Alibaba Cloud
Apsara Stack Enterprise

Operations and Maintenance Guide - Analytics and Artificial Intelligence

Version: 1911, Internal: V3.10.0

Issue: 20200315

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, consequent**

ial, 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 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 contact Alibaba Cloud directly if you discover any errors in this document.**

Document conventions

Style	Description	Example
	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.
	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.
	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.
	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]`

Style	Description	Example
{} or {a|b}	**This format is used for a required value, where only one item can be selected.**	`switch {active|stand}`

Contents

1 Apsara Big Data Manager (ABM) platform

1.1 What is Apsara Big Data Manager?

Apsara Big Data Manager (ABM) is an operations and maintenance platform tailored for big data products.

Currently, ABM supports the following products:

- **MaxCompute**
- **DataWorks**
- **StreamCompute**
- **Quick BI**
- **DataHub**
- **Machine Learning Platform for AI**

ABM supports operations and maintenance of big data products from perspectives such as business, services, clusters, and hosts. You can also install patches for big data products, customize alert configurations, and view O&M history through the ABM console.

ABM allows on-site Apsara Stack engineers to manage big data products with ease. For example, they can view performance metrics in real time, modify runtime configurations, and check and handle alerts in a timely manner.

1.2 Common operations

The data tables and legends in the Apsara Bigdata Manager (ABM) console facilitate operations. This topic uses MaxCompute as an example to describe the common operations.

Search for a project quickly

You can quickly search for a project based on the 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 Quick Search field, enter 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.

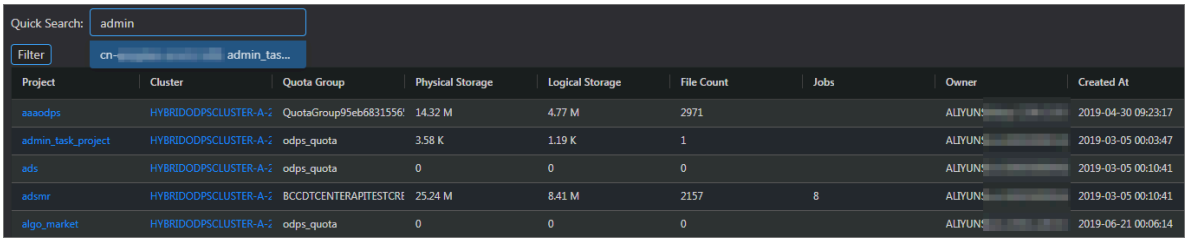

Example:

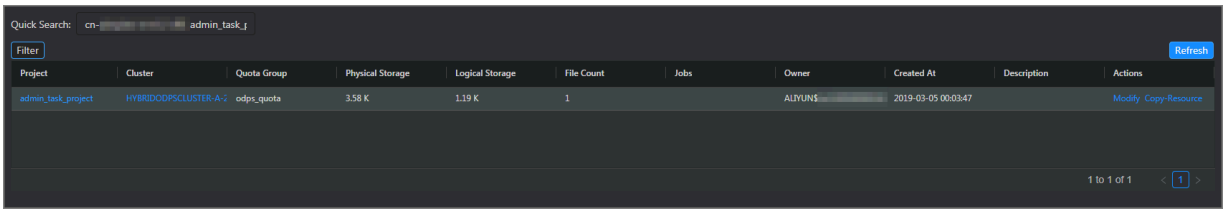

Filter projects

You can set filter conditions for multiple columns at the same time to quickly filter the projects you want.

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.

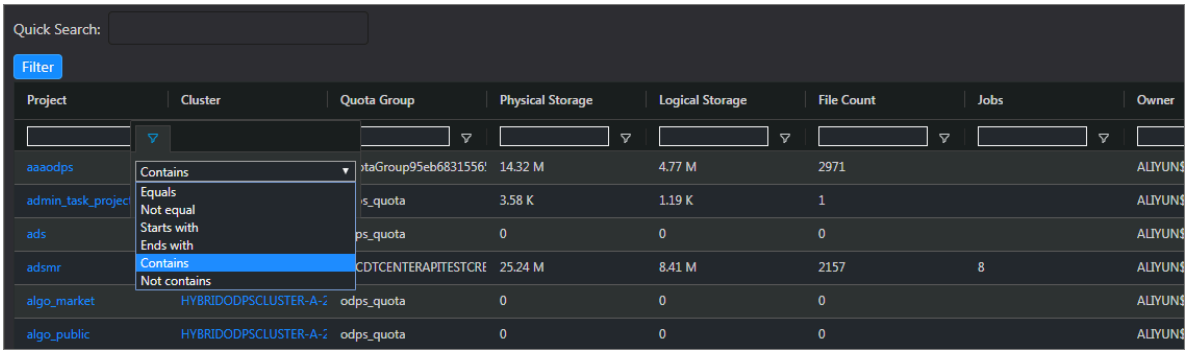

Optional filtering methods include:

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

4. After selecting the filtering method, enter the filter condition. The projects that meet the filter condition are automatically filtered.

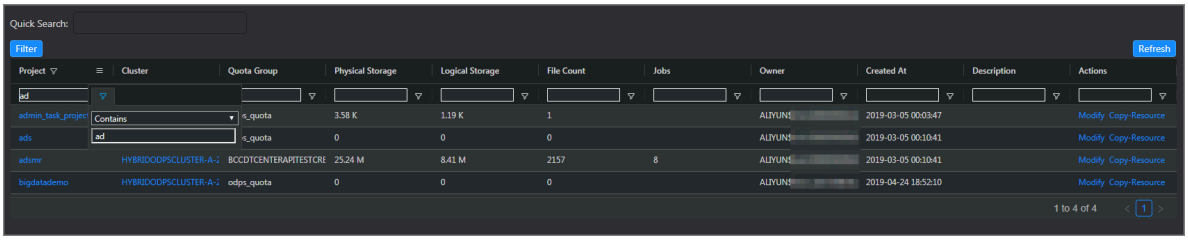

5. If the filtering result is not accurate, you can continue performing this operation on other columns.

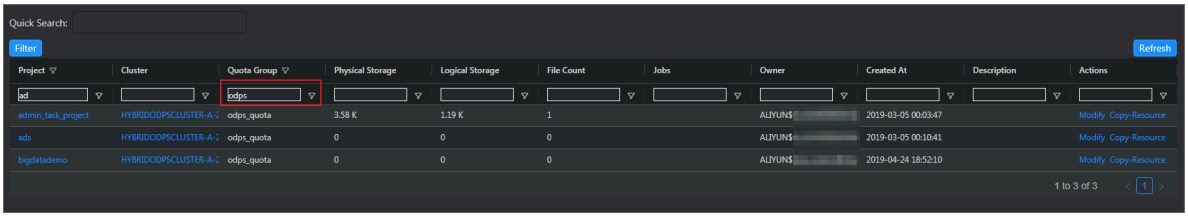

After you set the filter conditions for the projects, the Filter button is highlighted. If you need to cancel filtering, click the highlighted Filter button.

Search for items

You can search for items in a table by column, which is similar to filtering projects. For example, follow these 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.

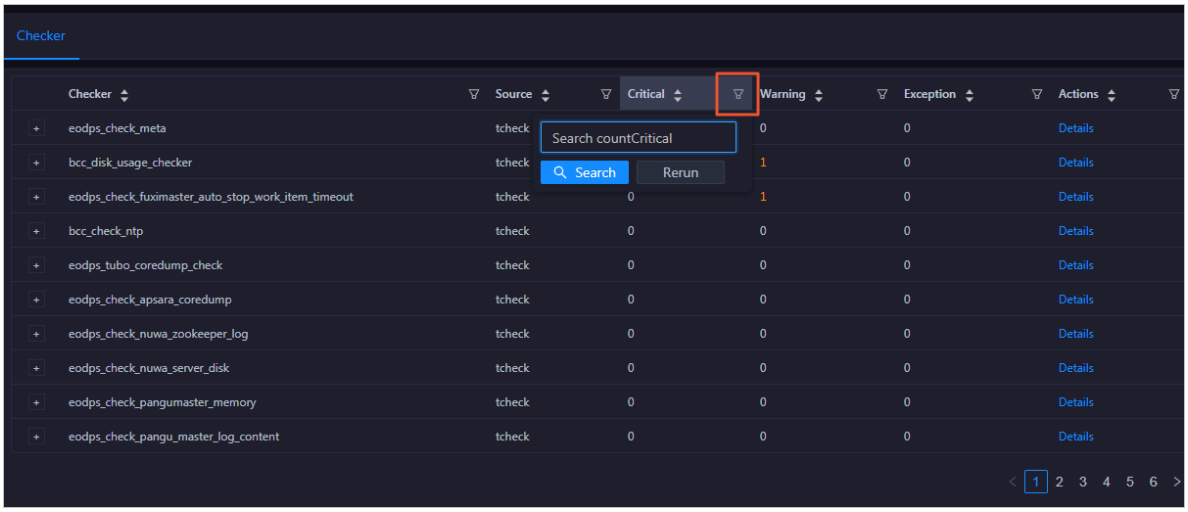

3. Click Search. The checkers that meet the requirements appear.
4. If the search result is not accurate, you can continue performing this operation on other columns.

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.

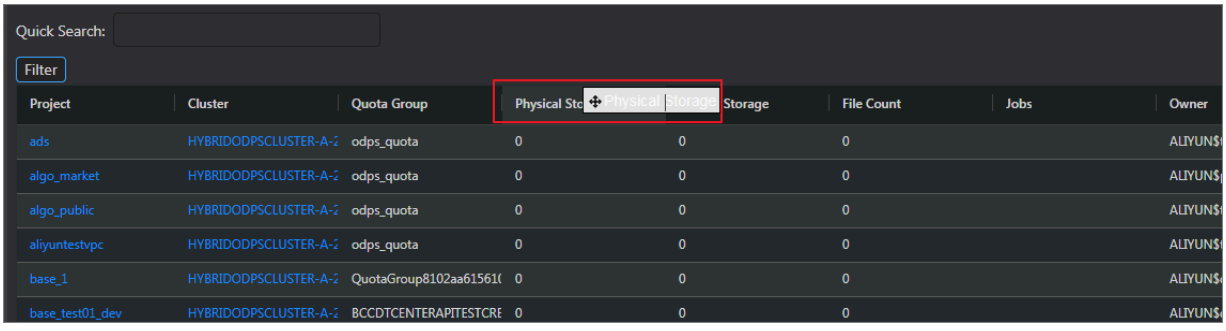

You can click ☰ in a column heading to customize the column.

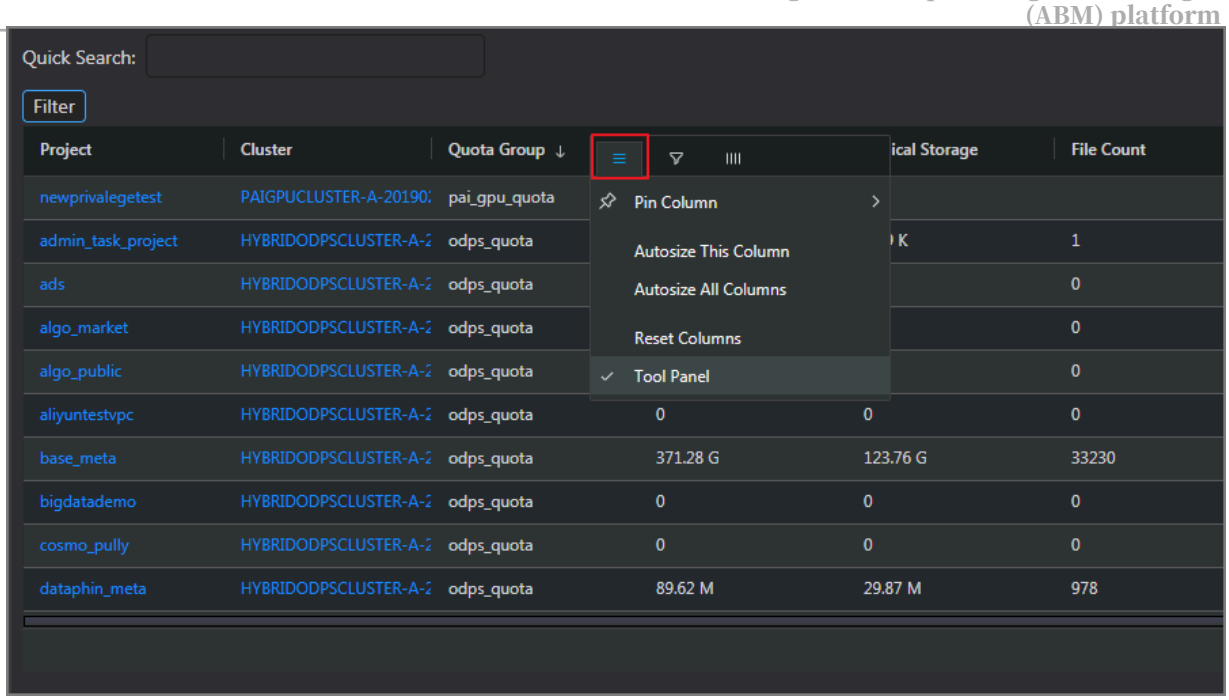

- **Pin Column: allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.**
- **Autosize This Column: allows you to adjust the width of a column automatically.**
- **Autosize All Columns: allows you to adjust the width of all columns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

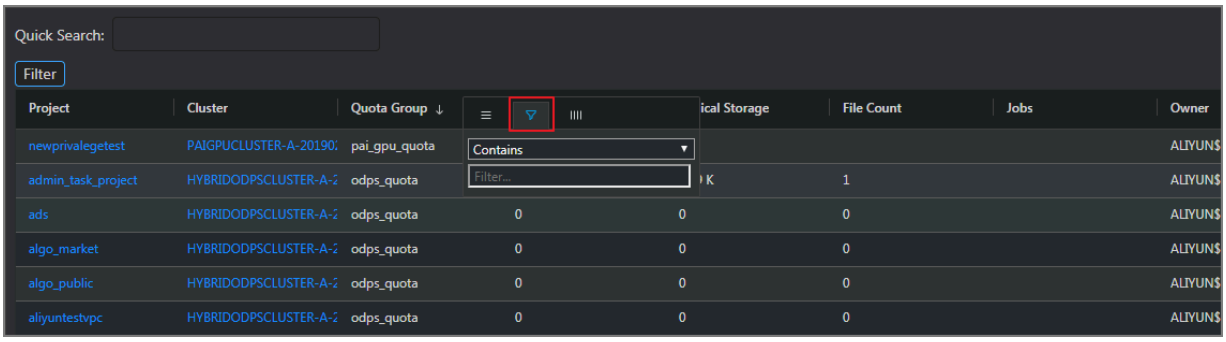

Click in a column heading and select the columns to be displayed.

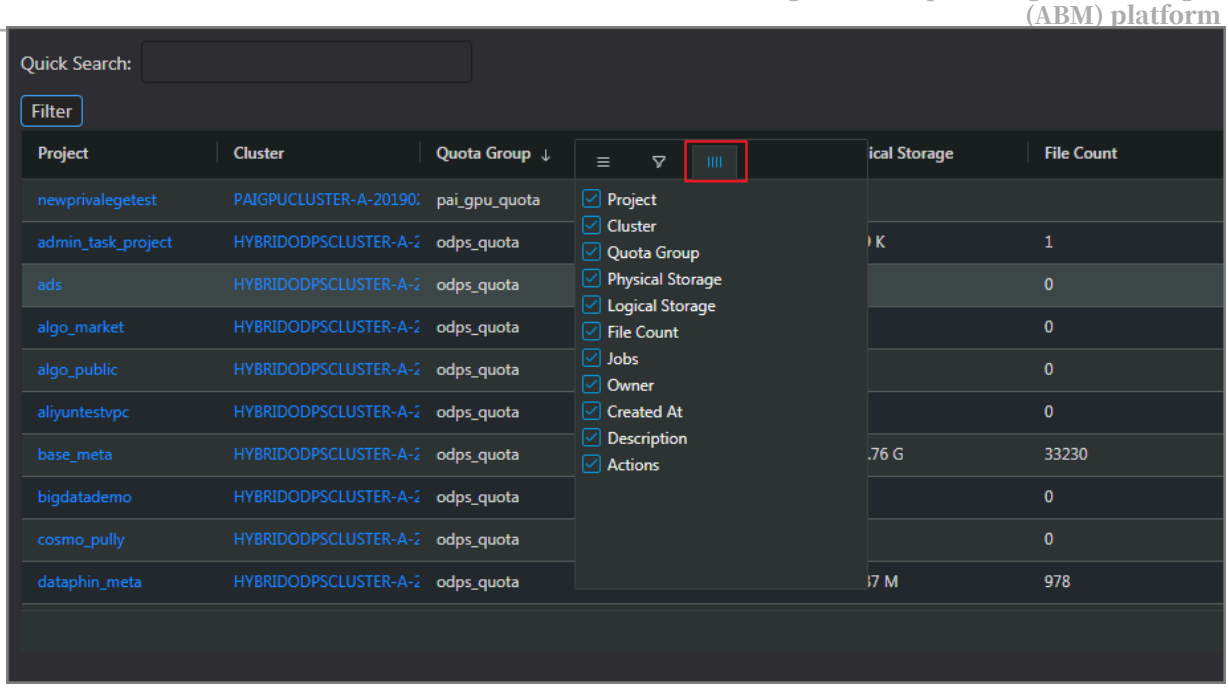

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

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

On the Project List page, click ☰ in a column heading and then select Tool Panel.

The tool panel is then attached to the right of the list.

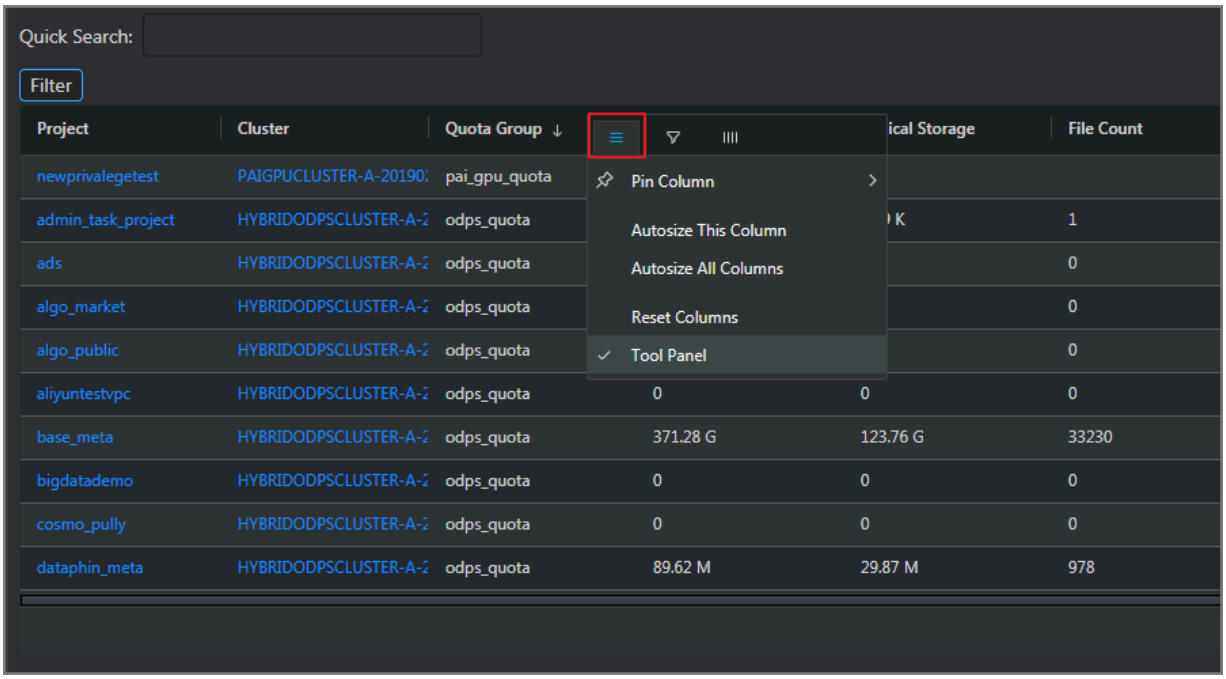

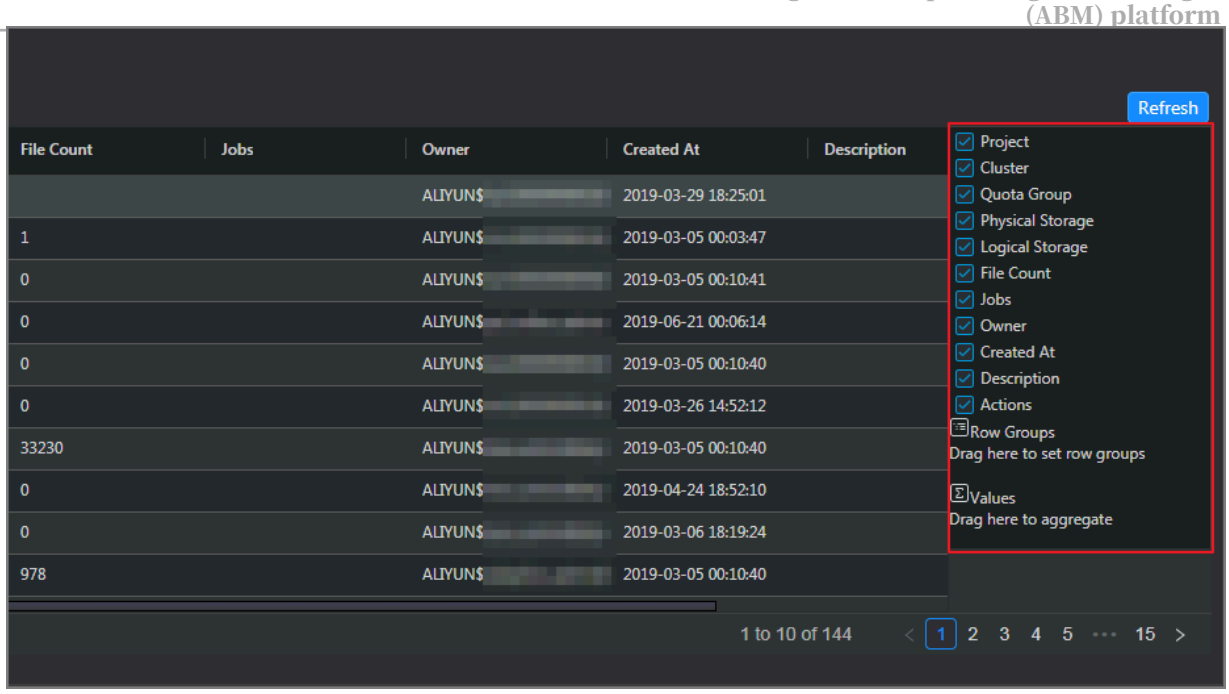

Sort projects based on a column

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.

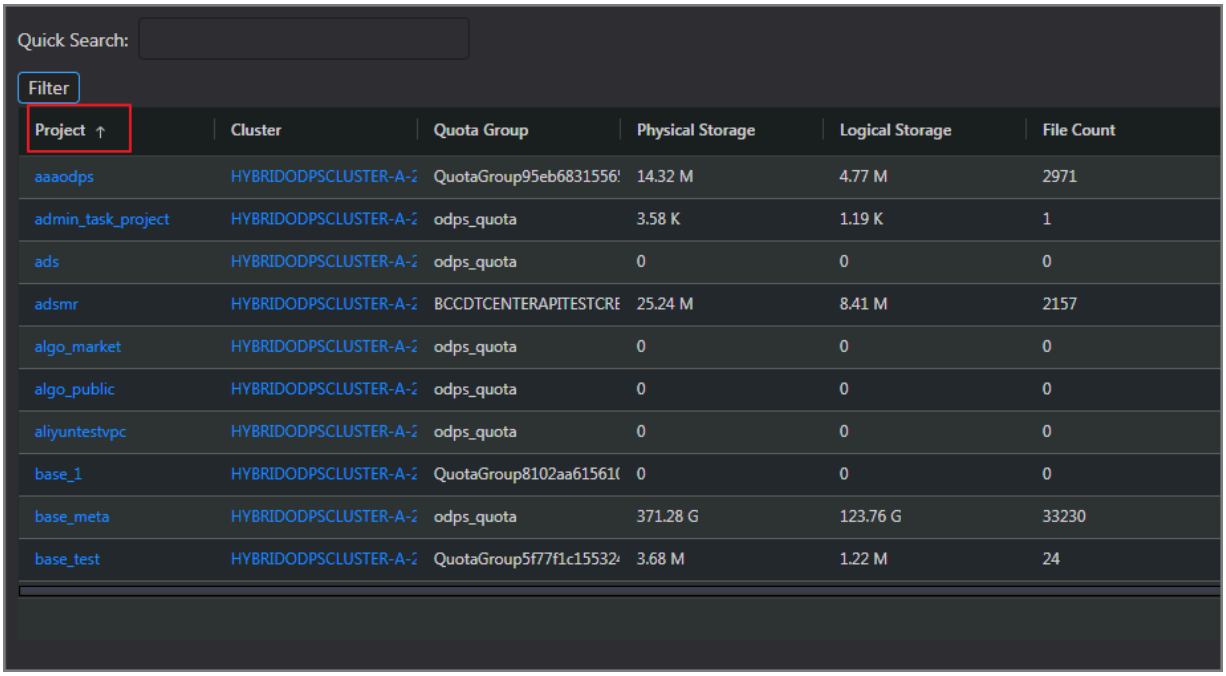

Sort items based on a column

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.**

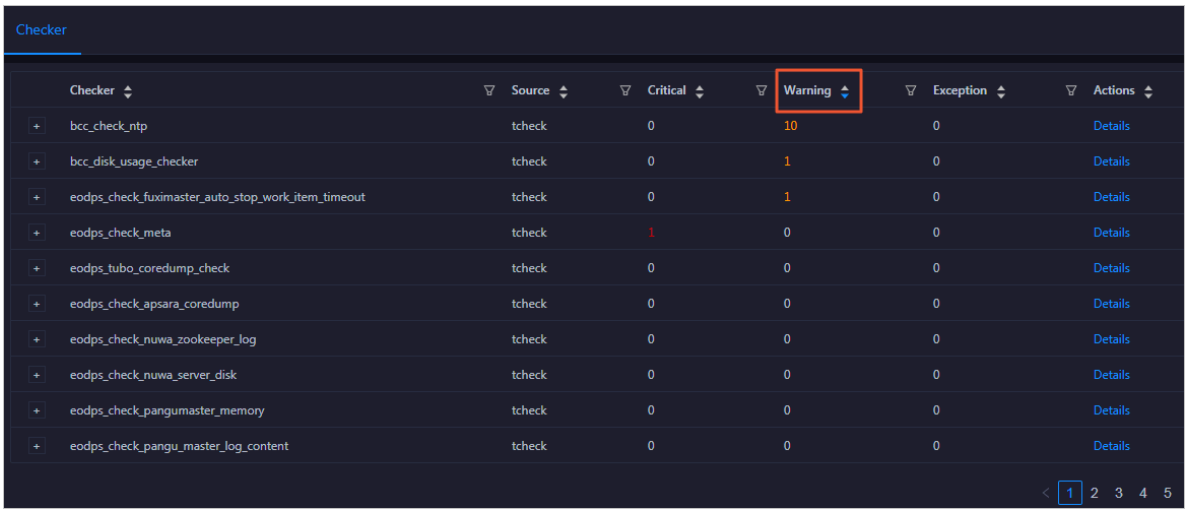

The highlighted up arrow indicates that the checkers are sorted in ascending order. The highlighted down arrow indicates that the checkers are sorted in descending order.

Trend chart 1

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, of the selected cluster.

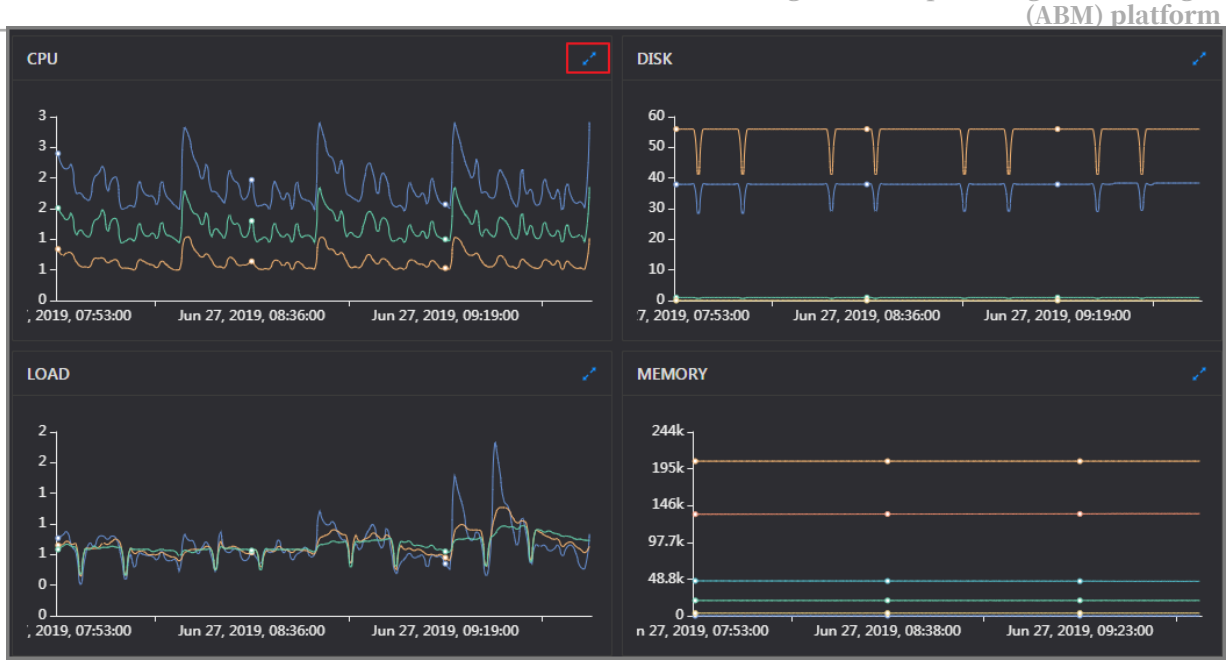

Take CPU as an example. The trend chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the specified cluster over time in different colors.

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

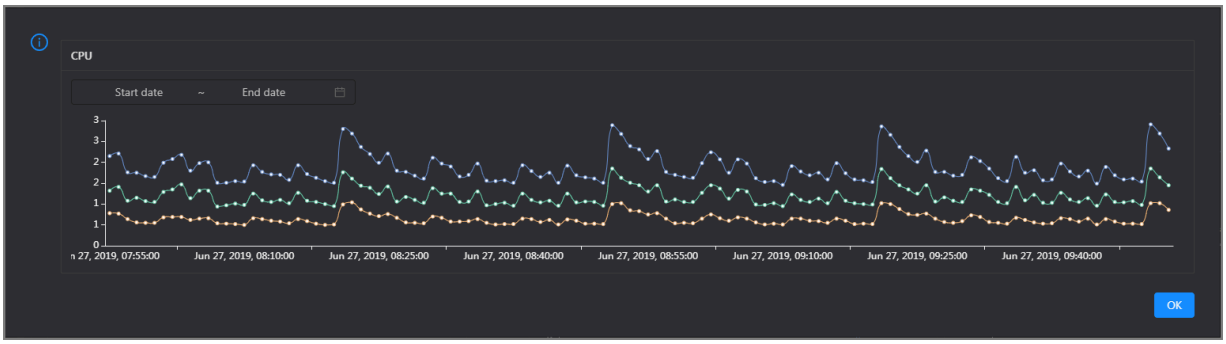

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

1.3 Quick start

1.3.1 Log on to the ABM console

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

Context

- **ASO access address in the format of http://*region-id*.aso.*intranet-domain-id*.com.**
- **Google Chrome browser (recommended).**

Procedure

1. **Open the browser.**
2. **Enter the ASO access address http://*region-id*.aso.*intranet-domain-id*.com in the address bar and then press Enter.**

   Figure 1-1: Log on to ASO

   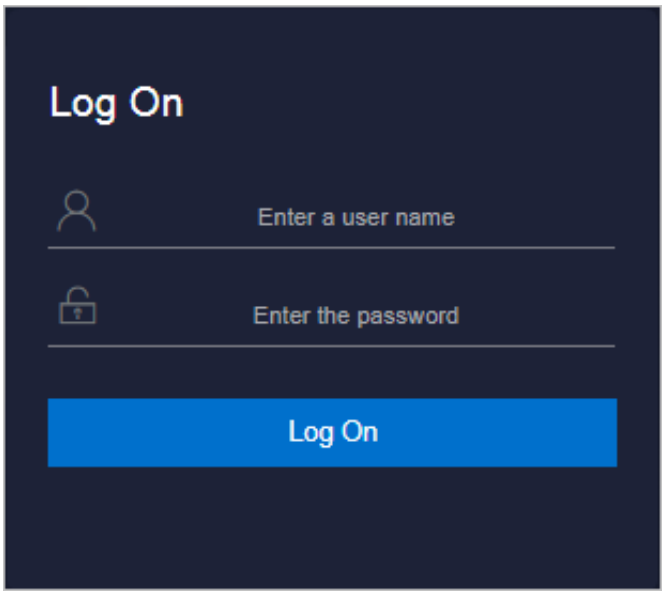

   Note:
   You can select the language from the drop-down list in the upper-right corner to change the language of ASO.
3. **Enter the correct username and password.**
   - **The system has three default users:**
     - **Security officer: manages other users or roles.**
     - **Auditor officer: views audit logs.**
     - **System administrator: used for other functions except those of the security officer and auditor officer.**
   - **You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or**

a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%), which meets the minimum complexity requirements.

4. Click Log On to log on to ASO.
5. In the left-side navigation pane, choose Products > Apsara Bigdata Manager to log on to the ABM console.

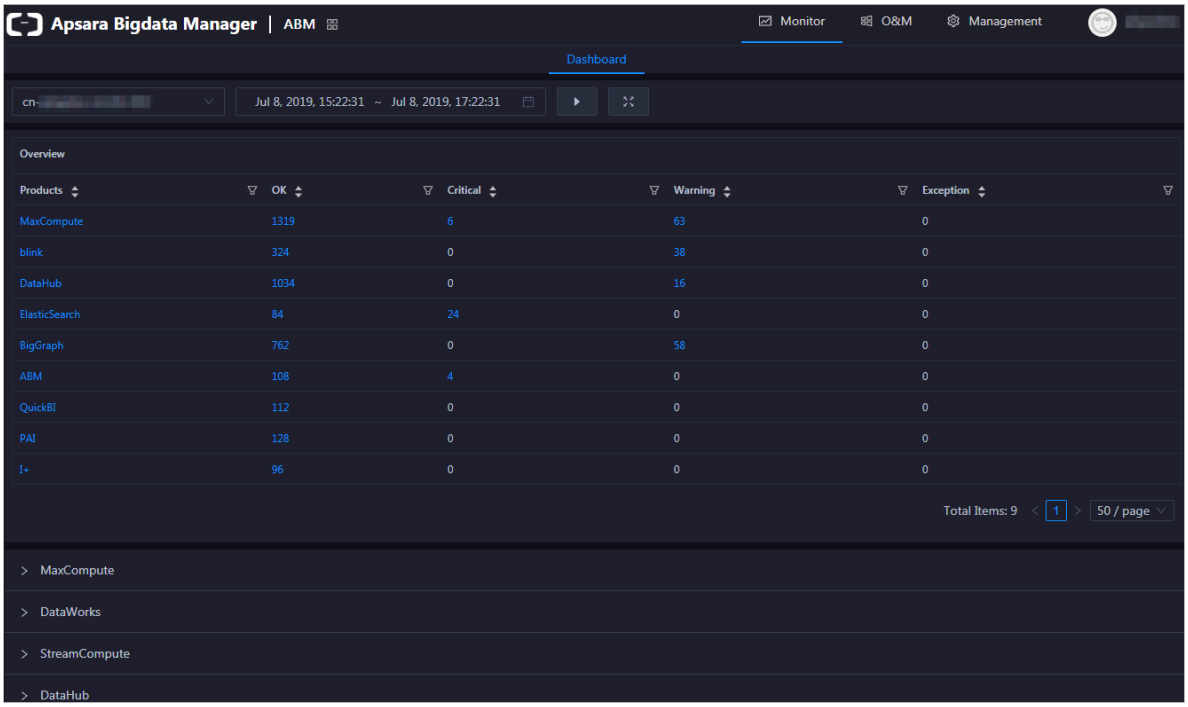

1.3.2 Set the theme of the console

You can set the theme of the Apsara Bigdata 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.

Theme	Description
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.

1.3.3 View the dashboard

The Apsara Bigdata Manager (ABM) dashboard displays key operation metrics of MaxCompute, DataWorks, Realtime Compute, and DataHub. It also provides information about alerts for all big data services, so that you can understand the overall running status of the big data services.

Prerequisites

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

Background

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*.

   Log on to the ABM console. The Dashboard page appears. To return to the Dashboard page from any other page, click in the upper-left corner and then click ABM.

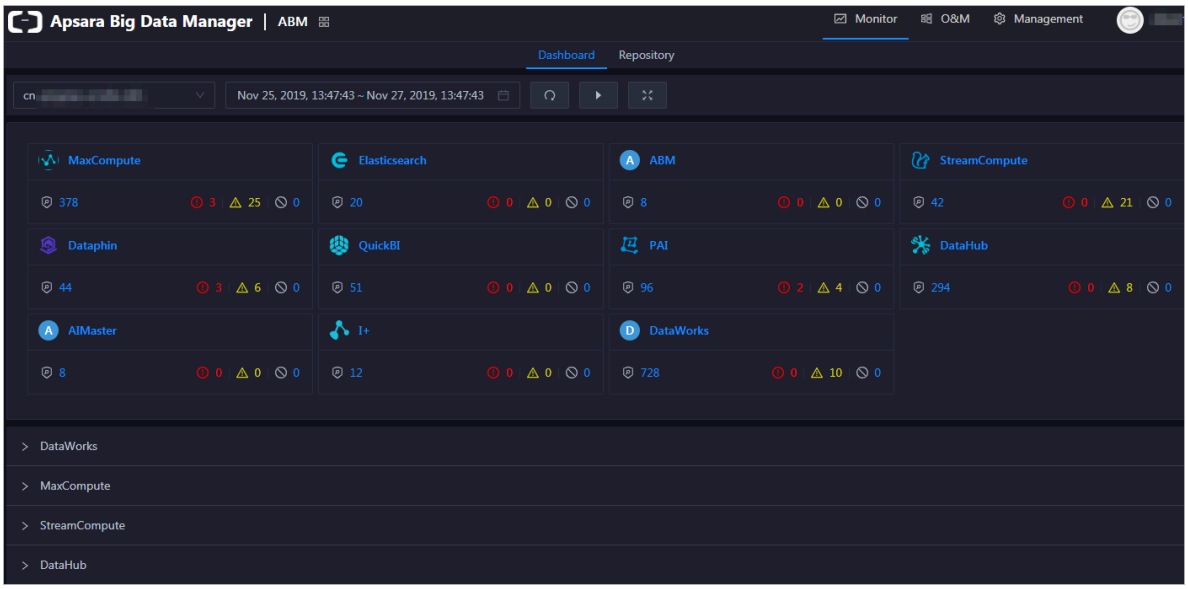

2. View and clear service alerts.

   In the overview section, you can view the number of alerts for all big data services. Pay attention to the Critical and Warning alerts. These alerts must be cleared in a timely manner.

   a. On the Dashboard page, click the number of Critical or Warning alerts of a service in the overview section. The Health Status page under Clusters of the service appears.

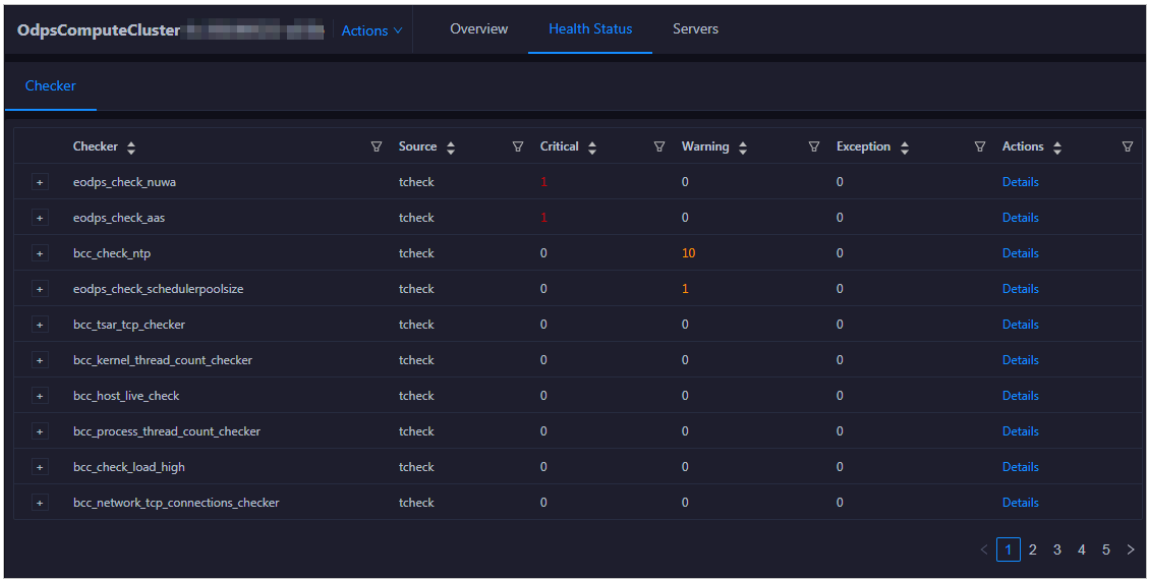

   On the Health Status page, you can view all the checkers of the service.

   b. Click Details in the Actions column of a checker with alerts. In the Details dialog box that appears, view the details of the checker and the scheme to clear the alerts. Follow the steps in the scheme to clear the alerts.

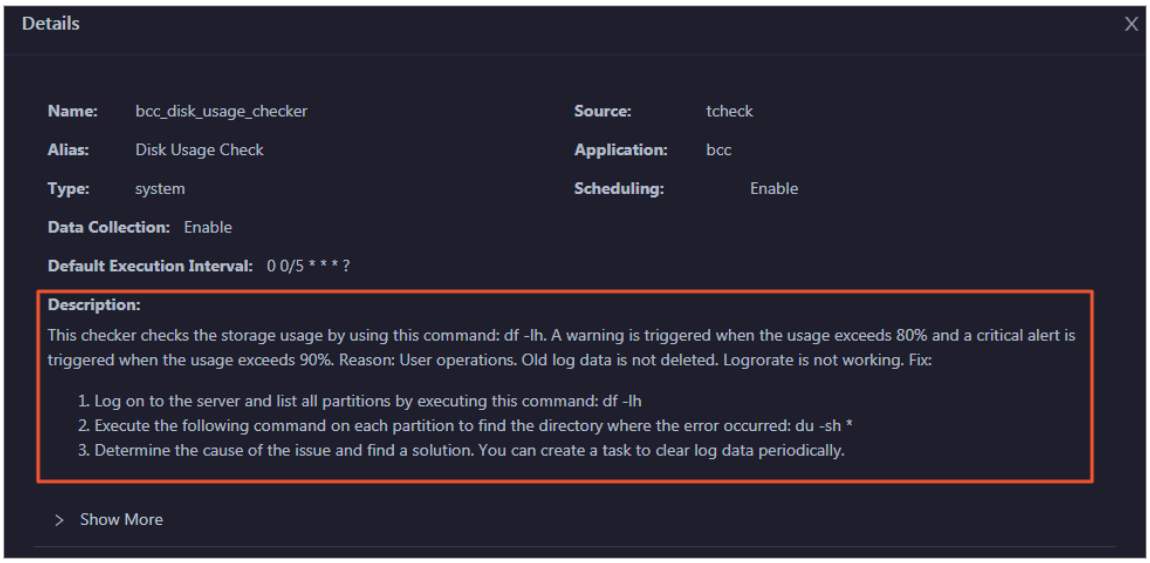

   c. Log on to the hosts with alerts if necessary for related operations.

Click + to expand a checker with alerts, and then click the Log On icon next to the name of a host with alerts.

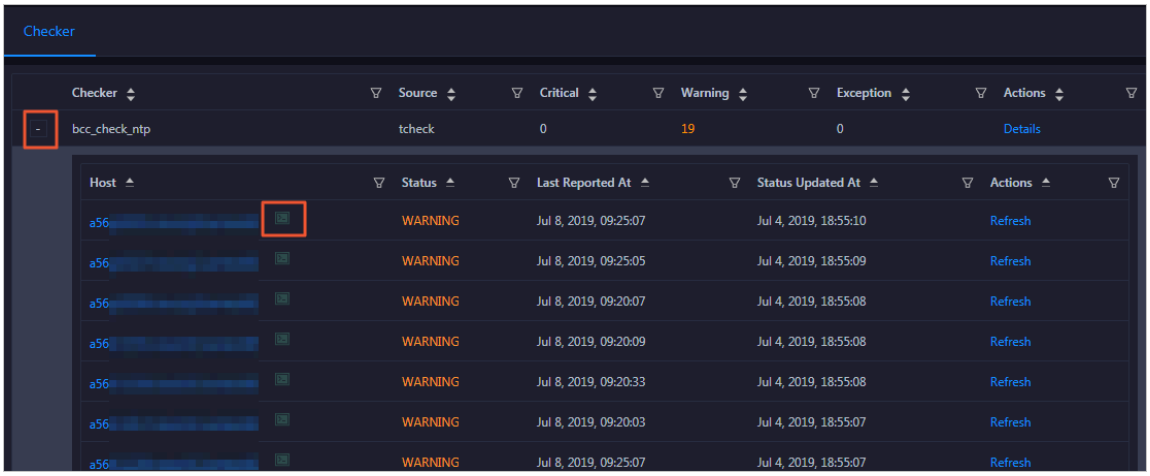

d. On the TerminalService page that appears, click the hostname on the left to log on to the host.

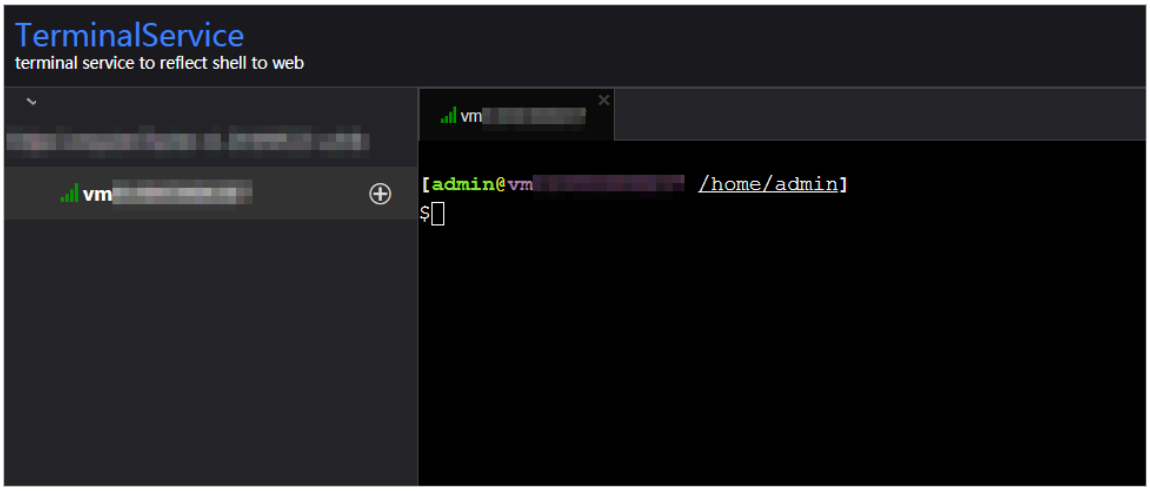

3. On the Dashboard page, click MaxCompute in the overview section to view relevant metrics.

In the MaxCompute section, you can view the job running 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 usage, and physical CPU usage.

4. On the Dashboard page, click DataWorks in the overview section to view relevant metrics.

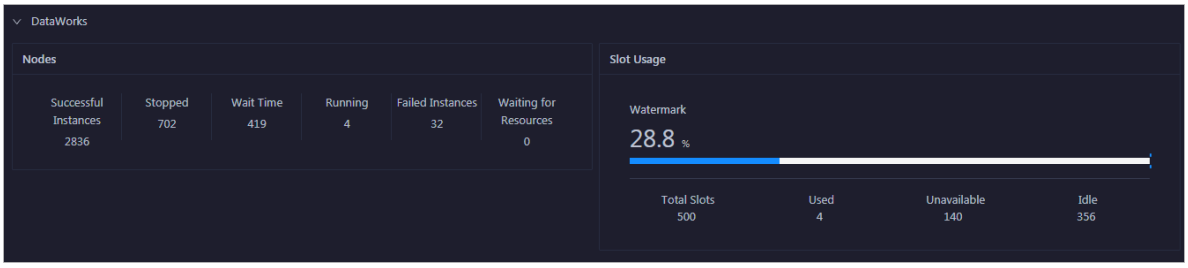

In the DataWorks section, you can view the node scheduling and slot usage of the DataWorks cluster. You can also view the trend chart of the total number of finished tasks.

5. On the Dashboard page, click StreamCompute in the overview section to view relevant metrics.

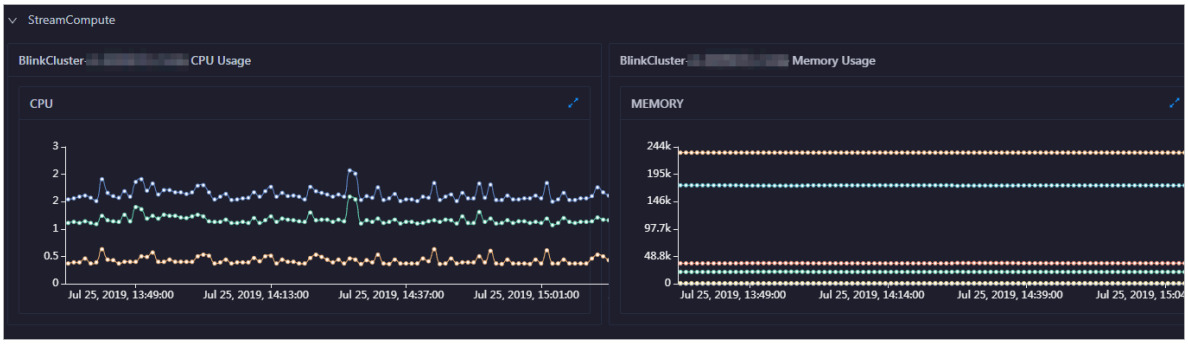

In the StreamCompute section, you can view the trend charts of the transactions per second (TPS), failover rate, CPU usage, and memory usage for the Realtime Compute cluster.

6. On the Dashboard page, click DataHub in the overview section to view relevant metrics.

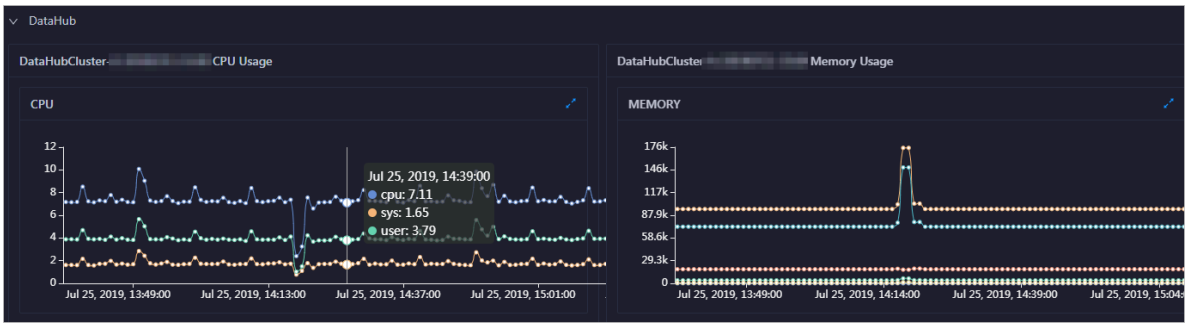

In the DataHub section, you can view the trend charts of the latency, records, queries per second (QPS), and throughput of the read and write operations. You can also view the trend charts of CPU and memory usage for the DataHub cluster.

1.3.4 View the cluster running status

Apsara Bigdata 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.

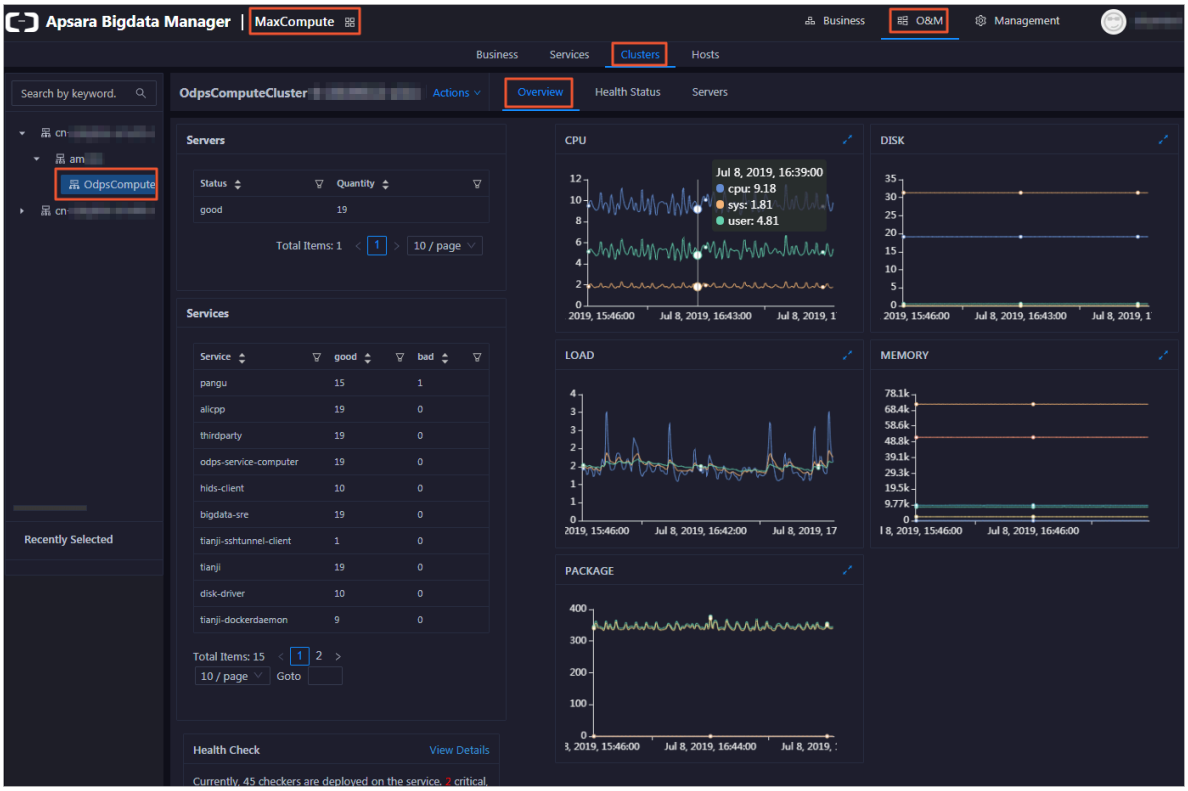

On the Overview page, you can view the host status, service status, health check result, and health check history of the selected cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster.

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*.

1.3.5 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 Bigdata 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.

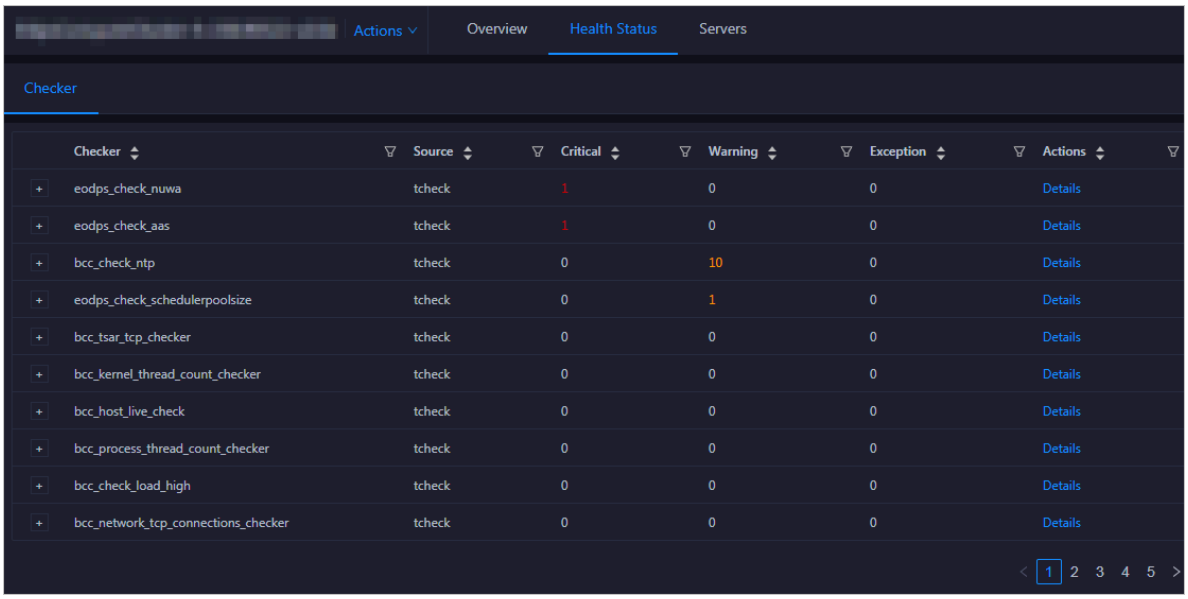

4. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.

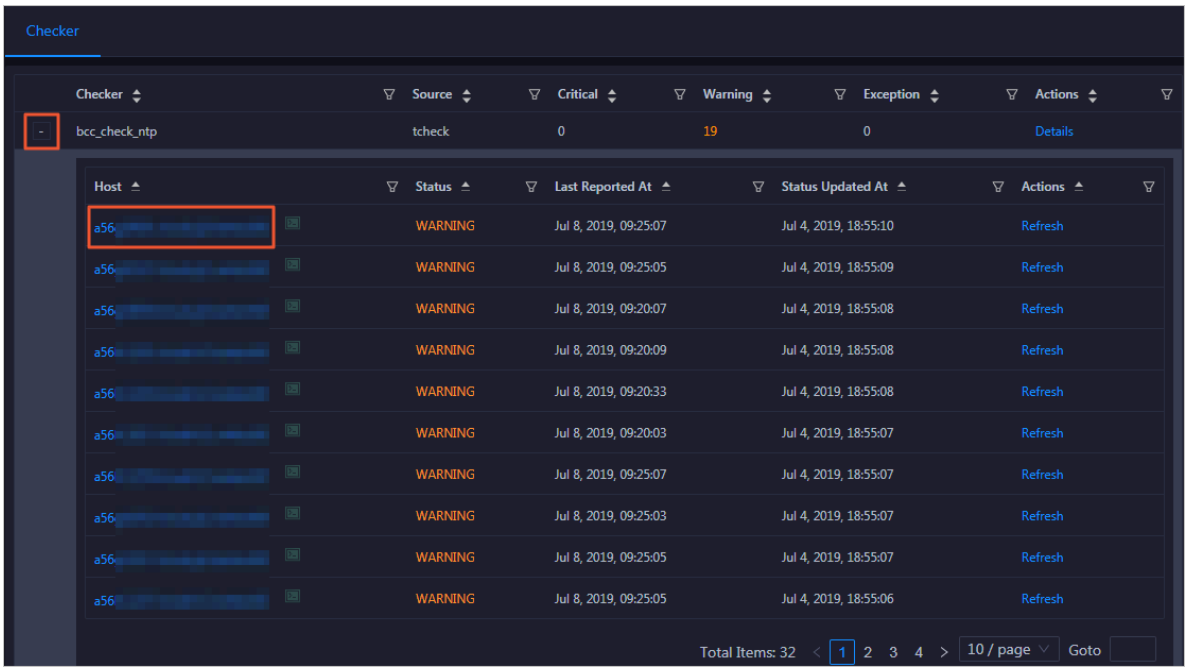

5. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.

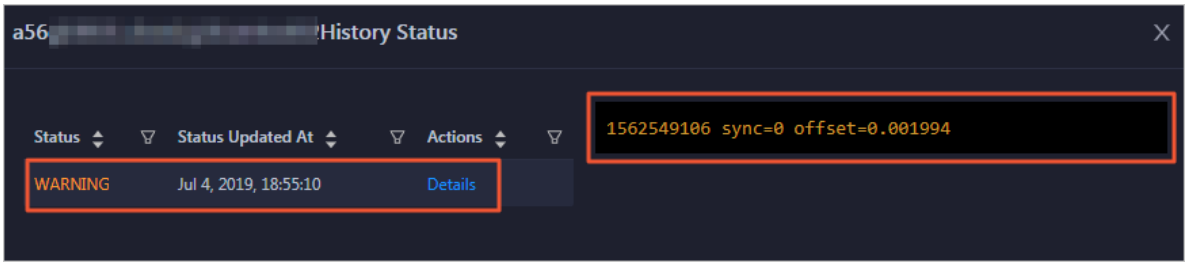

6. On the Health Status page, click Details in the Actions column of the checker to view the schemes to clear the alerts.

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.

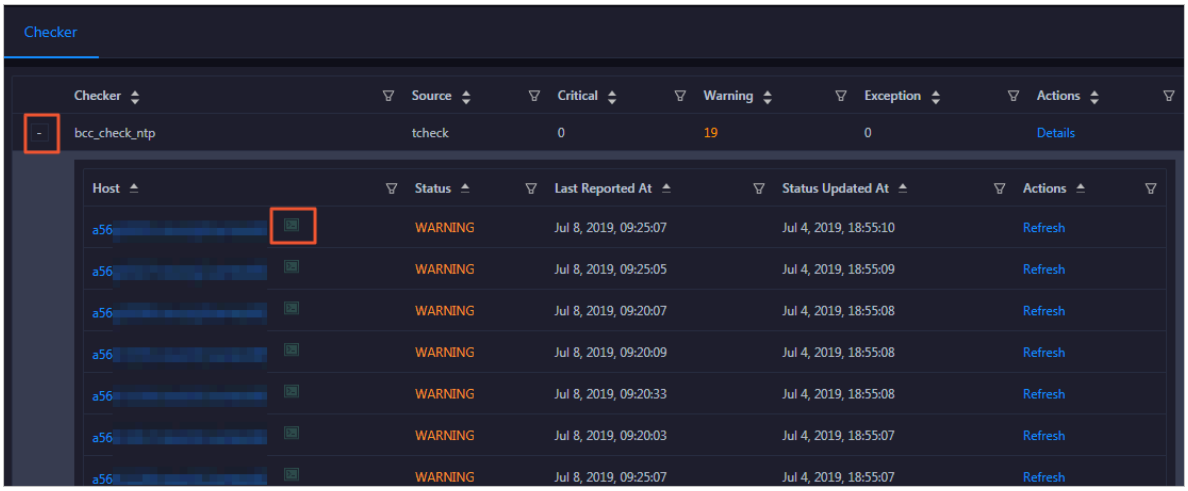

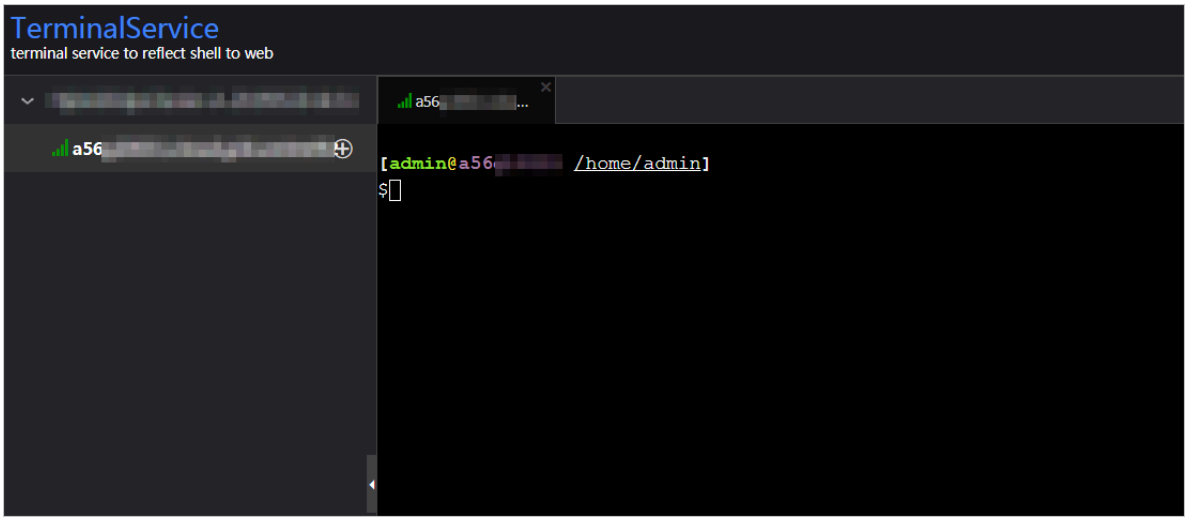

8. **After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.**

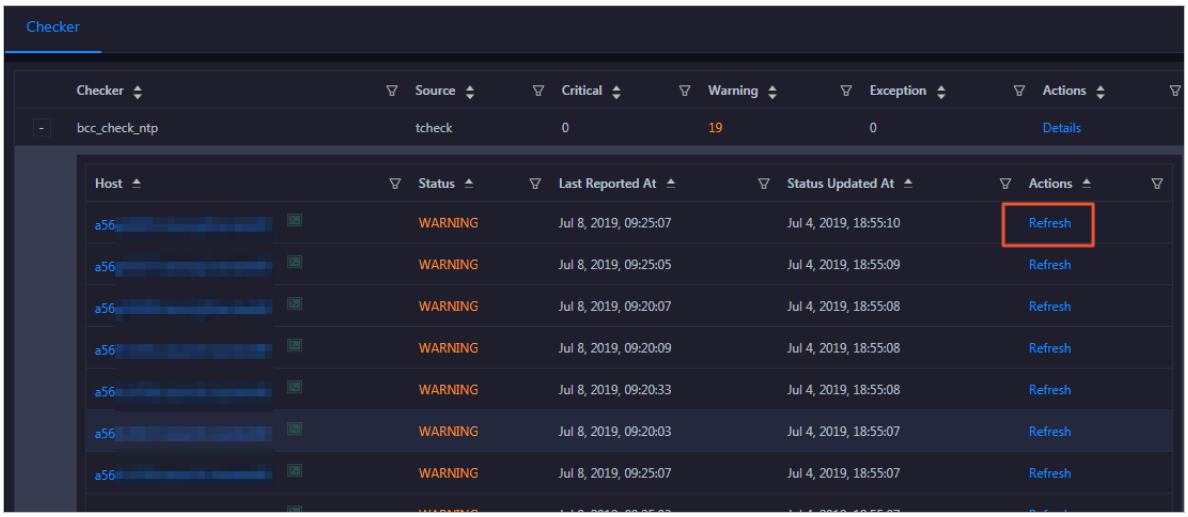

1.4 ABM

1.4.1 ABM dashboard

The Apsara Bigdata Manager (ABM) dashboard displays key operation metrics of MaxCompute, DataWorks, Realtime Compute, and DataHub. It also provides information about alerts for all big data services, so that you can understand the overall running status of the big data services. In addition, the dashboard supports automatic data refresh and full-screen display.

Entry

Log on to the ABM console. **The Dashboard page appears. To return to the Dashboard page from any other page, click**  **in the upper-left corner and then click ABM.**

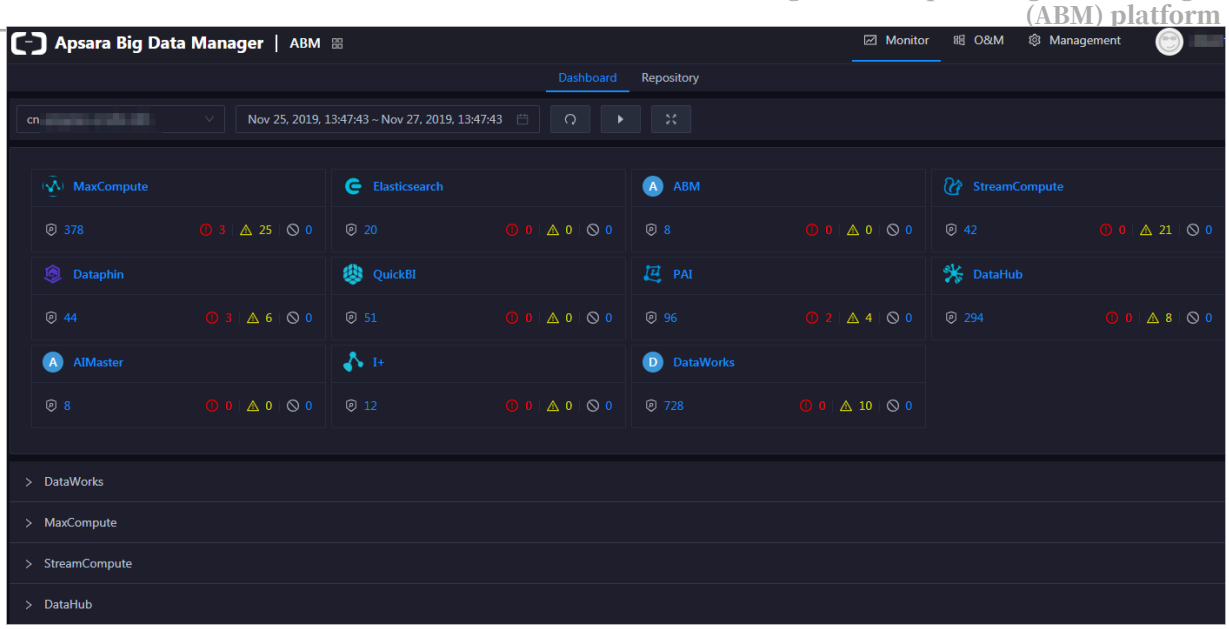

On the Dashboard page, you can select a region from the region drop-down list in the upper-left corner. In this way, you can view the cluster running status of each big data service in the specified region.

View and clear alerts of various services

In the overview section, you can view the respective number of Critical, Warning, and Exception alerts for each big data service. If a service has alerts, especially the Critical and Warning alerts, clear the alerts in a timely manner.

1. On the Dashboard page, click the number of Critical or Warning alerts of a service in the overview section. The Health Status page under Clusters of the service appears.

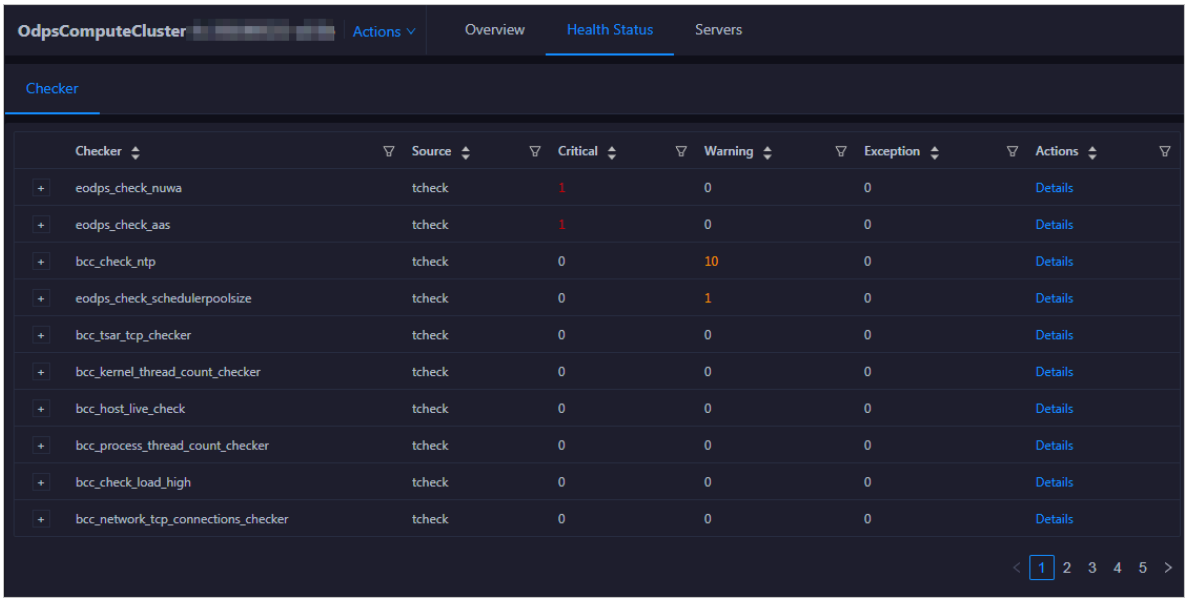

On the Health Status page, you can view all the checkers of the service.

2. Click Details in the Actions column of a checker with alerts. In the Details dialog box that appears, view the details of the checker and the scheme to clear the alerts. Follow the steps in the scheme to clear the alerts.

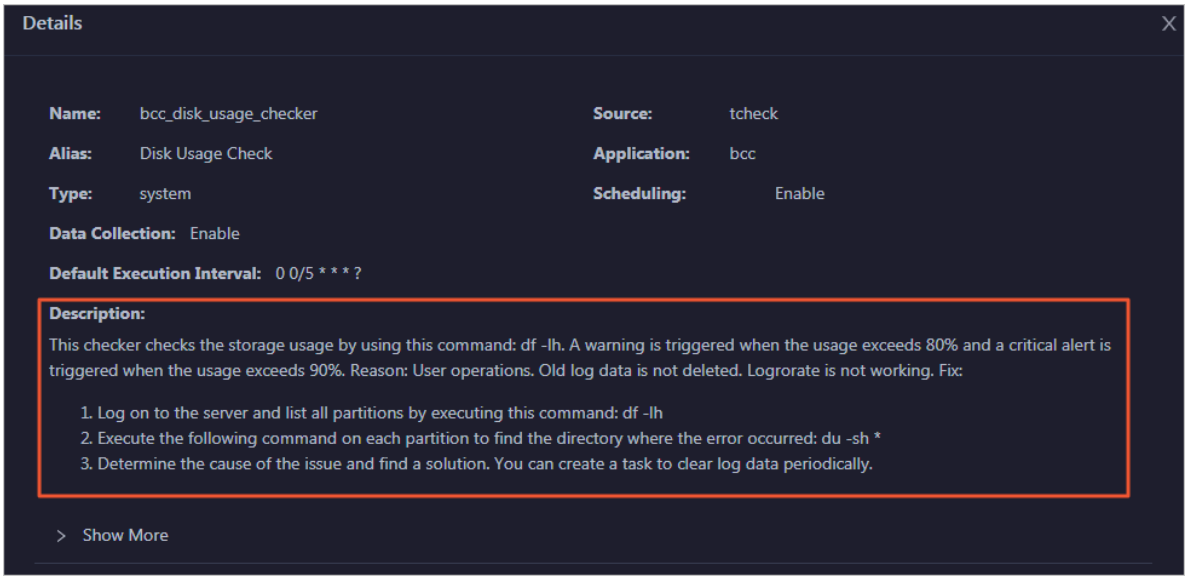

3. **Log on to the hosts with alerts if necessary for related operations.**

   Click + to expand a checker with alerts, and then 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.

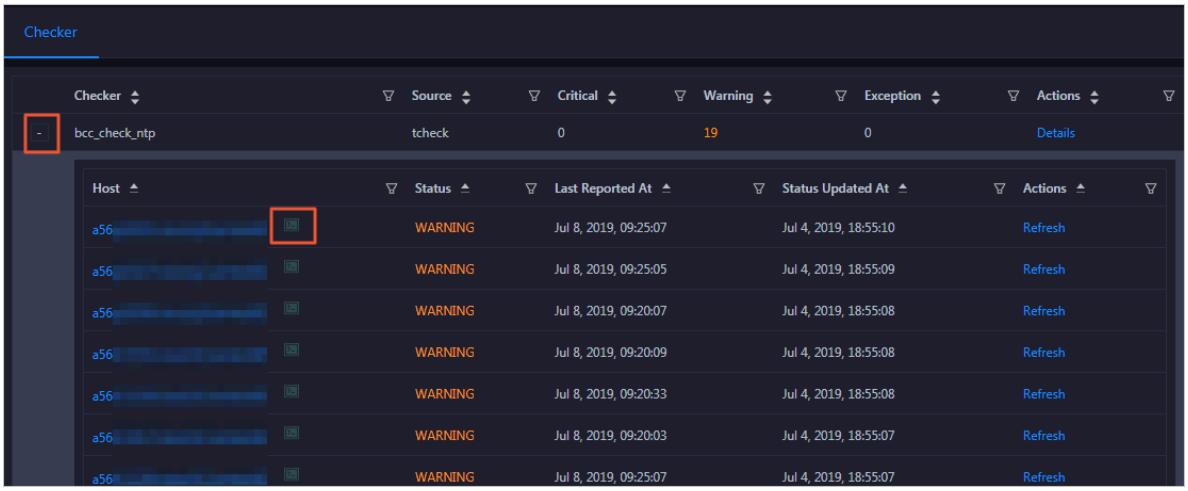

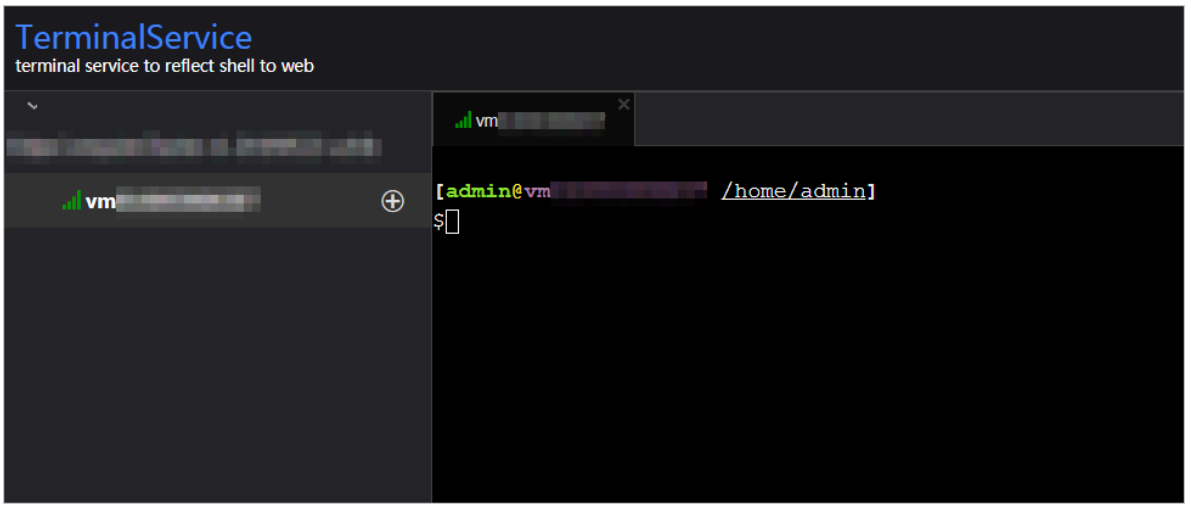

View key operation metrics of MaxCompute

The ABM dashboard displays key operation metrics of MaxCompute. On the Dashboard page, click MaxCompute in the overview section to view the metrics.

In the MaxCompute section, you can view the job running 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 usage, and physical CPU usage. For more information about the MaxCompute operation metrics, see *Cluster overview*.

View key operation metrics of DataWorks

The ABM dashboard displays key operation metrics of DataWorks. On the Dashboard page, click DataWorks in the overview section to view the metrics.

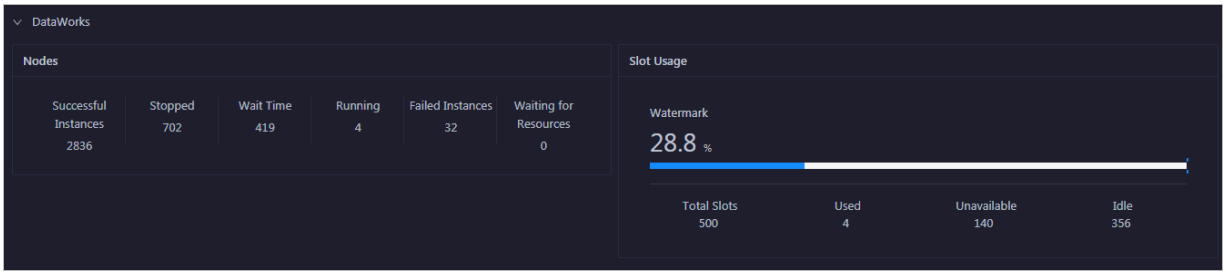

In the DataWorks section, you can view the node scheduling and slot usage of the DataWorks cluster. You can also view the trend chart of the total number of finished tasks. For more information about the DataWorks operation metrics, see *Service overview*.

View key operation metrics of Realtime Compute

The ABM dashboard displays key operation metrics of Realtime Compute. On the Dashboard page, click StreamCompute in the overview section to view the metrics.

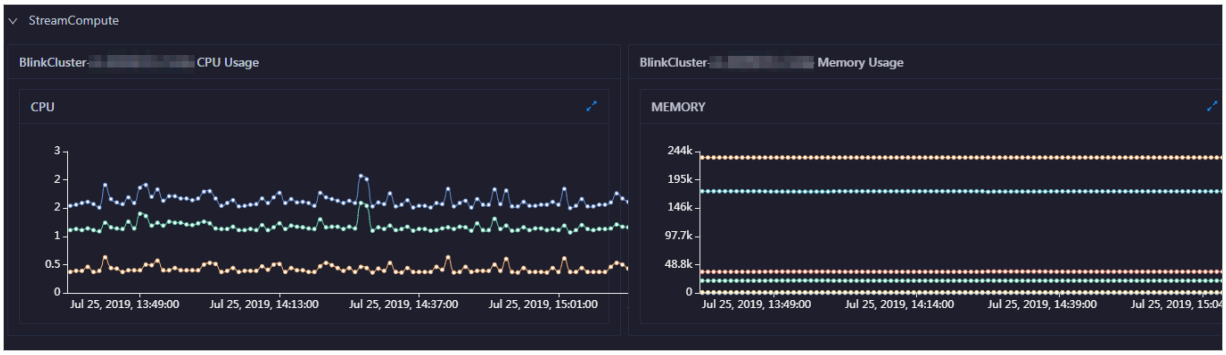

In the StreamCompute section, you can view the trend charts of the transactions per second (TPS), failover rate, CPU usage, and memory usage for the Realtime Compute cluster. For more information about the Realtime Compute operation metrics, see *Cluster overview*.

Enable or disable automatic refresh

By default, the Dashboard page does not refresh automatically and displays the statistics of the last two days. You can specify a time period to view the correspond ing statistics. If auto refresh is enabled, the Dashboard page refreshes automatica lly based on the specified interval and displays the latest data.

1. **On the Dashboard page, click ▶ at the top.**
2. **In the dialog box that appears, set the Refreshing every and Refreshing range parameters.**

   The Refreshing range parameter specifies the time period of the trend charts, such as the trend charts of CPU and memory usage, for each cluster.
3. **Click OK.**

   When automatic refresh is enabled, the ▶ icon is replaced with the ⏸ icon.

   The system automatically refreshes all data on the dashboard according to the specified time interval.

   To disable automatic refresh, click ⏸.

Display the dashboard in full-screen mode

The dashboard provides a full-screen display feature for you to view the running status of big data services clearly.

At the top of the Dashboard page, click to display the Dashboard page in full-screen mode.

1.4.2 ABM repository

The Repository page in the Apsara Bigdata 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.**

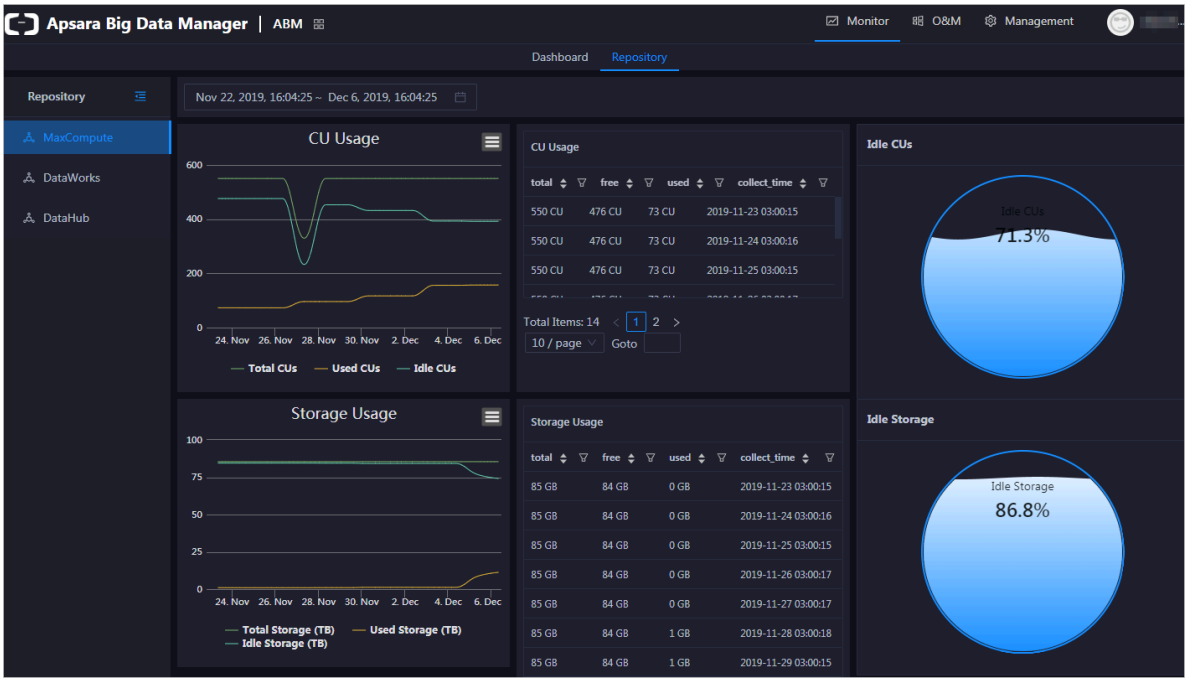

View the resource usage in MaxCompute

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.

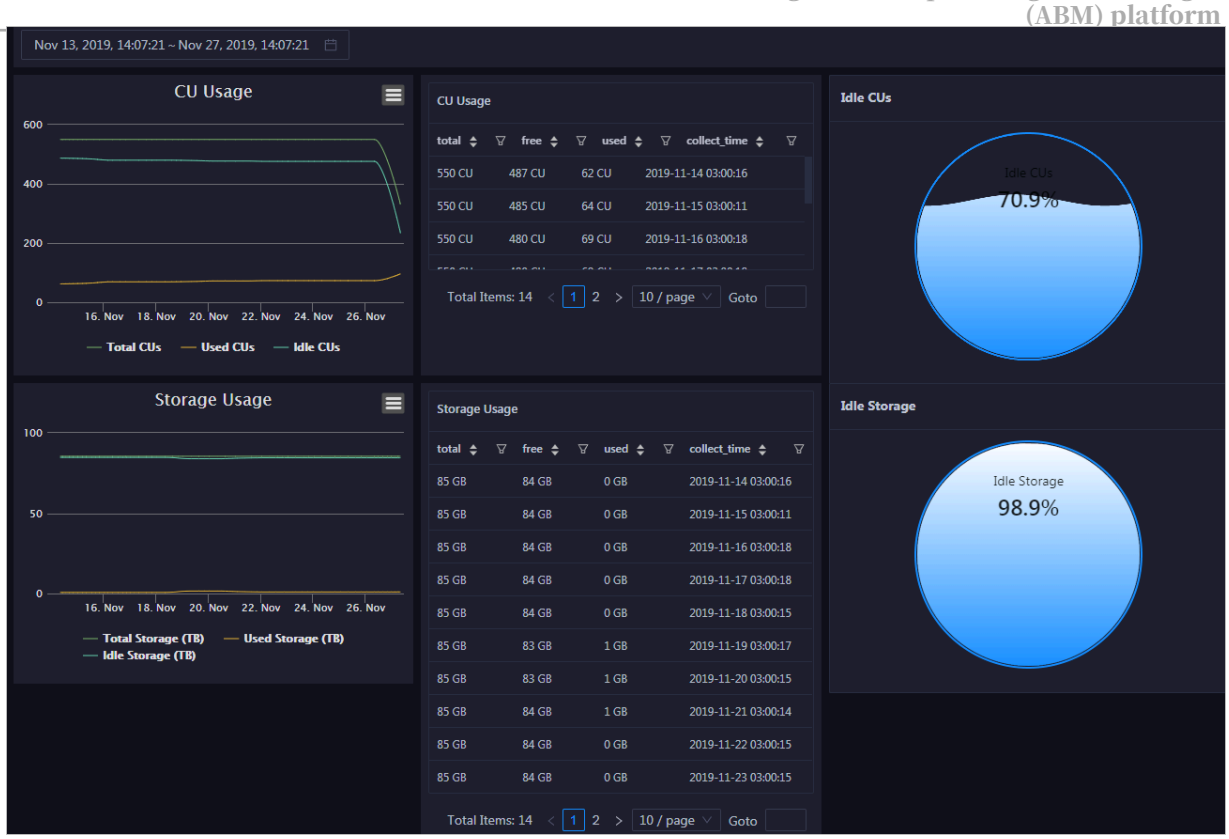

For MaxCompute, the Repository page displays the trend charts of CU and storage usage, records of CU and storage usage, and proportions of idle CUs and storage.

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.

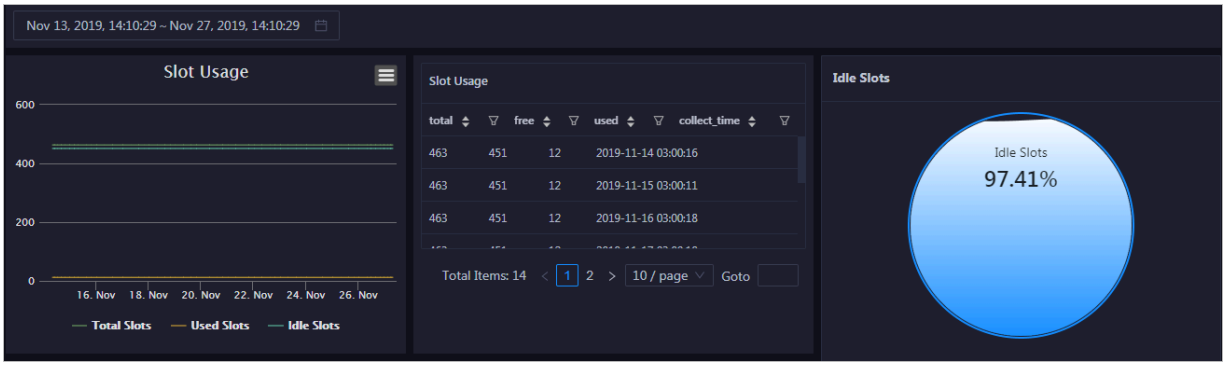

For DataWorks, the Repository page displays the trend chart of slot usage, records of slot usage, and proportion of idle slots.

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.

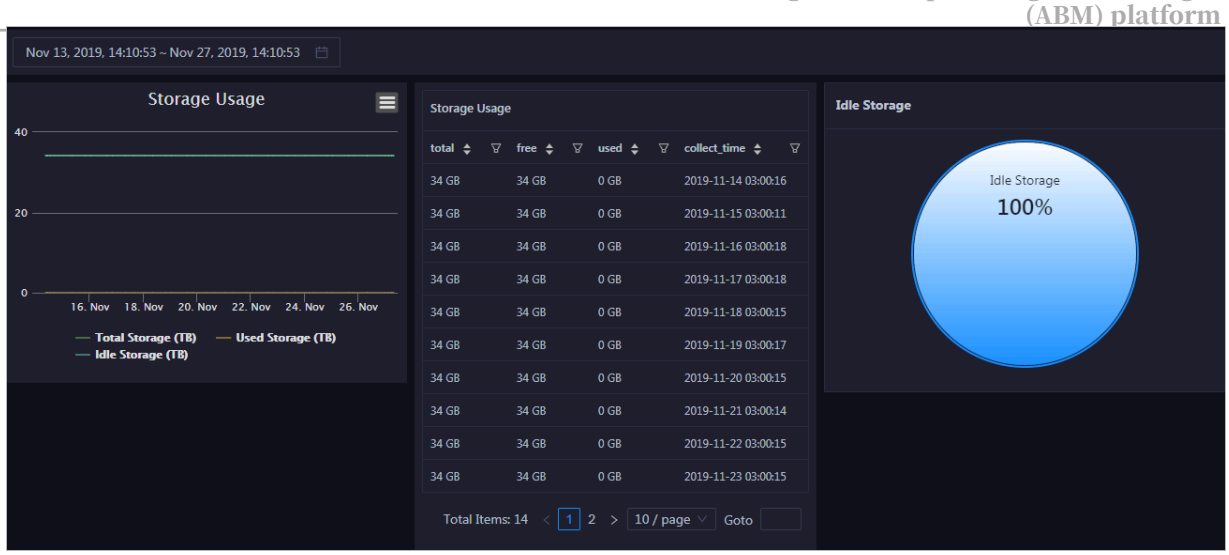

For DataHub, the Repository page displays the trend chart of storage usage, records of storage usage, and proportion of idle storage.

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*.

1.4.3 ABM O&M overview

This topic describes the features of Apsara Bigdata Manager (ABM) O&M and how to access the ABM O&M page.

Modules

ABM O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Services	Overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service in a cluster.
	Server	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Clusters	Overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmissi on, TCP connection, and root disk usage for a cluster.

Module	Feature	Description
	Health Status	Displays 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.
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.
	Health Status	Displays the checkers of the selected host, including the checker details, check results, check history, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click ABM.

3. On the page that appears, click O&M in the upper-right corner. The Services page appears.

The O&M page includes three modules, namely, Services, Clusters, and Hosts.

1.4.4 Service O&M

1.4.4.1 Service overview

The Overview page lists all Apsara Bigdata 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.

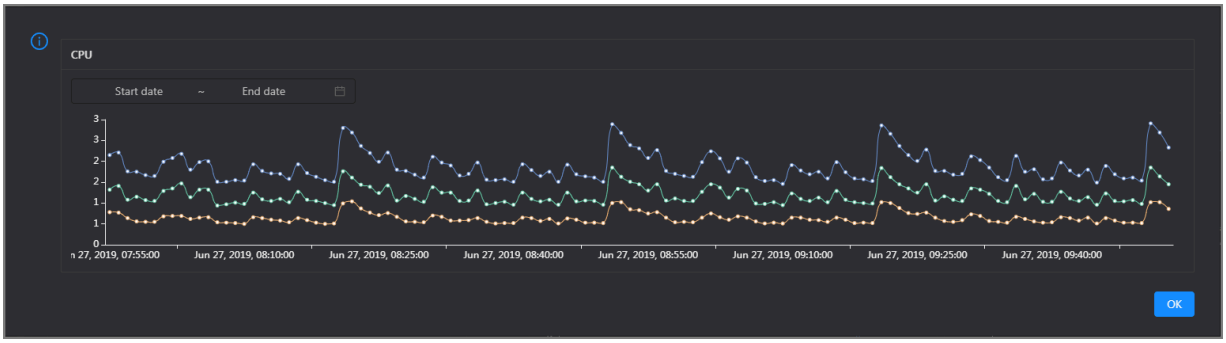

DISK

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.

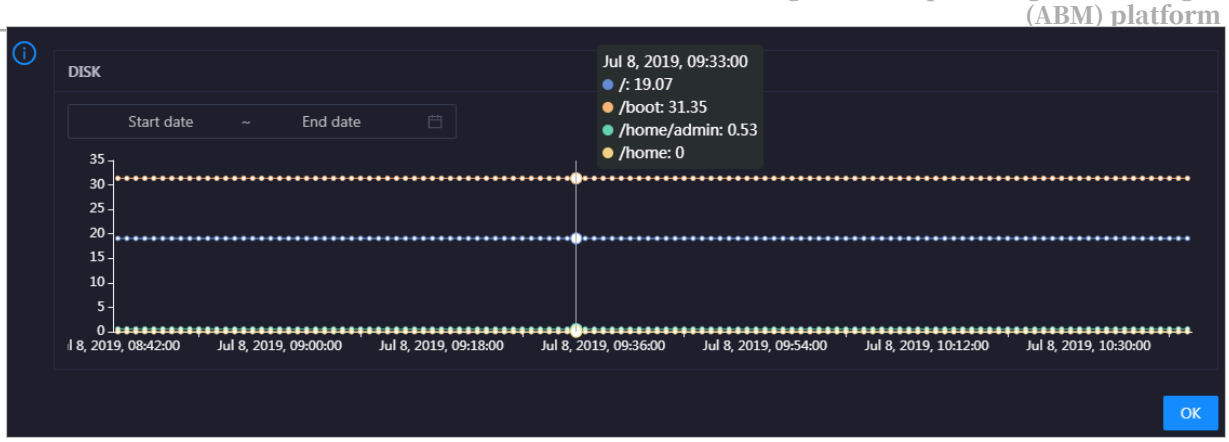

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

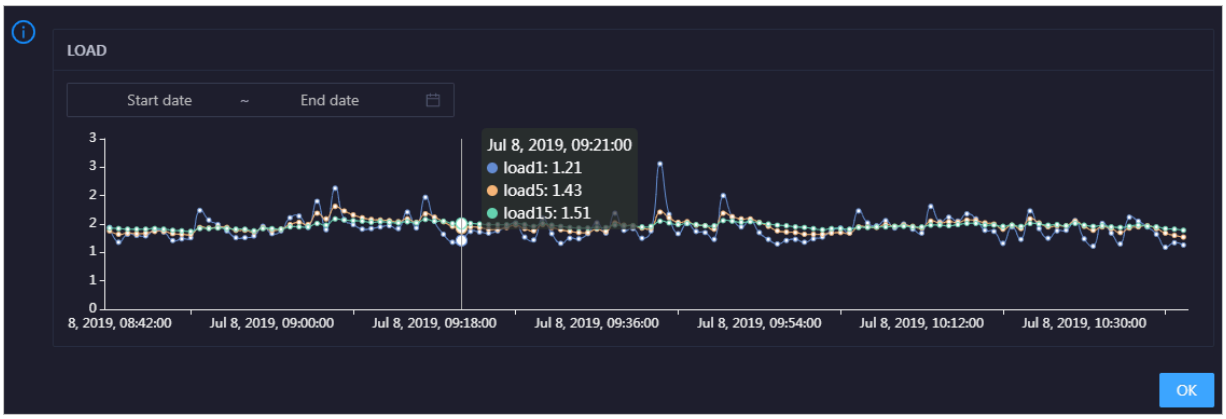

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

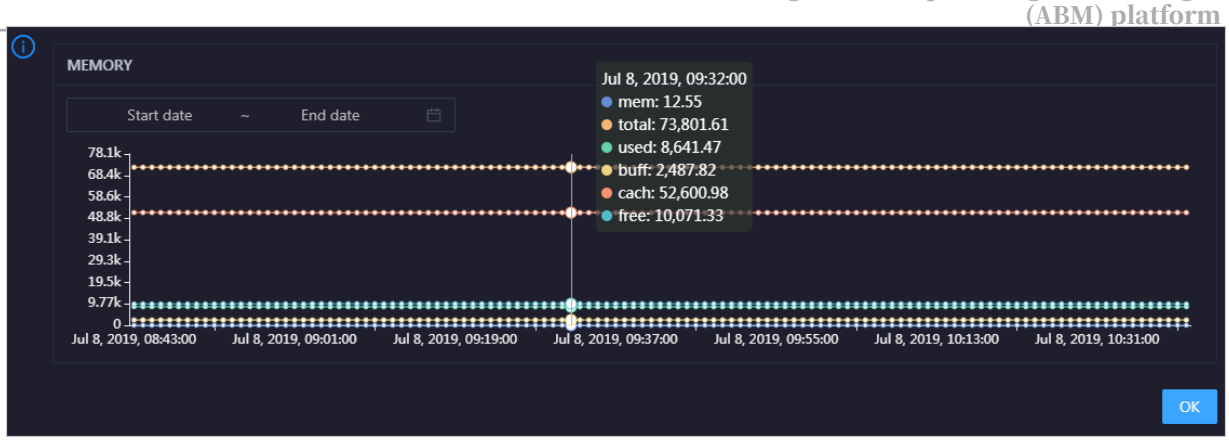

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

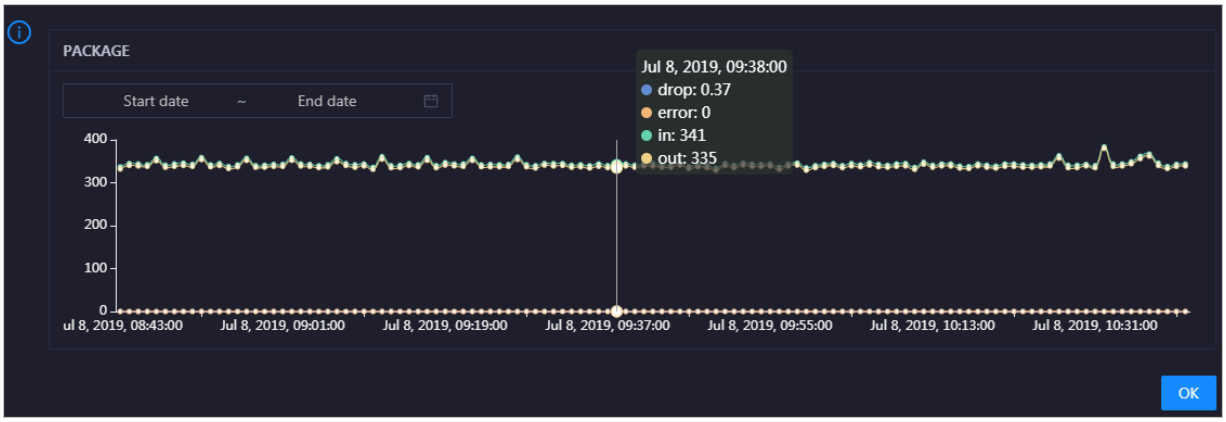

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

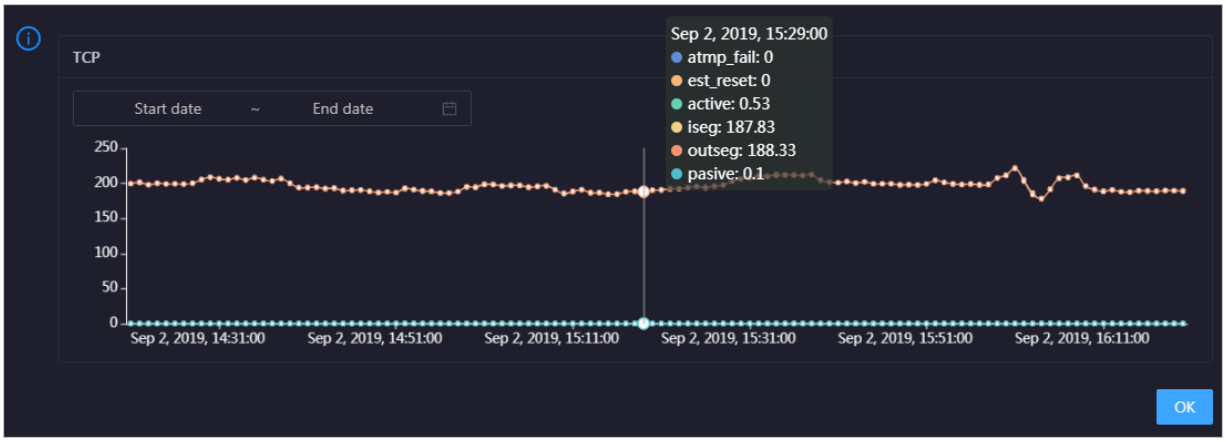

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

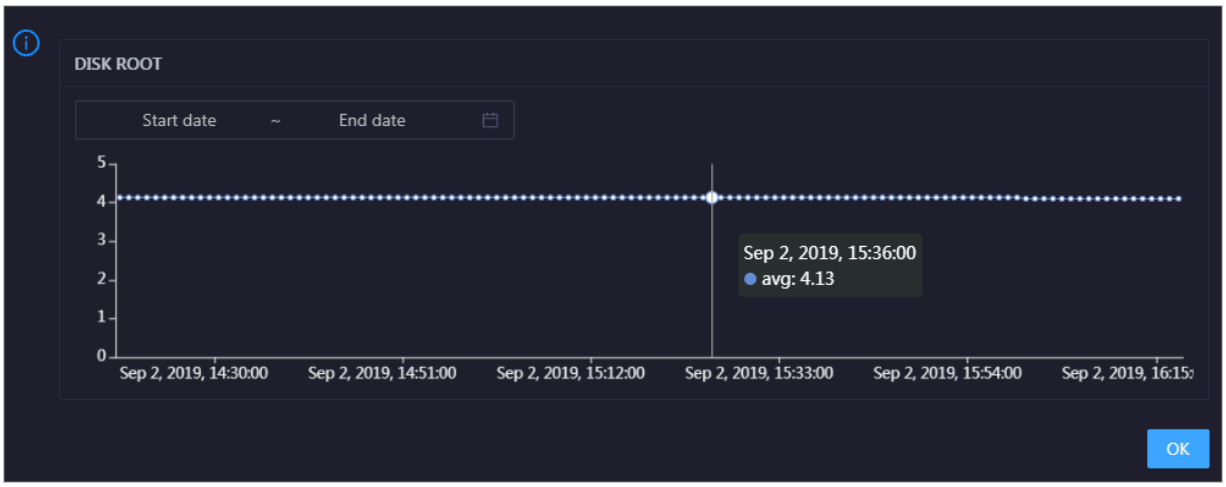

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

1.4.4.2 Service hosts

Apsara Bigdata 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.

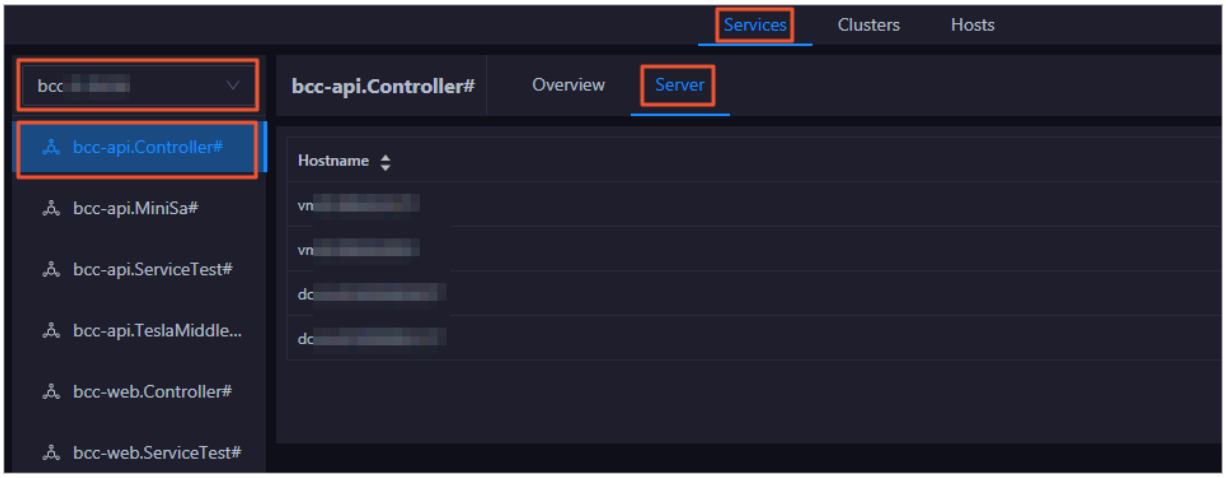

On the Server page, you can view the hosts where the selected service is run.

1.4.5 Cluster O&M

1.4.5.1 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, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster. The trend charts are described as follows:

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 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.

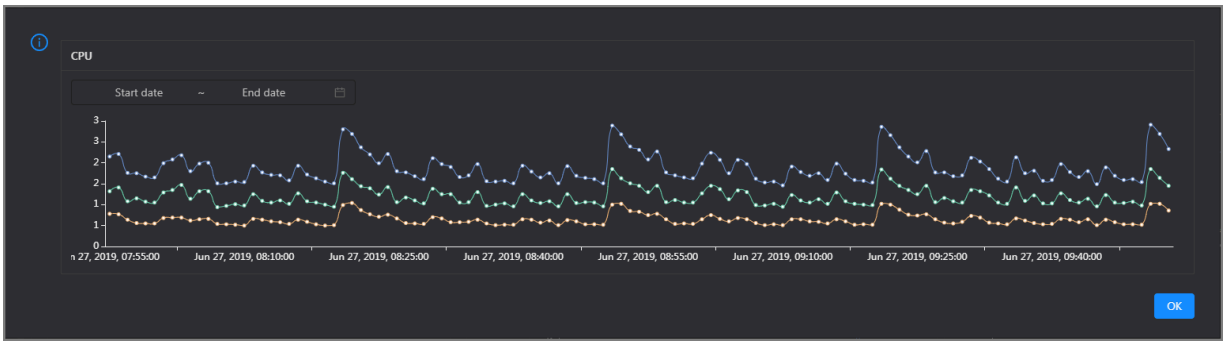

DISK

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

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

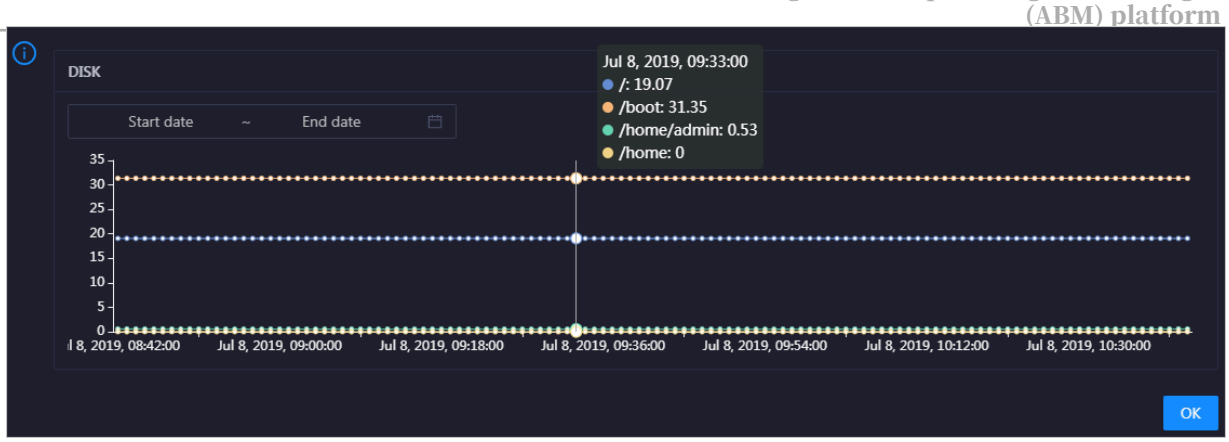

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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 cluster over time in different colors.

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

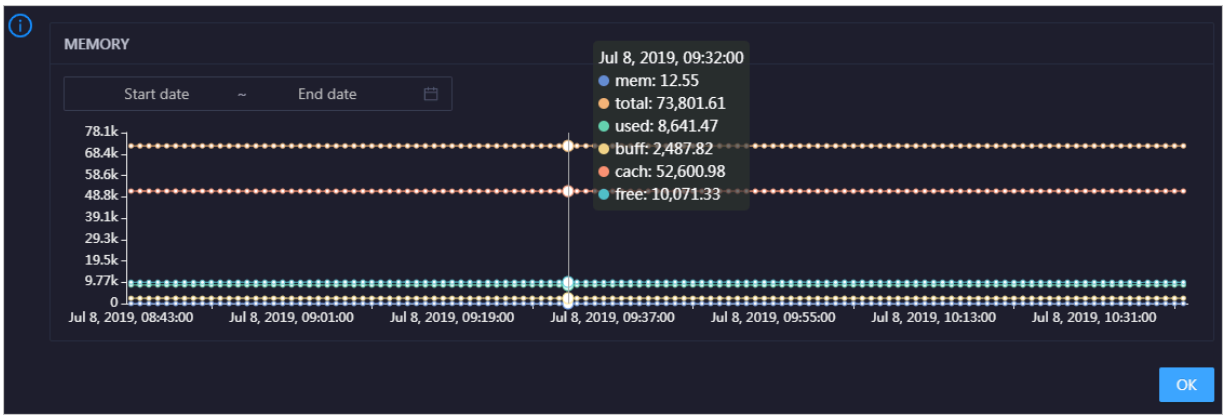

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

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

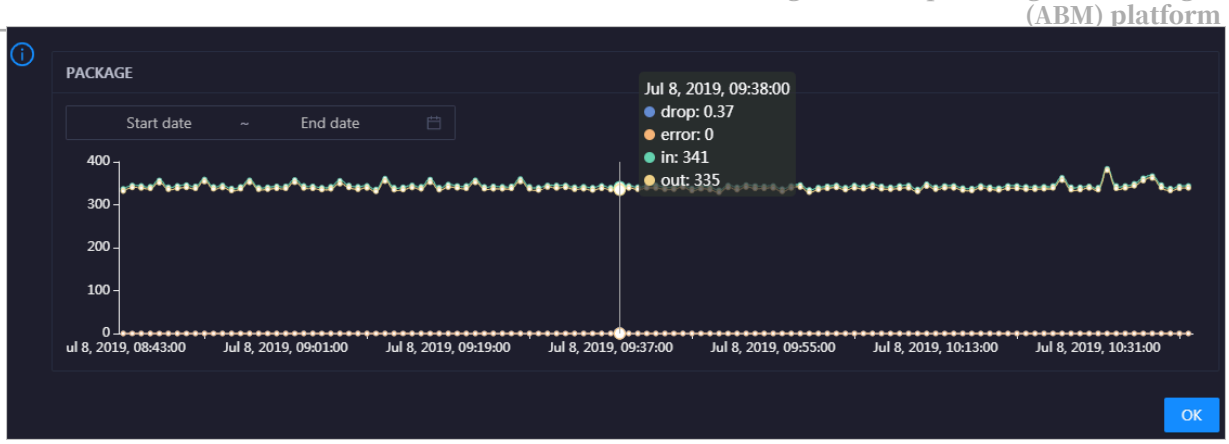

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

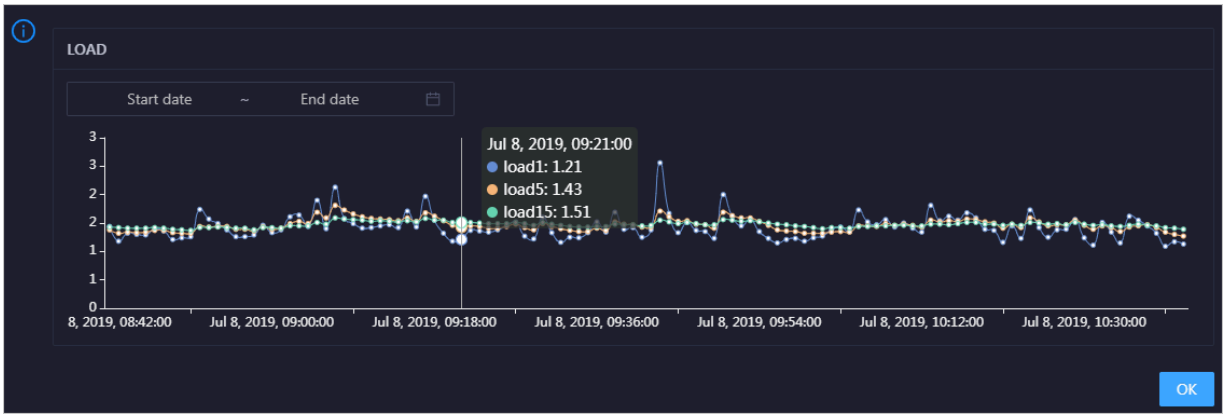

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

1.4.5.2 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.

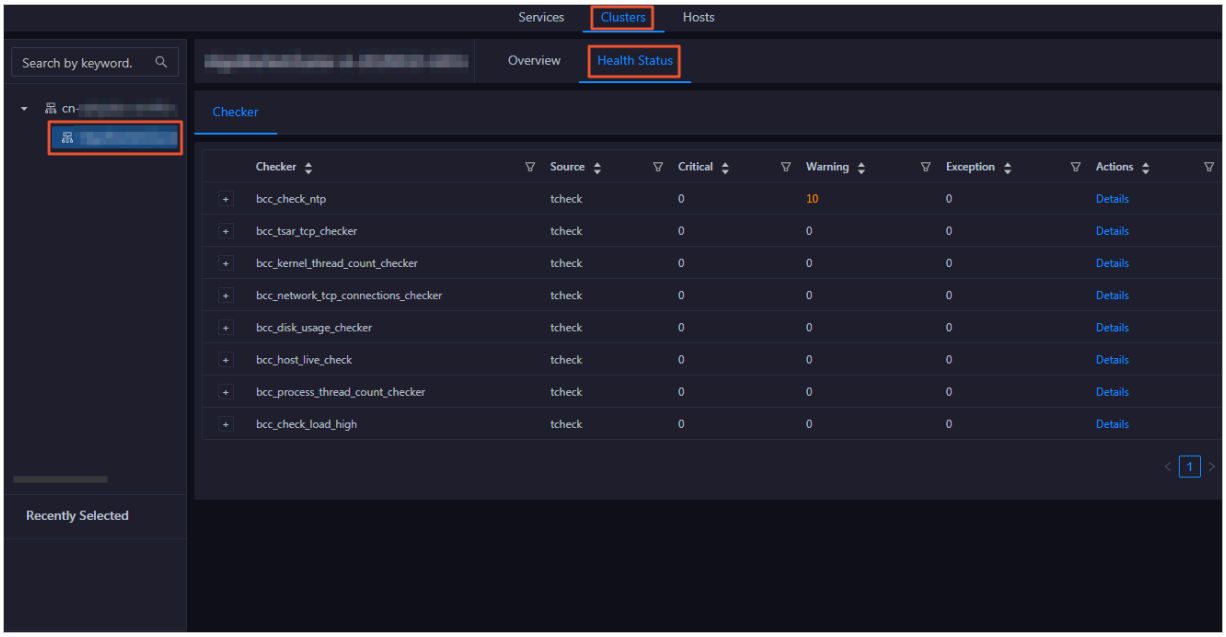

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

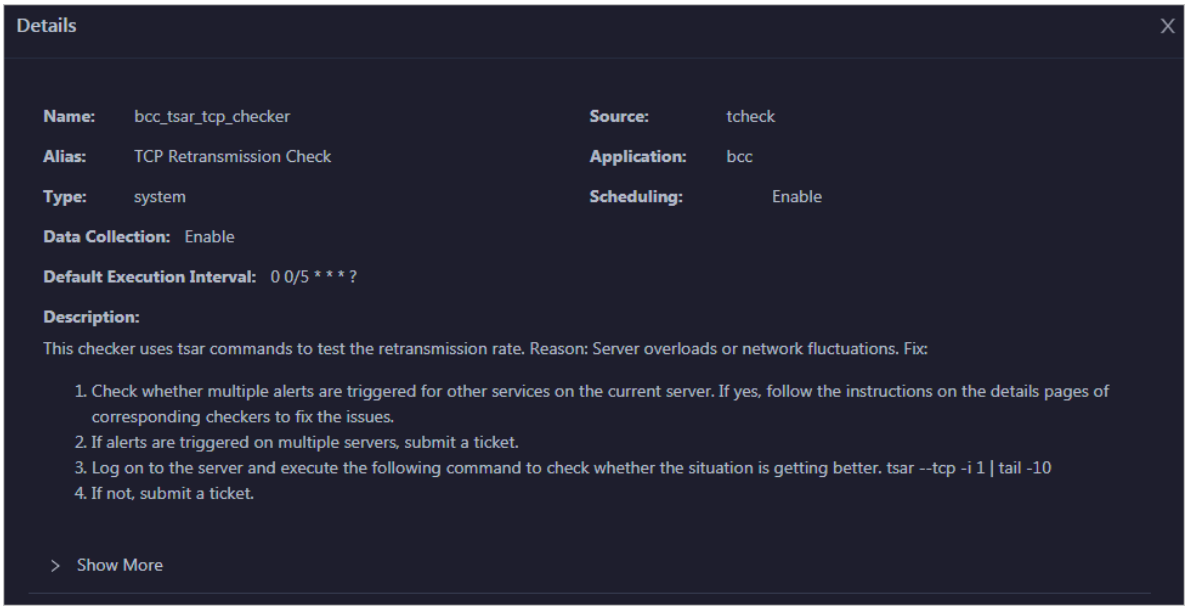

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

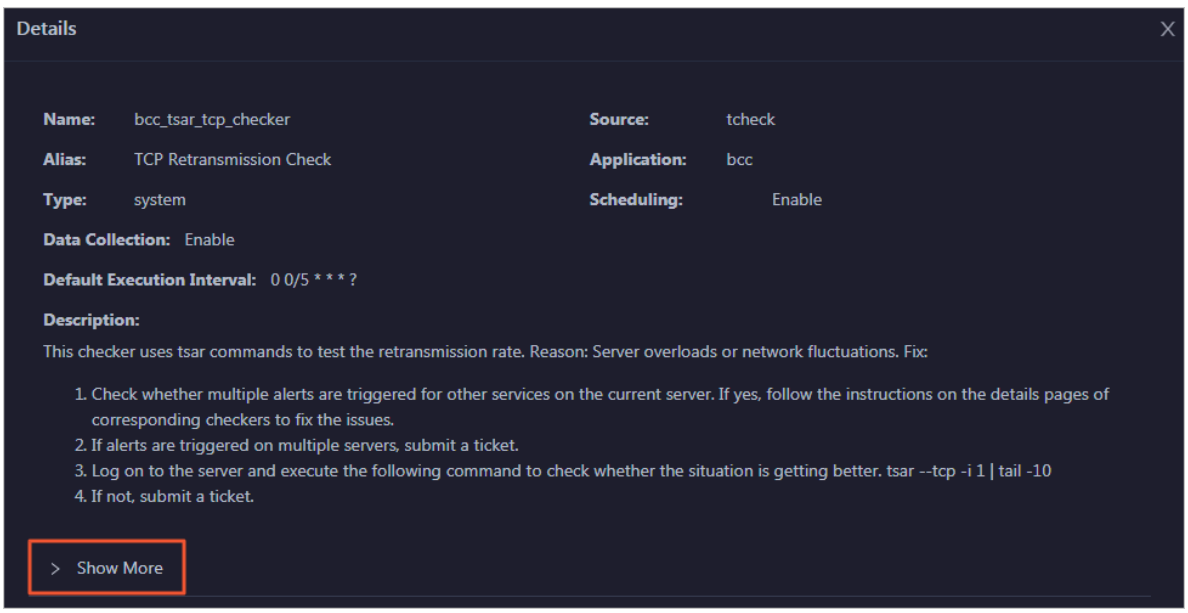

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View hosts with alerts and the alert causes

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

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

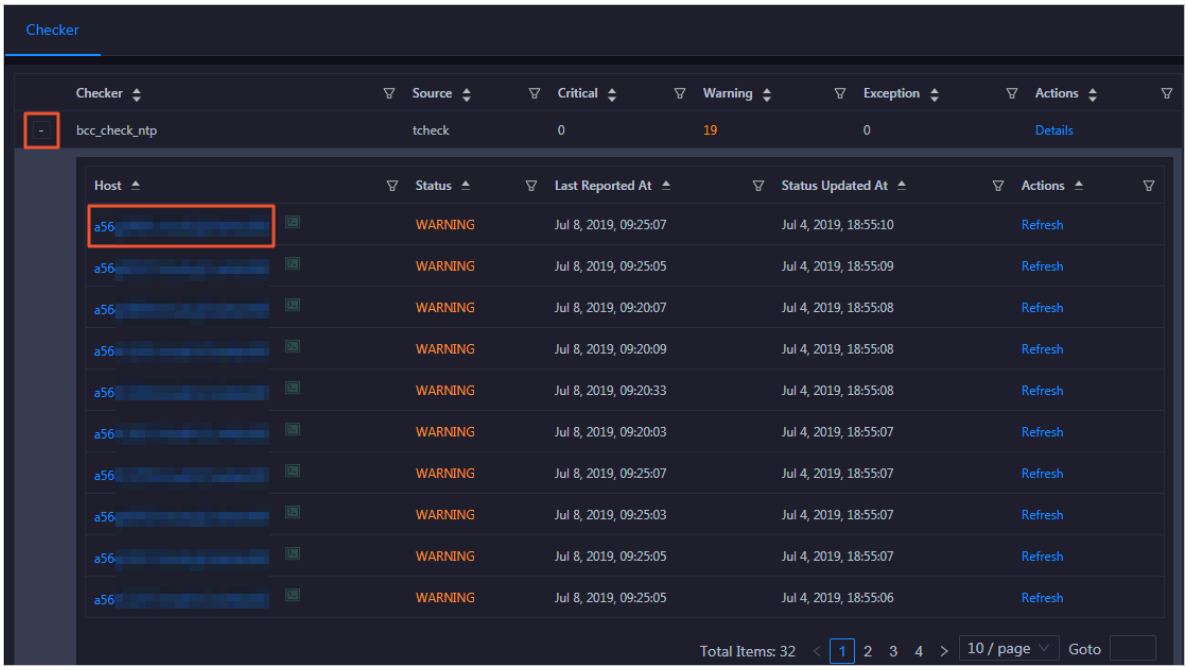

2. **Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

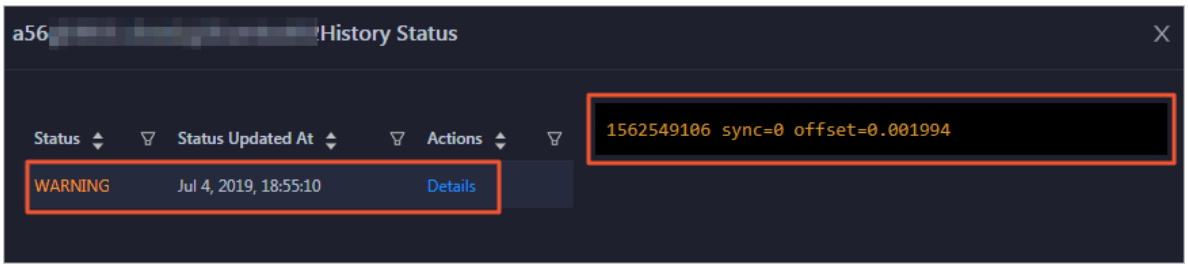

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

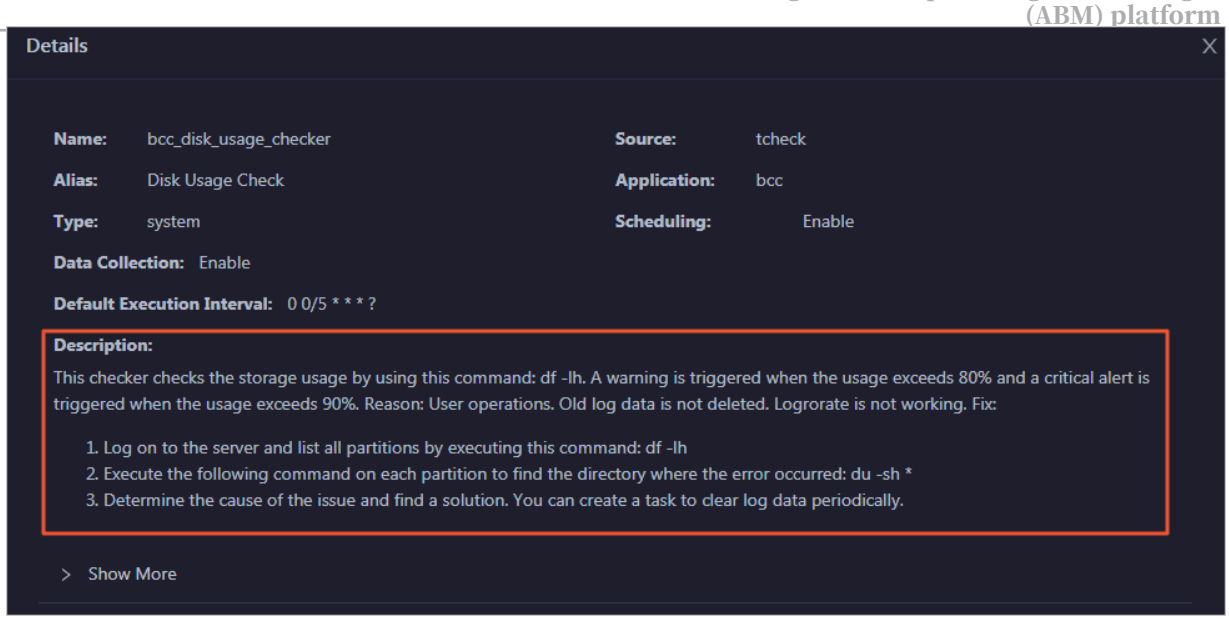

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

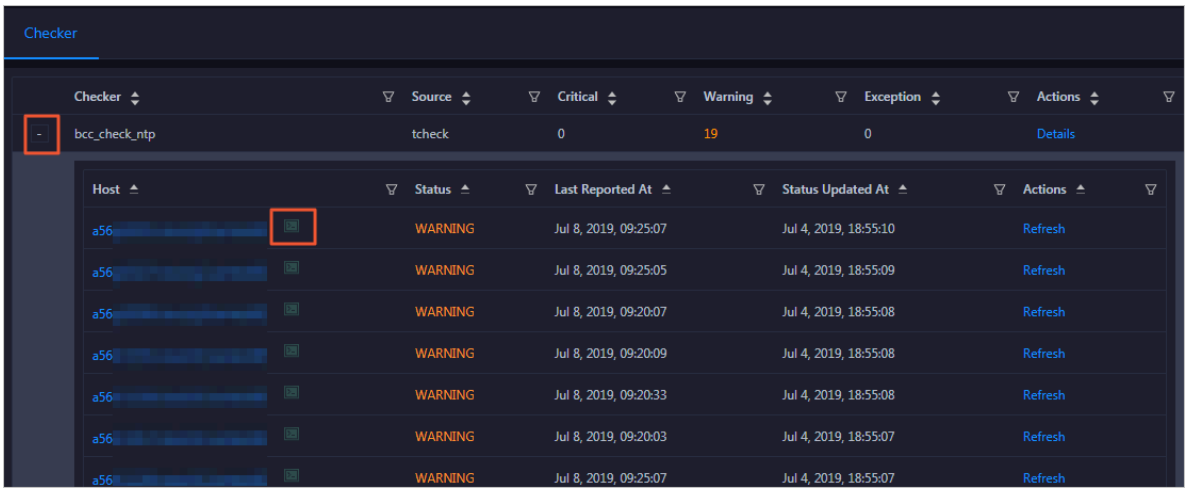

2. Click the Log On icon of a host. The TerminalService page appears.

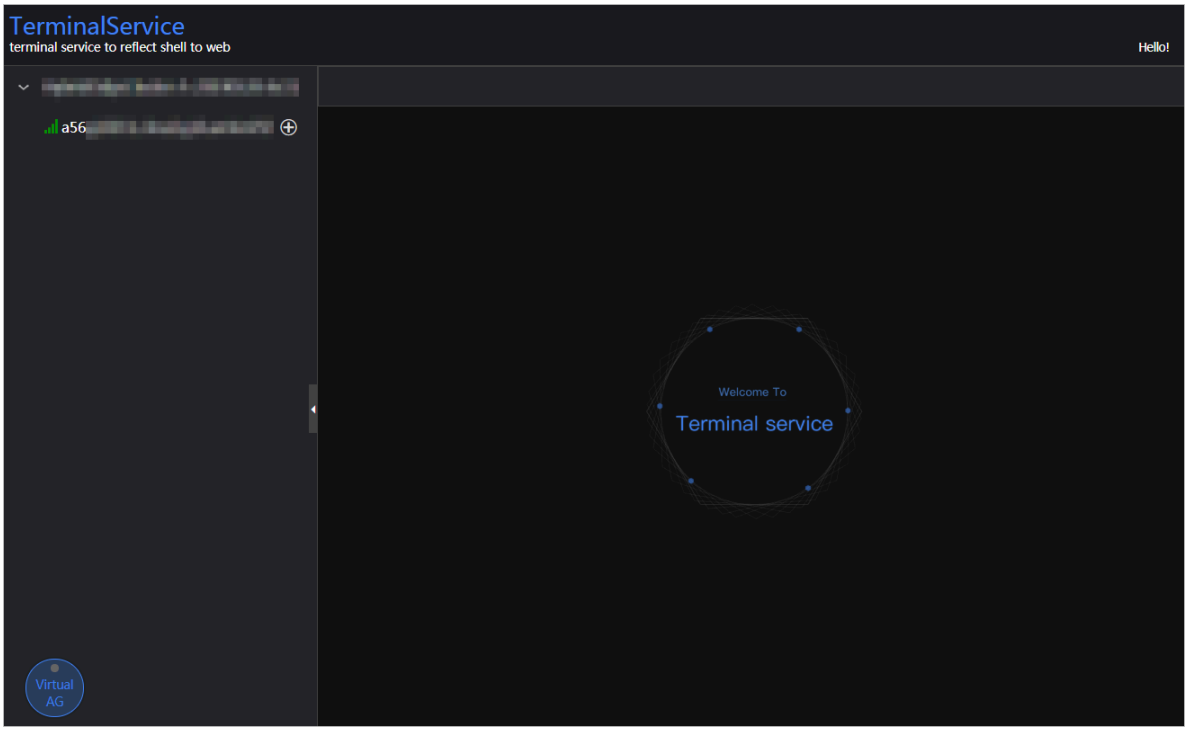

3. On the TerminalService page, click the hostname on the left to log on to the host.

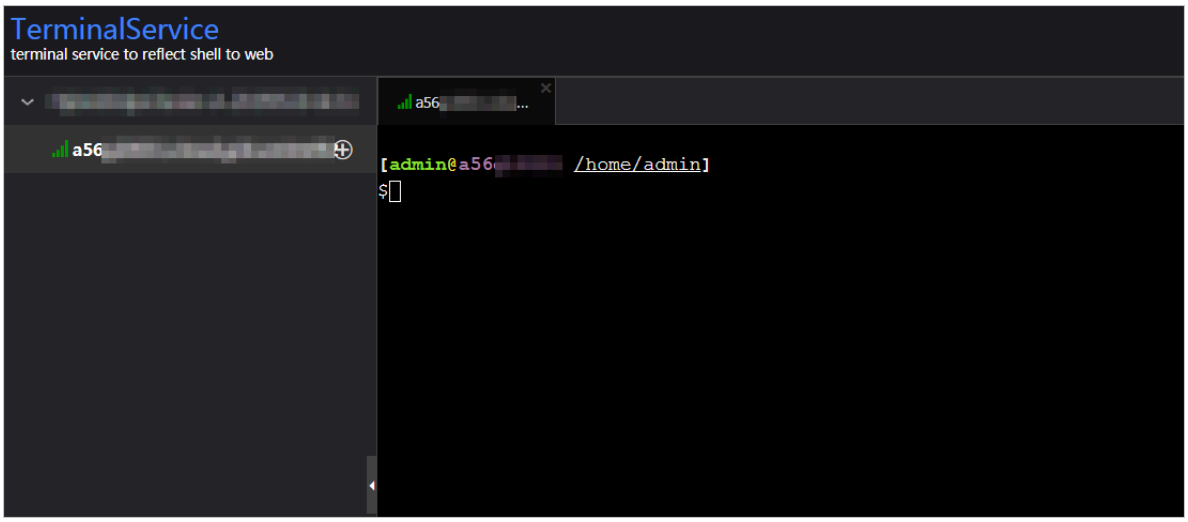

Run a checker again

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. In this way, you can check whether the alert is cleared.

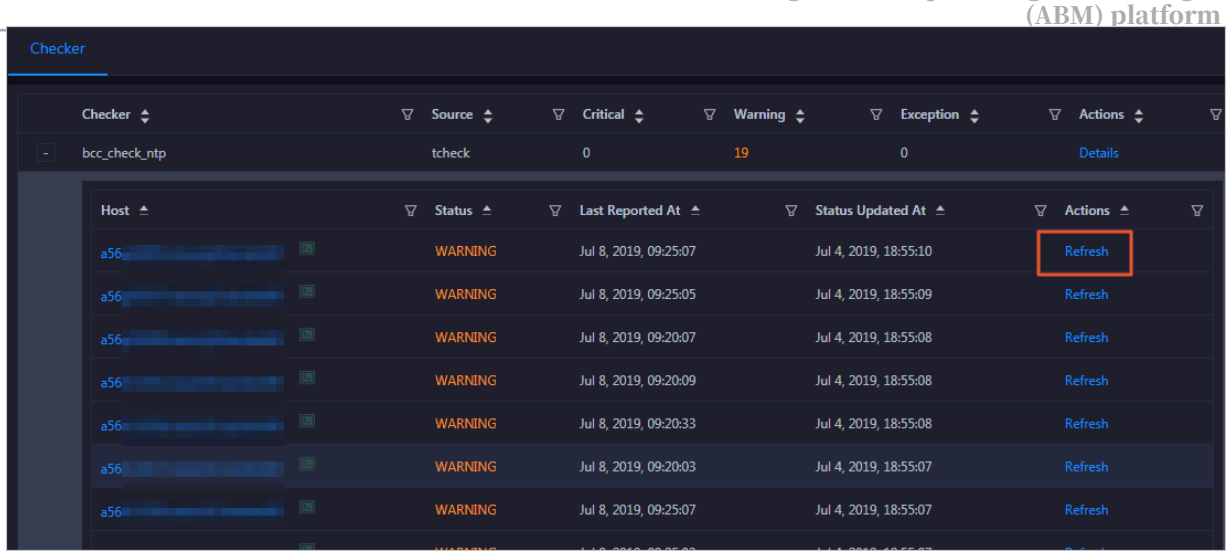

1.4.6 Host O&M

1.4.6.1 Host overview

The host overview page displays the overall running information about a host in an Apsara Bigdata 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.

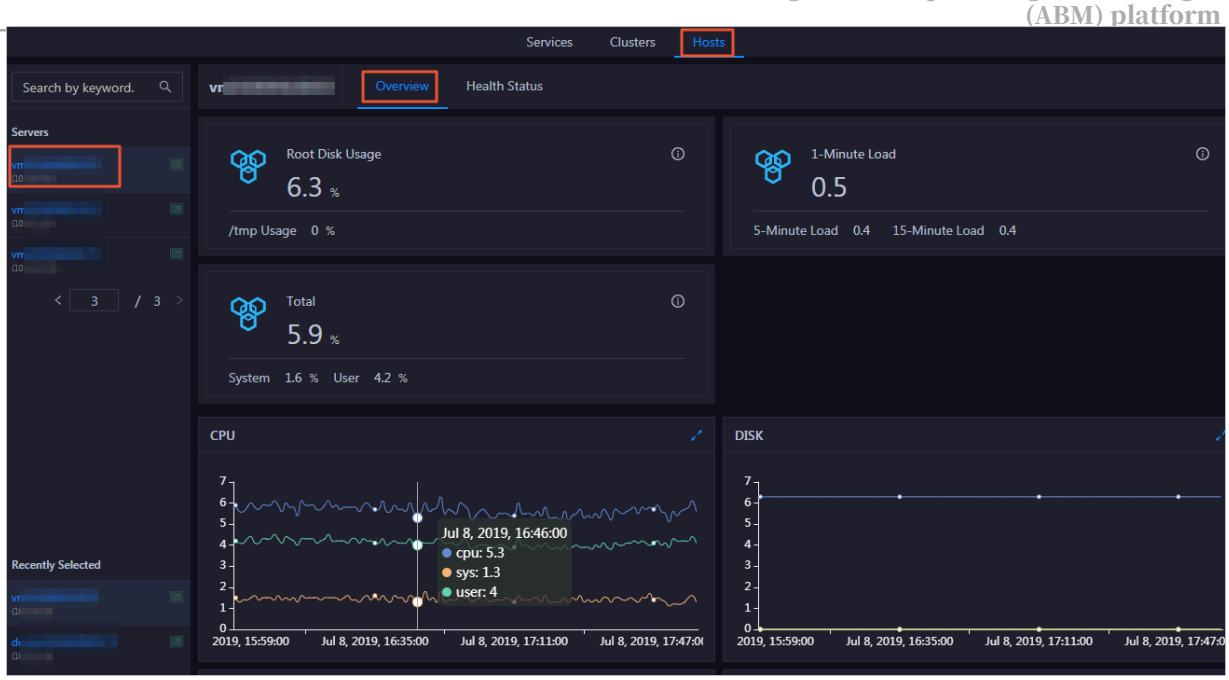

Root Disk Usage, Total, and 1-Minute Load

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.

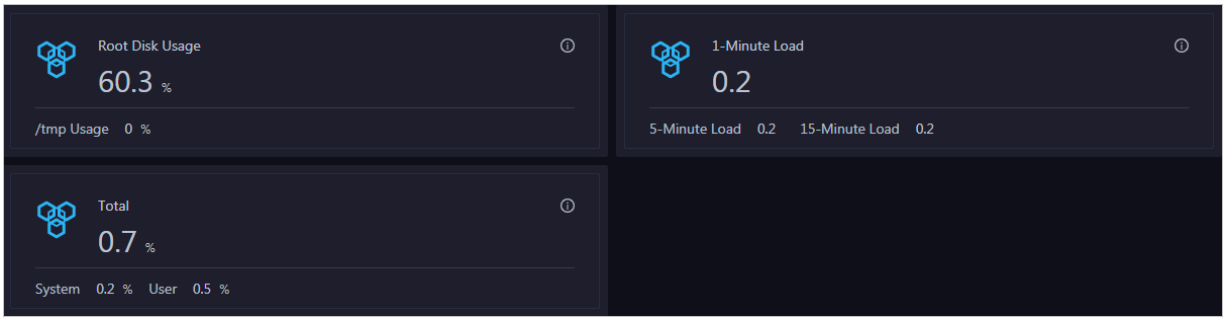

CPU

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 host over time in different colors.

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

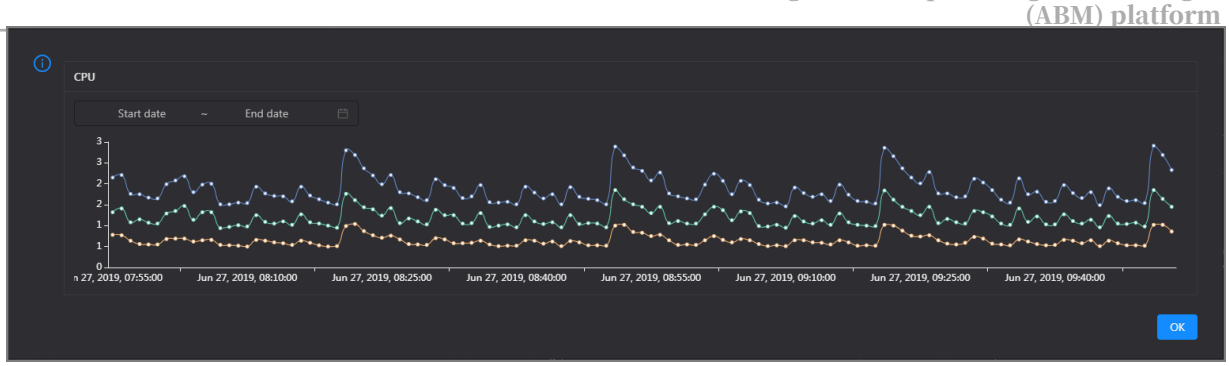

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

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

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

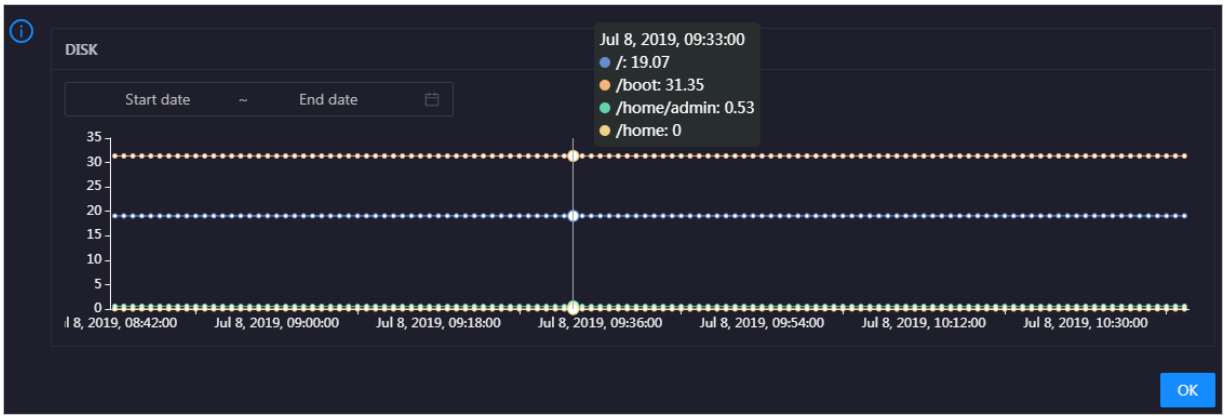

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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 host over time in different colors.

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

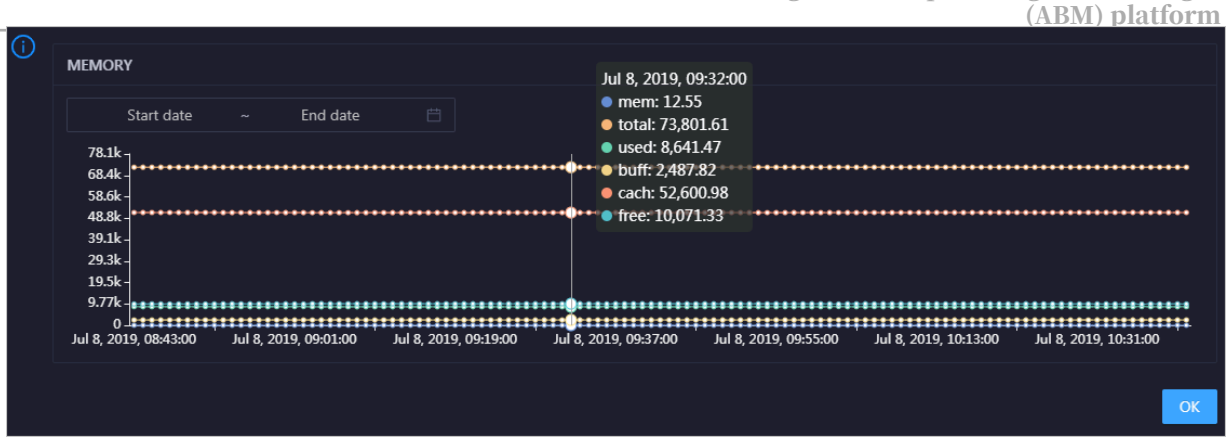

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

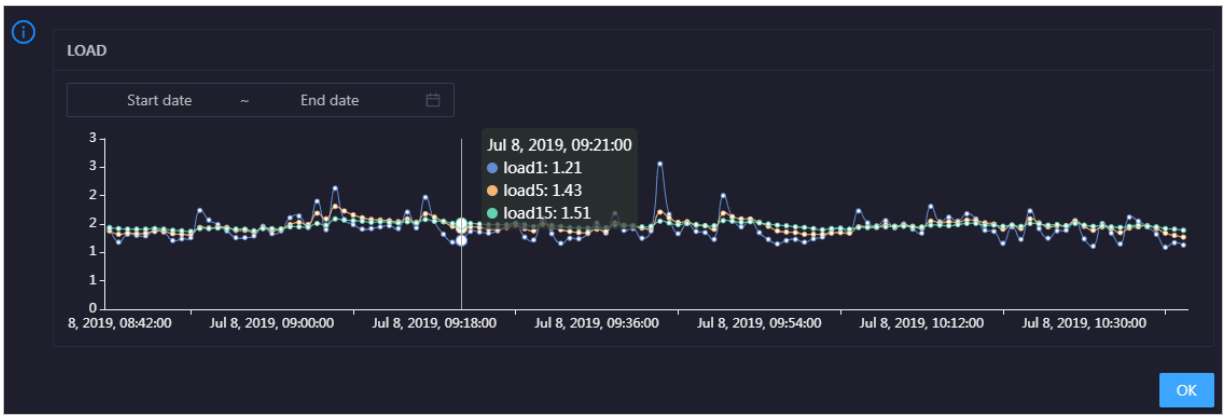

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 host over time in different colors. These trend lines reflect the data transmissi on status of the host.

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

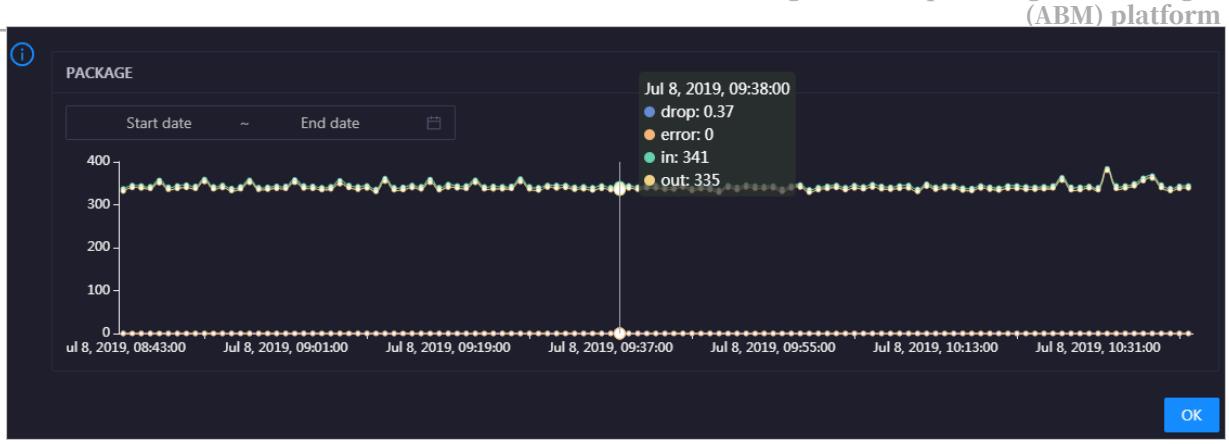

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

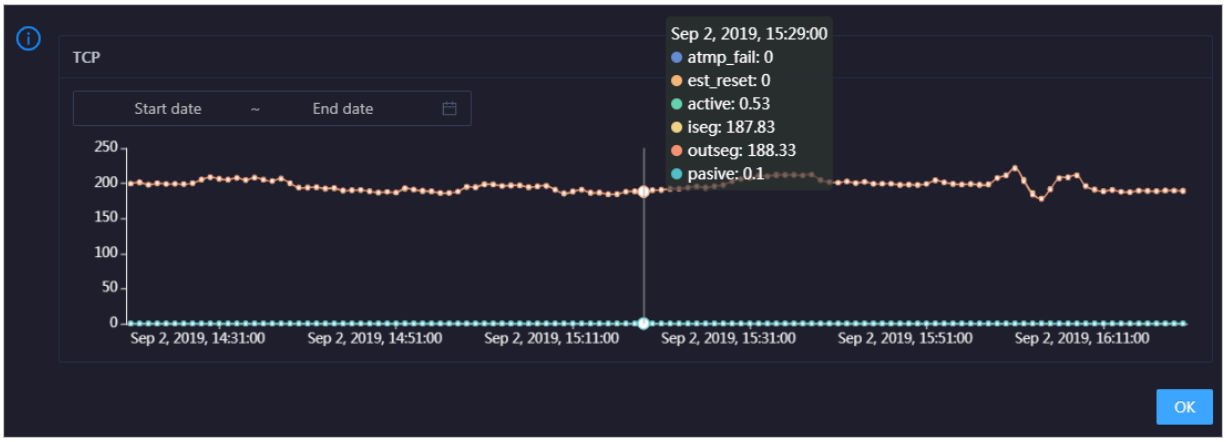

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

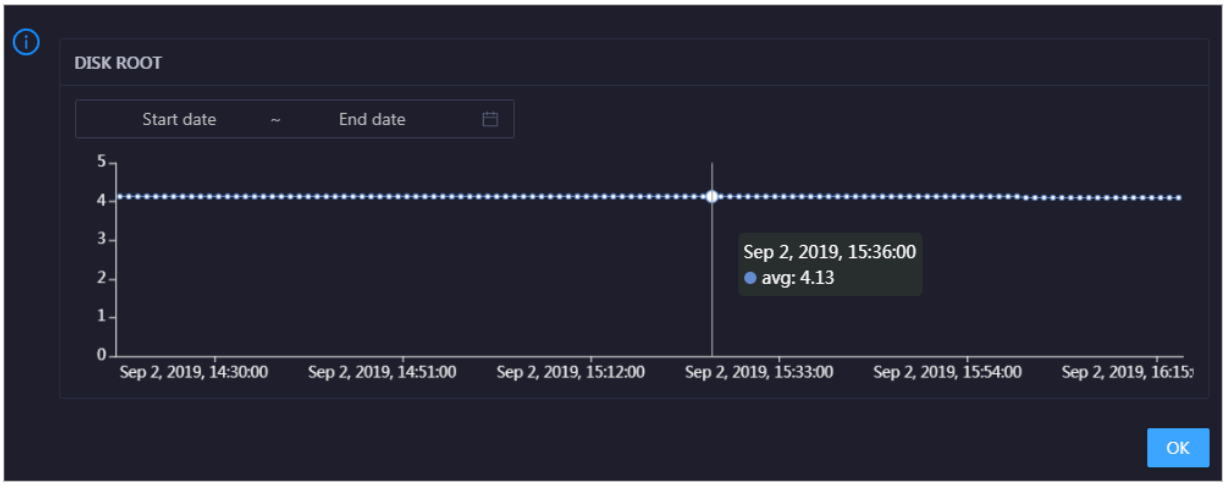

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

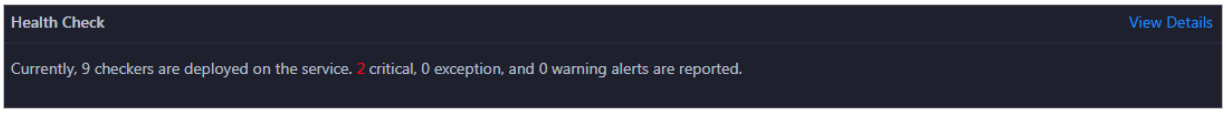

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

Health Check History

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

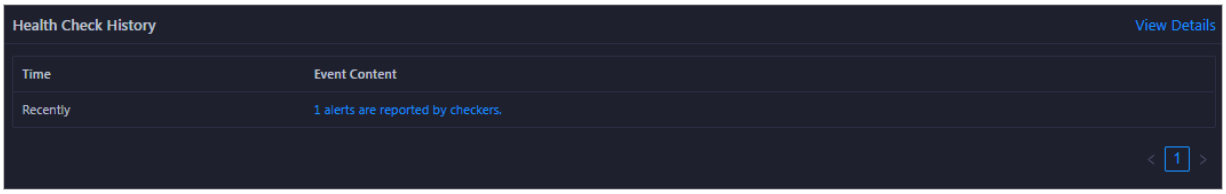

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

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

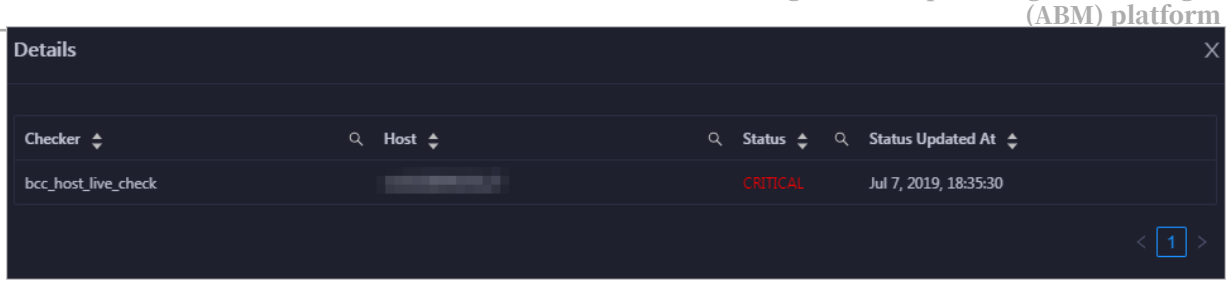

1.4.6.2 Host health

On the host health status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to the host.

Entry

At the top of the O&M page, click the Hosts tab. On the page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.

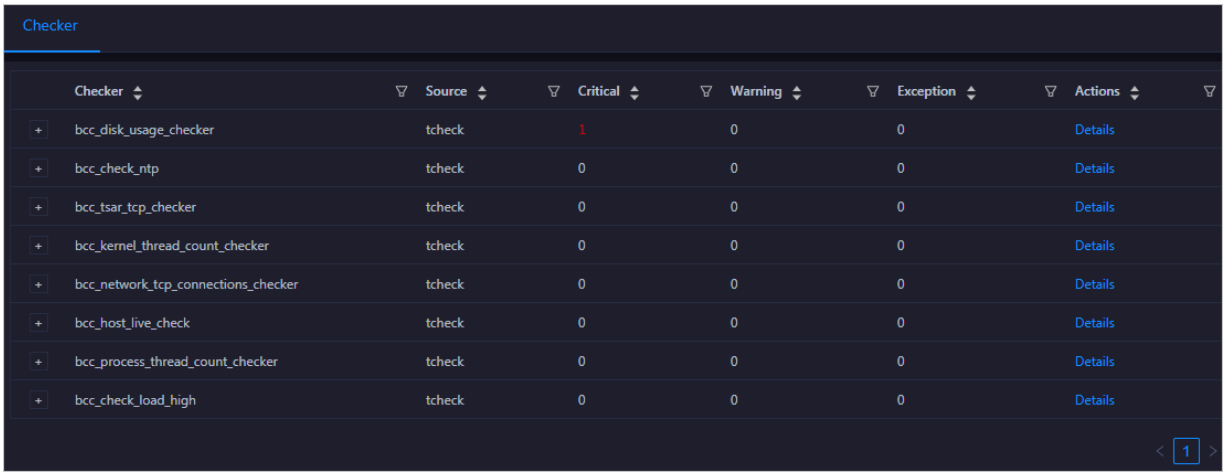

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

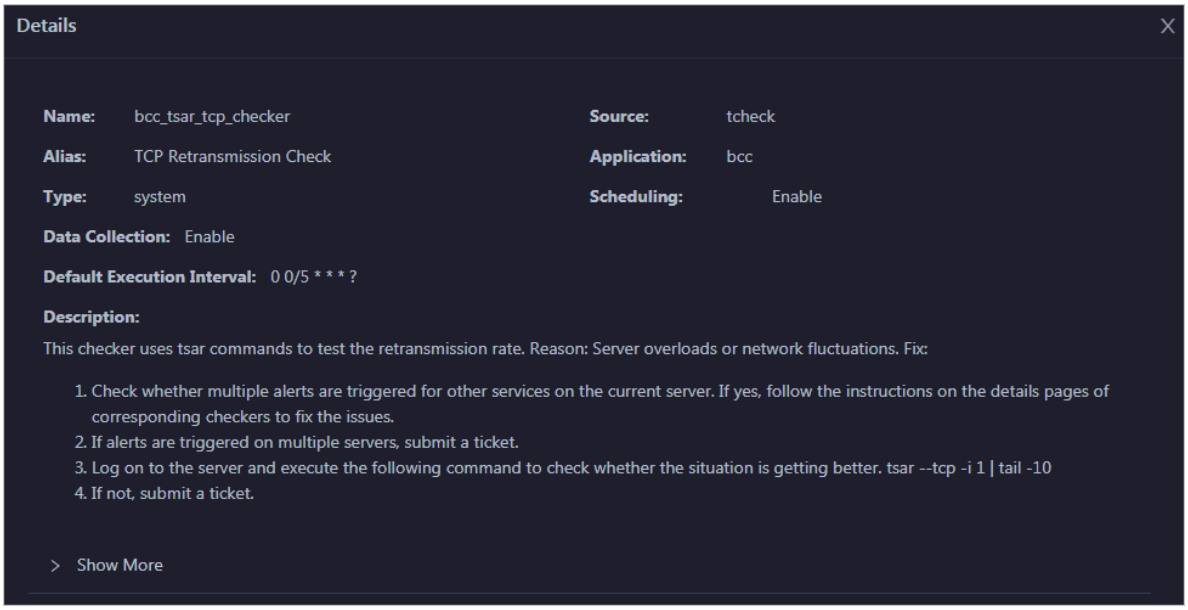

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

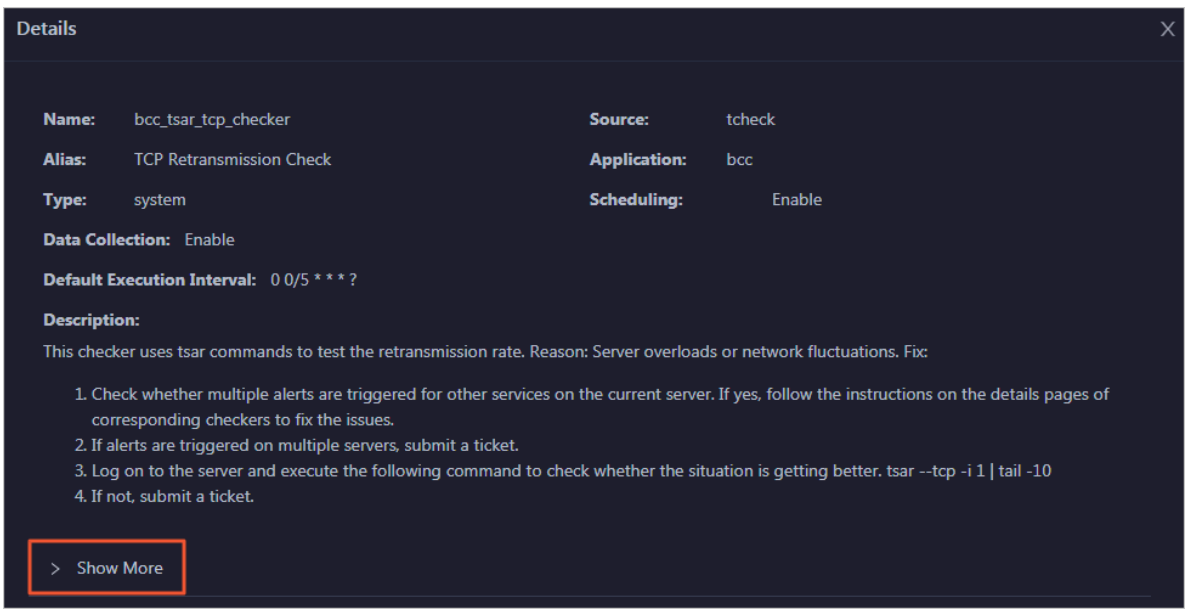

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

1. **On the Health Status page, click + to expand a checker with alerts.**

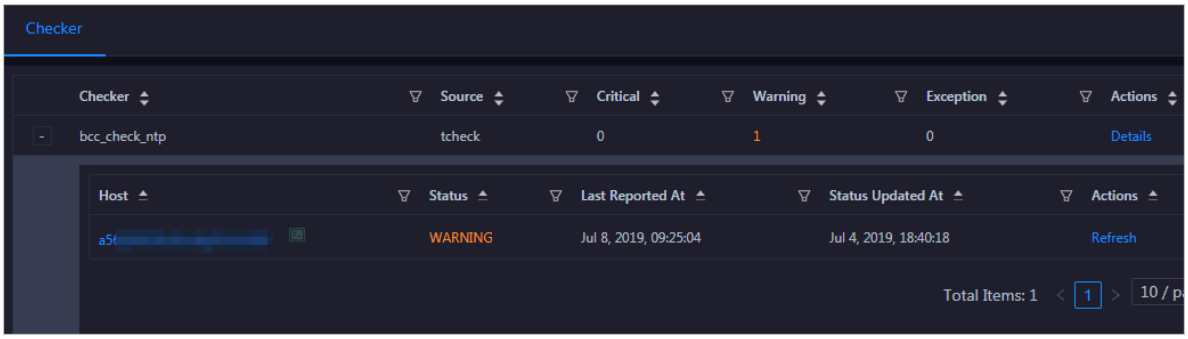

2. **Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

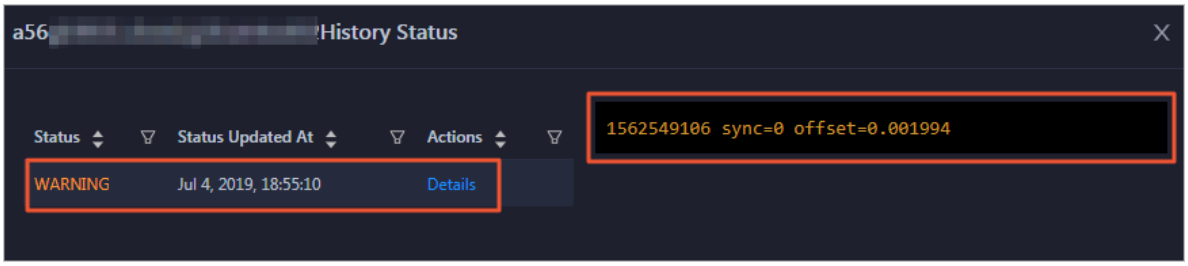

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

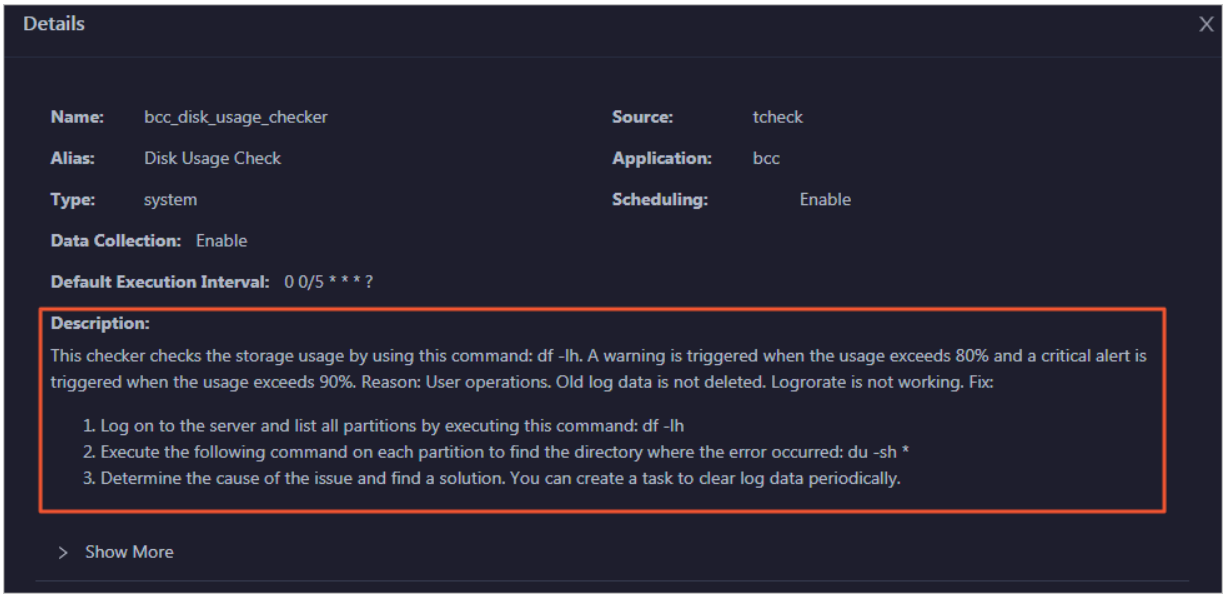

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

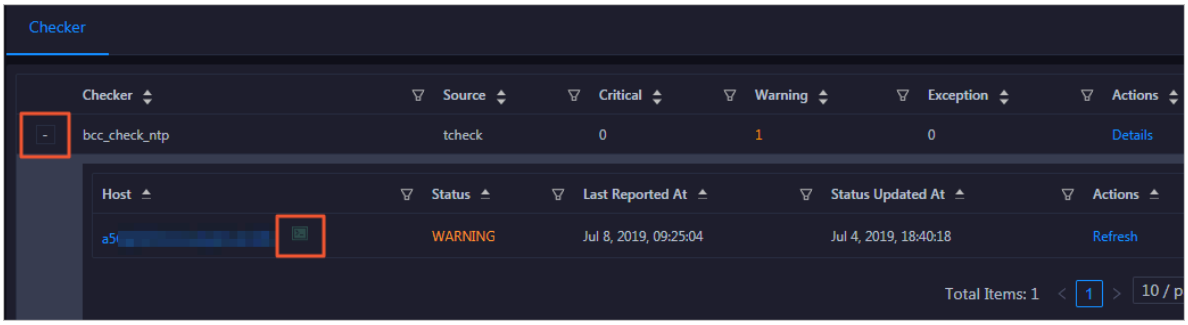

2. **Click the Log On icon of a host. The TerminalService page appears.**

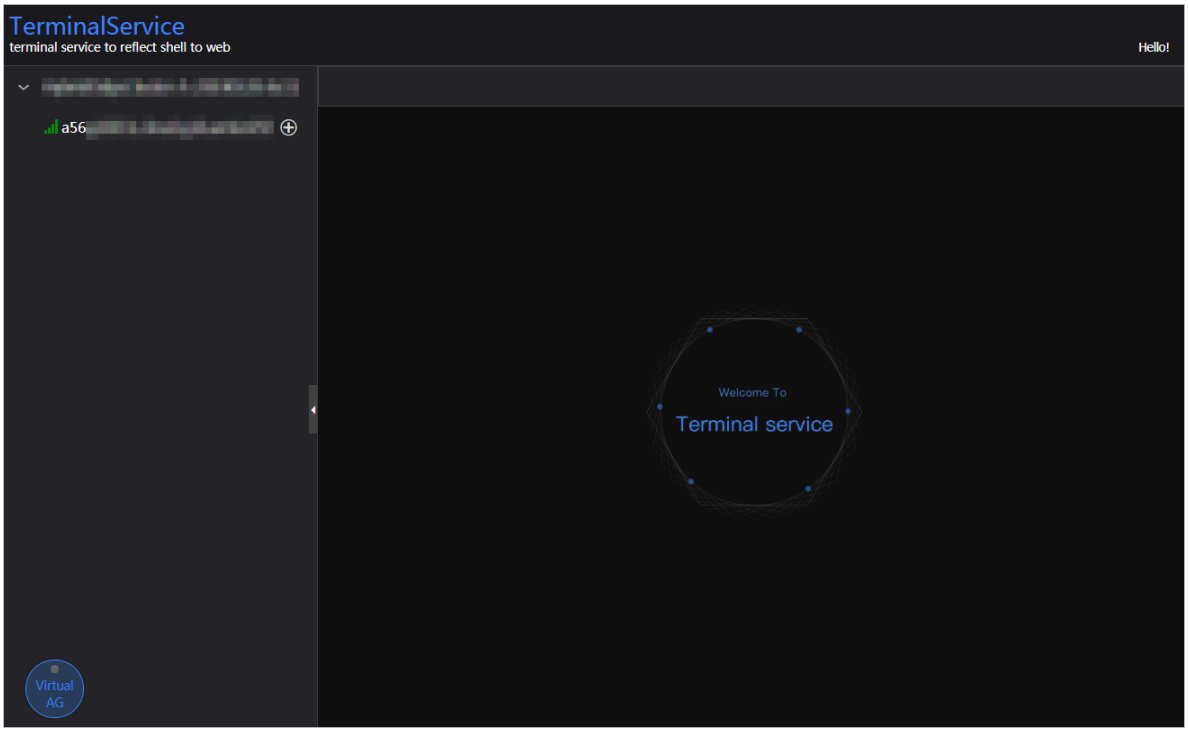

3. On the TerminalService page, click the hostname on the left to log on to the host.

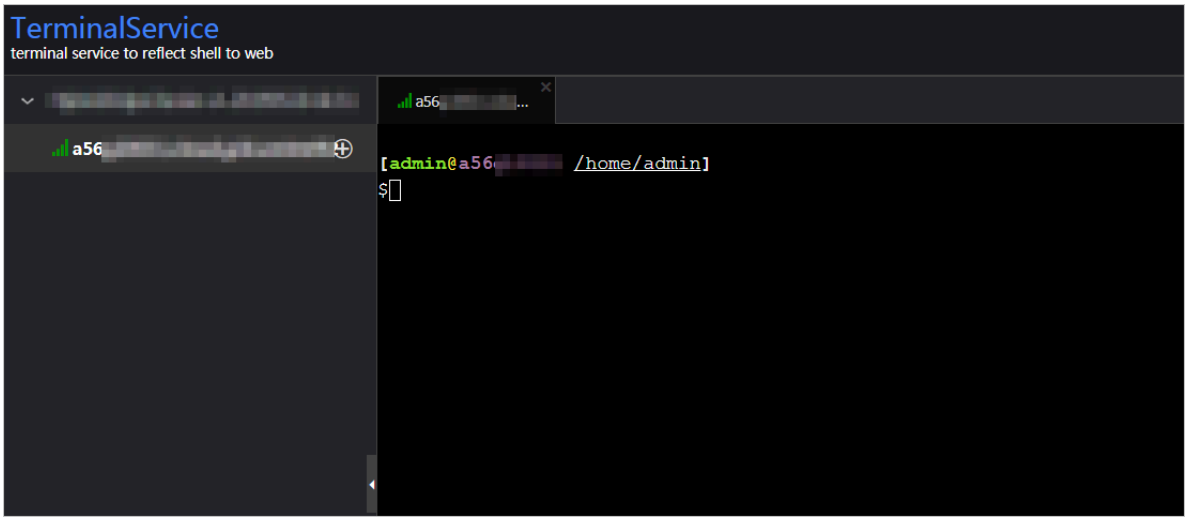

Run a checker again

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. In this way, you can check whether the alert is cleared.

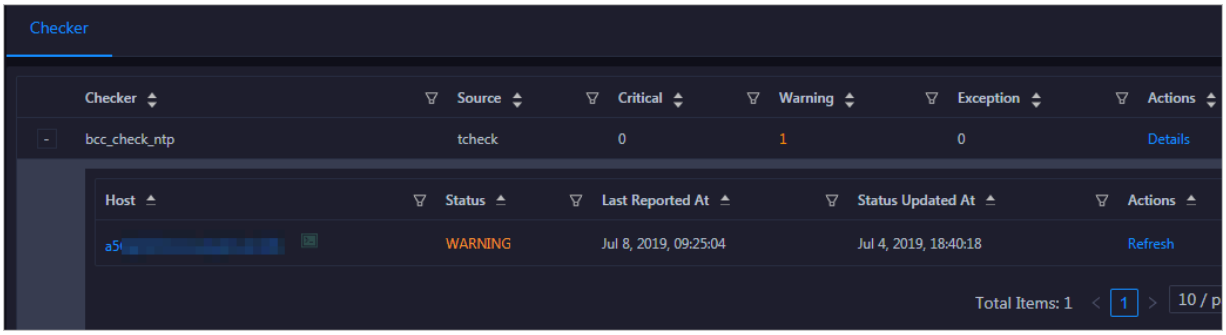

1.5 MaxCompute

1.5.1 MaxCompute workbench

In the Apsara Bigdata Manager (ABM) console, you can view your MaxCompute projects and project details, including the project overview, jobs, storage, configurations, quota groups, tunnels, resource analysis, and cross-cluster replication, on the MaxCompute workbench. You can also modify the current configuration of a project.

Entry

1. *Log on to the ABM console*.

2. Click in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page, click Business in the upper-right corner. The My Projects page under Workbench appears.

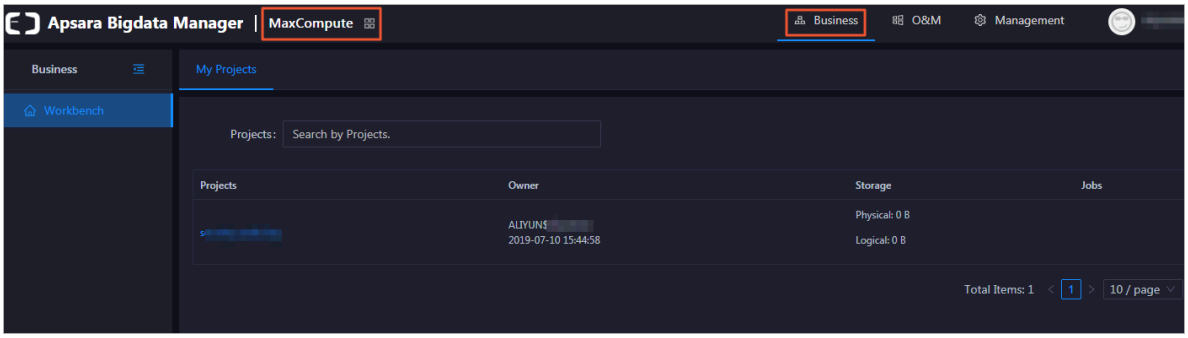

Overview

On the Overview page, you can view the following information about the selected project:

- Basic information about the project, such as the default quota group, creator, creation time, service, and region.
- Trend charts that display the trend lines of requested and used CPU and memory resources by minute over time in different colors.
- Trend charts that display the trend lines of CPU and memory usage by day over time in different colors.

On the My Projects page, click the name of a project in the project list. The Overview page appears.

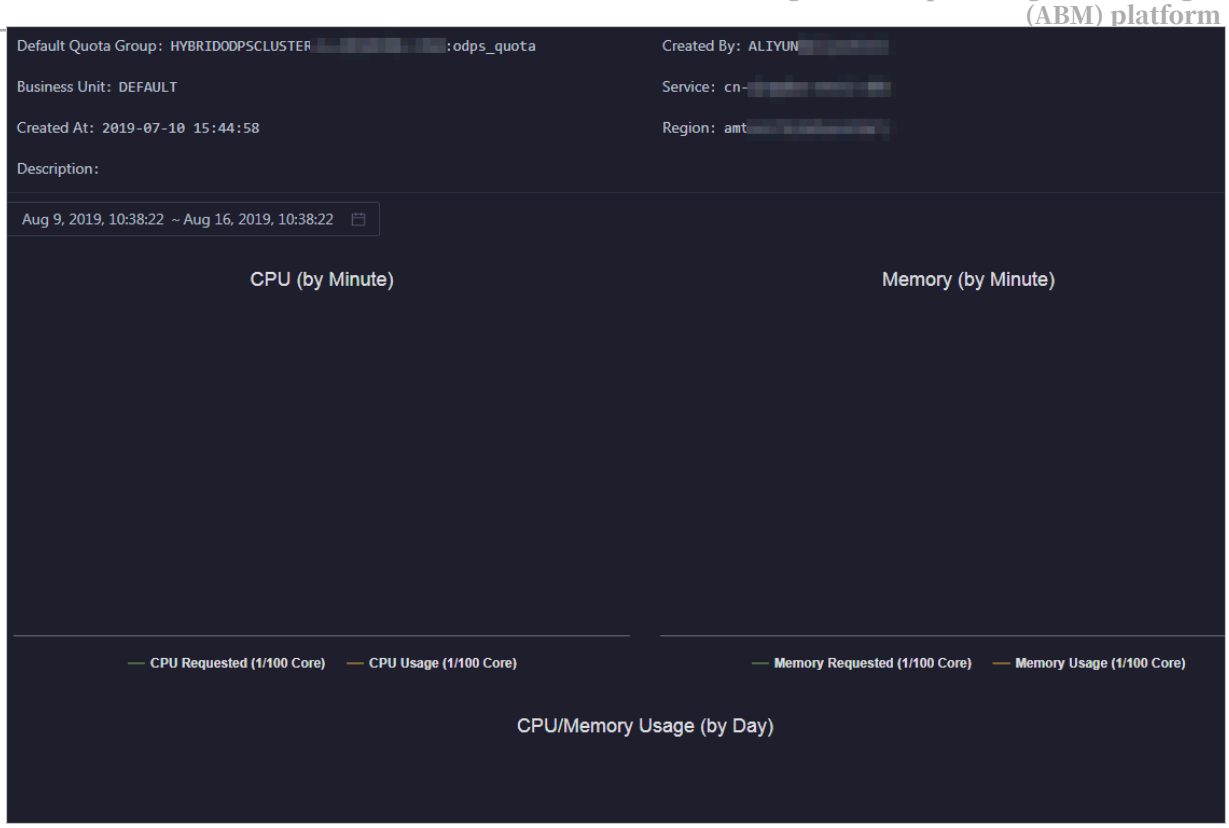

Jobs

On the Jobs page, you can view job snapshots by day over the last week. Detailed information about a job snapshot includes the job ID, project, quota group, submitter, running duration, minimum and maximum CPU usage, minimum and maximum memory usage, DataWorks node, running status, start time, priority, and type. You can also view the operational logs of a job to locate job failures.

On the My Projects page, click the name of a project in the project list and then click the Jobs tab. The Jobs page for the project appears.

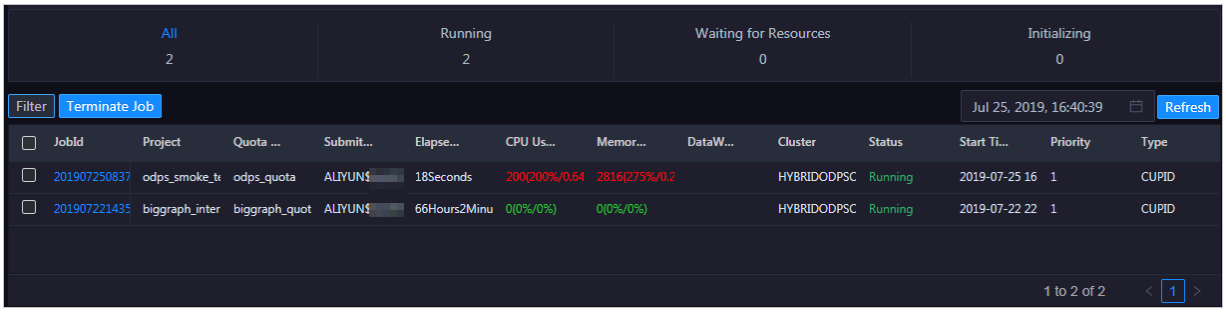

You can perform the following operations on jobs:

- **Customize columns or sort job snapshots based a specified column. For more information, see *Common operations*.**

- View operational logs of jobs or terminate jobs. For more information, see *Job snapshots*.

Storage

On the Storage page, you can view the storage usage, used storage space, storage quota, and available storage space. You can also view a trend chart that displays the trend lines of storage usage, the number of Apsara Distributed File System files, the number of tables, the number of partitions, and idle storage by day over time in different colors.

On the My Projects page, click the name of a project in the project list and then click the Storage tab. The Storage page for the project appears.

Configuration

On the Configuration page, you can configure the general, sandbox, SQL, MR, access control, and resource recycling properties of the project. You can also configure encryption algorithms for the project.

On the My Projects page, click the name of a project in the project list and then click the Configuration tab. The Properties page for the project appears.

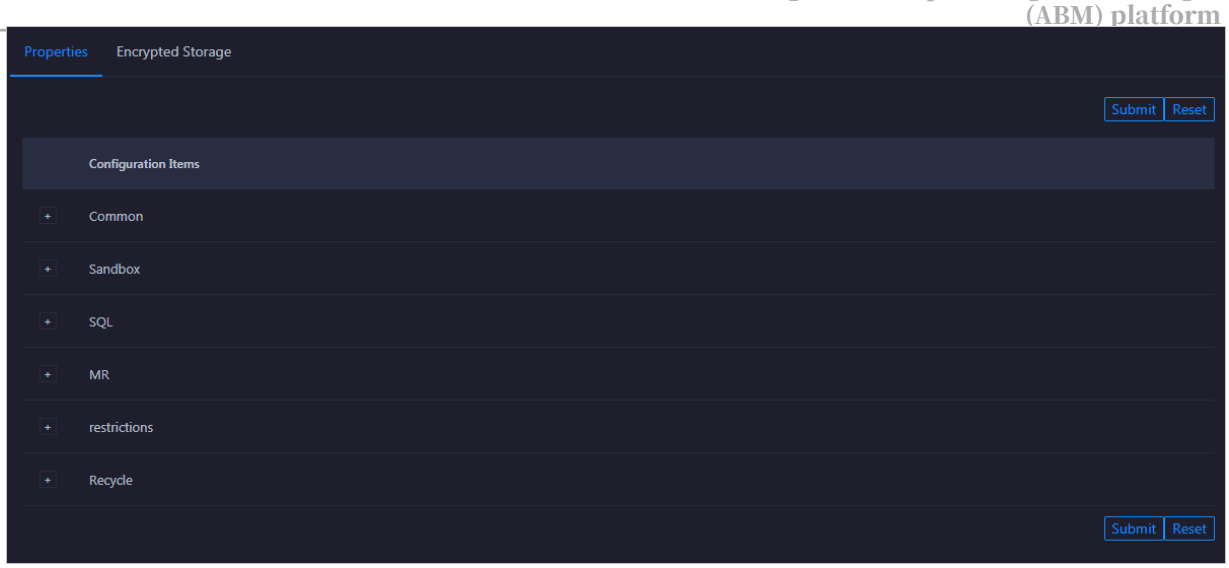

On the Properties page, you can view and modify each configuration item. To restore all configuration items to the default settings, click Reset.

On the Encrypted Storage page, you can configure the RC4 and AESCTR encryption algorithms.

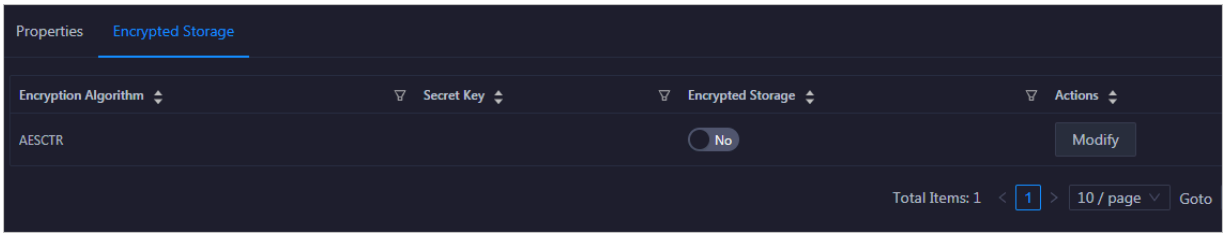

Quota Groups

On the Quota Groups page, you can view the quota groups of the project and the details of each quota group.

On the My Projects page, click the name of a project in the project list and then click the Quota Groups tab. The Quota Groups page for the project appears.

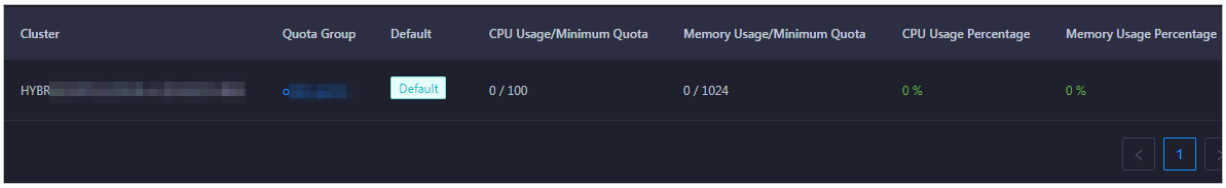

To view detailed information about a quota group, click the quota group name. For more information, see *View quota group details*.

Tunnel

On the Tunnel page, you can view the tunnel throughput of the project in the unit of bytes per minute. The Tunnel Throughput chart displays the trend lines of inbound traffic and outbound traffic over time in different colors.

On the My Projects page, click the name of a project in the project list and then click the Tunnel tab. The Tunnel page for the project appears.

Resource Analysis

On the Resource Analysis page, you can view the resource usage for the project from different dimensions, including tables, tasks, execution time, start time, and engines. For more information, see *Resource analysis*.

On the My Projects page, click the name of a project in the project list and then click the Resource Analysis tab. The Tables page appears.

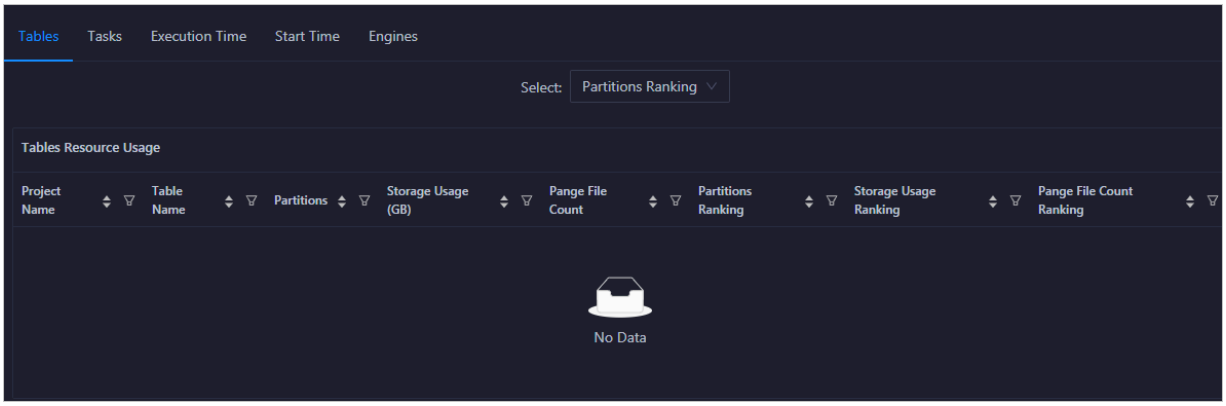

Cross-Cluster Replication

ABM allows you to view the projects with the cross-cluster replication feature enabled and the details and status of cross-cluster replication on the Cross-Cluster Replication page.

When you deploy multiple clusters for using MaxCompute, MaxCompute projects may be mutually dependent. In this case, data maybe directly read between projects. MaxCompute regularly scans tables or partitions that are directly read by other tables or partitions. If the duration of direct data reading reaches the specified upper limit, MaxCompute adds the tables or partitions to the cross-cluster replication list.

Assume that project 1 of cluster A depends on table 1 of project 2 of cluster B. In this case, project 1 directly reads data from table 1. If the duration of direct data

reading reaches the specified upper limit, MaxCompute adds table 1 to the cross-cluster replication list.

The Cross-Cluster Replication page consists of the Details and Configuration tabs.

- **Details: displays the information about the tables that support cross-cluster synchronization, including the project name, cluster name, table name, partition , storage space, number of files, and the cluster to which the data is synchronized .**
- **Configuration: displays the configuration of the tables that support cross-cluster synchronization, including the table name, priority, the cluster to which the data is synchronized, and lifecycle. You can also view the progress of cross-cluster replication for a table.**

1.5.2 Business O&M

1.5.2.1 Business O&M overview

This topic describes the features of MaxCompute business O&M and how to access the MaxCompute business O&M page.

Modules

- **Projects:**
  - **Project List: displays all projects in the MaxCompute cluster and the project details. You can filter, query, and sort projects. You can also modify the quota group of a project. If zone-disaster recovery is enabled, you can set resource replication parameters and determine whether to enable resource replication for a project.**
  - **Authorize Package for Metadata Warehouse: allows you to authorize members of a project to access the metadata warehouse.**
  - **Encryption at Rest: allows you to encrypt the data stored in MaxCompute projects.**
  - **Disaster Recovery: allows you to view the cluster status when zone-disaster recovery is enabled for MaxCompute. You can enable switchover between the primary and standby clusters. You can also determine whether to run scheduled tasks to synchronize resources between the primary and standby clusters.**

- **Jobs: displays the information about jobs deployed in the MaxCompute cluster. You can filter and search for these jobs. You can also view the operational logs, terminate a running job, and collect job logs.**
- **Business Optimization:**
  - **File Merging: allows you to create file merge tasks for clusters and projects. You can also filter merge tasks and view records of the tasks.**
  - **File Archiving: allows you to create file archiving tasks for clusters and projects. You can also filter archiving tasks and view records of the tasks.**
  - **Resource Analysis: allows you to view the resource usage of the cluster from different dimensions.**

Entry

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **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.**

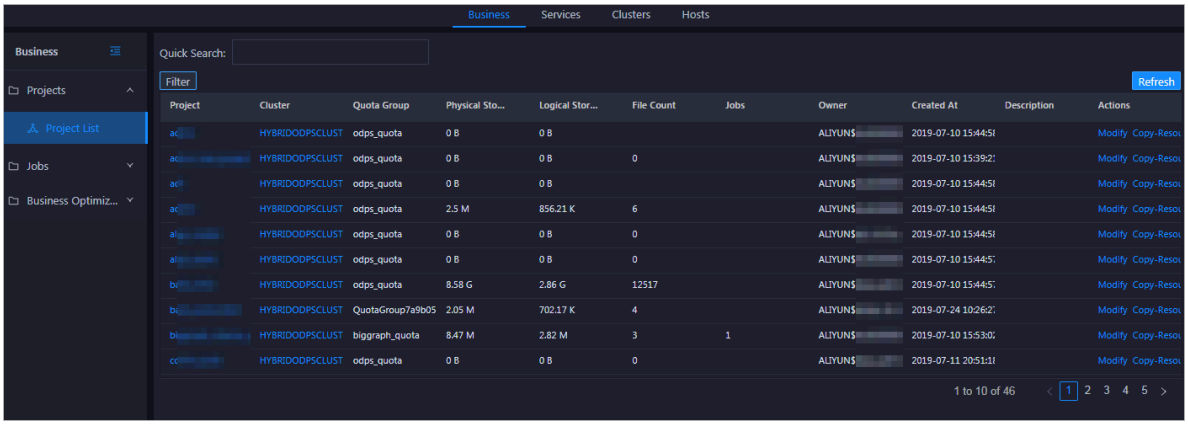

1.5.2.2 Project management

1.5.2.2.1 Project list

The Project List page displays all projects and project details in the MaxCompute cluster. You can filter, query, and sort projects. You can also modify the quota group

of a project. If zone-disaster recovery is enabled, you can set resource replication parameters and determine whether to enable resource replication for a project.

Entry

On the Business page, choose Projects > Project List in the left-side navigation pane to view projects in the cluster.

On the Project List page, you can view detailed information about all projects in the cluster. For example, you can view the name, cluster, storage, file quantity, running job quantity, owner, creation time, quota group, and description of a project.

Facilitate information retrieval

You can filter and query projects. You can also customize columns or sort projects based a specified column. For more information, see *Common operations*.

View project details

On the Project List page, you can click the name of a project to view its detailed information, including the project overview, jobs, storage, configurations, quota groups, tunnels, resource analysis, and cross-cluster replication. For more information, see *MaxCompute workbench*. You can also grant the permission of accessing the metadata warehouse to the project or encrypt data of the project. For more information, see *Project authorization for accessing the metadata warehouse* and *Storage encryption*.

Modify a project

You can modify the quota group and default cluster of a project.

1. On the Project List page, find the project to be modified and click Modify in the Actions column. In the Modify Project dialog box that appears, set relevant parameters.

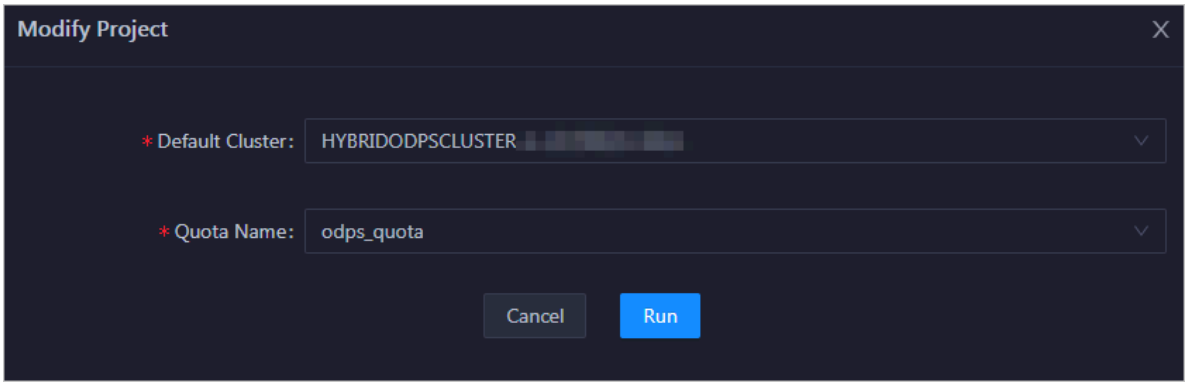

The parameters are described as follows:

- Default Cluster: the default cluster of the project. If the project belongs to multiple clusters, you can select a cluster from the drop-down list to serve as the default cluster.
- Quota Name: the quota group to which the project belongs. To change the quota group, select the specified quota group from the drop-down list.

2. Click Run. A message appears, indicating that the action has been submitted.

Configure resource replication

The resource replication feature can be configured only in zone-disaster recovery scenarios. In other scenarios, you can only view the settings. In zone-disaster recovery scenarios, you can determine whether to enable the resource replication feature for a project in the primary cluster. If the resource replication feature is enabled for a project, you can set data synchronization rules for the project to regularly synchronize data such as data tables to the standby cluster.

1. On the Project List page, find a specified project and click Copy-Resource in the Actions column. In the Copy Resource dialog box that appears, set relevant parameters.

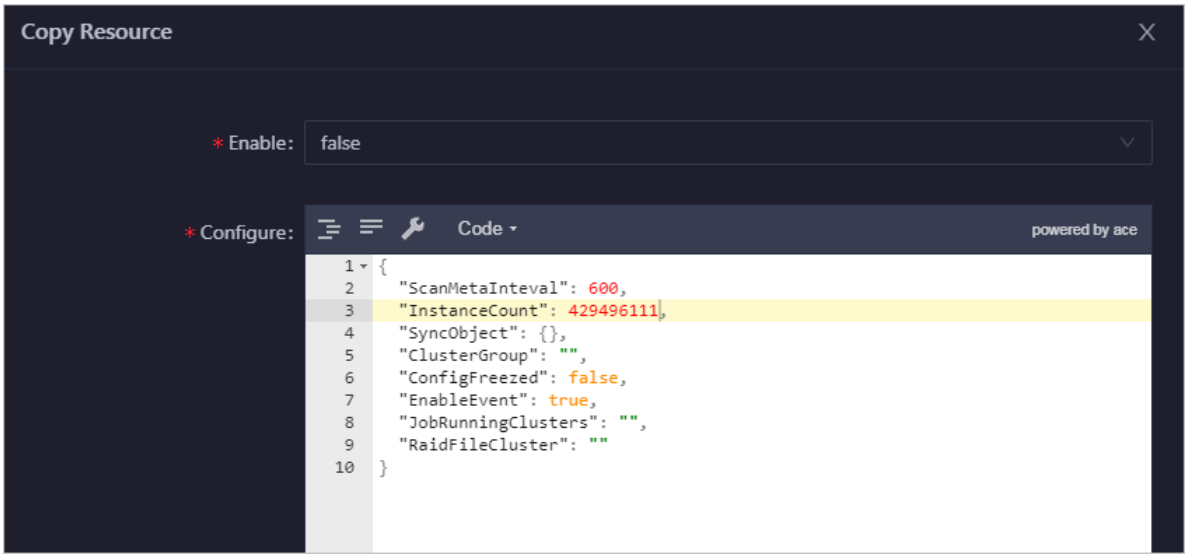

The parameters are described as follows:

- Enable: specifies whether to enable the resource replication feature. A value of true indicates that the resource replication feature is enabled. A value of false indicates that the resource replication feature is disabled.
- Configure: the data synchronization rules of a project. Generally, the default settings are used. If you need to modify the settings, consult second-line O&M engineers.

2. After modifying the code in Configure, click Compare Versions to view the highlighted differences between the code of the current version and that of the earlier version.
3. Click Run. A message appears, indicating that the action has been submitted.

1.5.2.2.2 Storage encryption

On the Encryption at Rest page, you can specify whether to encrypt the data stored in MaxCompute projects.

Prerequisites

If MaxCompute V3.8.0 or later is newly deployed, it supports storage encryption by default. If MaxCompute is upgraded from an earlier version to V3.8.0 or later, it does not support storage encryption by default. If you want to enable storage encryption, complete the configuration in the MaxCompute cluster.

Context

After storage encryption is enabled for a project, the feature cannot be disabled. After storage encryption is enabled, all data written to the project will be encrypted , but historical data cannot be automatically encrypted. To encrypt historical data, you can create rules and configure tasks.

Understand the concepts of rules and tasks in Apsara Bigdata Manager (ABM) before encrypting historical data for a MaxCompute project. You can create a rule to specify the time period of historical data to be encrypted in the specified project . After the rule is created, the system obtains the data in the specified time period every day after the data is exported from the metadata warehouse. You can create only one rule every day. If multiple rules are created on a single day, the new rule overwrites the old rule and takes effect. Each rule takes effect only once. You can create a key rotate task to encrypt the selected historical data.

Procedure

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click MaxCompute.
3. 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.
4. On the Project List page, click the name of the target project to go to the project details page.
5. On the project details page, click the Encryption at Rest tab. The Encrypt tab appears.
6. Enable storage encryption.

   After storage encryption is enabled, all data written to the project will be encrypted.

   a. On the Encrypt tab, click Modify in the Actions column. In the Configure Encrypted Storage dialog box that appears, set the region and project parameters.

      Encryption Algorithm: You can select AESCTR or RC4 from the drop-down list.

   b. Click Run.

      After storage encryption is enabled, the switch in the Encrypted Storage column is turned on.

7. To encrypt historical data or encrypted data, follow these steps:

   a. Create a rule.

      On the Create Rule tab, click OK in the Actions column of a time period in the Create Rule section. In the Create Rule dialog box that appears, click Run and then click OK. The new rule appears in the rule list.

      The available time periods include Last Three Months, Last Six Months, Three Months Ago, Six Months Ago, and All.

   b. Create a key rotate task.

      On the Configure Task tab, click Add a key rotate task. In the Edit Key Rotate Task dialog box that appears, set relevant parameters and click Run.

Parameter	Description
Region	The region where the cluster of the project of which data is to be encrypted resides. Select a region from the drop-down list.
Project Name	The name of the project of which data is to be encrypted.
Start Timestamp	The start time of the task.
Ended At	The end time of the task.
Priority	The priority of the task. A smaller value indicates a higher priority.
Enabled	Specifies whether the task is enabled.
Bandwidth Limit	Specifies whether to limit the concurrency of key rotate tasks for the project. · Yes: indicates that key rotate tasks cannot be run simultaneously. · No: indicates that key rotate tasks can be run simultaneously.
Maximum Concurrent Tasks	The maximum number of key rotate tasks that can be run at the same time in the cluster of the selected project. This parameter is valid only when Bandwidth Limit is set to No.

Parameter	Description
Maximum Number of Running Jobs	The maximum number of jobs that can be run at the same time in the cluster of the selected project. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the cluster of the selected project, not only the key rotate tasks.
Merge Parameters	`{ "odps.merge.cross.paths": "true", "odps.idata.useragent": "odps encrypt key rotate via force mergeTask", "odps.merge.max.filenumber.per.job": "10000000", "odps.merge.max.filenumber.per.instance": "10000", "odps.merge.failure.handling": "any", "odps.merge.maintain.order.flag": "true", "odps.merge.smallfile.filesize.threshold": "4096", "odps.merge.quickmerge.flag": "true", "odps.merge.maxmerged.filesize.threshold": "4096", "odps.merge.force.rewrite": "true", "odps.merge.restructure.action": "hardlink" }`

8. Optional: View the history of encrypting data in the project.

   On the Historical Queries tab, select a date from the Date drop-down list. Then, you can view the information about the storage encryption on the specified date.

1.5.2.2.3 Project authorization for accessing the metadata warehouse

On the Authorize Package for Metadata Repository page, you can grant the permission of accessing the metadata warehouse to projects and project members.

Prerequisites

- If MaxCompute V3.8.1 or later is newly deployed, the package of the metadata warehouse is installed by default. Then, you can directly grant the permission of accessing the metadata warehouse in Apsara Bigdata Manager (ABM). If you upgrade MaxCompute from an earlier version to V3.8.1 or later, the package of the metadata warehouse is not installed by default. In this case, you need to manually install the package of the metadata warehouse in the MaxCompute cluster before granting the permission of accessing the metadata warehouse.
- A project is created in DataWorks.

Context

To allow a project to access the resources of the metadata warehouse, grant the corresponding permission to the project and install the package to the project in ABM. When you install the package, ABM retrieves the authentication information , such as the AccessKey, of the project from DataWorks. If the project is created in MaxCompute, an error message is returned during installation.

Procedure

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **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.**
4. **On the Project List page, click the name of the target project to go to the project details page.**
5. **On the project details page, click the Configuration tab. On the page that appears, click the Authorize Package for Metadata Repository tab.**
6. **Click Authorize in the Actions column. In the Authorize Package dialog box that appears, click Run. A message appears, indicating that the permission of accessing the metadata warehouse is granted to the project.**
7. **After the permission is granted, click Install in the Actions column. In the Install Package dialog box that appears, click Run. A message appears, indicating that the package of the metadata warehouse is installed.**

   After the package is installed, the switch in the Authorized column is turned on.

1.5.2.2.4 Disaster recovery

When the primary MaxCompute cluster fails, you can quickly switch services form the primary cluster to the standby cluster in the Apsara Bigdata Manager (ABM) console to restore services. This topic describes the elements on the Disaster Recovery page and the prerequisites and procedure for switchover. Only zone-disaster recovery is supported.

Entry

> **Note:**
> **The disaster recovery feature is available only when zone-disaster recovery is enabled.**

On the Business page, choose Projects > Disaster Recovery in the left-side navigation pane to view information about the primary and standby MaxCompute clusters.

The Disaster Recovery page consists of the following elements:

- **Primary and Standby sections: The Primary section displays the name of the primary cluster and the number of projects with the primary cluster as the default cluster. The Standby section displays the name of the standby cluster and the number of projects with the standby cluster as the default cluster.**

  

  Note:
  The total number of projects in the primary cluster is the same as that in the standby cluster.
- **Resource Synchronization Status: specifies whether to run the scheduled task to synchronize resources between the primary and standby MaxCompute clusters. When you turn on this switch, resources are scheduled between the primary and standby clusters every 30 minutes.**
- **View Resource Synchronization History: allows you to view the execution history of the scheduled resource synchronization task.**

Prerequisites for switchover

- **A zone-disaster recovery environment is built.**
- **Your ABM account is granted the required permissions to perform O&M operations on MaxCompute and can be used to log on to the ABM console.**
- **The VIP address of the current ABM cluster has been switched to the standby ABM cluster. For more information, see** *Switch the VIP address the current ABM cluster to the standby ABM cluster***.**
- **You have disabled the scheduled task for synchronizing resources between the primary and standby MaxCompute clusters. For more information, see** *Enable or disable the resource synchronization between primary and standby MaxCompute clusters***.**
- **The Business Continuity Management Center (BCMC) switchover of MaxCompute has been completed. The services on which MaxCompute depends, including AAS, Table Store, and MiniRDS, are running properly.**

Enable or disable the resource synchronization between primary and standby MaxCompute clusters

When the resource synchronization feature is enabled, the scheduled resource synchronization task is run to synchronize resources, such as a compiled JAR package, between the primary and standby MaxCompute clusters every 30 minutes. You need to keep the scheduled resource synchronization task disabled until the switchover between the primary and standby clusters is completed.

1. **On the Business page, choose Projects > Disaster Recovery in the left-side navigation pane to view information about the primary and standby MaxCompute clusters.**

   If the Resource Synchronization Status switch is turned on, the scheduled resource synchronization task is enabled. If the switch is turned off, the task is disabled.
2. **Turn on or off the Resource Synchronization Status switch to enable or disable the scheduled resource synchronization task between the primary and standby clusters.**

Switch the VIP address the current ABM cluster to the standby ABM cluster

Before the switchover between the primary and standby MaxCompute clusters, you can execute the Change Bcc Dns-Vip Relation For Disaster Recovery scheme in ABM to replace the VIP address of the standby ABM cluster with that of the current ABM cluster.

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click Management in the upper-right corner. The Management appears.**
4. **Click Jobs in the left-side navigation pane, and click Job Management on the right side to go to the Schemes page.**
5. **In the scheme list, find the Change Bcc Dns-Vip Relation For Disaster Recovery scheme, and click Run in the Actions column. On the page that appears, set the following two parameters in Target Group:**

   Set the NowBccApiOneIp parameter to the IP address of any Docker container in a path of the current ABM cluster, for example, bcc > bcc-api > controller# >

#Docker#xx.xx.xx.xx. Set the NewBccApiOneIp parameter to the IP address of any Docker container in the same path of the standby ABM cluster.

6. **Click Run in the upper-right corner and confirm the risks of running the job.**
7. **Click Confirm. The job running page appears.**
8. **Click Start at the top of the page.**

Start switchover

After all the prerequisites for switchover are met, you can start MaxCompute switchover.

1. **On the Business page, choose Projects > Disaster Recovery in the left-side navigation pane to view information about the primary and standby MaxCompute clusters.**
2. **Click Start Switchover in the upper-right corner. A dialog box appears, asking you to confirm whether the VIP address of the standby ABM cluster has been replaced with that of the current ABM cluster.**

   If the VIP address of the standby ABM cluster has not been replaced with that of the current ABM cluster, click No and then replace the VIP address of the standby ABM cluster with that of the current ABM cluster. For more information, see *Switch the VIP address the current ABM cluster to the standby ABM cluster*.
3. **Click Yes. The Stop Resource Replication page appears.**

   Note:
   In this step, the scheduled resource synchronization task is automatically disabled.
4. **After resource replication is disabled, click Next Step. The Switch Control Cluster page appears.**

   In this step, the services are automatically switched from the primary cluster to the standby cluster, which takes about 30 seconds. You can determine whether the switchover is successful based on the values of the Current Primary Cluster

and Current Standby Cluster parameters on the page. When the primary and standby clusters are switched, the switchover is complete.

After the switchover, you need to perform the following operations:

a. Click Restart Standby Cluster. It takes about 20 seconds to restart the standby cluster. When the standby cluster is restarted, the value of the MaxCompute Cluster Status parameter changes from Abnormal to Normal.
b. Click Restart Frontend Server. It takes about 20 seconds to restart the front-end server. When the front-end server is restarted, a success message appears.
c. Click Test adminTask to check whether the MaxCompute service is normal. If the test is passed, the clusters are switched. The Next Step button becomes operable, and the Switching... message disappears.

5. Click Next Step. The Switch Computing Cluster page appears.

   In this step, the default computing cluster of the projects in the primary cluster is changed to the standby cluster, and that of the projects in the standby cluster is changed to the primary cluster. Each project has a switchover progress bar. If the progress bar of a project is highlighted, the switchover is complete.

   > **Note:**
   > If the computing cluster of a project fails to be switched, you can contact O&M engineers to locate the cause of the exception. If the project can be fixed, fix it and click Retry to continue the switchover. If the project is damaged or does not need to change the computing cluster, you can click Next Step after confirming that other projects have been switched.

6. Click Next Step. The Switch Replication Service to Standby Service page appears.
7. After the switchover is completed, click Next Step. The Collect Statistics about Unsynchronized Data page appears.

   This step takes some time, depending on the data volume. Wait until the step is completed. After the collection is completed, the system lists all projects with unsynchronized data. You can check the data that has not been synchronized.

   You must select the projects with unsynchronized data, and click Download Unsynchronized Data of Selected Projects to download the data to a local device so that you can manually fill in the missing data later based on the statistics. Only

after the unsynchronized data is downloaded does the Next Step button become operable.

> **Note:**
> If the unsynchronized data is abnormal, you can click Recollect Unsynchronized Data.

8. Click Next Step. The Repair Metadata page appears.

   In this step, the data in the primary and standby clusters becomes the same. Select all projects, click Repair Metadata of Selected Projects, and then wait for results.

   - If some projects fail to be fixed, click Download Last Execution Log and send the logs to O&M engineers to analyze the cause of the exception. After the exception is resolved, you can fix the projects again.
   - If you do not need to fix all projects, click Next Step after the necessary projects are fixed.

9. After the metadata is fixed, click Next Step. The Manually Fill in Missing Data page appears.

   In this step, you need to log on to the DataWorks console, and manually fill in the missing data according to the unsynchronized data downloaded in the Collect Statistics about Unsynchronized Data step. After filling in the missing data, select all projects and click Confirm Data Repair Complete. Then, the Next Step button becomes operable.

10. Click Next Step. The Repair Unsynchronized Resources page appears.

    In this step, it takes some time to count the projects to be fixed, depending on the data volume. Wait until the results appear. If some projects in the standby cluster are inconsistent with those in the primary cluster, you need to fill in the missing data manually. Otherwise, proceed to the next step.

11. After the unsynchronized resources are fixed, click Complete and Next. The Enable Resource Replication page appears.

    In this step, the scheduled resource synchronization task is automatically started.

12. After enabling resource replication, click Next Step. The Complete Wizard page appears.

13. Click Back in the upper-left corner. The primary and standby clusters have been switched.

1.5.2.3 Job management

1.5.2.3.1 Job snapshots

The Job Snapshots page allows you to manage the tasks created in MaxCompute and the merge tasks created in Apsara Bigdata Manager (ABM). You can also view the job details by using Logview, terminate a job, and collect job logs on this page.

View job snapshots

You can view job snapshots by day over the last week. Detailed information about a job snapshot includes the job ID, project, quota group, submitter, running duration , minimum and maximum CPU usage, minimum and maximum memory usage, DataWorks node, running status, start time, priority, and type. You can also view the operational logs of a job to locate job failures.

The Job Snapshots page supports multiple operations to facilitate information retrieval. For example, you can filter and sort job snapshots. For more information, see *Common operations*.

1. On the Business page, choose Jobs > Job Snapshots in the left-side navigation pane. The Job Snapshots page appears.

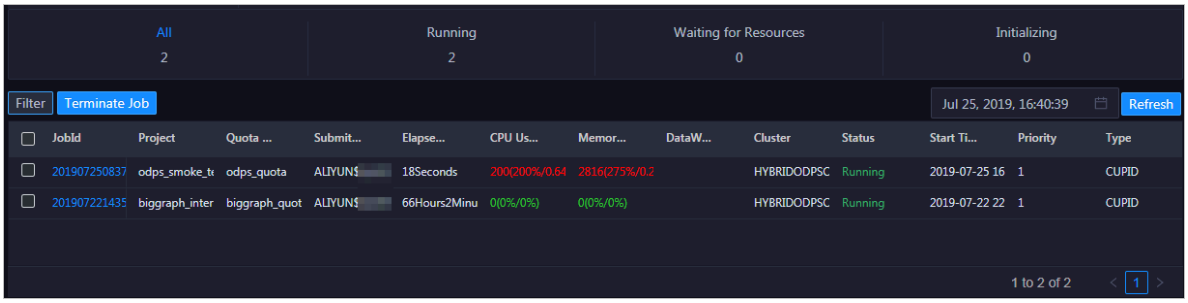

2. In the upper-right corner of the job snapshot list, select the date and time to view job snapshots by day over the last week.
3. Click All, Running, Waiting for Resources or Initializing to view job snapshots in the corresponding status on the specified date.
4. Click the job ID of a job snapshot, and then click DetailLogview. A dialog box appears, containing a link to Logview.

5. Click click here to access the Logview page and view detailed information about the job.

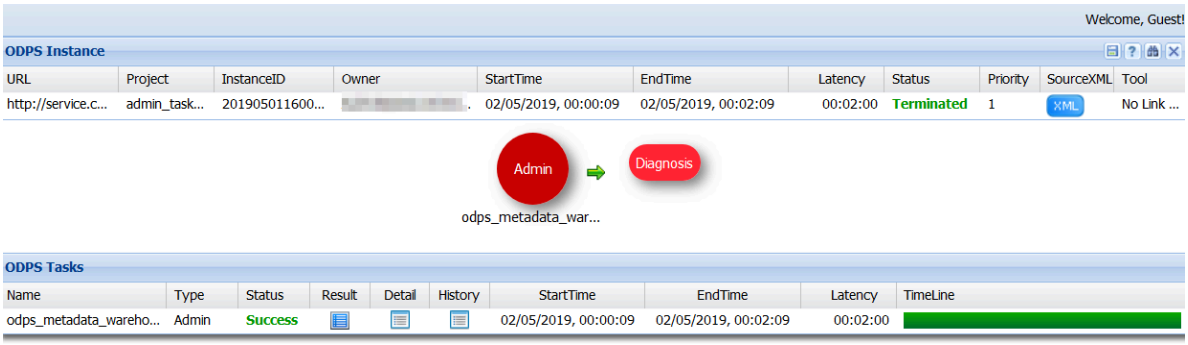

Terminate a job

1. On the Business page, choose Jobs > Job Snapshots in the left-side navigation pane. The Job Snapshots page appears.

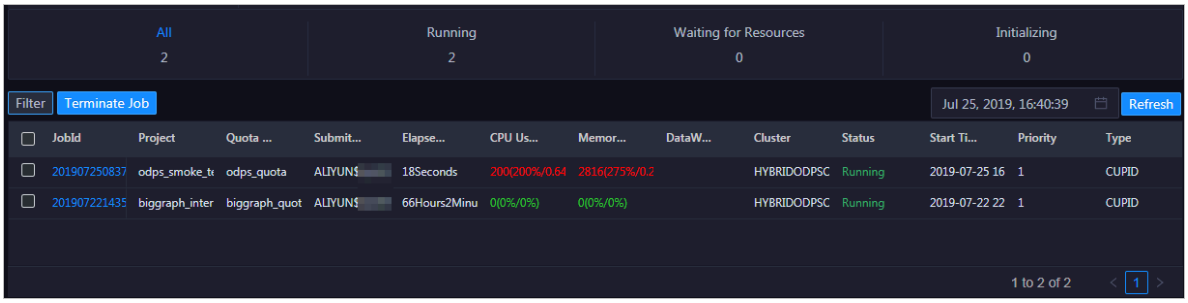

2. Select one or more jobs, and then click Terminate Job. In the dialog box that appears, view the information about the job or jobs to be terminated.

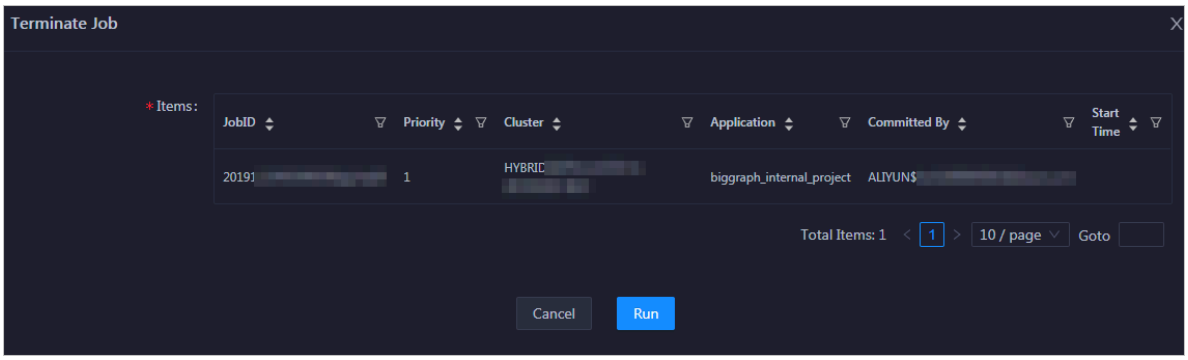

3. Click Run. A message appears, indicating the running result.

Collect job logs

When an exception occurs during job running, you can locate and analyze the issue by collecting job logs.

1. On the Business page, choose Jobs > Job Snapshots in the left-side navigation pane. The Job Snapshots page appears.

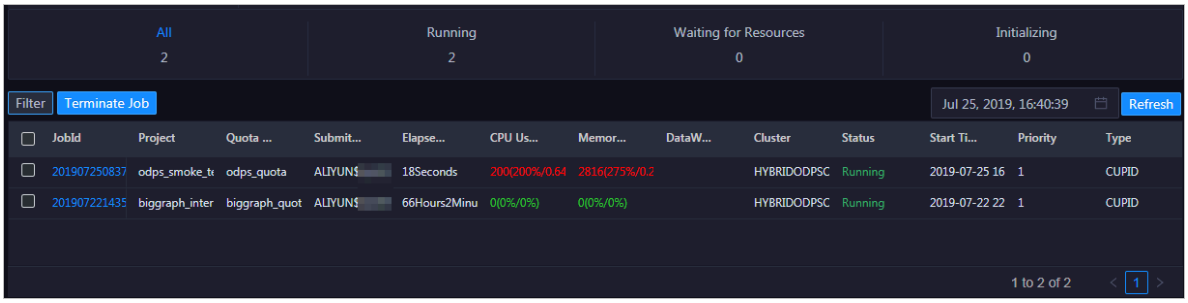

2. In the upper-left corner of the Job Snapshots page, choose Actions > Collect Job Logs.
3. In the Collect Job Logs dialog box that appears, set the parameters.

   The following table describes the required parameters.

Parameter	Description
Target Service	The target service from which you want to collect job logs. Select a target service from the drop-down list.
intanceid	Optional. The ID of the job instance.
requestid	Optional. The ID of the request returned when the job execution fails. If the value you specify is not a request ID, job logs that contain the corresponding value will be collected.
Time Period	The time period in which you want to collect job logs.
Time Interval	Optional. The time interval for collecting job logs. Unit: hours.
Degree of Concurrency	The maximum number of nodes from which you can collect job logs at the same time.

4. Click Run to start job log collection.
5. View the execution status and progress of job log collection.

   In the upper-left corner of the Job Snapshots page, click Actions, and then click Execution History next to Collect Job Logs to view the execution status and history of job log collection.

   In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful. You can click Details in the Details column of a task in the Running state to view the execution progress.

6. **View the path for storing the job logs.**

   You can click Details in the Details column of a task in the Success state to view the execution details. In the Steps section, click the name of the node to show detailed information, and then click View Details in the Actions column to view the path for storing the job logs.

1.5.2.4 Business optimization

1.5.2.4.1 File merging

Excessive small files in a MaxCompute cluster occupy a lot of memory resources. Apsara Bigdata Manager (ABM) allows you to merge small files in clusters and projects to release memory resources occupied by excessive small files.

Create a merge task for a cluster

When excessive small files exist in most projects of a MaxCompute cluster, you can create a merge task to merge the small files in the cluster in a unified manner.

1. **On the Business page, choose Business Optimization > File Merging in the left-side navigation pane. The Merge Tasks tab appears.**

2. In the Merge Tasks for Clusters section, click Create Merge Task, and set relevant parameters in the dialog box that appears.

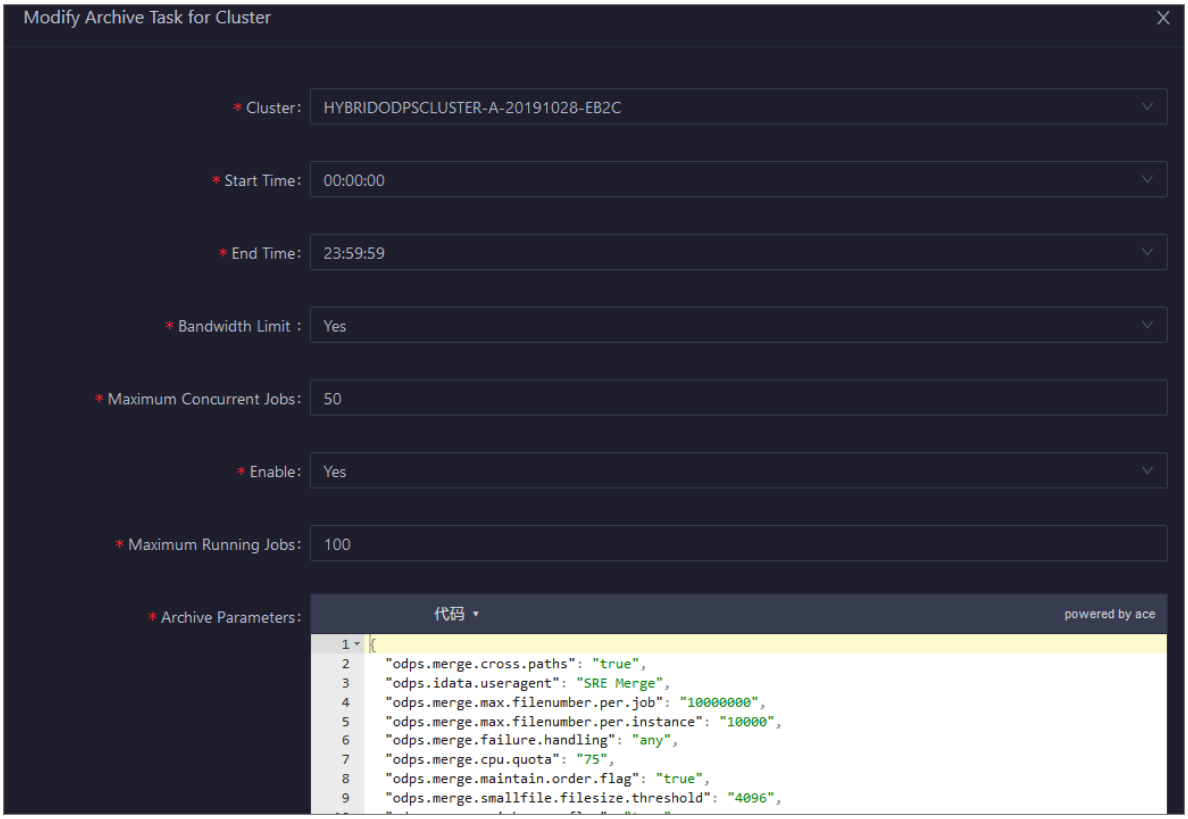

The following table describes the required parameters.

Parameter	Description
Cluster	The cluster in which you want to run the merge task. Select a cluster from the drop-down list.
Start Time	The start time of the merge task.
End Time	The end time of the merge task.
Bandwidth Limit	Specifies whether to limit the concurrency of merge tasks for the cluster. · Yes: indicates that merge tasks cannot be run simultaneously. · No: indicates that merge tasks can be run simultaneously.
Maximum Concurrent Tasks	The maximum number of merge tasks that can be run at the same time in the selected cluster. This parameter is valid only when Bandwidth Limit is set to No.
Enabled	Specifies whether the merge task is enabled.

Parameter	Description
Merge Parameters	The parameter configuration for the merge task. You can use the following default configuration: `{ "odps.idata.useragent": "SRE Merge", "odps.merge.cpu.quota": "75", "odps.merge.quickmerge.flag": "true", "odps.merge.cross.paths": "true", "odps.merge.smallfile.filesize.threshold": "4096", "odps.merge.maxmerged.filesize.threshold": "4096", "odps.merge.max.filenumber.per.instance": "10000", "odps.merge.max.filenumber.per.job": "10000000", "odps.merge.maintain.order.flag": "true", "odps.merge.failure.handling": "any" }`
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the selected cluster. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the selected cluster, not only the merge tasks.

3. Click Compare Versions to view the differences between the modified values and the original parameter values.
4. Click Run. A message appears, indicating that the action has been submitted.

   After the merge task is created, it appears in the list of merge tasks for clusters.

Create a merge task for a project

When excessive small files exist in only a few projects of a MaxCompute cluster, you can create a merge task to merge the small files in each project.

1. On the Business page, choose Business Optimization > File Merging in the left-side navigation pane. The Merge Tasks tab appears.

2. In the Merge Tasks for Projects section, click Create Merge Task, and set relevant parameters in the dialog box that appears.

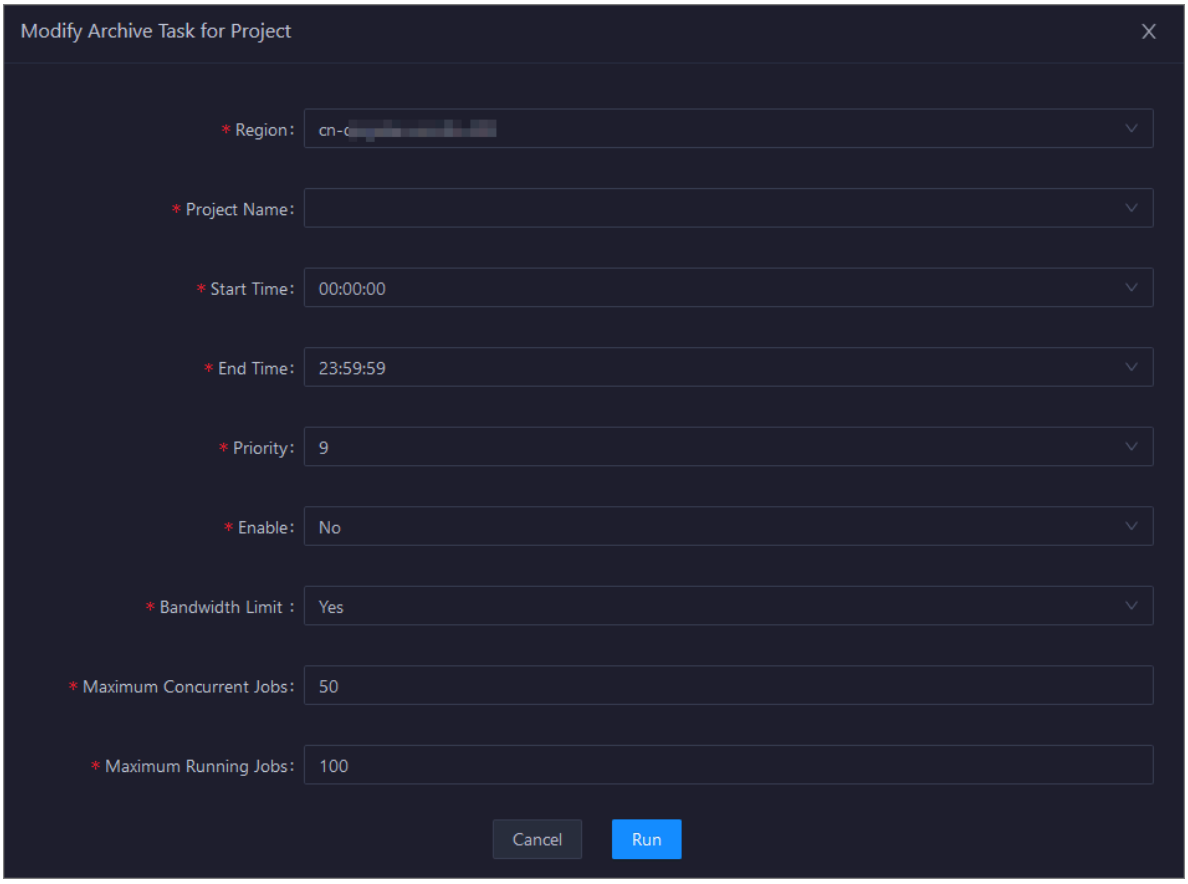

The following table describes the required parameters.

Parameter	Description
Region	The region where the cluster of the selected project resides. Select a region from the drop-down list.
Project Name	The name of the project in which you want to run the merge task. Select a project from the drop-down list.
Start Time	The start time of the merge task.
Priority	The priority of the merge task. A smaller value indicates a higher priority.
End Time	The end time of the merge task.
Enabled	Specifies whether the merge task is enabled.

Parameter	Description
Bandwidth Limit	Specifies whether to limit the concurrency of merge tasks for the project. · Yes: indicates that merge tasks cannot be run simultaneously. · No: indicates that merge tasks can be run simultaneously.
Maximum Concurrent Tasks	The maximum number of merge tasks that can be run at the same time in the cluster of the selected project. This parameter is valid only when Bandwidth Limit is set to No.
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the cluster of the selected project. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the cluster of the selected project, not only the merge tasks.

3. Click Compare Versions to view the differences between the modified values and the original parameter values.
4. Click Run. A message appears, indicating that the action has been submitted.

   After the merge task is created, it appears in the list of merge tasks for projects.

View merge task statistics

On the Business page, choose Business Optimization > File Merging in the left-side navigation pane, and then click the Historical Statistics tab to view the historical statistics of merge tasks for clusters and projects.

```
Merge task chart
```

The trend chart for merge tasks displays statistics on the execution of all merge tasks for each day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks, number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

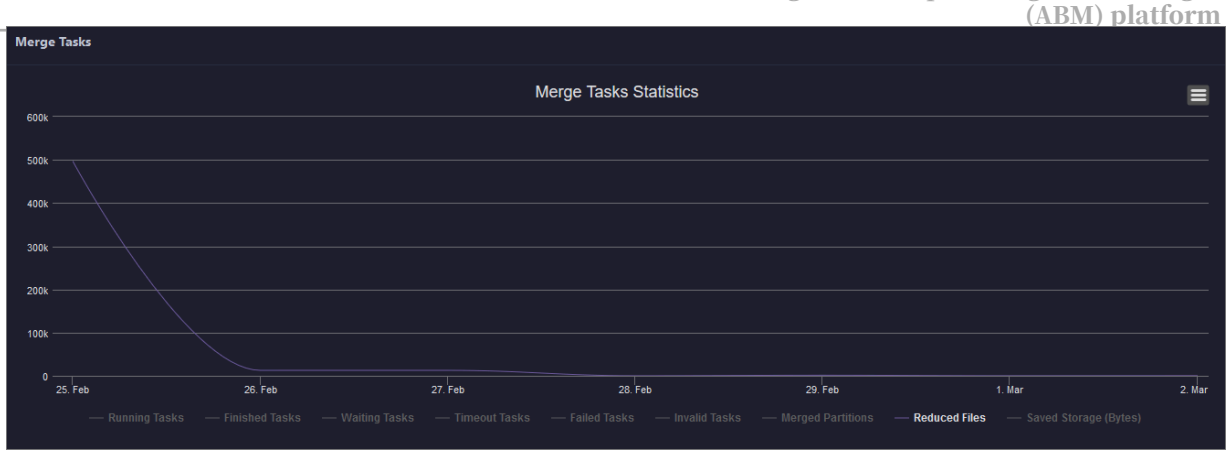

`Cluster statistics and project statistics`

The two tables list statistics on the execution of merge tasks for clusters and projects for a specified day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks , number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

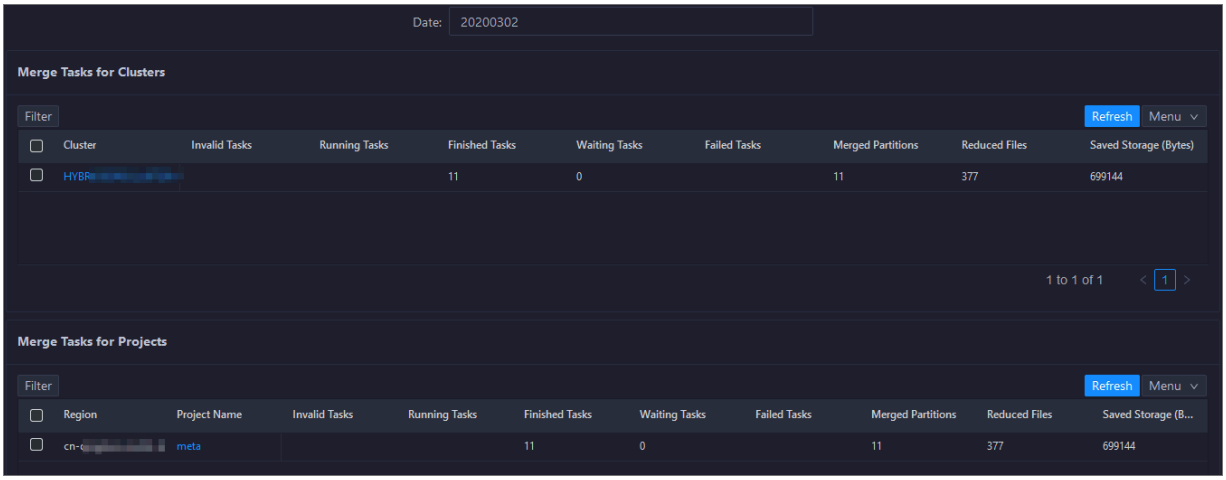

1.5.2.4.2 File archiving

In the Apsara Bigdata Manager (ABM) console, you can create archive tasks to compress idle files in MaxCompute clusters and projects to save storage space for the clusters.

Definition

ABM sorts the tables or partitions created more than 90 days ago in a cluster by storage space, and then compresses the first 100,000 tables or partitions.

Create an archive task for a cluster

When excessive idle files exist in most projects of a MaxCompute cluster, you can create an archive task to compress the idle files in the cluster in a unified manner.

1. **On the Business page, choose Business Optimization > File Archiving in the left-side navigation pane. The Archive Tasks tab appears.**
2. **In the Archive Tasks for Clusters section, click Create Archive Task, and set relevant parameters in the dialog box that appears.**

   The following table describes the required parameters.

Parameter	Description
Cluster	The cluster in which you want to run the archive task. Select a cluster from the drop-down list.
Start Time	The start time of the archive task.
End Time	The end time of the archive task.
Bandwidth Limit	Specifies whether to limit the concurrency of archive tasks for the cluster. · Yes: indicates that archive tasks cannot be run simultaneously. · No: indicates that archive tasks can be run simultaneously.
Maximum Concurrent Jobs	The maximum number of archive tasks that can be run at the same time in the selected cluster. This parameter is valid only when Bandwidth Limit is set to No.
Enabled	Specifies whether the archive task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the selected cluster. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the selected cluster, not only the archive tasks.

Parameter	Description
Archive Parameters	The parameter configuration for the archive task. You can use the following default configuration: `{ "odps.idata.useragent": "SRE Archive", "odps.oversold.resources.ratio": "100", "odps.merge.quickmerge.flag": "true", "odps.merge.cross.paths": "true", "odps.merge.smallfile.filesize.threshold": "4096", "odps.merge.maxmerged.filesize.threshold": "4096", "odps.merge.max.filenumber.per.instance": "10000", "odps.merge.max.filenumber.per.job": "10000000", "odps.merge.maintain.order.flag": "true", "odps.sql.hive.compatible": "true", "odps.merge.compression.strategy": "normal", "odps.compression.strategy.normal.compressor": "zstd", "odps.merge.failure.handling": "any", "odps.merge.archive.flag": "true" }`

3. Click Compare Versions to view the differences between the modified values and the original parameter values.
4. Click Run. A message appears, indicating that the action has been submitted.

   After the archive task is created, it appears in the list of archive tasks for clusters.

Create an archive task for a project

When excessive idle files exist in only a few projects of a MaxCompute cluster, you can create an archive task to compress the idle files in each project.

> **Note:**
> If the tables or partitions of a project are not ranked top 100,000 in the cluster of the project, the archive task cannot compress the idle files in the project.

1. On the Business page, choose Business Optimization > File Archiving in the left-side navigation pane. The Archive Tasks tab appears.

2. In the Archive Tasks for Projects section, click Create Archive Task, and set relevant parameters in the dialog box that appears.

   The following table describes the required parameters.

Parameter	Description
Region	The region where the cluster of the selected project resides. Select a region from the drop-down list.
Project Name	The name of the project in which you want to run the archive task. Select a project from the drop-down list.
Start Time	The start time of the archive task.
Priority	The priority of the archive task. A smaller value indicates a higher priority.
End Time	The end time of the archive task.
Bandwidth Limit	Specifies whether to limit the concurrency of archive tasks for the project. · Yes: indicates that archive tasks cannot be run simultaneously. · No: indicates that archive tasks can be run simultaneously.
Maximum Concurrent Jobs	The maximum number of archive tasks that can be run at the same time in the cluster of the selected project. This parameter is valid only when Bandwidth Limit is set to No.
Enabled	Specifies whether the archive task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the cluster of the selected project. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the cluster of the selected project, not only the archive tasks.

3. Click Compare Versions to view the differences between the modified values and the original parameter values.
4. Click Run. A message appears, indicating that the action has been submitted.

   After the archive task is created, it appears in the list of archive tasks for projects .

View archive task statistics

On the Business page, choose Business Optimization > File Archiving in the left-side navigation pane, and then click the Historical Statistics tab to view the historical statistics of archive tasks for clusters and projects.

`Archive task chart`

The trend chart for archive tasks displays statistics on the execution of all archive tasks for each day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks, number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

`Cluster statistics and project statistics`

The two tables list statistics on the execution of archive tasks for clusters and projects for a specified day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks , number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

1.5.2.4.3 Resource analysis

Apsara Bigdata Manager (ABM) allows you to analyze the resources for MaxCompute clusters from multiple dimensions so that you can better understand the data storage in MaxCompute. The dimensions include tables, tasks, execution time, start time, and engines.

Tables

From this dimension, you can view the detailed information about all tables in each project, including the number of partitions, storage space, number of Apsara Distributed File System files, ranking of the number of partitions, and ranking of the number of Apsara Distributed File System files. You can also sort tables based on the number of partitions, storage space, or number of Apsara Distributed File System files.

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane. The Tables tab appears.

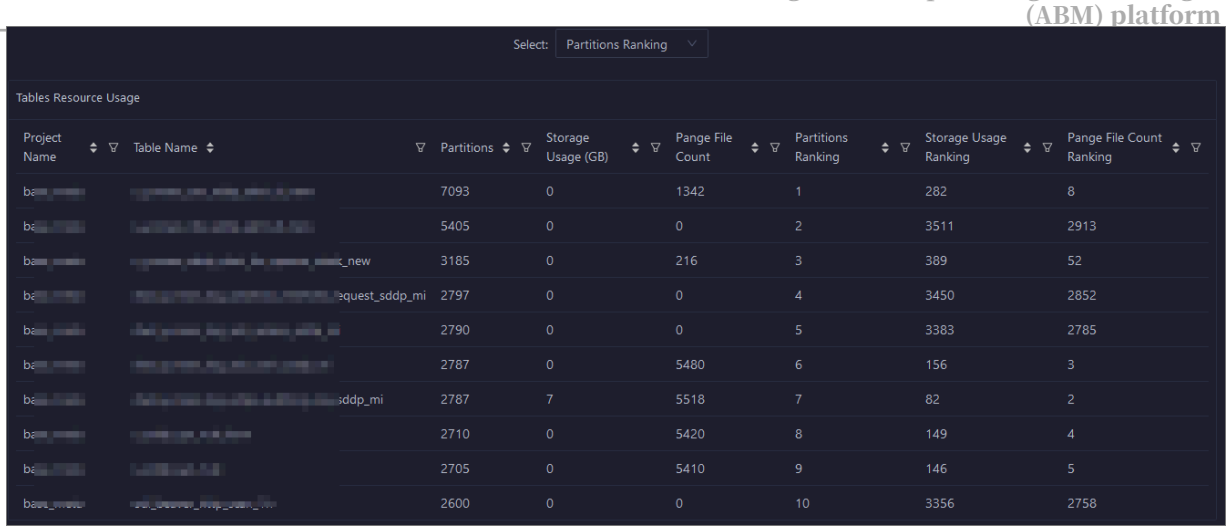

Projects

From this dimension, you can view the detailed information about storage for each project, including the number of Apsara Distributed File System files, storage space , CU usage, memory usage, number of tasks, number of tables, idle storage, and daily and weekly increases of these items.

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Projects tab. The Projects tab appears.

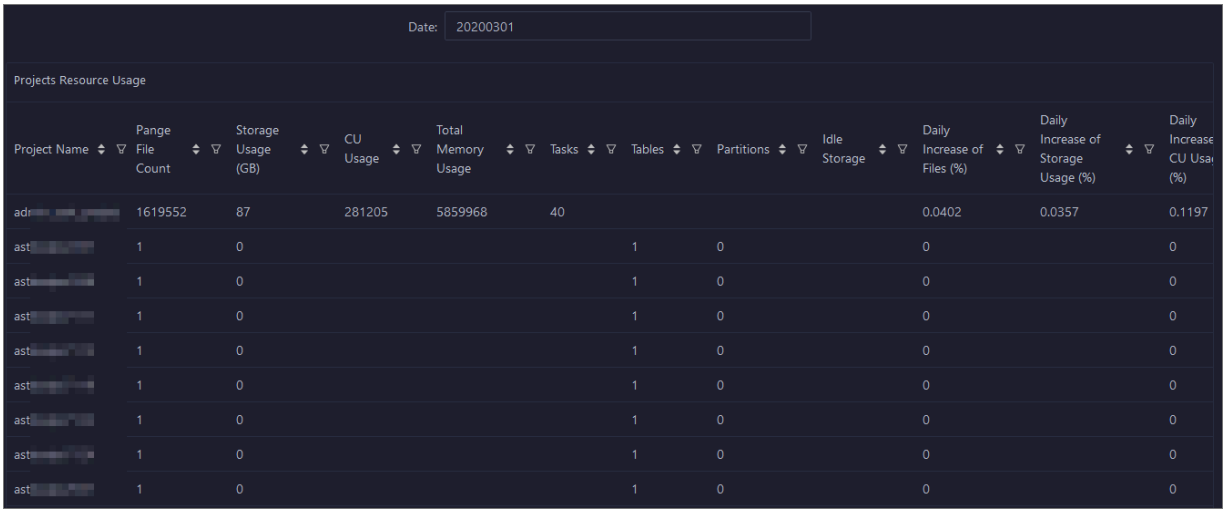

Tasks

From this dimension, you can view the detailed information about tasks in each project, including the ID of the task instance, running status, CU usage, start time, end time, execution time, ranking of CU usage, and SQL statements.

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Tasks tab. The Tasks tab appears.

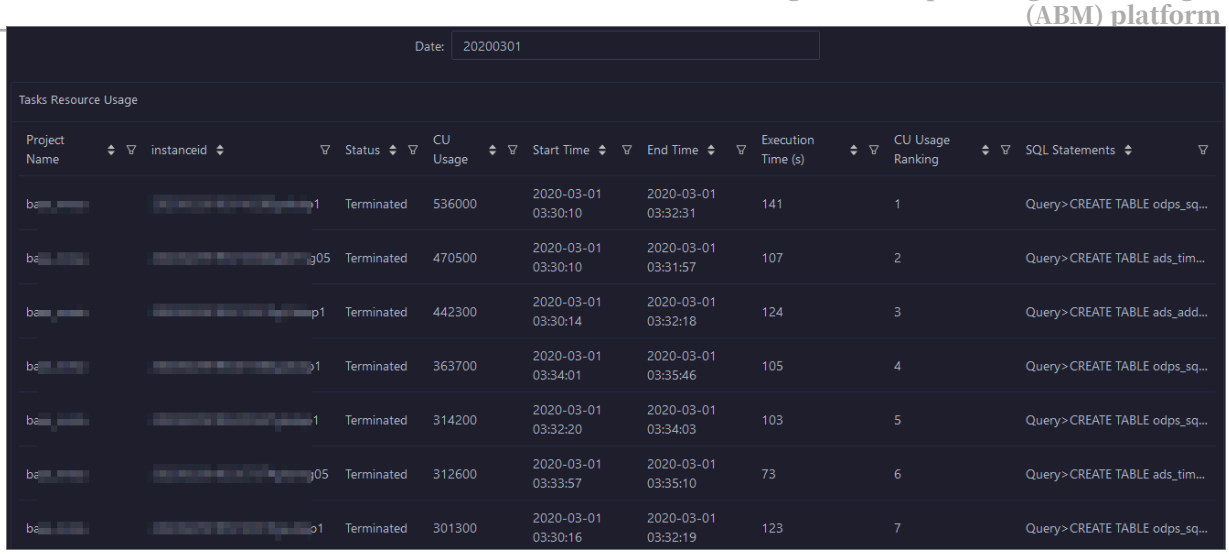

Project Name	instanceid	Status	CU Usage	Start Time	End Time	Execution Time (s)	CU Usage Ranking	SQL Statements
ba	1	Terminated	536000	2020-03-01 03:30:10	2020-03-01 03:32:31	141	1	Query>CREATE TABLE odps_sq...
ba	g05	Terminated	470500	2020-03-01 03:30:10	2020-03-01 03:31:57	107	2	Query>CREATE TABLE ads_tim...
ba	p1	Terminated	442300	2020-03-01 03:30:14	2020-03-01 03:32:18	124	3	Query>CREATE TABLE ads_add...
ba	1	Terminated	363700	2020-03-01 03:34:01	2020-03-01 03:35:46	105	4	Query>CREATE TABLE odps_sq...
ba	1	Terminated	314200	2020-03-01 03:32:20	2020-03-01 03:34:03	103	5	Query>CREATE TABLE odps_sq...
ba	05	Terminated	312600	2020-03-01 03:33:57	2020-03-01 03:35:10	73	6	Query>CREATE TABLE ads_tim...
ba	o1	Terminated	301300	2020-03-01 03:30:16	2020-03-01 03:32:19	123	7	Query>CREATE TABLE odps_sq...

Execution time

From this dimension, you can view the numbers of tasks whose execution time is within 5 minutes, within 15 minutes, within 30 minutes, within 60 minutes, and over 60 minutes respectively in each project. The execution time chart displays the trend lines of the numbers of tasks with different execution time by day in different colors.

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Execution Time tab. The Execution Time tab appears.

Date	Less than 5 Minutes	Less than 15 Minutes	Less than 30 Minutes	Less than 60 Minutes	More than 60 Minutes
20200301	34679	1	0	0	0
20200229	34992	1	0	0	0
20200228	34753	1	0	0	0
20200227	34457	1	0	0	0
20200226	26242	3	0	0	0
20200225	31435	47	3	0	0
20200224	34305	3	0	0	0

Start time

From this dimension, you can view the numbers of tasks started in different time periods for each project. The time interval is set to 30 minutes. The task chart

displays the trend line of the number of tasks started in the specified time period by day.

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Start Time tab. The Start Time tab appears.

Engines

From this dimension, you can view the trend lines of performance statistics of tasks for each project, including CPU usage (cost_cpu), memory usage (cost_mem), execution time (cost_time), input in the unit of bytes (input_bytes), input per CU in the unit of bytes (input_bytes_per_cu), number of input records (input_reco rds), number of input records per CU (input_records_per_cu), output in the unit of bytes (output_bytes), output per CU in the unit of bytes (output_bytes_per_cu), number of output records (output_records), and number of output records per CU (output_records_per_cu).

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Engines tab. The Engines tab appears.

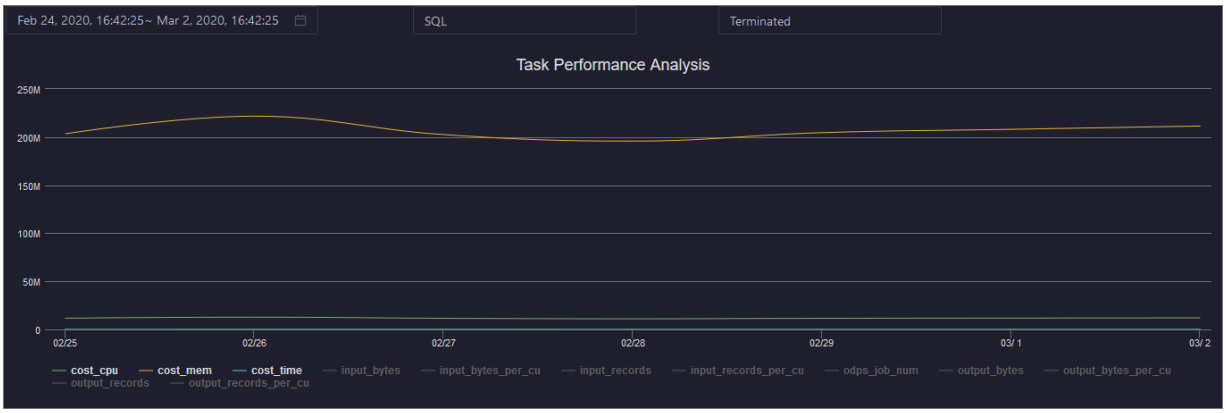

1.5.3 Service O&M

1.5.3.1 Control service O&M

1.5.3.1.1 Control service O&M overview

This topic describes the features of control service O&M and how to access the control service O&M page.

Modules

- **Overview page: displays the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.**
- **Health Status page: displays all checkers for the control service, 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.**
- **Instances page: displays information about the service roles, including the host, service status, CPU application, and memory application of each service role.**
- **Configuration page: provides the access to configuring global computing, cluster -level computing, computing scheduling, and cluster endpoints.**
- **Start Service Role or Stop Service Role action: allows you to enable or disable the control service roles of MaxCompute and view the execution history. You can also locate the failure cause when service role disabling or enabling fails.**
- **Start Admin Console action: allows you to start AdminConsole.**
- **Collect Service Logs action: allows you to collect service logs for the specified time period. This helps you locate the failure cause.**

Entry

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.**

4. **On the Services page, click Control in the left-side navigation pane. The Overview page for the control service appears.**

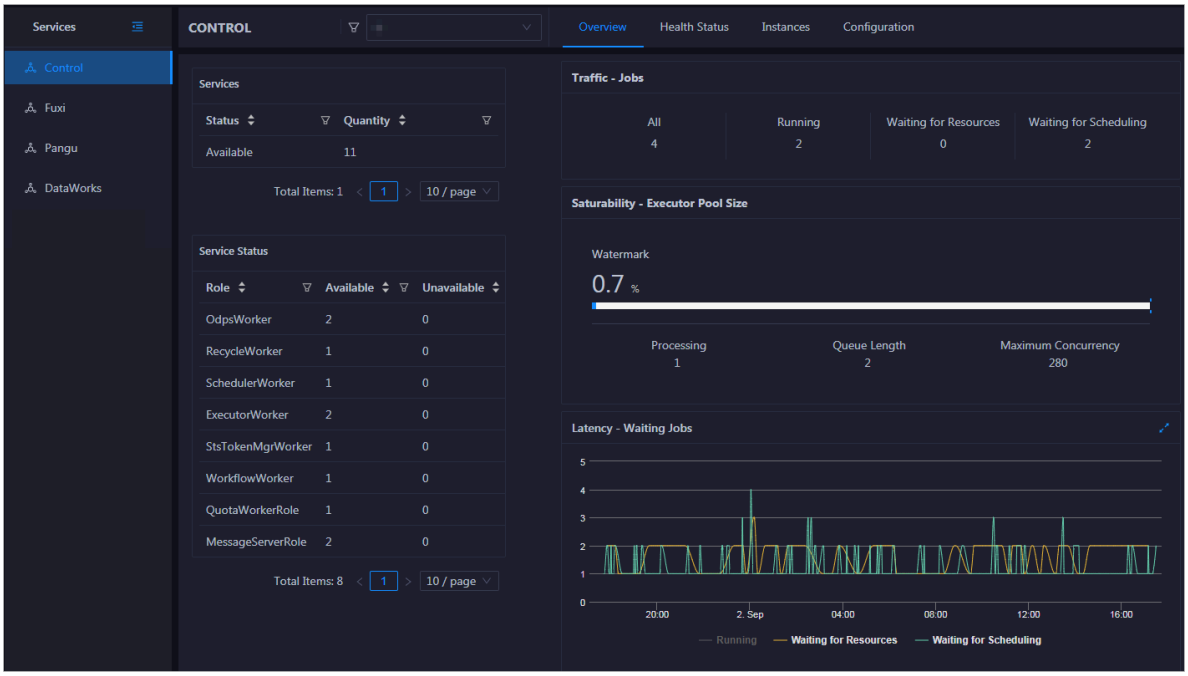

1.5.3.1.2 Control service overview

The Overview page displays the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

Entry

On the Services page, click Control in the left-side navigation pane. The Overview page for the control service appears.

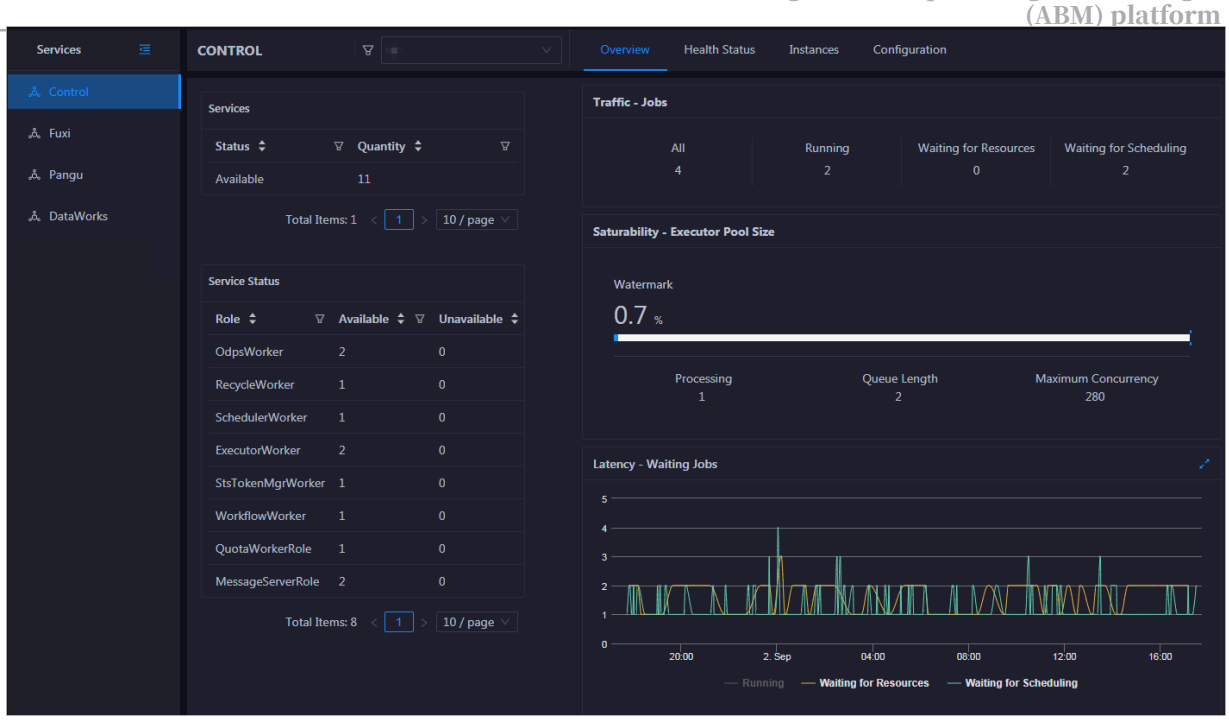

On the Overview page, you can view the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

Services

This section displays the numbers of available services and unavailable services respectively.

Service Status

This section displays all control service roles. You can also view the numbers of available and unavailable services respectively for each service role.

Traffic - Jobs

This section displays the total number of jobs in the cluster, and the numbers of running jobs, jobs waiting for resources, and jobs waiting for scheduling respective ly.

Saturability - Executor Pool Size

The section displays information about the thread pool, including the resource usage, number of jobs being processed, queue length, and maximum concurrency.

Latency - Waiting Jobs

This section displays the trend chart of jobs. The chart displays the trend lines of the numbers of running jobs, jobs waiting for resources, and jobs waiting for scheduling in different colors.

1.5.3.1.3 Control service health

On the Health Status page for the control service, 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 Services page, click Control in the left-side navigation pane, and then click the Health Status tab.

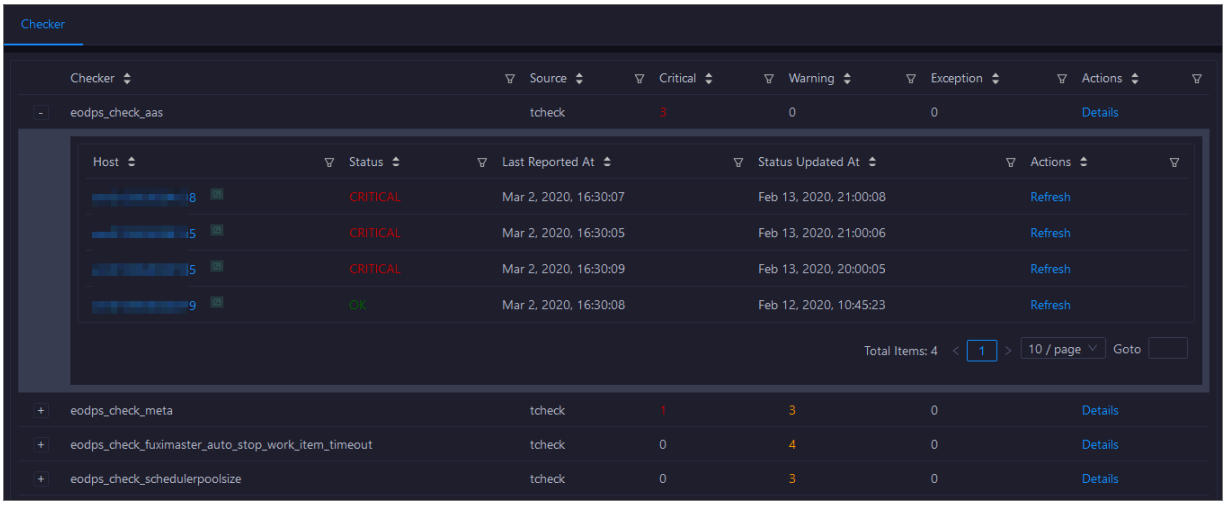

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

Supported operations

On the 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. For more information, see *Cluster health*.

1.5.3.1.4 Control service instances

The Instances page displays information about the service roles, including the host, service status, CPU application, and memory application of each service role.

Entry

On the Services page, click Control in the left-side navigation pane, and then click the Instances tab.

The Instance page displays information about the service roles, including the host, service status, CPU application, and memory application of each service role.

Other operations

You can filter or sort service roles by column to facilitate information display. For more information, see *Common operations*.

1.5.3.1.5 Control service configuration

The Configuration page under Control is the access to configuring global computing, cluster-level computing, computing scheduling, and cluster endpoints. If you need to modify the configurations of the control service, submit a ticket to apply for technical support, and then modify the configurations carefully under the guidance of technical support engineers.

On the Services page, click Control in the left-side navigation pane, and then click the Configuration tab.

The Configuration page consists of the following tabs:

- **Computing: provides the global computing configuration, cluster-level computing configuration, and compute scheduling configuration features.**
- **Tunnel Routing Address: provides the cluster endpoint configuration feature.**

1.5.3.1.6 Metadata warehouse for the control service

This topic describes how to view the complete time and status of the output tasks of metadata warehouse and the trend chart of the consumed time for running tasks in MaxCompute.

The metadata warehouse in MaxCompute regularly runs data output tasks every day. Apsara Bigdata Manager (ABM) obtains the status of output tasks every 30 minutes. If an output task of the metadata warehouse is not completed within 24 hours, the output task is regarded as a failure.

On the Services page, click Control in the left-side navigation pane, and then click the Metadata Repository tab.

The Metadata Repository page displays the throughput of metadata warehouse and the trend chart of consumed time for running tasks. The time displayed in the Completed At column indicates the time when the output task is completed. The time displayed in the Collected At column indicates the last time when ABM collects the status of the output task.

1.5.3.1.7 Disable or enable a control service role

Apsara Bigdata Manager (ABM) allows you to disable or enable control service roles of MaxCompute and view the execution history. You can also locate the failure cause when service role disabling or enabling fails.

Disable a service role

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.**
4. **On the Services page, click Control in the left-side navigation pane. Click Actions in the upper-left corner, and then click Stop Service Role.**
5. **In the dialog box that appears, select a service role to be disabled, and then click Run. A message appears, indicating that the action has been submitted.**

6. **Click Actions in the upper-left corner, and then click Execution History next to Stop Service Role to check whether the action is successful in the execution history.**

   The current status, submission time, start time, end time, and operator of each action are recorded in the execution history.
7. **Click Details to view the execution details.**

   On the execution details page, you can view the job name, execution status, execution steps, script, and parameter settings. You can also download the execution details to your local device.

Enable a service role

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.**
4. **On the Services page, click Control in the left-side navigation pane. Click Actions in the upper-left corner, and then click Start Service Role.**
5. **In the dialog box that appears, select a service role to be enabled, and then click Run. A message appears, indicating that the action has been submitted.**
6. **Click Actions in the upper-left corner, and then click Execution History next to Start Service Role to check whether the action is successful in the execution history.**

   The current status, submission time, start time, end time, and operator of each action are recorded in the execution history.
7. **Click Details to view the execution details.**

   On the execution details page, you can view the job name, execution status, execution steps, script, and parameter settings. You can also download the execution details to your local device.

Locate the failure cause

This section uses service role enabling as an example to describe how to locate the failure cause.

1. **In the execution history dialog box, click Details in the Details column of the task to view the details.**

2. In the dialog box that appears, click View Details for a failed step to locate the failure cause.

   You can also view the parameter settings, outputs, error messages, script, and execution parameters to locate the failure cause.

1.5.3.1.8 Start AdminConsole

AdminConsole is a management platform of MaxCompute. It is disabled by default. Apsara Bigdata Manager (ABM) allows you to quickly start AdminConsole to better manage MaxCompute clusters.

Prerequisites

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

Step 1: Start AdminConsole

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Control in the left-side navigation pane.
5. On the Control page, click Actions in the upper-left corner, and then click Start Admin Console.
6. In Start Admin Console dialog box that appears, click Run. A message appears, indicating that the action has been submitted.

Step 2: View the execution status or progress

1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Start Admin Console to view the execution history.

   In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful.
2. You can click Details in the Details column of a task in the RUNNING state to view the execution progress.

Step 3: Optional. Locate the failure cause

If the status of the task is FAILED, you can view the execution logs to locate the failure cause.

1. **On the Control page, click Actions in the upper-left corner, and then click Execution History next to Start Admin Console to view the execution history.**
2. **In the execution history dialog box, click Details in the Details column of the task to view the details.**
3. **On the Servers tab of the failed step, click View Details in the Actions column of a failed server. The Execution Output tab appears in the Execution Details section. You can view the output to locate the failure cause.**

1.5.3.1.9 Collect service logs

Apsara Bigdata Manager (ABM) allows you to collect service logs for the specified time period. This helps you locate the failure cause.

Prerequisites

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

Step 1: Collect service logs

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.**
4. **On the Services page, click Control in the left-side navigation pane.**
5. **On the Control page, click Actions in the upper-left corner, and then click Collect Service Logs.**
6. **In the Collect Service Logs dialog box that appears, set the parameters.**

   The following table describes the required parameters.

Parameter	Description
Target Service	The target service from which you want to collect service logs. Select a target service from the drop-down list.
Time Period	The time period in which the job logs that you want to collect are generated.

Parameter	Description
Degree of Concurrency	The maximum number of nodes from which you can collect service logs at the same time.
Hostname	The name of the host. Separate multiple hostnames with commas (,).

7. Click Run. A message appears, indicating that the action has been submitted.

Step 2: View the execution status or progress

1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Collect Service Logs to view the execution history.

   In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful.
2. You can click Details in the Details column of a task in the RUNNING state to view the execution progress.

Step 3: Optional. Locate the failure cause

If the status of the task is FAILED, you can view the execution logs to locate the failure cause.

1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Collect Service Logs to view the execution history.
2. In the execution history dialog box, click Details in the Details column of the task to view the details.
3. On the Servers tab of the failed step, click View Details in the Actions column of a failed server. The Execution Output tab appears in the Execution Details section. You can view the output to locate the failure cause.

1.5.3.2 Job Scheduler O&M

1.5.3.2.1 Job Scheduler O&M overview

This topic describes the O&M features of Job Scheduler and how to access the Job Scheduler O&M page.

Modules

- **Overview page: displays the key operation metrics of Job Scheduler, including the service overview, service status, resource usage, and compute node overview . You can also view the trend charts of CPU and memory usage on this page.**
- **Health Status page: displays all checkers for Job Scheduler, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.**
- **Quotas page: allows you to view, add, or modify Job Scheduler quota groups.**
- **Instances page: displays the information about the master nodes and service roles of Job Scheduler and allows you to restart the master nodes.**
- **Compute Nodes page: displays all Job Scheduler compute nodes and allows you to add compute nodes to or remove compute nodes from the blacklist or read-only list.**
- **Enable SQL Acceleration or Disable SQL Acceleration action: allows you to enable or disable SQL acceleration for Job Scheduler.**
- **Restart Fuxi Master Node action: allows you to restart the primary and secondary master nodes for Job Scheduler.**

Entry

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.**

4. On the Services page, click Fuxi in the left-side navigation pane and then select a cluster. The Overview page for the selected cluster appears.

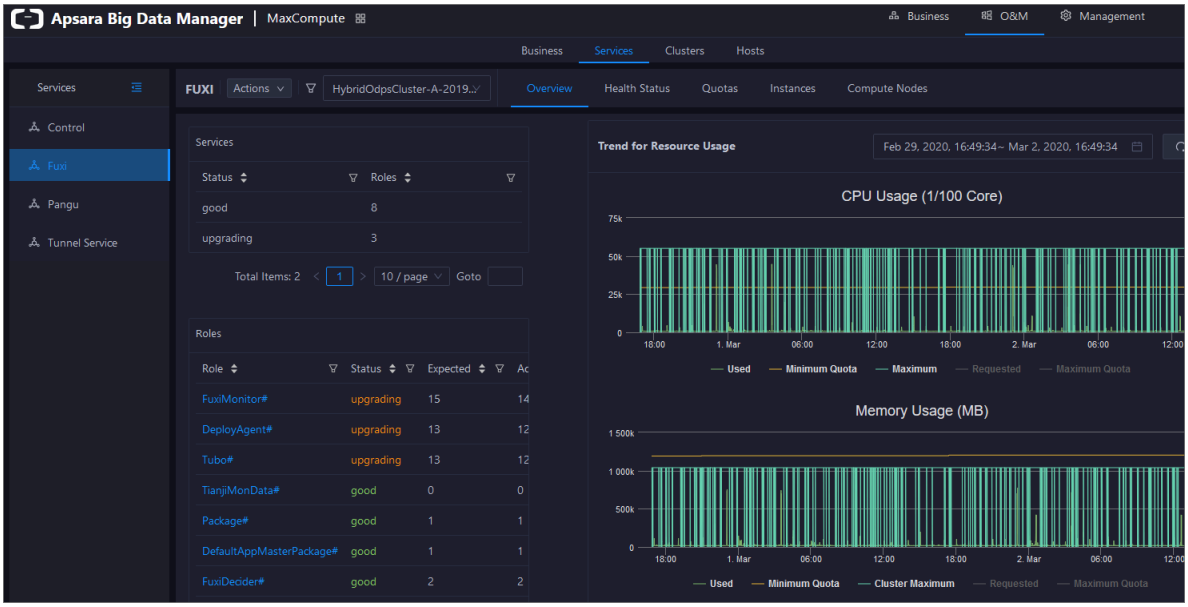

1.5.3.2.2 Job Scheduler overview

The Overview page displays the key operation metrics of Job Scheduler, including the service overview, service status, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.

Entry

1. On the Services page, click Fuxi in the left-side navigation pane.

2. Select a cluster from the drop-down list, and then click the Overview tab. The Overview page for the selected cluster appears.

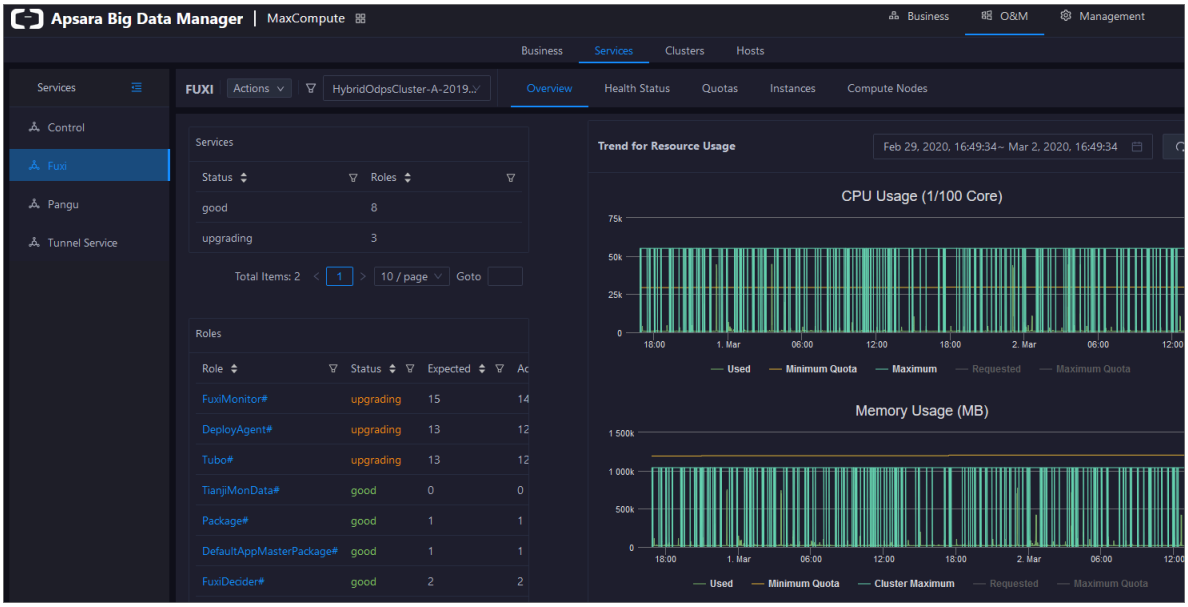

The Overview page displays the key operation metrics of Job Scheduler, including the service overview, service status, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

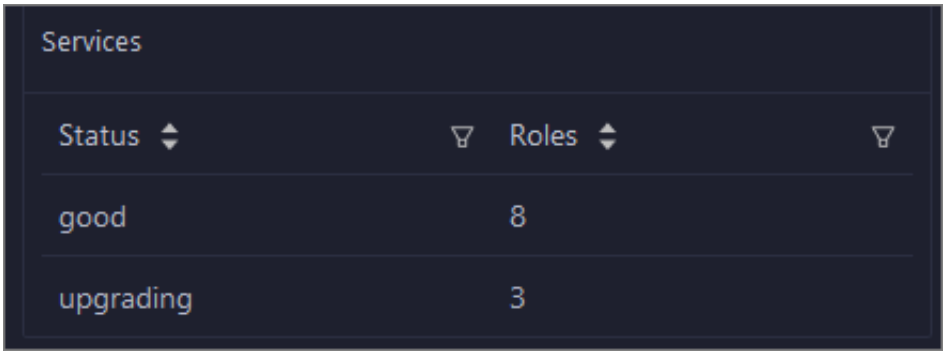

Roles

This section displays all Job Scheduler service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

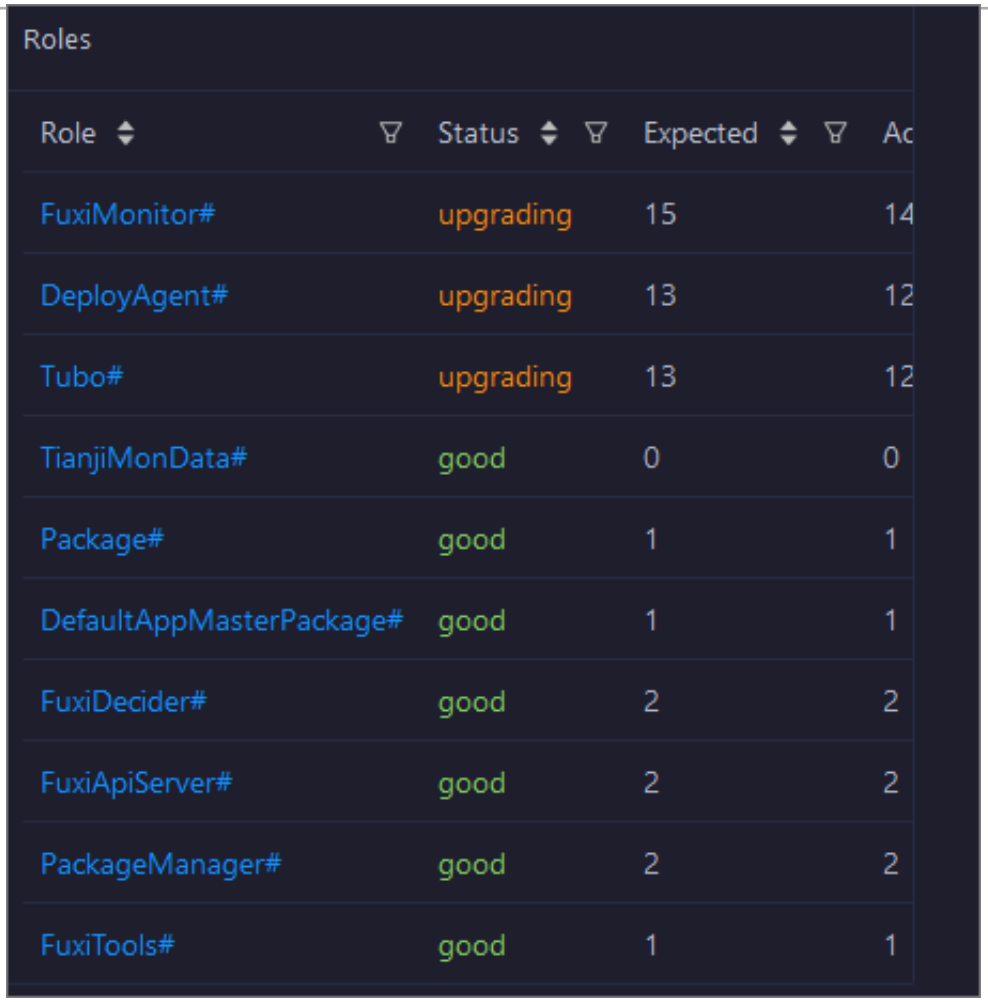
Roles

Role	Status	Expected	Ac
FuxiMonitor#	upgrading	15	14
DeployAgent#	upgrading	13	12
Tubo#	upgrading	13	12
TianjiMonData#	good	0	0
Package#	good	1	1
DefaultAppMasterPackage#	good	1	1
FuxiDecider#	good	2	2
FuxiApiServer#	good	2	2
PackageManager#	good	2	2
FuxiTools#	good	1	1

Click the name of a service role to go to the Apsara Infrastructure Management Framework console to view details.

Saturability - Resource Usage

This section displays the usage and allocation of CPU and memory resources.

- **CPU (Core): displays the CPU usage, the total number of CPU cores, the number of available CPU cores, and the CPU cores for SQL acceleration.**
- **Memory (Bytes): displays the memory usage, the total memory size, the available memory size, and the size of memory for SQL acceleration.**

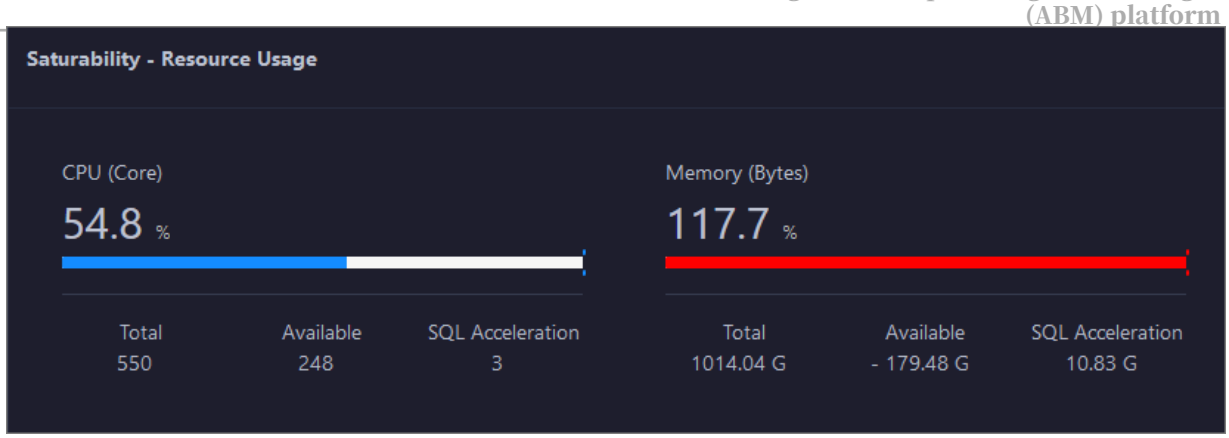

CPU Usage (1/100 Core) and Memory Usage (MB)

This section displays the trend charts of CPU and memory usage for Job Scheduler . Each trend chart displays the trend lines of the used quota, minimum quota , maximum cluster quota, requested quota, and maximum quota over time in different colors.

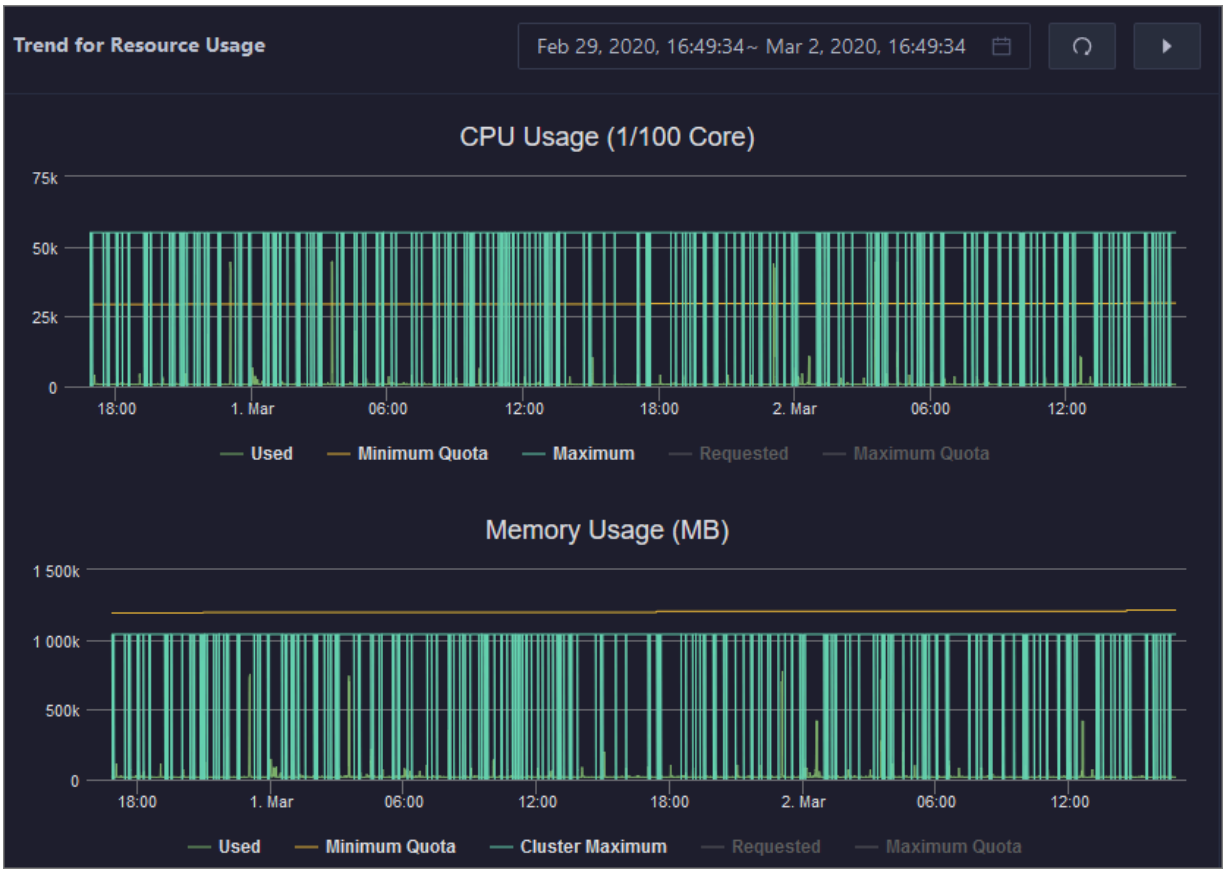

Compute Nodes

This section displays the details of Job Scheduler compute nodes, including the online rate, the total number of compute nodes, the number of online compute nodes, and the number of compute nodes in the blacklist.

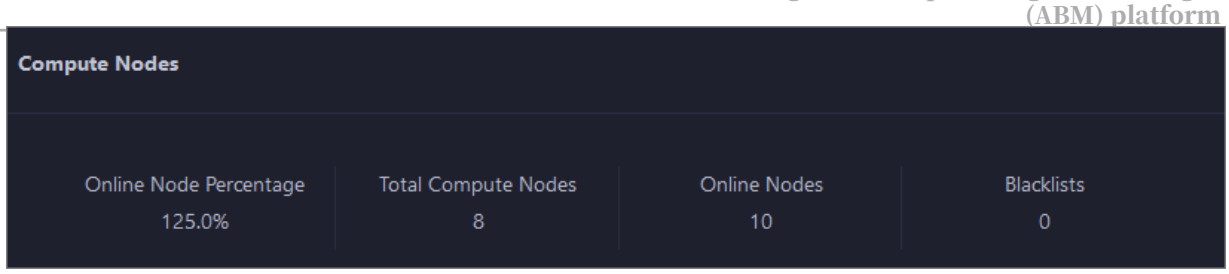

1.5.3.2.3 Job Scheduler health

On the Health Status page for Job Scheduler, you can view all checkers of Job Scheduler, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. **On the Services page, click Fuxi in the left-side navigation pane.**
2. **Select a cluster from the drop-down list, and then click the Health Status tab. The Health Status page for Job Scheduler appears.**

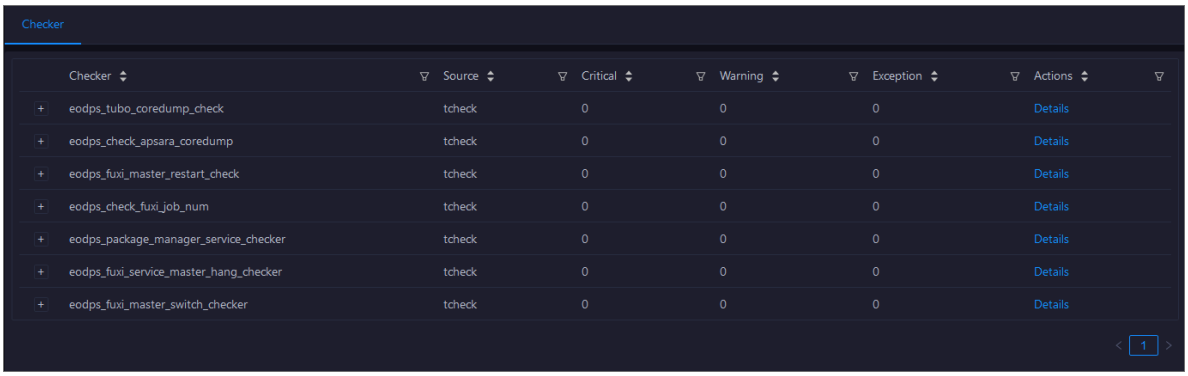

On the Health Status page, you can view all checkers of the Job Scheduler service and the check results for all hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

Supported operations

On the 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. For more information, see *Cluster health*.

1.5.3.2.4 Job Scheduler quota management

You can view, add, or modify Job Scheduler quota groups on the Quotas page. A quota group is used to allocate computing resources, including CPU and memory resources, to MaxCompute projects.

Entry

1. **On the Services page, click Fuxi in the left-side navigation pane.**
2. **Select a cluster from the drop-down list, and then click the Quotas tab. The Quotas page for the selected cluster appears.**

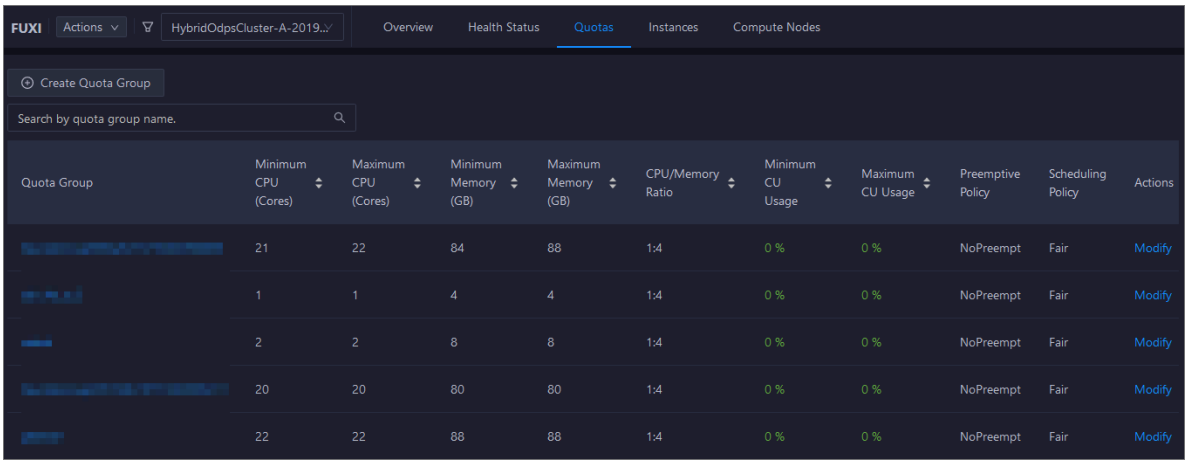

The Quotas page lists the existing Job Scheduler quota groups.

Add a quota group

1. **On the Quotas page, click Create Quota Group in the upper-left corner.**
2. **In the Quota Group dialog box that appears, set the parameters as instructed.**

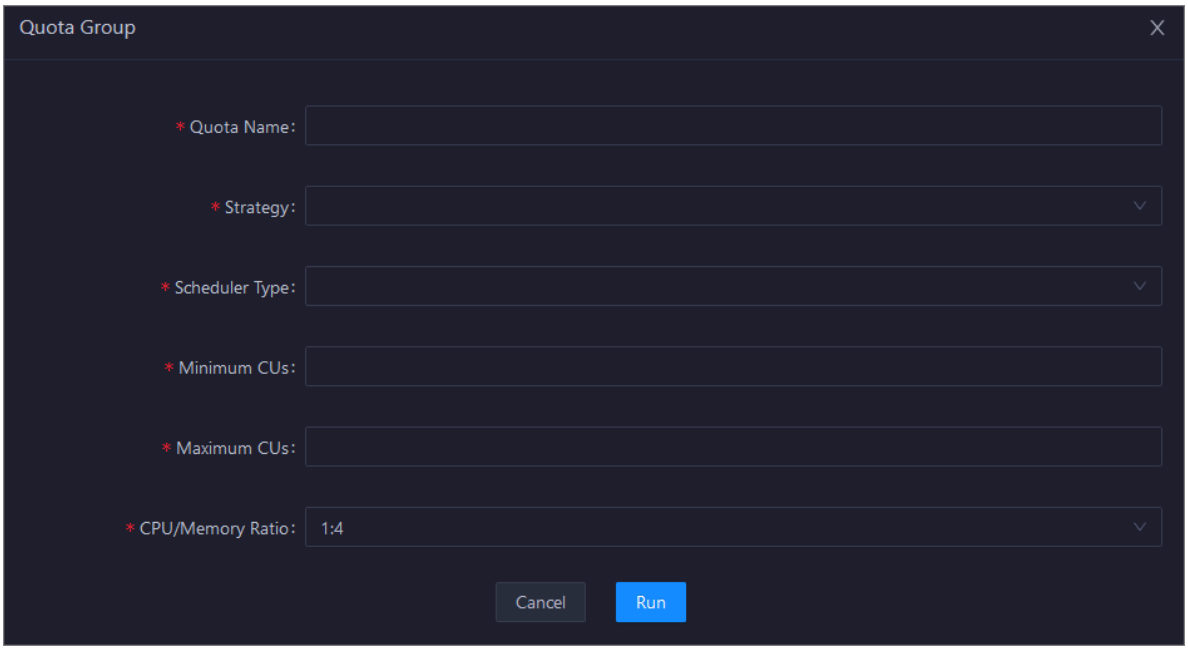

3. **Click Run. A message appears, indicating that the action has been submitted.**

   After the quota group is created, it appears in the quota group list.

View quota group details

Click the name of a quota group to view the details. The Resource Usage tab displays the trend charts of CPU and memory usage. The Applications page displays the projects that use the quota group resources.

Figure 1-2: Resource usage

Figure 1-3: Applications

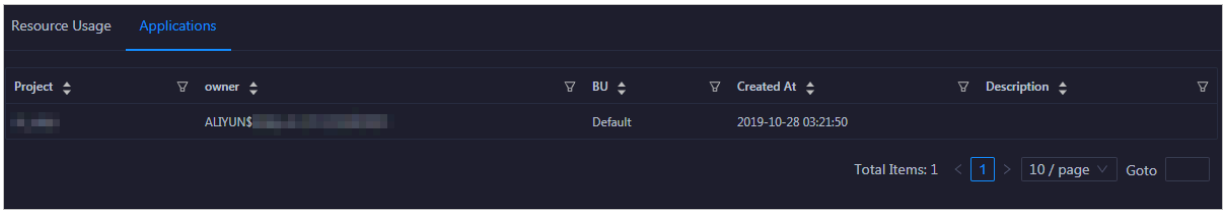

Modify a quota group

1. **On the Quotas page, find the quota group you want to modify, click Modify in the Actions column, and then modify parameters as instructed in the dialog box that appears.**
2. **Click Run. A message appears, indicating that the action has been submitted.**

   You can check whether the quota group is successfully modified in the quota group list after the configuration is completed.

1.5.3.2.5 Job Scheduler instances

This topic describes how to view the information about the master nodes and server roles of Job Scheduler and how to restart the master nodes.

Entry

1. **On the Services page, click Fuxi in the left-side navigation pane.**
2. **Select a cluster from the drop-down list, and then click the Instances tab. The Instances page for Job Scheduler appears.**

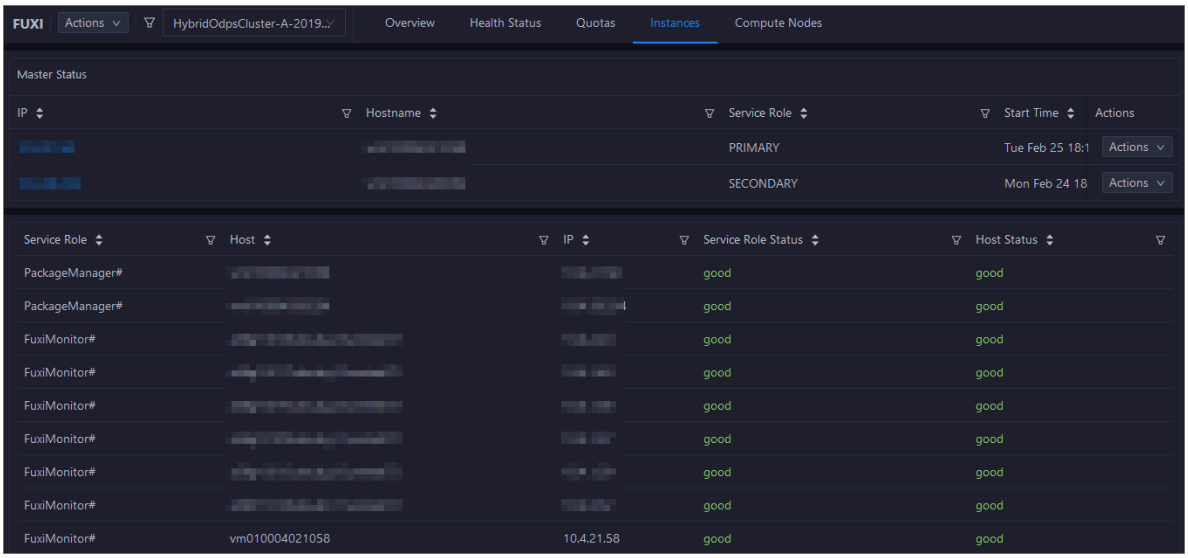

The Instances page displays the information about the master nodes and service roles of Job Scheduler. The information about the master nodes includes the IP address, hostname, service role, and start time. The information about the service roles includes the service role name, service role host, service role status, and host status.

Supported operations

You can filter or sort service roles by column to facilitate information retrieval. For more information, see *Common operations*.

You can restart the master nodes of Job Scheduler. For more information, see *Restart a master node of Job Scheduler*.

1.5.3.2.6 Job Scheduler compute nodes

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including the total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and

whether it is active. In addition, you can add compute nodes to or remove compute nodes from the blacklist or read-only list on the Compute Nodes page.

Entry

1. **On the Services page, click Fuxi in the left-side navigation pane.**
2. **Select a cluster from the drop-down list, and then click the Compute Nodes tab. The Compute Nodes page for Job Scheduler appears.**

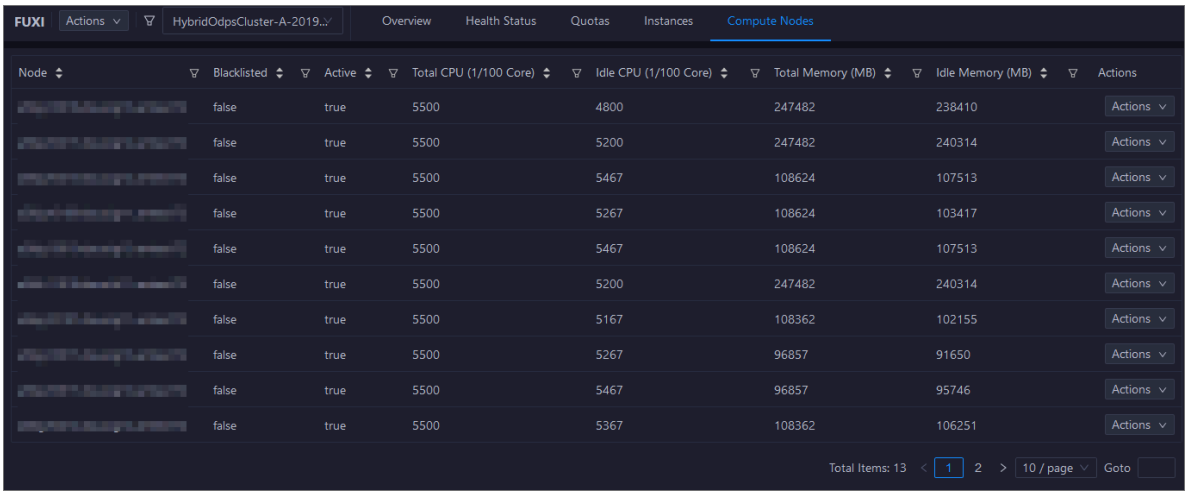

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including the total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active.

Blacklist and read-only setting

You can add compute nodes to or remove compute nodes from the blacklist or read-only list. To add compute nodes to the blacklist, follow these steps:

1. **On the Compute Nodes page, click Actions for the target compute node and then select Add to Blacklist.**

2. In the dialog box that appears, click Run. A message appears, indicating that the action has been submitted.

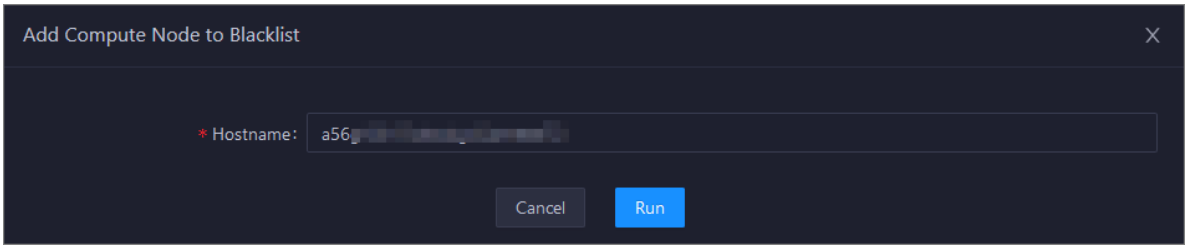

The value of the Hostname parameter is automatically filled. You do not need to specify a value for this parameter.

You can check whether a compute node is added to the blacklist in the compute node list after the configuration is completed.

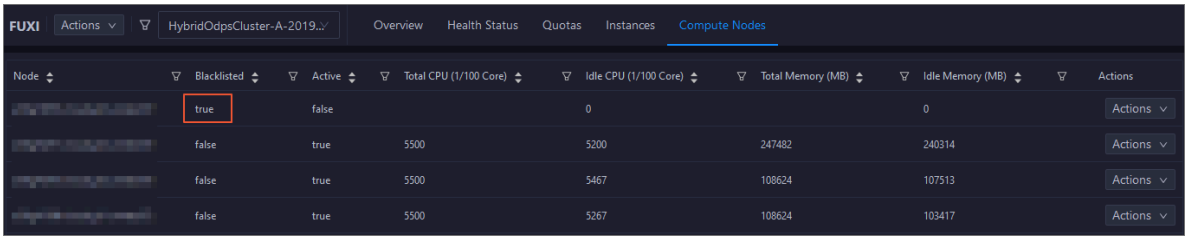

1.5.3.2.7 Enable or disable SQL acceleration

You can enable or disable SQL acceleration for Job Scheduler in the Apsara Bigdata Manager (ABM) console. Enabling SQL acceleration can greatly increase the speed of running SQL statements in Job Scheduler, but it consumes more computing resources.

Enable SQL acceleration

1. On the Services page, click Fuxi in the left-side navigation pane and then select a cluster.

2. **Click Actions next to FUXI in the upper-left corner, and then click Enable SQL Acceleration.**

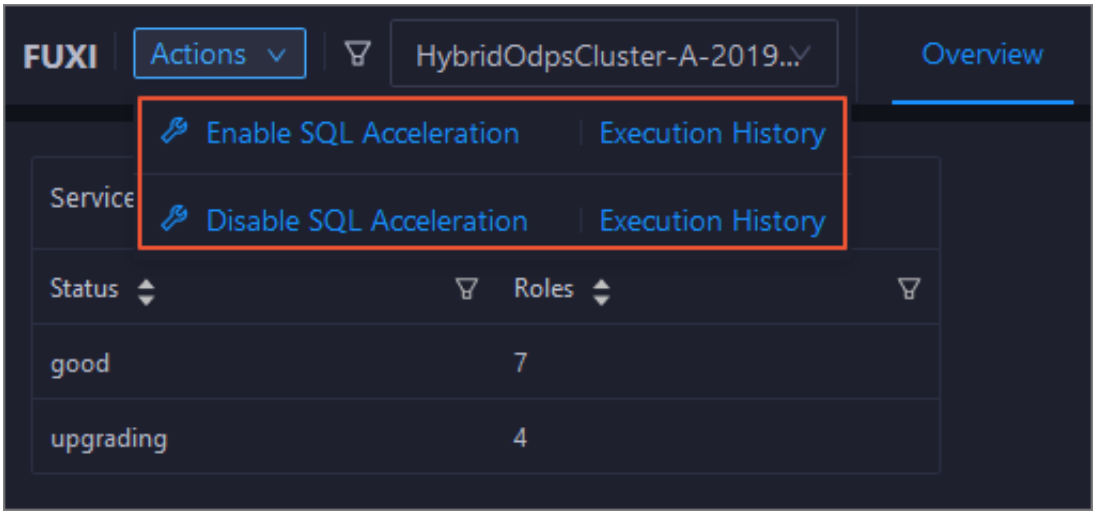

3. **In the dialog box that appears, set the WorkerSpans parameter.**

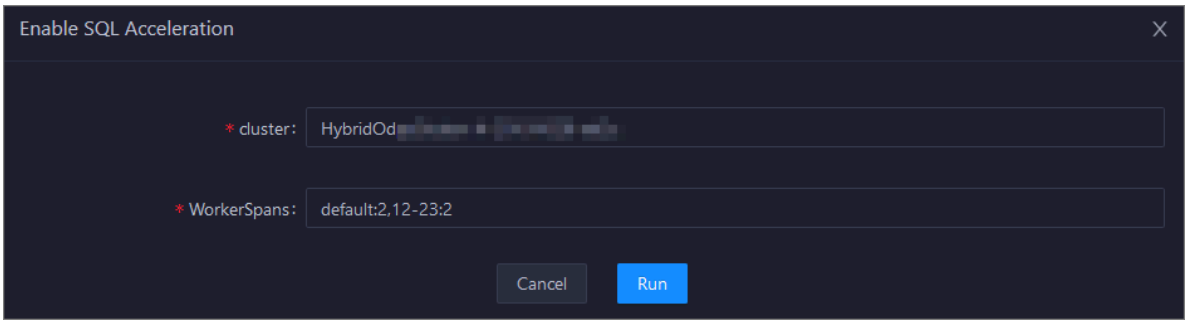

WorkerSpans: the default resource quota of the cluster and the resource quota within the specified period. Default value: default:2,12-23:2. The default value indicates that the default resource quota is 2 and the resource quota for the period from 12:00 to 23:00 is also 2. You can set the resource quota as needed. For example, during business peak hours, you can set this parameter to default:2,12-23:4 to increase the resource quota.

4. **Click Run. A message appears, indicating that the action has been submitted.**

Disable SQL acceleration

1. **On the Services page, click Fuxi in the left-side navigation pane and then select a cluster.**

2. **Click Actions next to FUXI in the upper-left corner, and then click Disable SQL Acceleration.**

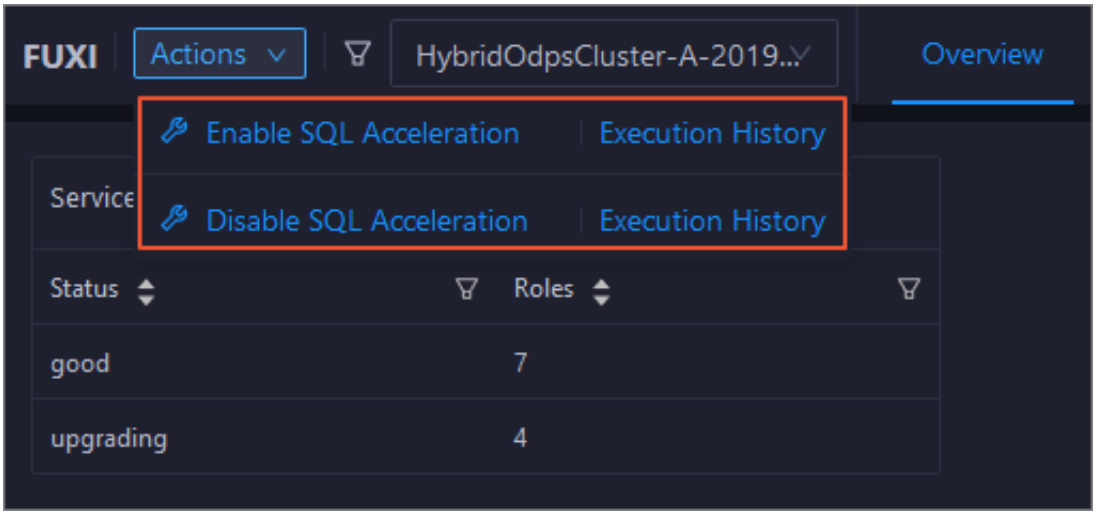

3. **In the dialog box that appears, click Run. A message appears, indicating that the action has been submitted.**

View the execution history of enabling or disabling SQL acceleration

After you submit the action of enabling or disabling SQL acceleration, you can check whether the current action is completed by viewing the execution history. The system executes the action as a job. It provides execution records and logs for each execution so that you can locate faults encountered during the execution of the job. To view the execution history of enabling SQL acceleration, follow these steps:

1. **On the Services page, click Fuxi in the left-side navigation pane and then select a cluster.**

2. **Click Actions next to FUXI in the upper-left corner, and then click Execution History next to Enable SQL Acceleration.**

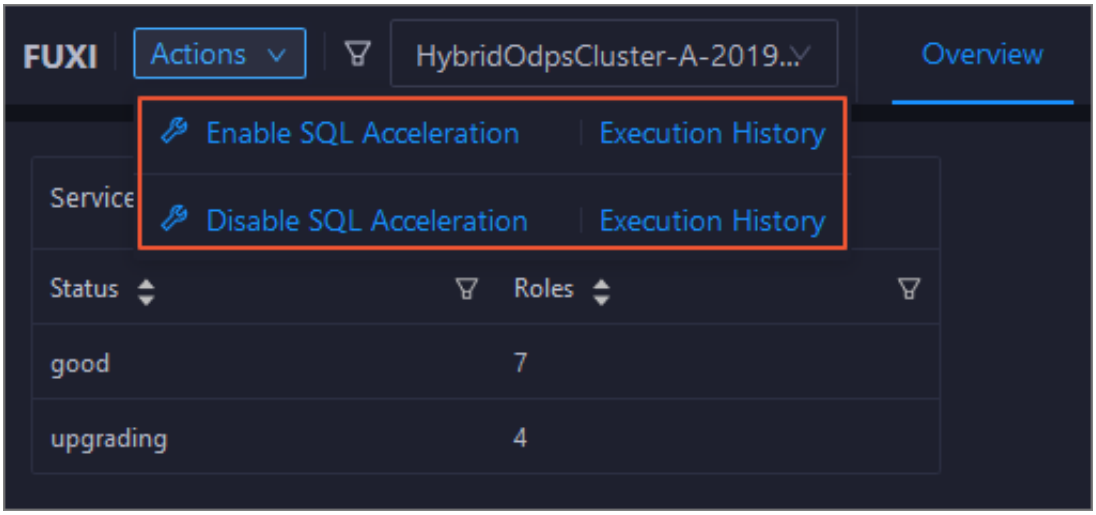

3. **In the dialog box that appears, view the execution history of enabling SQL acceleration.**

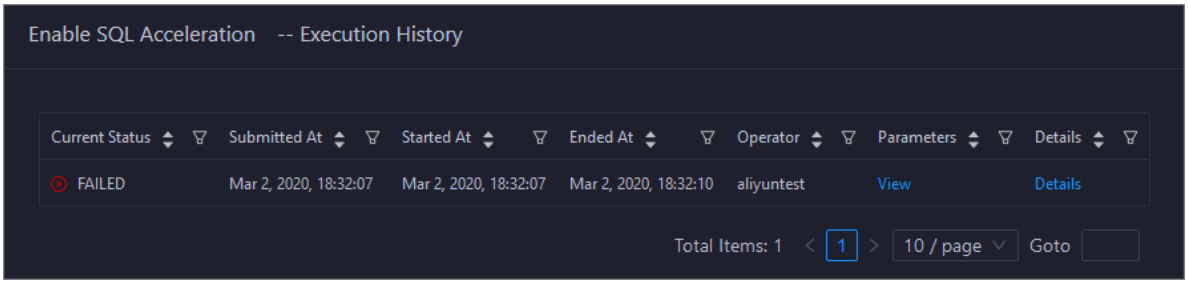

The current status, submission time, start time, end time, and operator of each action are recorded in the execution history.

4. If the execution fails, click Details to go to the Jobs page to locate the failure cause.

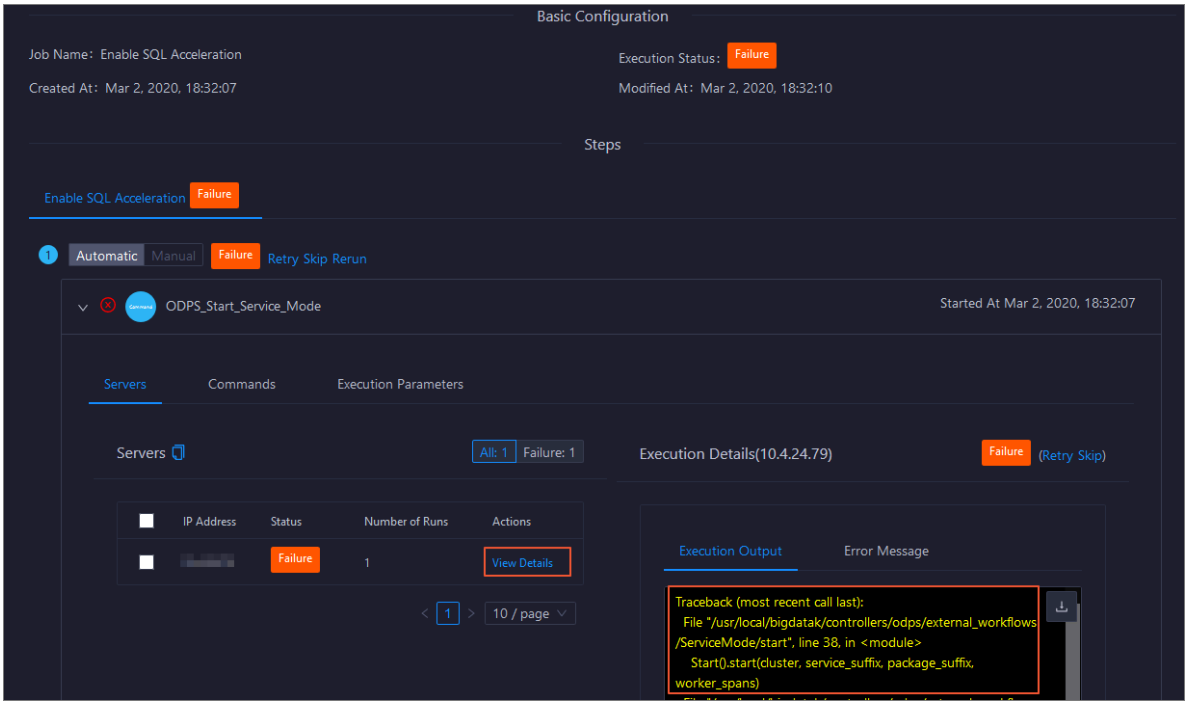

1.5.3.2.8 Restart a master node of Job Scheduler

Job Scheduler is the resource management and task scheduling system of the Apsara system. Apsara Bigdata Manager (ABM) allows you to quickly restart the primary and secondary master nodes of Job Scheduler. Cluster services are not affected during the restart process.

Prerequisites

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

Step 1: Restart a master node of Job Scheduler

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Fuxi in the left-side navigation pane, and then click the Instances tab.
5. On the Instances page, click Actions and select Restart Fuxi Master Node in the Actions column of a primary or secondary master node.

6. **In the Restart Fuxi Master Node dialog box that appears, click Run. A message appears, indicating that the action has been submitted.**

Step 2: View the execution status or progress

1. **On the Fuxi page, click Actions in the upper-left corner, and then click Execution History next to Restart Fuxi Master Node to view the execution history.**

   In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful.
2. **You can click Details in the Details column of a task in the RUNNING state to view the execution progress.**

Step 3: Optional. Locate the failure cause

If the status of the task is FAILED, you can view the execution logs to locate the failure cause.

1. **On the Fuxi page, click Actions in the upper-left corner, and then click Execution History next to Restart Fuxi Master Node to view the execution history.**
2. **In the execution history dialog box, click Details in the Details column of the task to view the details.**
3. **On the Servers tab of the failed step, click View Details in the Actions column of a failed server. The Execution Output tab appears in the Execution Details section. You can view the output to locate the failure cause.**

1.5.3.3 Apsara Distribute File System O&M

1.5.3.3.1 Apsara Distribute File System O&M overview

This topic describes the O&M features of Apsara Distributed File System and how to access the Apsara Distributed File System O&M page.

Modules

- **Overview page: displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, storage usage, and storage node overview. You can also view the trend charts of storage usage and file count on this page.**

- **Health Status page: displays all checkers for Apsara Distributed File System, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.**
- **Instances page: displays the information about the master nodes and service roles of Apsara Distributed File System. On this page, you can change the primary master node or run a checkpoint on a master node of Apsara Distributed File System.**
- **Storage Nodes page: displays the information about the storage nodes of Apsara Distributed File System. On this page, you can set the status of a storage node to Disabled or Normal. In addition, you can set the status of a disk on a storage node to Normal or Error.**
- **Change Primary Master Node action: allows you to change the primary master node of Apsara Distributed File System in a cluster.**
- **Run Checkpoint on Master Node action: allows you to run checkpoints on master nodes of Apsara Distributed File System to write memory data into disks.**
- **Empty Recycle Bin action: allows you to empty the recycle bin of Apsara Distributed File System.**
- **Enable Data Rebalancing or Disable Data Rebalancing action: allows you to enable or disable the data rebalancing feature of Apsara Distributed File System.**

Entry

1. *Log on to the ABM console*.
2. **Click  in the upper-left corner, and then click MaxCompute.**
3. **On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.**

4. On the Services page, click Pangu in the left-side navigation pane and then select a cluster. The Overview page for Apsara Distributed File System appears.

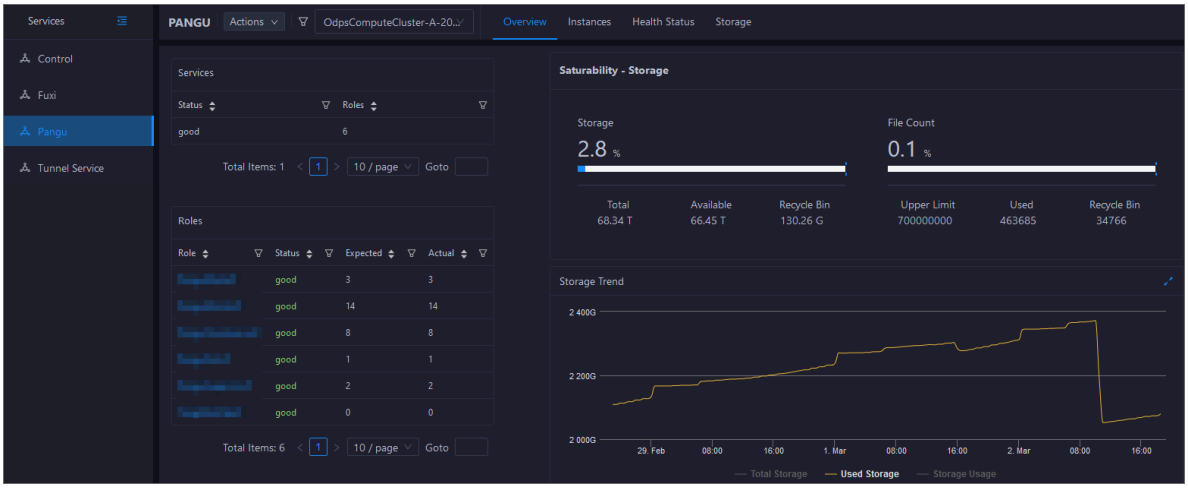

1.5.3.3.2 Apsara Distributed File System overview

On the Overview page for Apsara Distributed File System, you can view the key operation metrics, including the service overview, service status, storage usage, and storage node overview. You can also view the trend charts of storage usage and file count on this page.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.
2. Select a cluster from the drop-down list, and then click the Overview tab. The Overview page for Apsara Distributed File System appears.

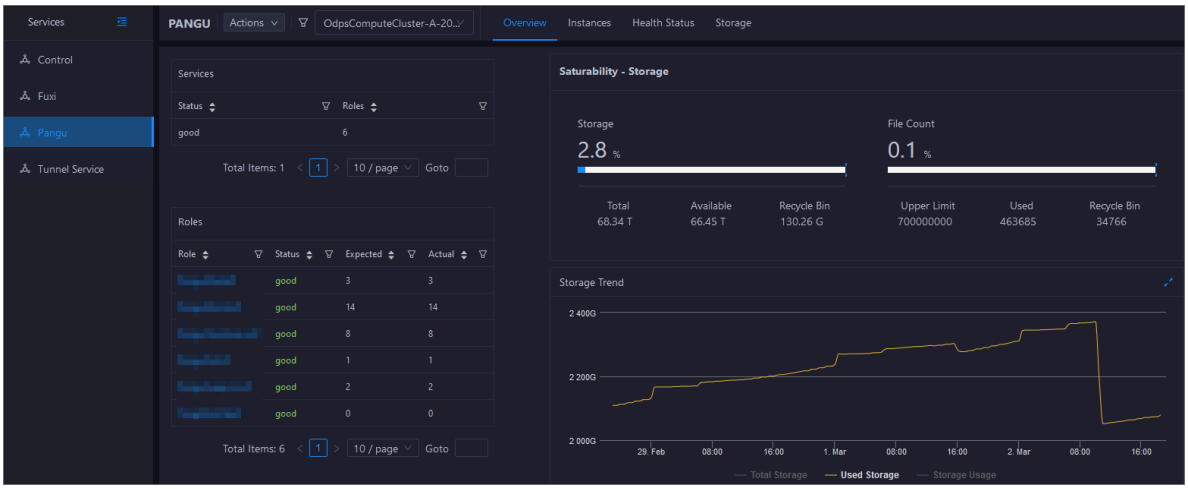

The Overview page displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, health check result, health check history, storage usage, and storage node overview. You can also view the trend charts of storage usage and file count on this page.

Services

This section displays the status of Apsara Distributed File System and the number of service roles.

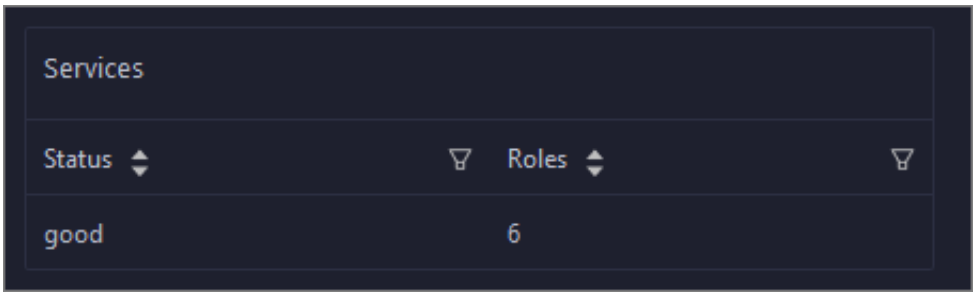

Roles

This section displays all Apsara Distributed File System service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Saturability - Storage

This section displays the storage usage (Storage) and file count (File Count).

- **Storage: displays the storage usage, total storage size, available storage size, and recycle bin size.**
- **File Count: displays the file count usage, maximum number of files, number of existing files, and number of files in the recycle bin.**

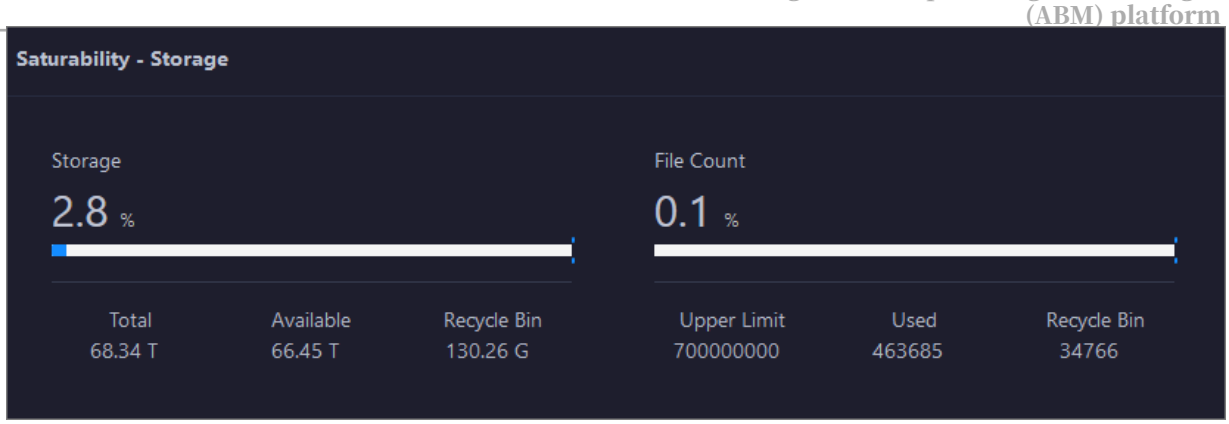

Storage Trend and File Count Trend

This section displays the storage usage and file count charts. The storage usage chart displays the trend lines of the total storage size, used storage size, and storage usage over time in different colors. The file count chart displays the trend line of the file count.

Click in the upper-right corner of the chart to zoom in the chart. The following figure shows an enlarged storage usage chart.

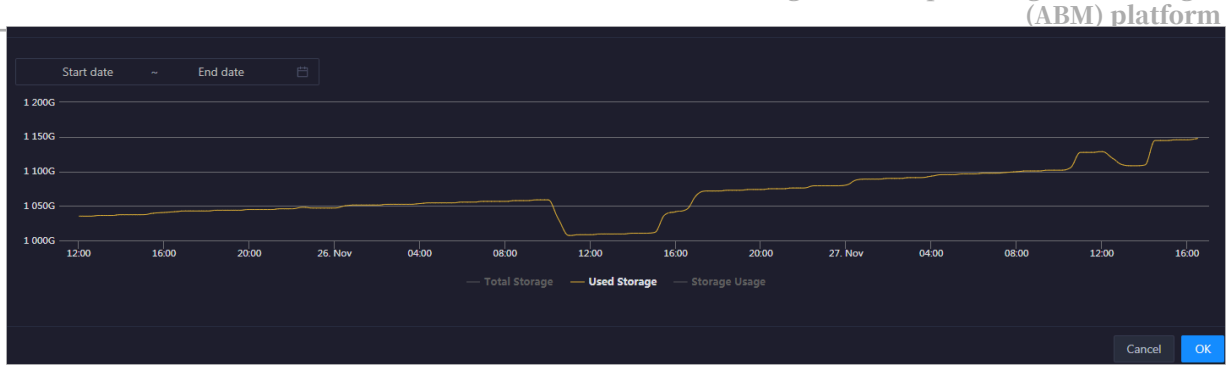

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

Storage Nodes

This section displays the information about the storage nodes of Apsara Distributed File System, including the respective number of data nodes, normal nodes, disks, and normal disks. You can also view the faulty node percentage and faulty disk percentage in this section.

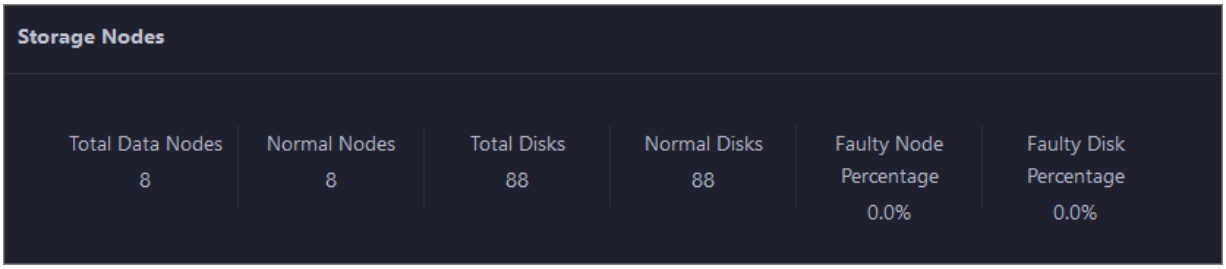

1.5.3.3.3 Apsara Distributed File System instances

This topic describes how to view the information about the master nodes and services roles of Apsara Distributed File System. This topic also describes how to change the primary master node or run a checkpoint on a master node of Apsara Distributed File System.

Entry

1. **On the Services page, click Pangu in the left-side navigation pane.**

2. **Select a cluster from the drop-down list, and then click the Instances tab. The Instances page for Apsara Distributed File System appears.**

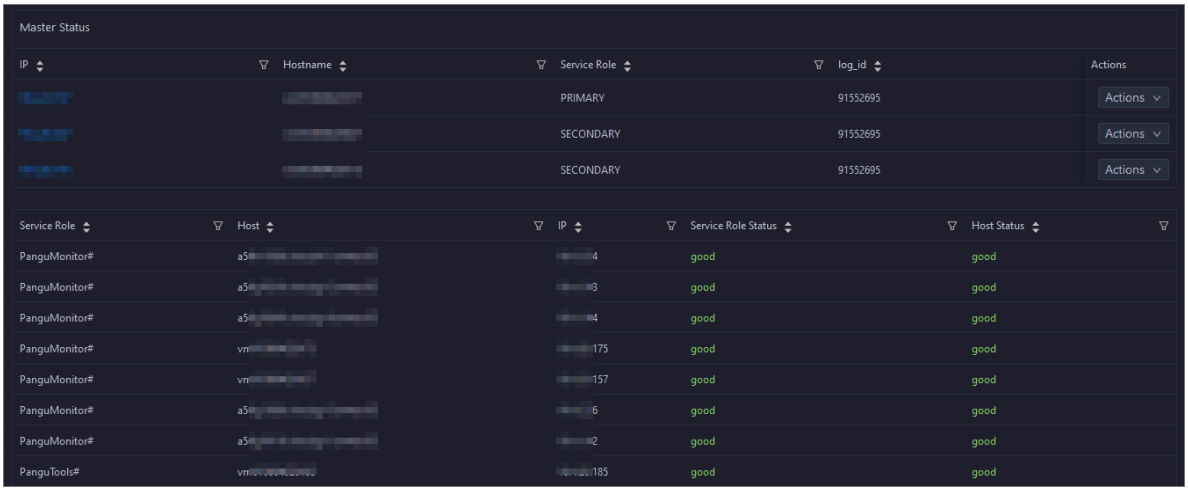

The Instances page displays the information about the master nodes and service roles of Apsara Distributed File System. The information about the master nodes includes the IP address, hostname, service role, and log ID. The information about the service roles includes the service role name, service role host, service role status, and host status.

Supported operations

You can filter or sort service roles by column to facilitate information retrieval. For more information, see *Common operations*.

You can change the primary master node or run a checkpoint on a master node of Apsara Distributed File System. For more information, see *Change the primary master node of Apsara Distributed File System* **and** *Run a checkpoint on a master node of Apsara Distributed File System*.

1.5.3.3.4 Apsara Distributed File System health

On the Health Status page for Apsara Distributed File System, you can view all checkers of Apsara Distributed File System, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. **On the Services page, click Pangu in the left-side navigation pane.**

2. **Select a cluster from the drop-down list, and then click the Health Status tab. The Health Status page for Apsara Distributed File System appears.**

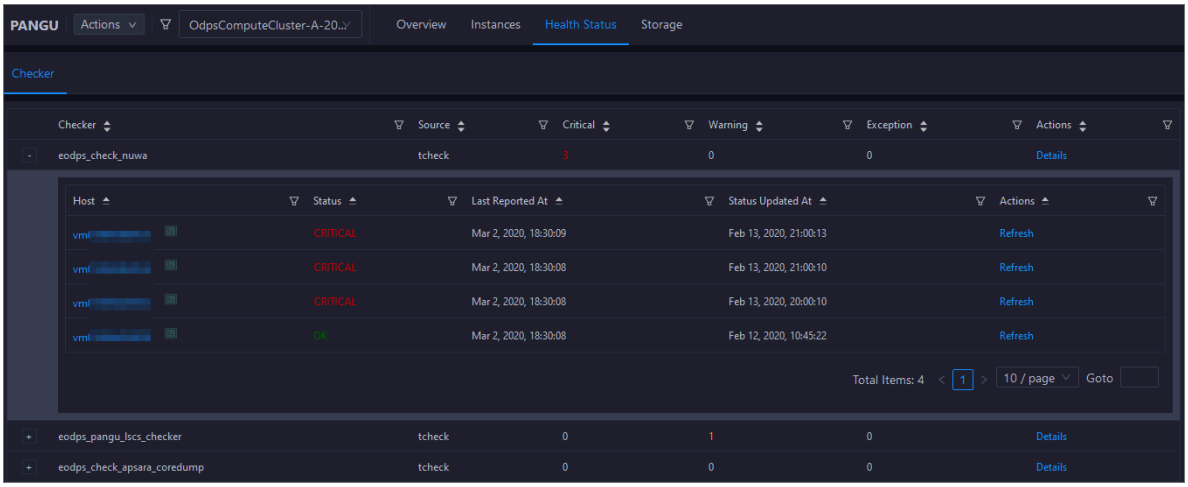

On the Health Status page, you can view all checkers of Apsara Distributed File System and the check results for all hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

Supported operations

On the 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. For more information, see *Cluster health*.

1.5.3.3.5 Apsara Distributed File System storage

This topic describes how to view the storage overview and storage node information of Apsara Distributed File System, and how to set the status of storage nodes and data disks.

Entry to the Storage Overview page

1. **On the Services page, click Pangu in the left-side navigation pane.**

2. **Select a cluster from the drop-down list, and then click the Storage tab. The Storage Overview page for Apsara Distributed File System appears.**

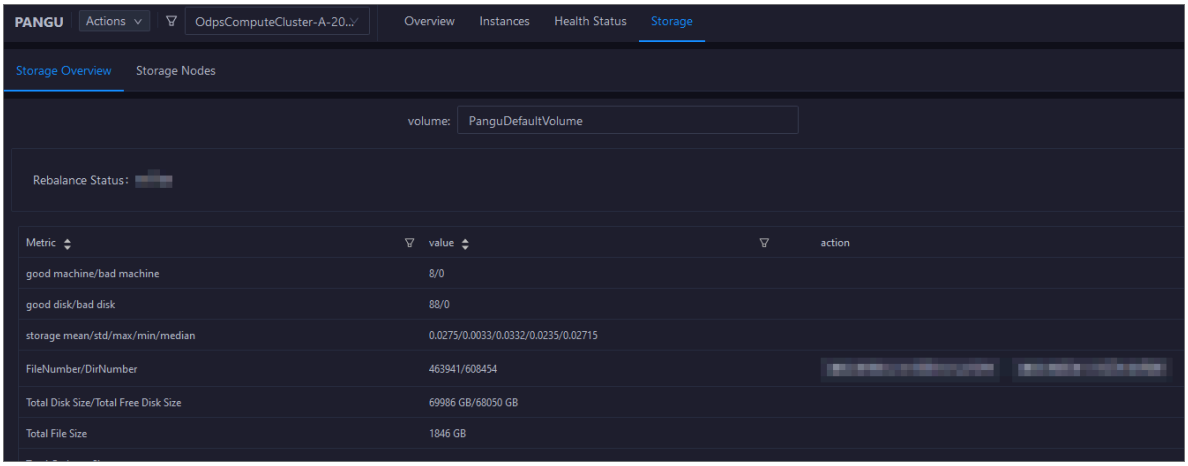

The Storage Overview page displays whether data rebalancing is enabled, key metrics and their values, suggestions to handle exceptions, and rack specifications of Apsara Distributed File System. The Storage Nodes page displays the information about all storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, time to live (TTL), and send buffer size. You can also set the status of storage nodes and data disks on this page.

Entry to the Storage Nodes page

1. **On the Services page, click Pangu in the left-side navigation pane.**
2. **Select a cluster from the drop-down list, and then click the Storage tab. The Storage Overview page for Apsara Distributed File System appears.**

3. Click the Storage Nodes tab. The Storage Nodes page appears.

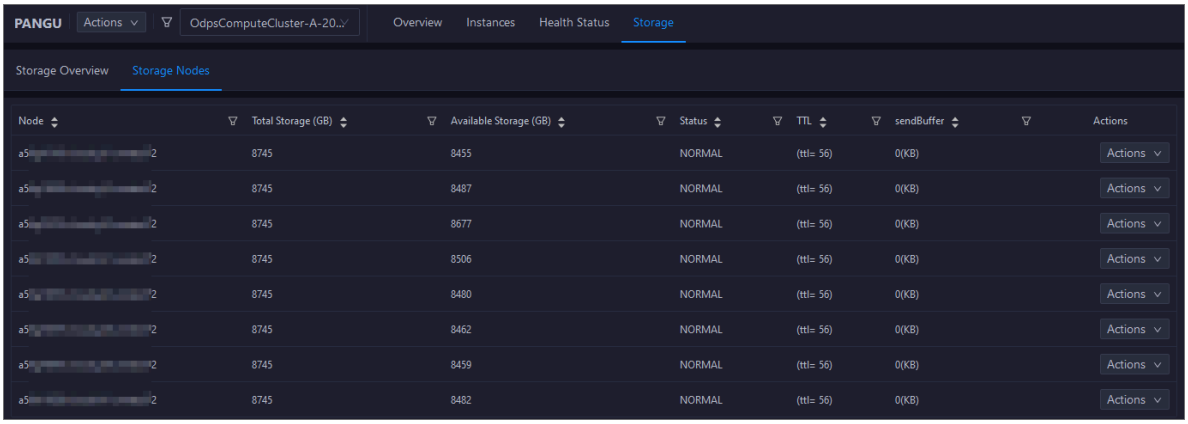

The Storage Nodes page displays the information about all storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, TTL, and send buffer size.

Set the storage node status

You can set the storage node status to Disabled or Normal. This section describes how to set the status of a storage node to Disabled.

1. On the Storage Nodes page, find the target storage node and choose Actions > Set Node Status to Disabled in the Actions column.
2. In the dialog box that appears, click Run. A message appears, indicating that the action has been submitted.

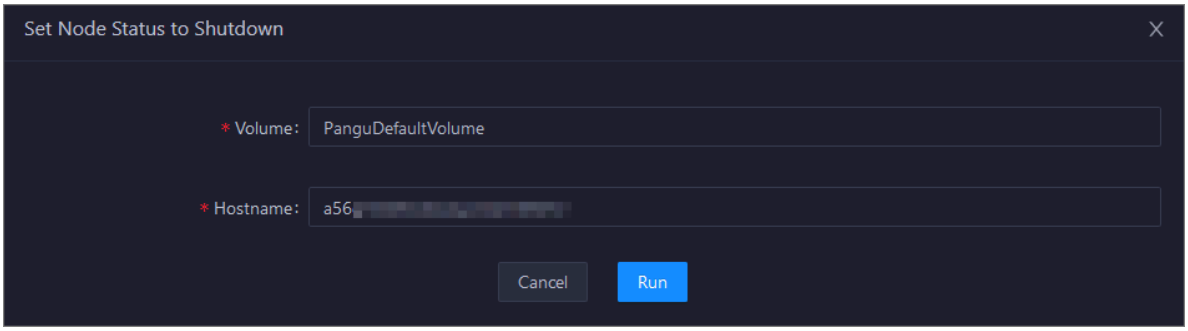

The values of the Volume and Hostname parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

You can check whether the status of storage node is changed in the storage node list.

Set the data disk status

You can set the data disk status to Error or Normal. This section describes how to set the status of a data disk to Error.

1. **On the Storage Nodes page, find the target storage node and choose Actions > Set Disk Status to Error in the Actions column.**
2. **In the dialog box that appears, set the Diskid parameter.**

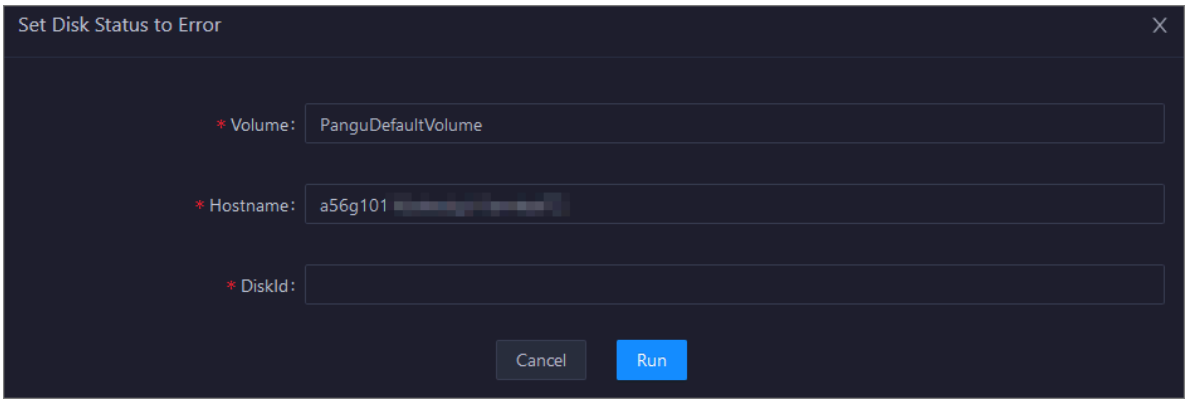

The values of the Volume and Hostname parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

3. **Click Run. A message appears, indicating that the action has been submitted.**

1.5.3.3.6 Change the primary master node of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to perform primary/secondary switchover on the master nodes of Apsara Distributed File System. After the primary/secondary switchover is completed, a secondary master node becomes the new primary master node, and the original primary master node becomes a new secondary master node.

Prerequisites

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

Background

A volume in Apsara Distributed File System is similar to a namespace in Hadoop Distributed File System (HDFS). The default volume is PanguDefaultVolume. Multiple volumes may exist if a cluster consists of numerous nodes. A volume has

three master nodes. One of the nodes serves as the primary master node, whereas the other two nodes serve as secondary master nodes.

Procedure

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click MaxCompute.
3. On the page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Pangu in the left-side navigation pane. Select a cluster from the drop-down list, and then click the Instances tab.
5. In the Master Status section of the Instances page, find the target master node and choose Actions > Change Primary Master Node in the Actions column. In the dialog box that appears, set required parameters.

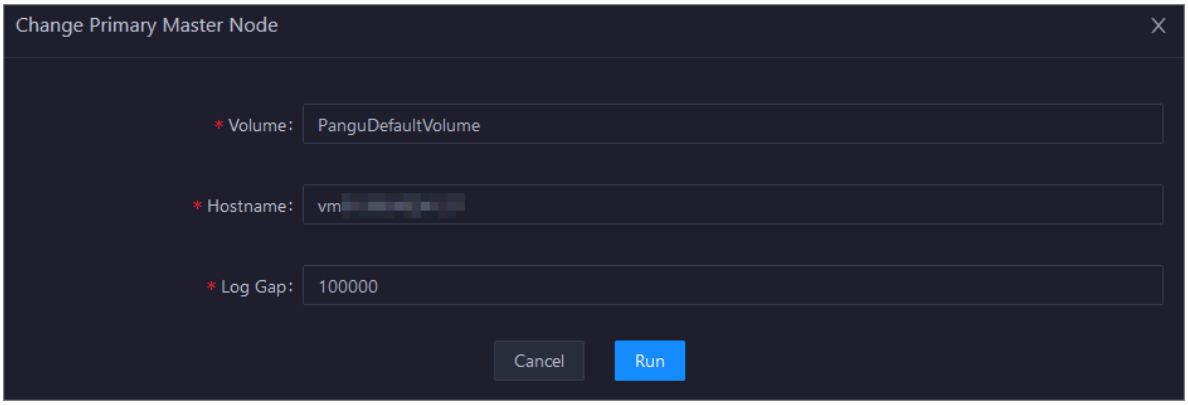

The parameters are described as follows:

- Volume: the volume whose primary master node needs to be changed. Default value: PanguDefaultVolume. If a cluster consists of multiple volumes, set this parameter to the name of the actual volume whose primary master node needs to be changed.
- Hostname: the hostname of the secondary master node that is to be the new primary master node.
- Log Gap: the maximum log number gap between the original primary and secondary master nodes to be switched. During the switchover, the system checks the log number gap between the original primary and secondary master nodes. If the gap is less than the specified value, switchover is allowed. Otherwise, you cannot change the primary master node. Default value: 100000.

6. Click Run. A message appears, indicating that the action has been submitted. Then, the Change Primary Master Node dialog box appears.

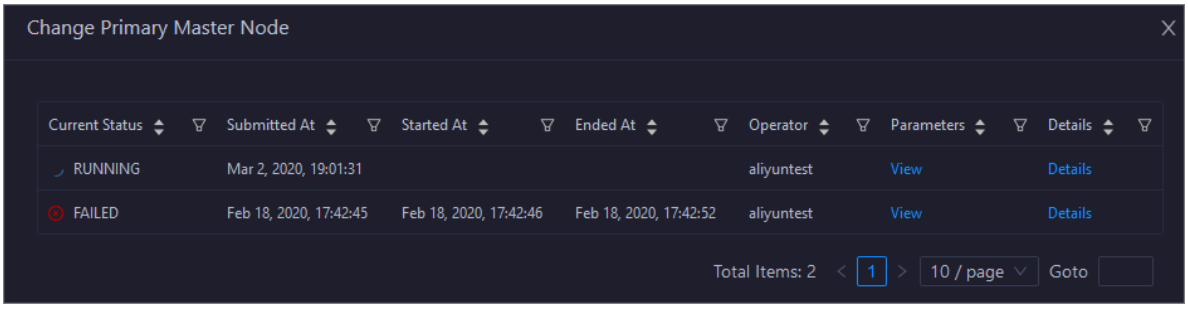

The Change Primary Master Node dialog box displays the switchover history. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

7. If the status is FAILED, click Details in the Details column to locate the failure cause.

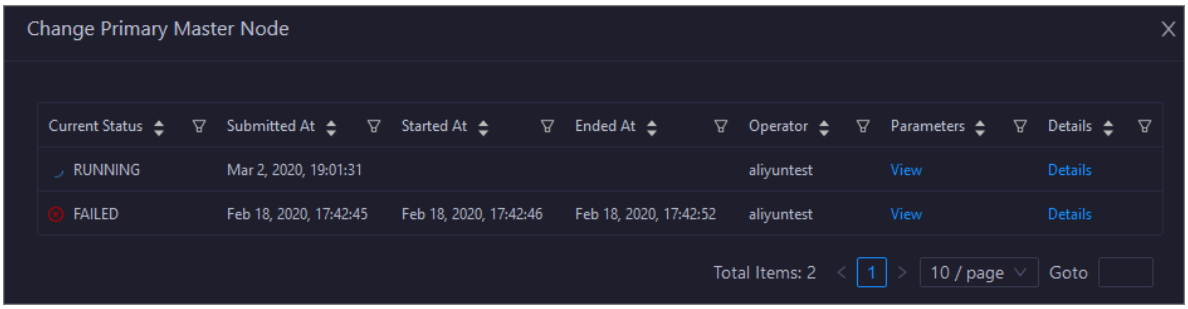

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.5.3.3.7 Empty the recycle bin of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to empty the recycle bin of Apsara Distributed File System.

Prerequisites

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

Procedure

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.

2. Choose Actions > Empty Recycle Bin in the upper-left corner.

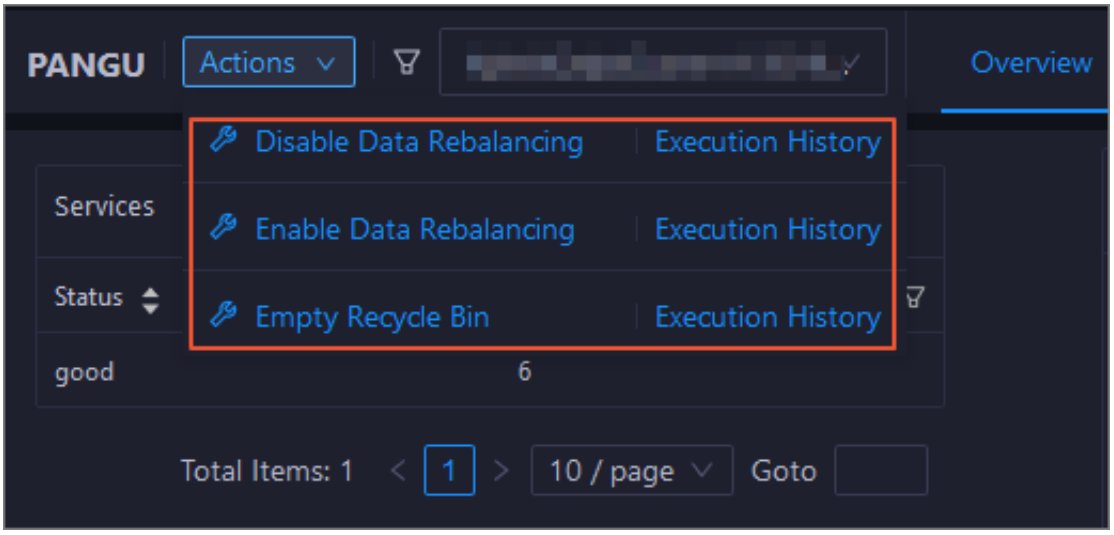

3. In the dialog box that appears, set the Volume parameter. The default value is PanguDefaultVolume.

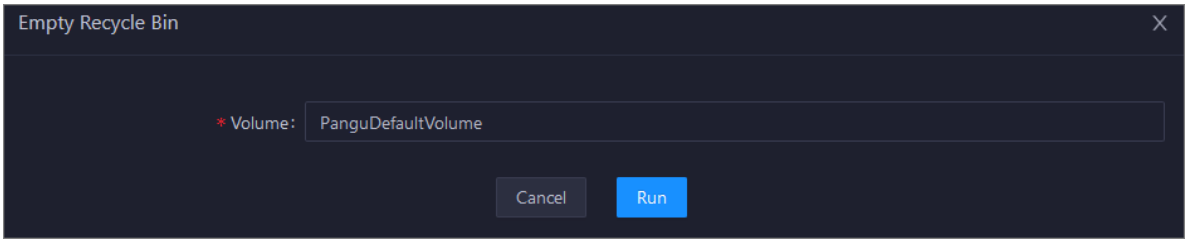

4. Click Run. A message appears, indicating that the action has been submitted.
5. View the execution status.

   Move the pointer over Actions in the upper-left corner, and then click Execution History next to Empty Recycle Bin to view the execution history.

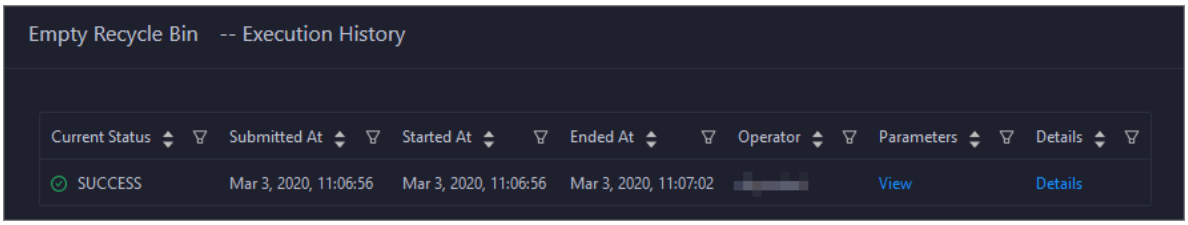

   In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. If the status is FAILED, click Details in the Details column to locate the failure cause.

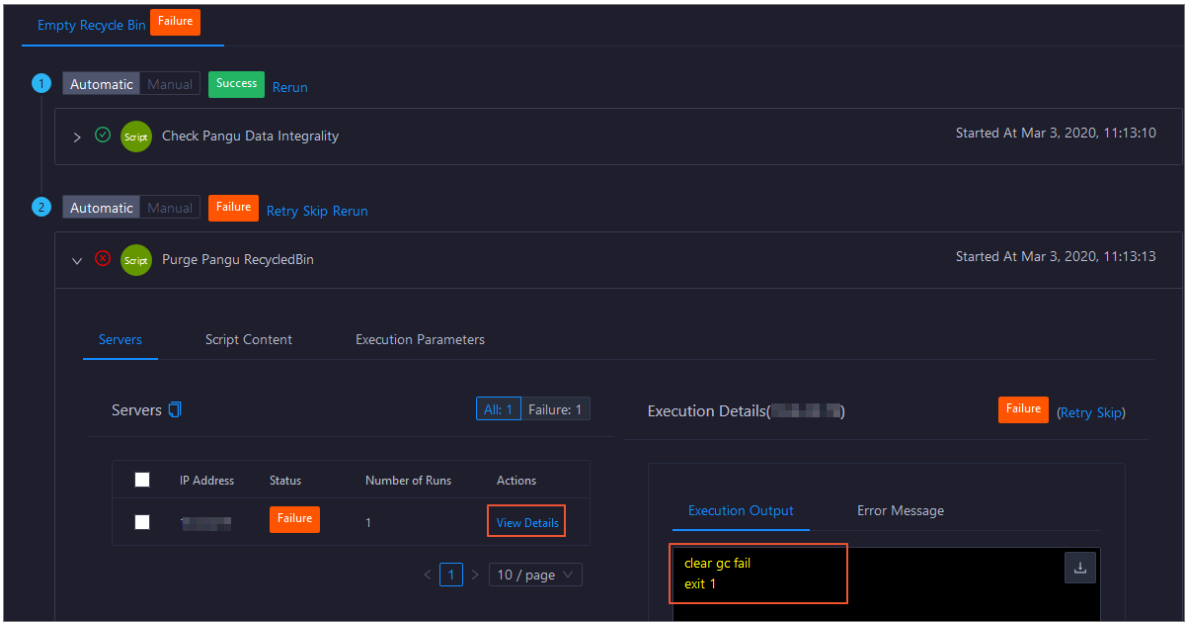

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.5.3.3.8 Enable or disable data rebalancing for Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to enable or disable data rebalancing for Apsara Distributed File System.

Prerequisites

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

Disable data rebalancing

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.

2. Choose Actions > Disable Data Rebalancing in the upper-left corner.

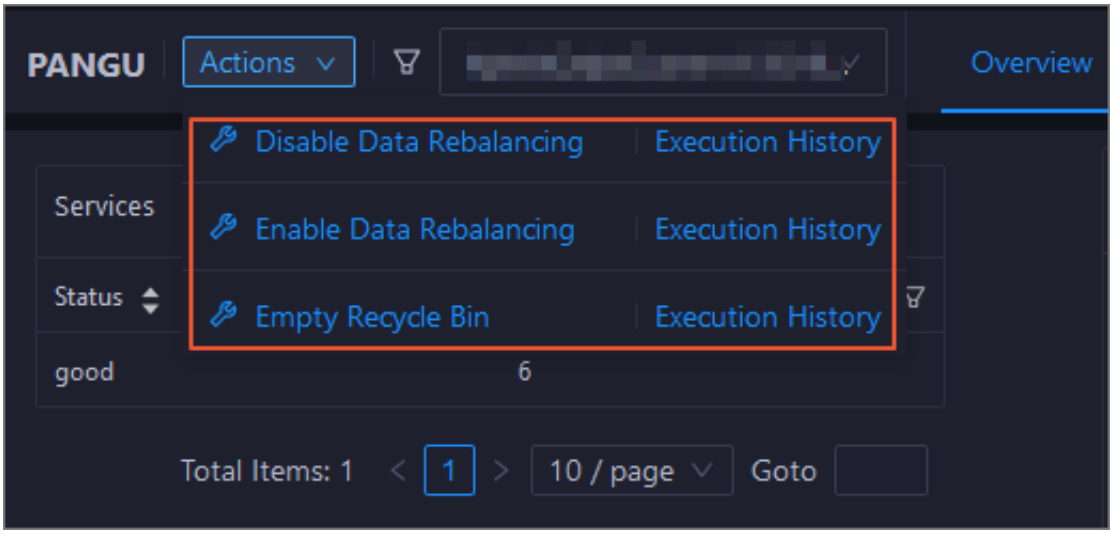

3. In the dialog box that appears, set the Volume parameter. The default value is PanguDefaultVolume.

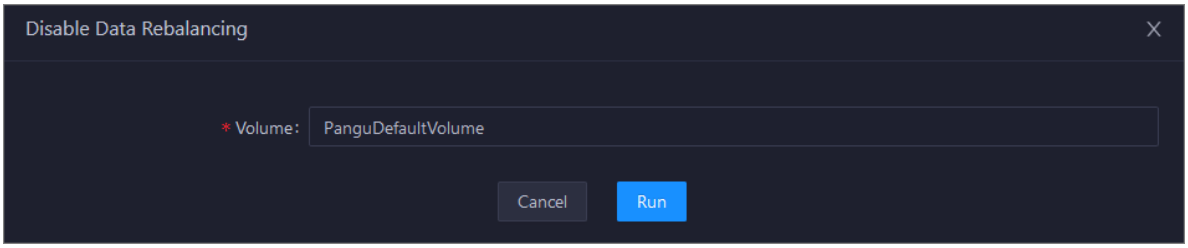

4. Click Run. A message appears, indicating that the action has been submitted.
5. View the execution status.

   Move the pointer over Actions, and then click Execution History next to Disable Data Rebalancing to view the execution history.

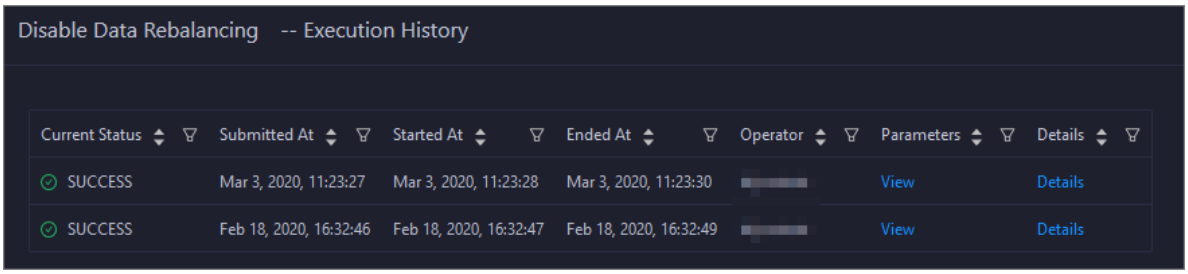

   In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.
6. If the status is FAILED, click Details in the Details column to locate the failure cause. For more information, see *Locate the failure cause*.

Enable data rebalancing

1. **On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.**
2. **Choose Actions > Enable Data Rebalancing in the upper-left corner.**

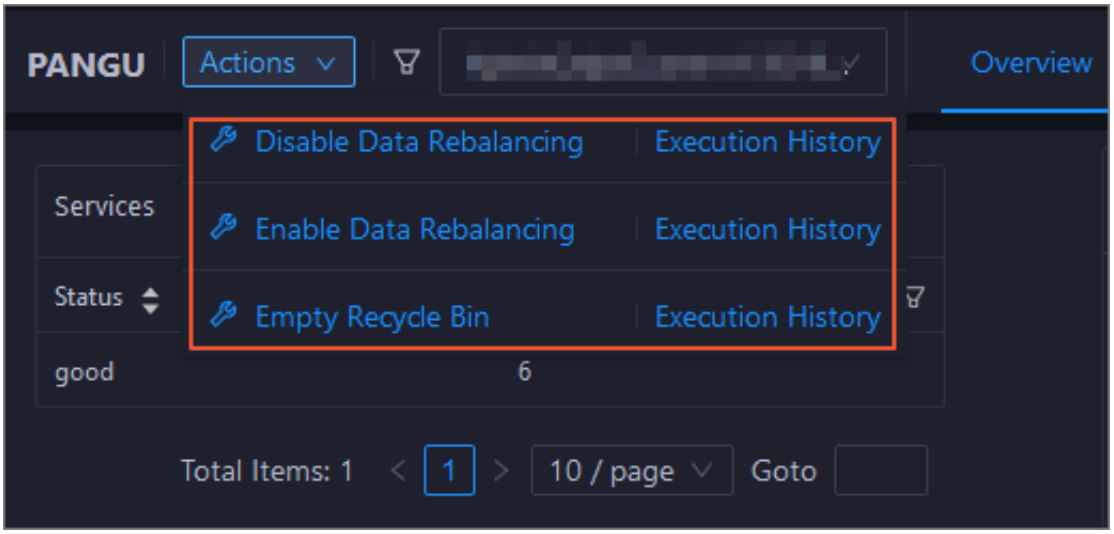

3. **In the dialog box that appears, set the Volume parameter. The default value is PanguDefaultVolume.**

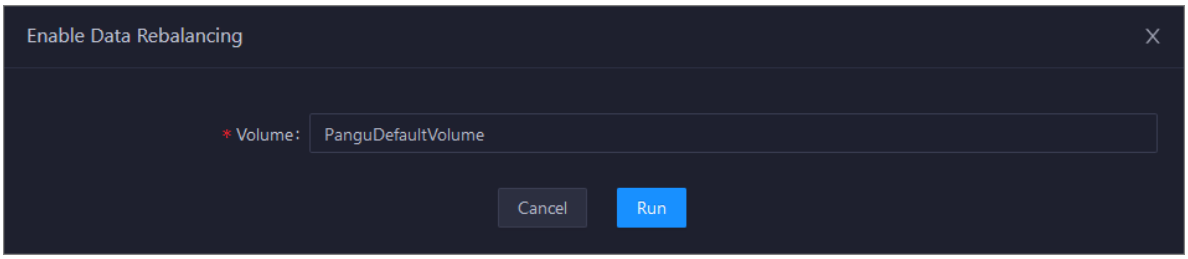

4. **Click Run. A message appears, indicating that the action has been submitted.**
5. **View the execution status.**

   Move the pointer over Actions, and then click Execution History next to Enable Data Rebalancing to view the execution history.

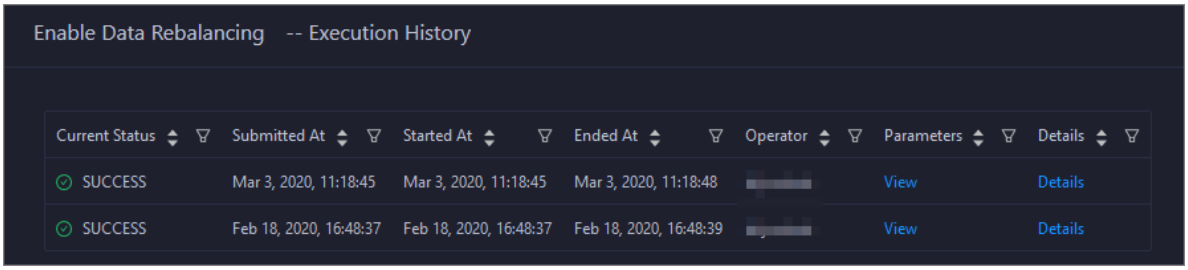

   In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. If the status is FAILED, click Details in the Details column to locate the failure cause. For more information, see *Locate the failure cause*.

Locate the failure cause

This section uses the procedure of locating the failure cause for enabling data reblancing as an example.

1. In the execution history dialog box, click Details in the Details column for a failed execution.
2. In the dialog box that appears, click View Details for a failed step to locate the failure cause.

   You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.5.3.3.9 Run a checkpoint on a master node of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to run checkpoints on master nodes of Apsara Distributed File System. This operation writes memory data into disks. When a failure occurs to Apsara Distributed File System, you can use checkpoints to restore data to the status before the failure. This guarantees data consistency.

Prerequisites

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

Procedure

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Pangu in the left-side navigation pane. Select a cluster from the drop-down list, and then click the Instances tab.

5. In the Master Status section of the Instances page, find the target master node and choose Actions > Run Checkpoint on Master Node in the Actions column. In the dialog box that appears, set the required parameter.

   Volume: the volume of the master node on which a checkpoint is run. Default value: PanguDefaultVolume.

6. Click Run. A message appears, indicating that the action has been submitted. Then, the Run Checkpoint on Master Node dialog box appears.

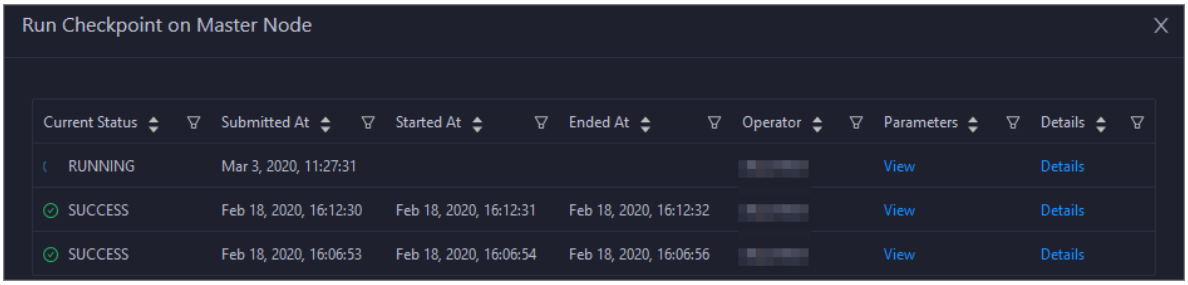

   The Run Checkpoint on Master Node dialog box displays the execution history of the checkpoint on the master node. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

7. If the status is FAILED, click Details in the Details column to locate the failure cause.

   You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.5.3.4 Tunnel service

1.5.3.4.1 Tunnel service O&M overview

This topic describes the concept and O&M features of the tunnel service, and how to access the tunnel service O&M page.

What is the tunnel service?

The tunnel service serves as the data tunnel of MaxCompute. You can use this service to upload data to or download data from MaxCompute.

Modules

- Overview page: displays the information about the tunnel service, including service overview and service status. On this page, you can also view the throughput trend chart.

- **Instances page: displays the information about the service roles of the tunnel service.**
- **Restart Tunnel Server action: allows you to restart one or more tunnel servers.**

Entry

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the page that appears, click O&M in the upper-right corner, and then click the Services tab.**
4. **On the Services page, click Tunnel Service in the left-side navigation pane. The Overview page for the tunnel service appears.**

1.5.3.4.2 Tunnel service overview

The Overview page for the tunnel service displays key operation metrics of the tunnel service, including service overview, service status, and throughput.

Entry

On the Services page, click Tunnel Service in the left-side navigation pane. The Overview page for the tunnel service appears.

The Overview page displays key operation metrics of the tunnel service, including service overview and service status. On this page, you can also view the throughput trend chart.

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

Roles

This section displays all service roles of the tunnel service and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Tunnel Throughput

This chart displays the trend lines of the inbound traffic and outbound traffic in the tunnel service by minute over time in different colors.

1.5.3.4.3 Tunnel service instances

The Instances page displays the information about the tunnel service roles, including the name, host, IP address, status, and host status.

Entry

On the Services page, click Tunnel Service in the left-side navigation pane. Then, click the Instances tab on the right. The Instances page for the tunnel service appears.

The Instances page displays the information about all tunnel service roles, including the name, host, IP address, status, and host status. The host statuses include good, bad, and upgrading.

Other operations

You can filter or sort service roles based on a column to facilitate information retrieval on the Instances page. For more information, see *Common operations*.

1.5.3.4.4 Restart tunnel servers

Apsara Bigdata Manager (ABM) allows you to restart tunnel servers for the corresponding service roles of the tunnel service.

Prerequisites

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

Context

You can restart one or more tunnel servers at a time on the Instances tab of the Tunnel Service page.

Step 1: Restart tunnel servers

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click MaxCompute.**
3. **On the page that appears, click O&M in the upper-right corner, and then click the Services tab.**
4. **On the Services page, click Tunnel Service in the left-side navigation pane, and then click the Instances tab.**
5. **On the Instances page, select one or more service roles for which you want to restart the tunnel servers, and then choose Actions > Restart Tunnel Server in the upper-left corner.**

6. **In the Restart Tunnel Server dialog box that appears, set relevant parameters.**

   The following table describes the required parameters.

Parameter	Description
Force Restart	Specifies whether to forcibly restart the tunnel server for the selected service role. Valid values: · no_force: does not forcibly restart the tunnel server. If a service role is in the running state, the corresponding tunnel server is not restarted. · force: forcibly restarts the tunnel server. A tunnel server is restarted no matter which state the corresponding service role is in.
Hostname	The hostname of the selected service role. Multiple hostnames are separated with commas (,). The value is automatically filled. You cannot specify a value for this parameter.

7. **Click Run. A message appears, indicating that the action has been submitted.**

Step 2: View the execution status or progress

1. **On the Overview tab or the Instances tab of the Tunnel Service page, click Actions in the upper-left corner. Then, click Execution History next to Restart Tunnel Server to view the execution history.**

   In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.
2. **Click Details in the Details column of a task in the RUNNING state to view the execution progress.**

Step 3: Optional. Locate the failure cause

If the status of the task is FAILED, you can view the execution logs to locate the failure cause.

1. **On the Overview tab or the Instances tab of the Tunnel Service page, click Actions in the upper-left corner. Then, click Execution History next to Restart Tunnel Server to view the execution history.**
2. **In the execution history dialog box, click Details in the Details column of the task to view the details.**

3. On the Servers tab of the failed step, click View Details in the Actions column of a failed server. The Execution Output tab appears in the Execution Details section. You can view the output to locate the failure cause.

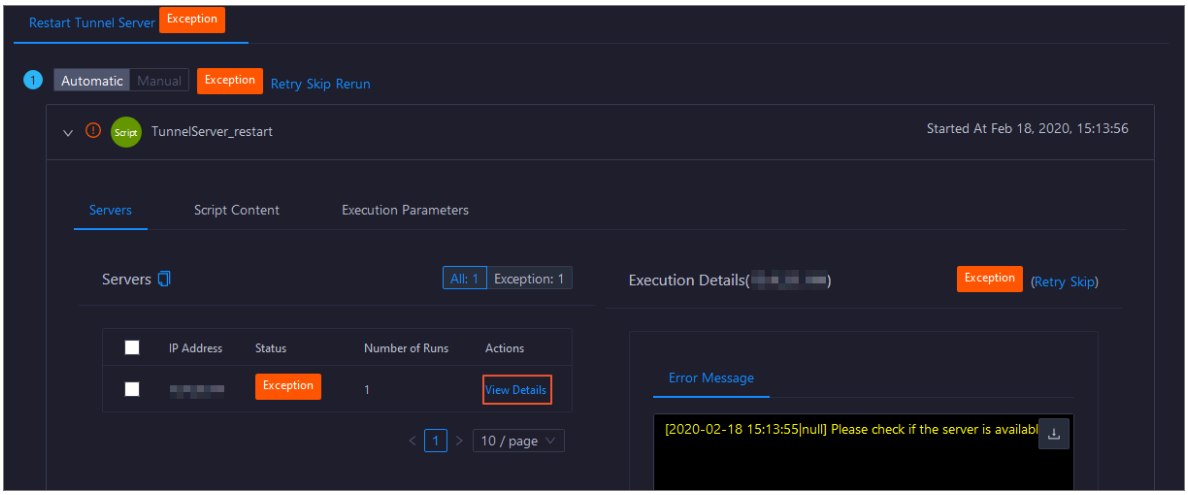

1.5.4 Cluster O&M

1.5.4.1 Cluster O&M overview

This topic describes the cluster O&M features of MaxCompute supported by Apsara Bigdata Manager (ABM) and how to access the MaxCompute cluster O&M page.

Cluster O&M features

ABM supports the following MaxCompute cluster O&M features:

- **Overview page: displays the overall running information about a cluster. On this page, you can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster. In the Log on section, you can click the name of the host whose role is pangu master, fuxi master, or odps ag to log on to the host.**
- **Health Status page: displays 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 .**
- **Servers page: displays the information about hosts in a cluster, including the hostname, IP address, role, type, CPU usage, memory usage, root disk usage, packet loss rate, and packet error rate.**

- **Cluster scaling action: allows you to scale out or scale in a MaxCompute cluster by adding or removing physical hosts.**

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click MaxCompute.
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.

1.5.4.2 Cluster overview

This topic describes how to access the Overview page of a MaxCompute cluster, the information on the page, and the operations that you can perform on this page.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click MaxCompute.
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 quickly log on to a host commonly used in MaxCompute cluster O&M. In addition, you can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster.

Log on

In this section, you can log on to a host commonly used in MaxCompute cluster O&M whose role is pangu master, fuxi master, or odps ag.

1. In the Log on section, click the name of a host in the Hostname column. The Hosts page for the host appears.

2. Click the Log On icon of the host in the upper-left corner. The TerminalService page appears.

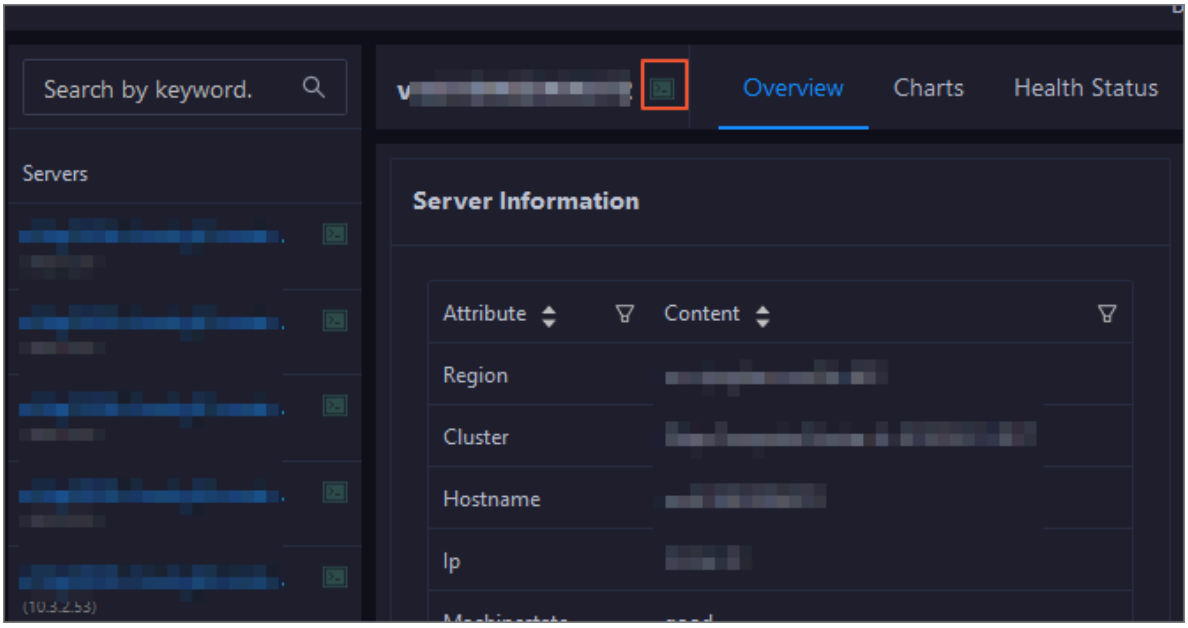

3. Click the hostname in the left-side navigation pane to log on to the host.

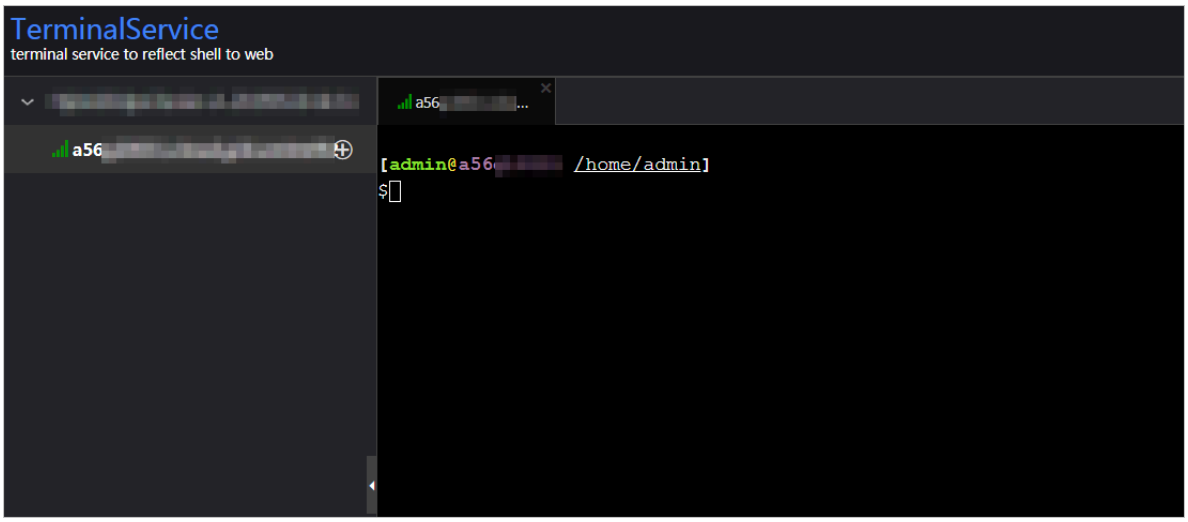

Servers

This section displays all host statuses and the number of hosts in each status. The host statuses include good and bad.

Services

This section displays all services deployed in the cluster and the respective number of services in the good and bad states.

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 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.

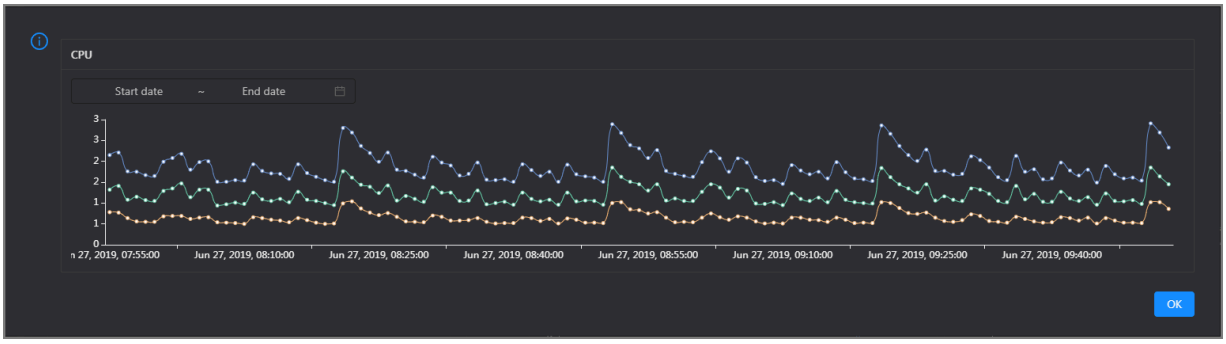

DISK

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

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

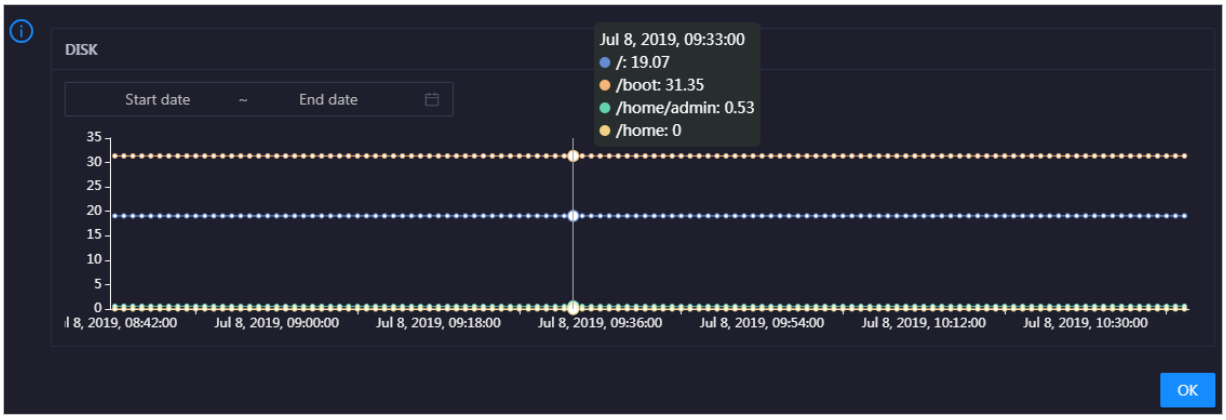

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

LOAD

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

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

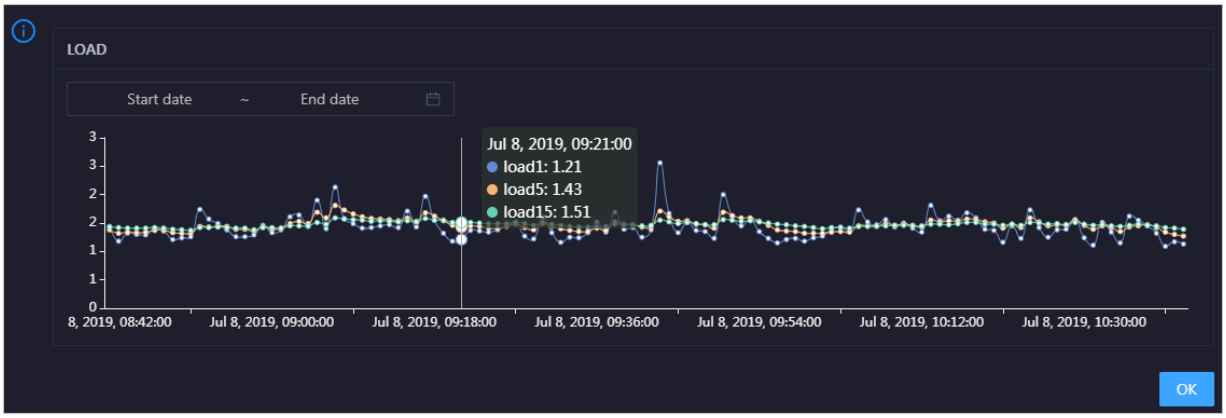

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

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 cluster over time in different colors.

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

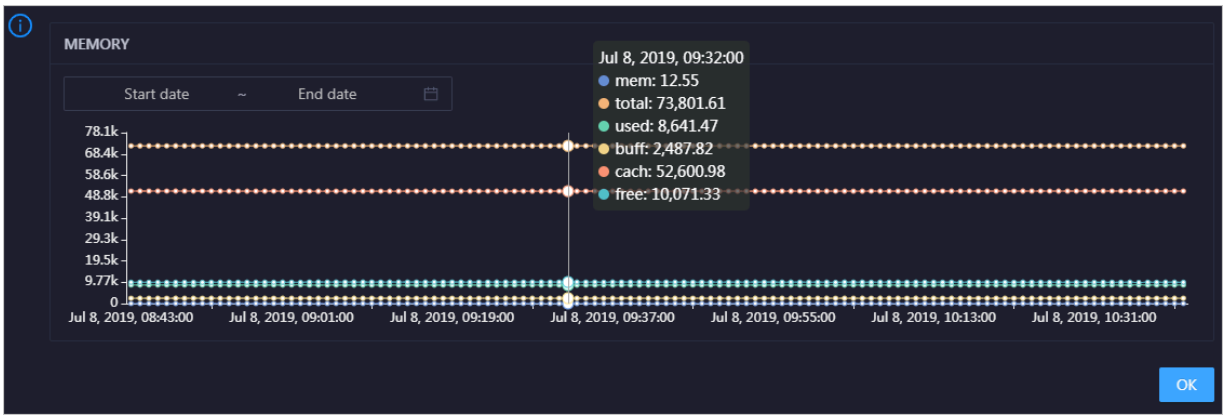

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

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

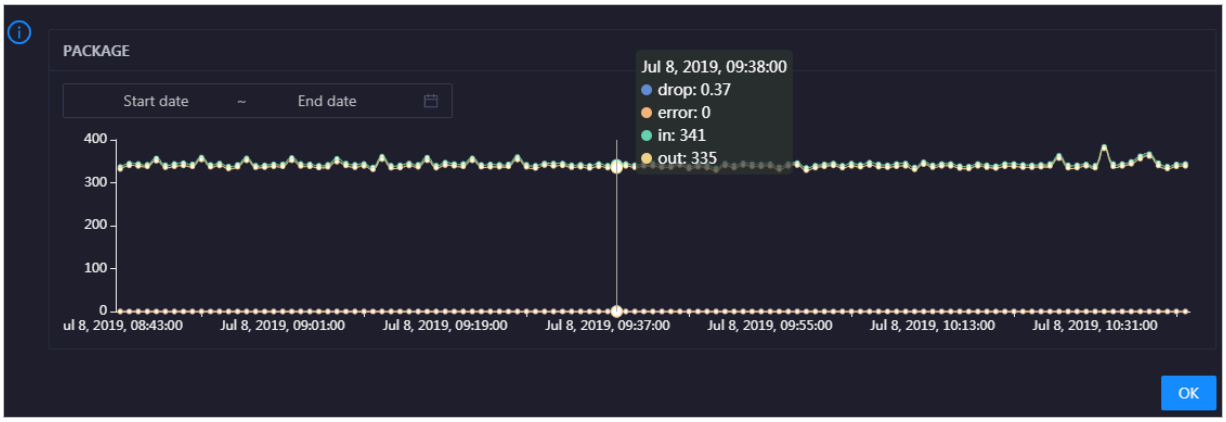

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

Health Check

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

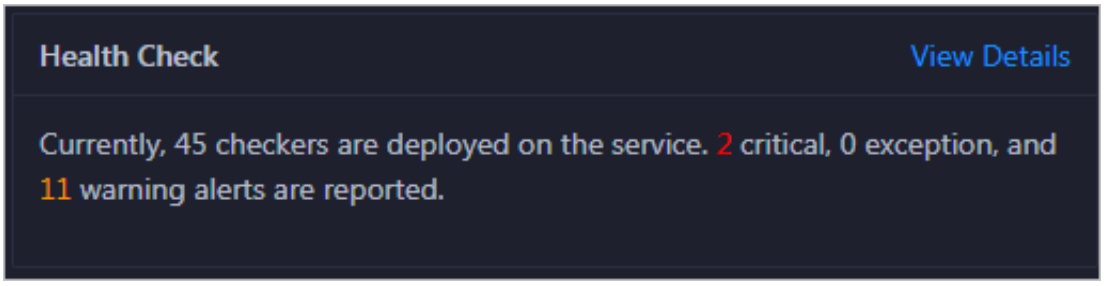

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Cluster health*.

Health Check History

This section displays a record of the health checks performed on the cluster. You can view the respective number of Critical, Warning, and Exception alerts for each health check.

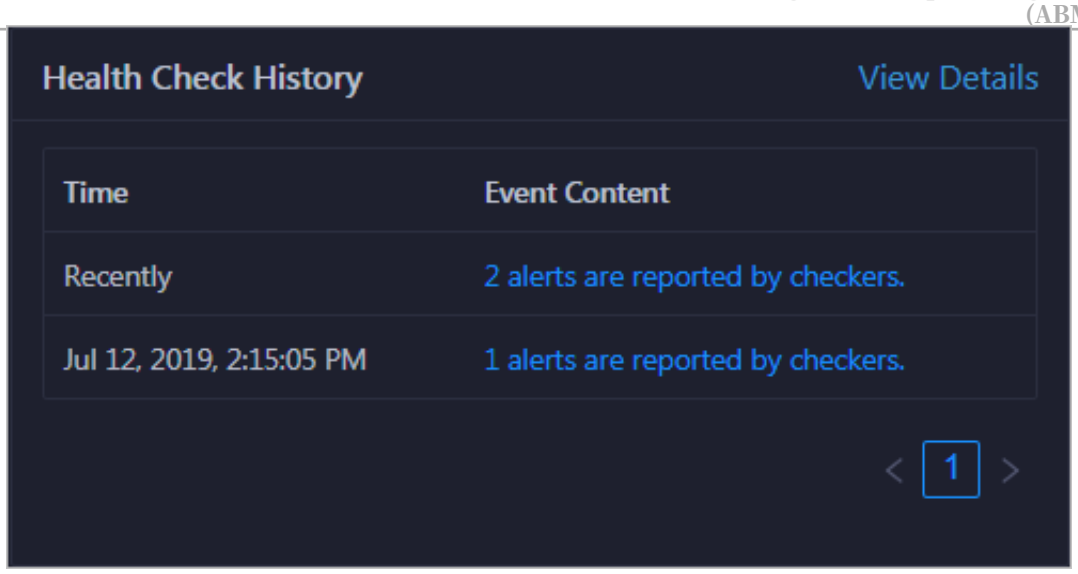

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Cluster health*.

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

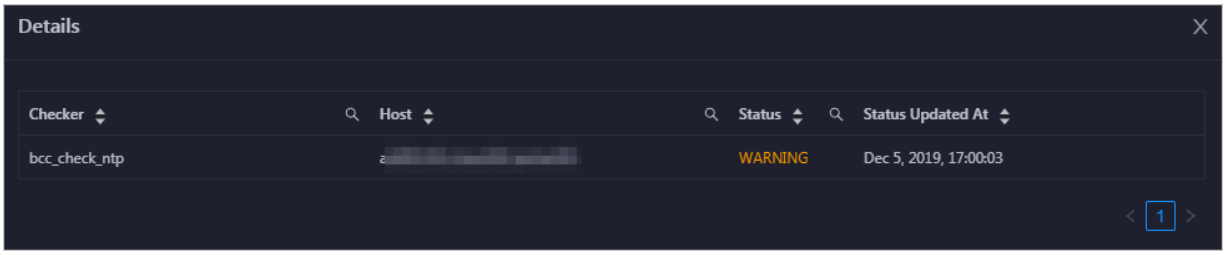

1.5.4.3 Cluster health

On the 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

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click MaxCompute.
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, and then click the Health Status tab. The Health Status page appears.

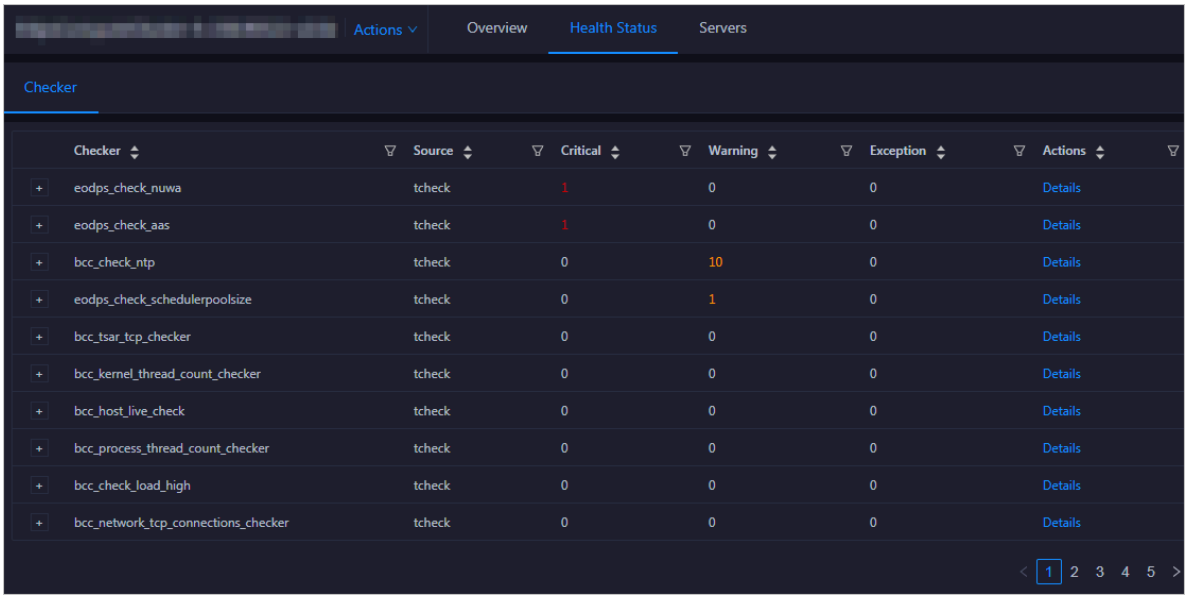

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

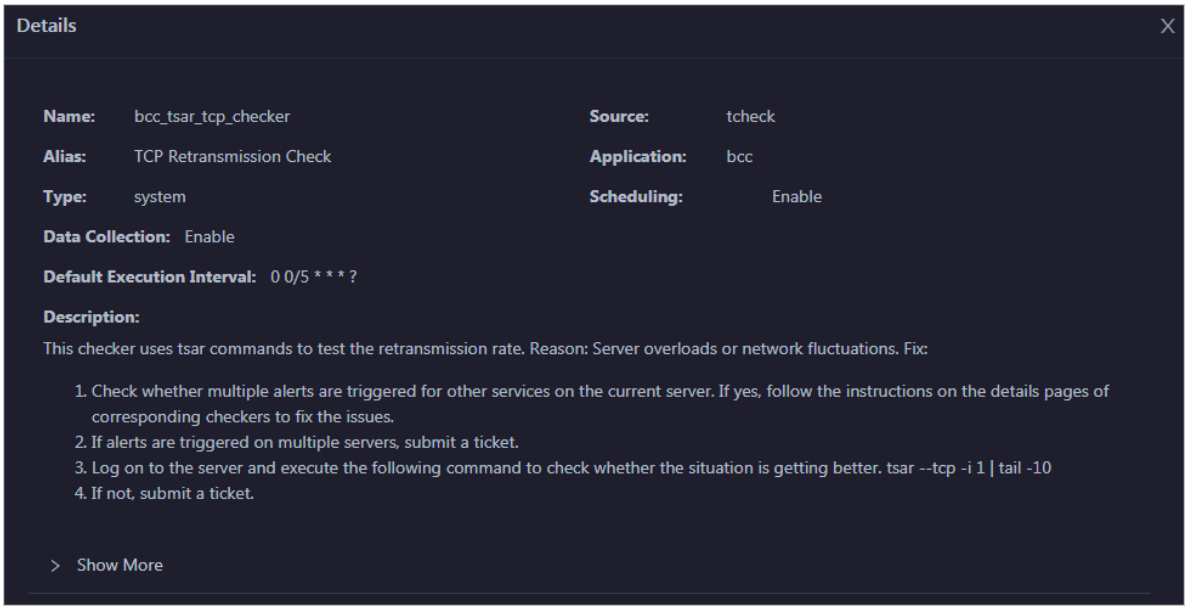

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

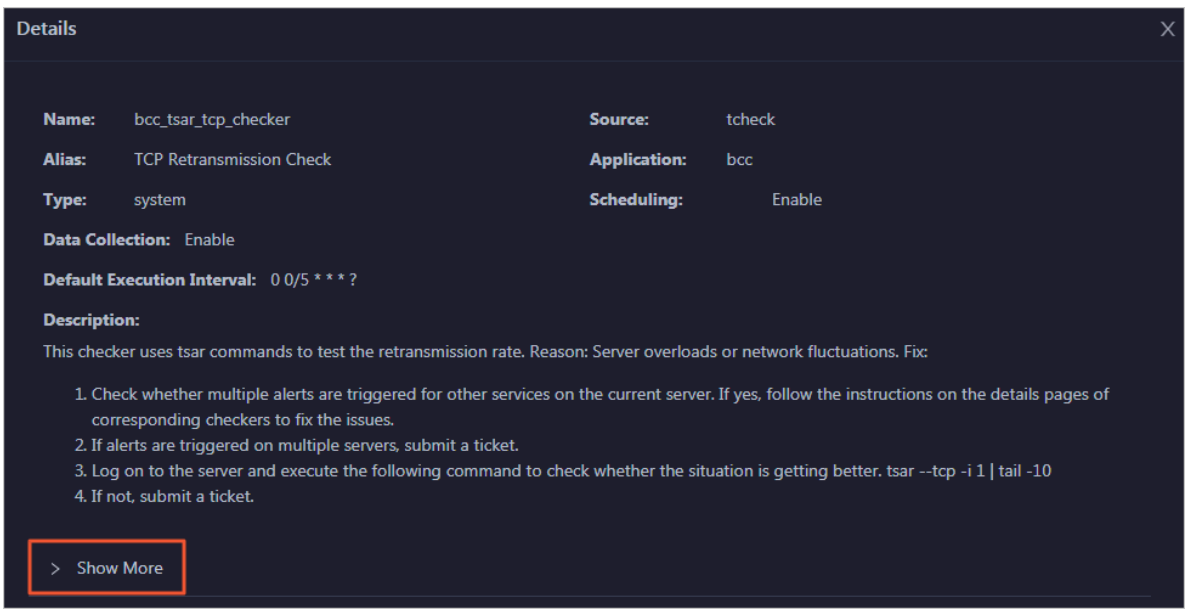

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View hosts with alerts and the alert causes

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

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

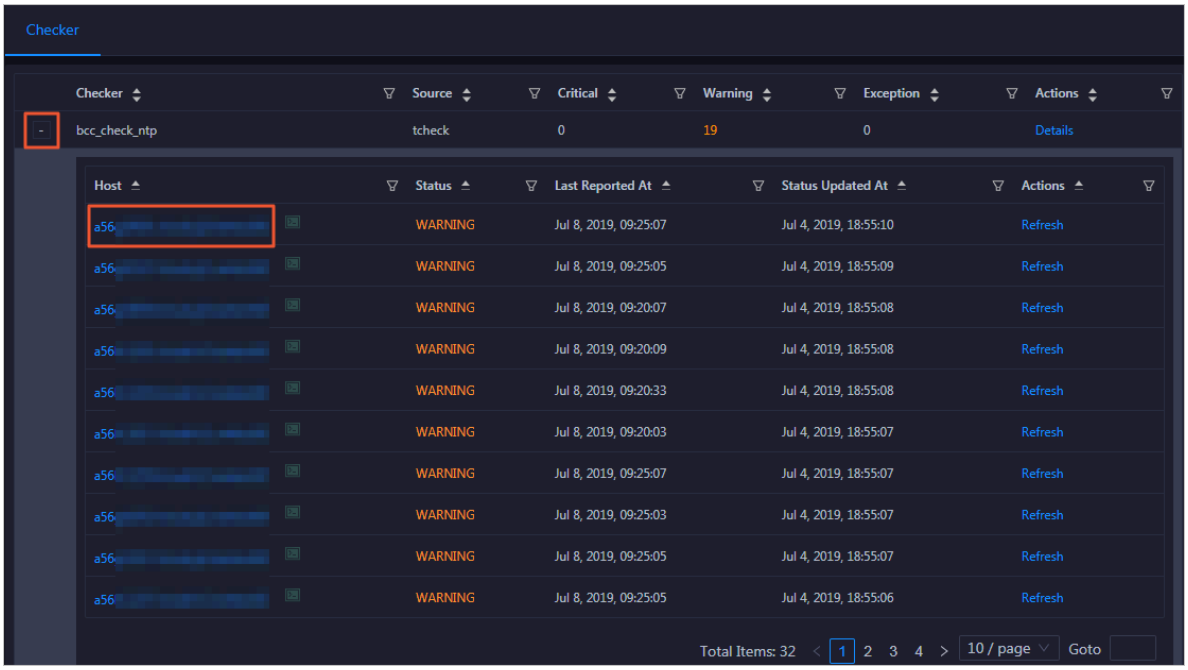

2. **Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

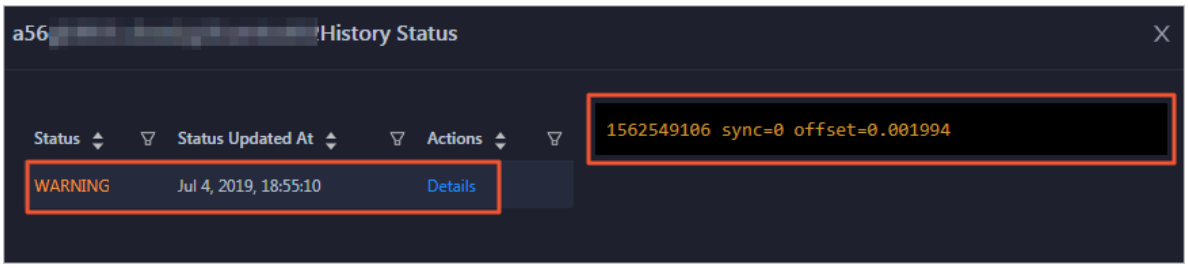

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

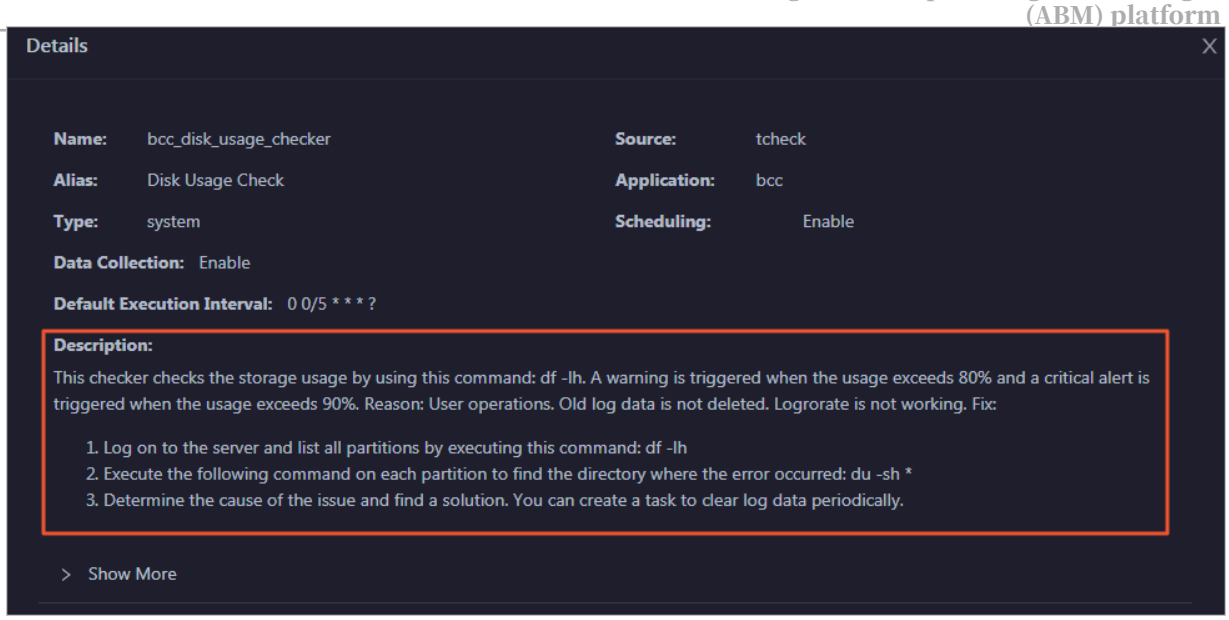

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

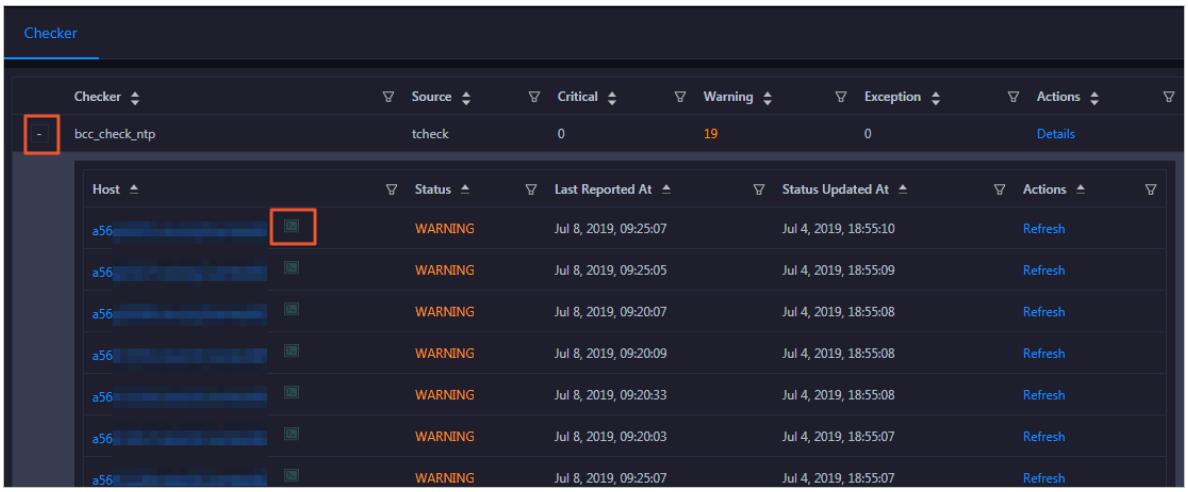

2. Click the Log On icon of a host. The TerminalService page appears.

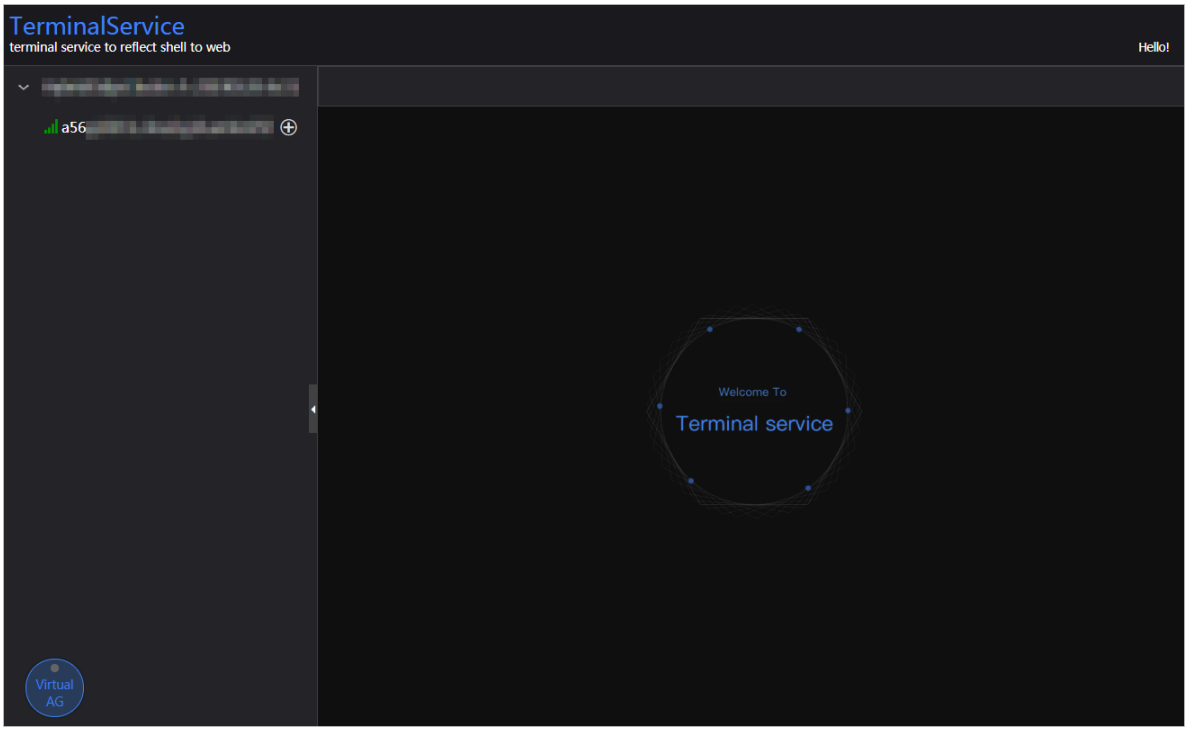

3. On the TerminalService page, click the hostname on the left to log on to the host.

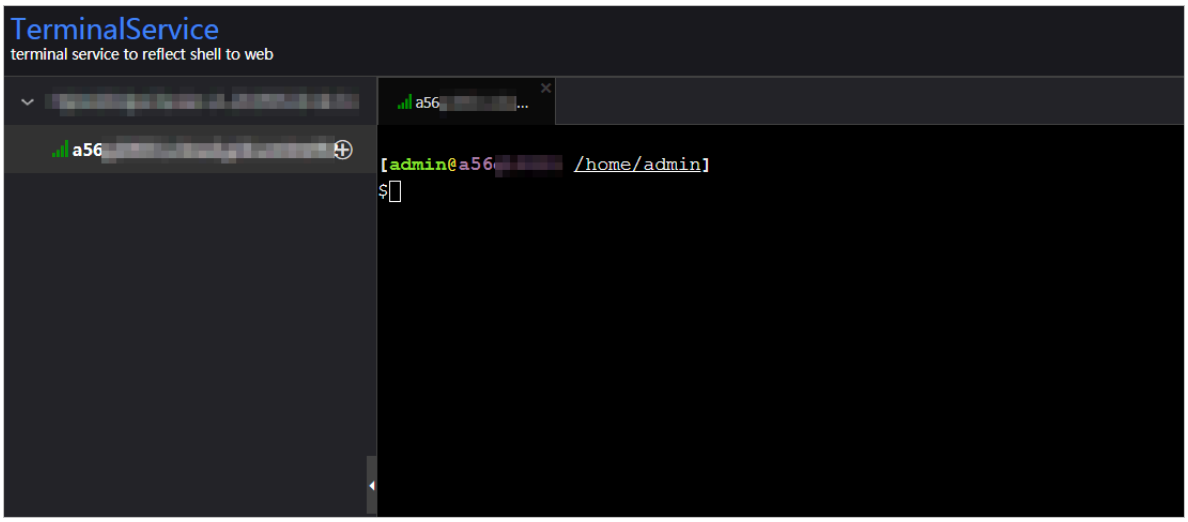

Run a checker again

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. In this way, you can check whether the alert is cleared.

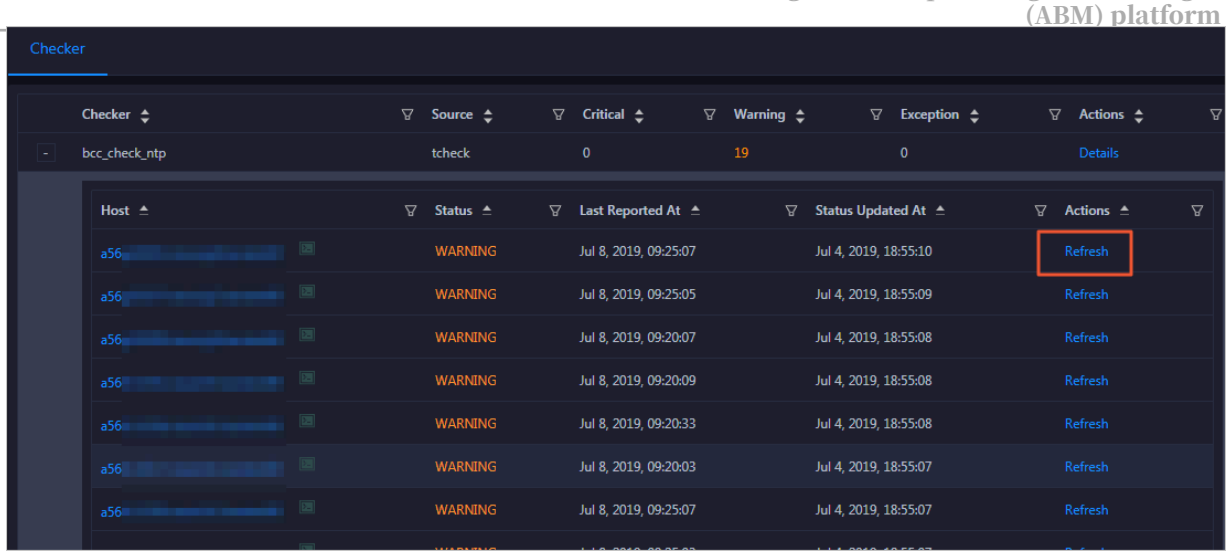

1.5.4.4 Cluster hosts

The cluster host page displays information about hosts, including the hostname, IP address, role, type, CPU usage, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

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

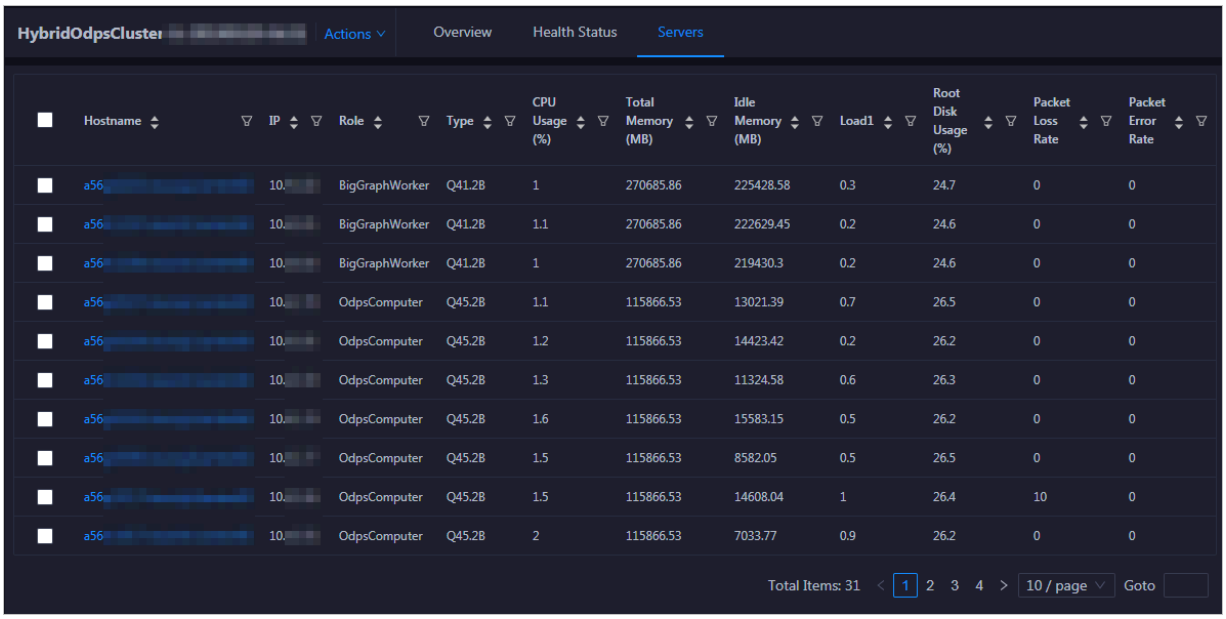

To view more information about a host, click the name of the host. The *Host overview* page appears.

1.5.4.5 Cluster scaling

Apsara Bigdata Manager (ABM) supports MaxCompute cluster scaling. To scale out a MaxCompute cluster, add physical hosts in the default cluster of Apsara

Infrastructure Management Framework to the MaxCompute cluster. To scale in a MaxCompute cluster, remove physical hosts from the MaxCompute cluster to the default cluster of Apsara Infrastructure Management Framework.

Background

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework can be considered as an available resource pool that provides resources for scaling out business clusters. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

In the ABM console, when you scale out a MaxCompute cluster, the system adds physical hosts in the default cluster to the MaxCompute cluster. When you scale in a MaxCompute cluster, the system removes physical hosts from the MaxCompute cluster to the default cluster.

Prerequisites

- **Scale-out: The physical host to be added to a MaxCompute cluster is an SInstance host in the default cluster of Apsara Infrastructure Management Framework.**
- **Scale-out: If you use a host as a template host for scale-out, the host is an SInstance host. You can log on to the admingateway host in the MaxCompute cluster to view SInstance hosts.**
- **Scale-in: The physical host to be removed from a MaxCompute cluster is an SInstance host. You can log on to the admingateway host in the MaxCompute cluster to view SInstance hosts.**

Scale out a MaxCompute cluster

You can add multiple hosts to a MaxCompute cluster at a time to scale out the cluster. To achieve this, you need to specify an existing host as the template host. When you scale out the MaxCompute cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. Log on to the admingateway host in the MaxCompute cluster. Run the `r ttrl` command to view SInstance hosts. For more information about how to log on to a host, see *Log on to a host*.

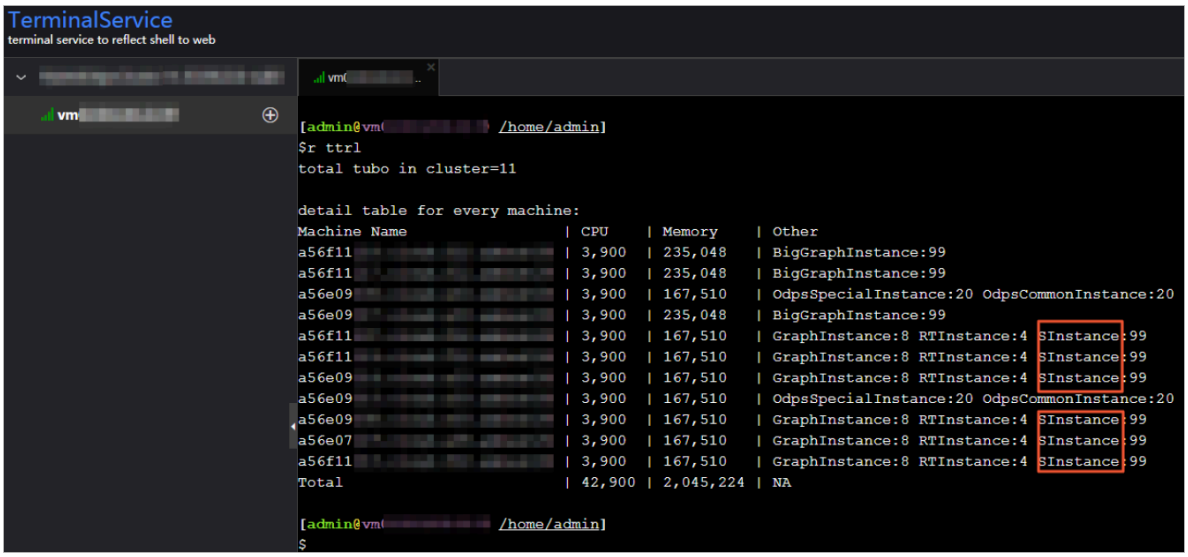

2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Servers tab. On the page that appears, select an SInstance host as the template host.

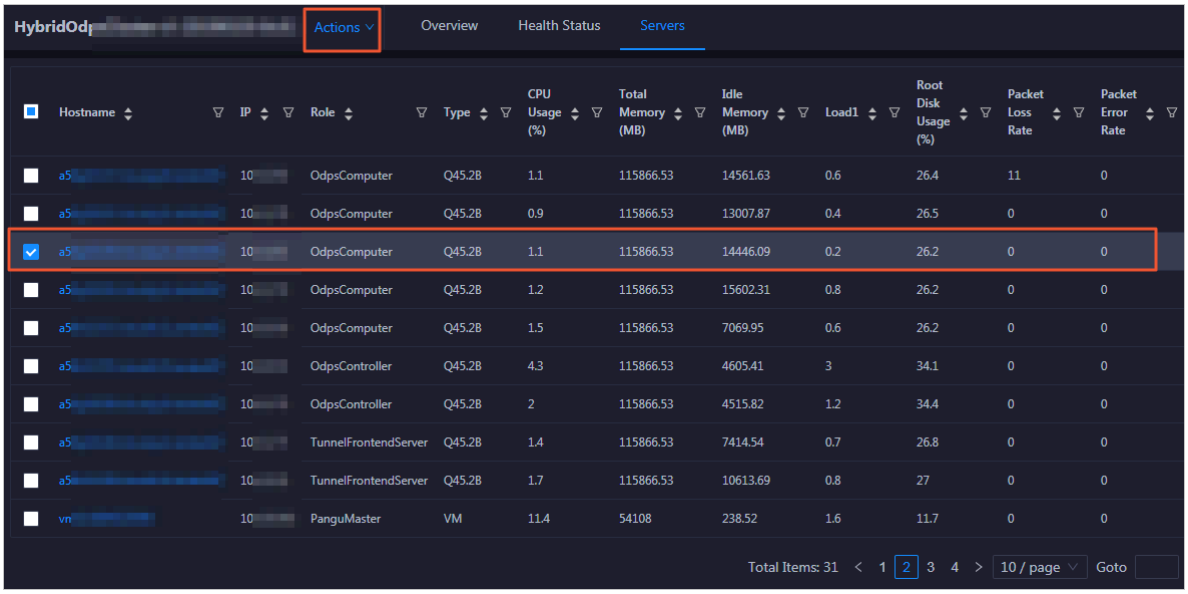

3. Choose Actions > Scale out Cluster in the upper-left corner. In the Scale out Cluster dialog box that appears, set relevant parameters.

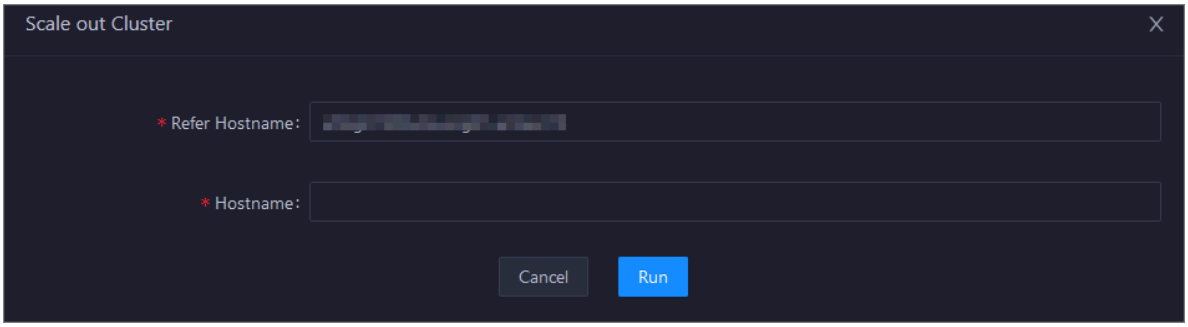

The parameters are described as follows:

- Refer Hostname: the name of the template host. By default, the name of the selected host is used.
- Hostname: the name of the host to be added to the MaxCompute cluster. The drop-down list displays all available hosts in the default cluster for scale-out. You can select one or more hosts from the drop-down list.

4. Click Run. A message appears, indicating that the action has been submitted.
5. View the scale-out status.

   Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale out Cluster to view the scale-out history.

   It may take some time for the cluster to be scaled out. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. If the status is RUNNING, click Details in the Details column to view the steps and progress of the scale-out.

7. If the status is FAILED, click Details in the Details column to locate the failure cause. For more information, see *Locate the failure cause*.

Scale in a MaxCompute cluster

You can remove multiple hosts from a MaxCompute cluster at a time to scale in the cluster.

1. Log on to the admingateway host in the MaxCompute cluster. Run the `r ttrl` command to view SInstance hosts. For more information about how to log on to a host, see *Log on to a host*.

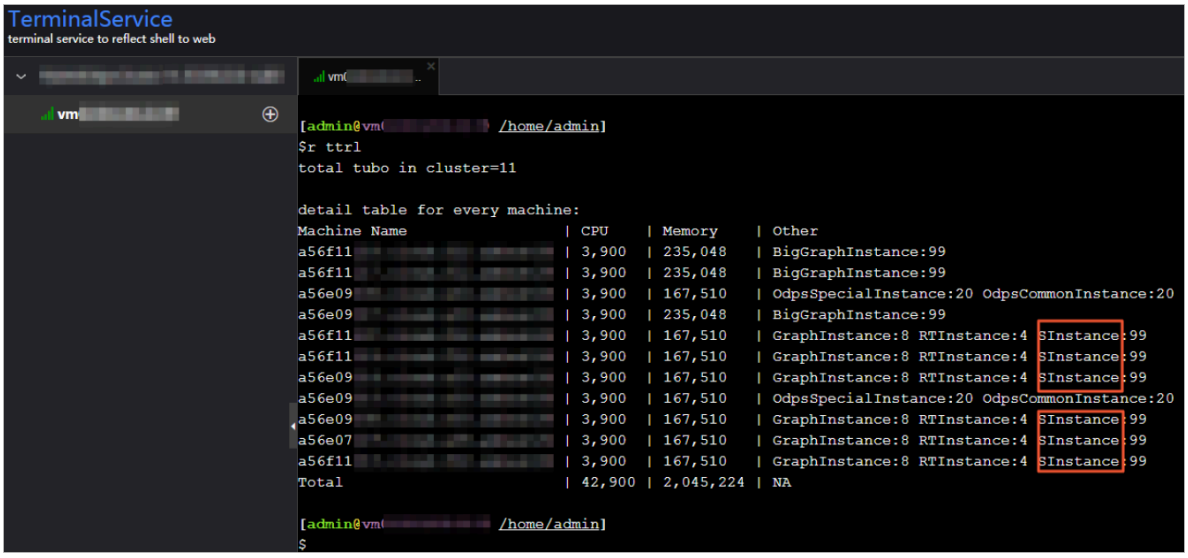

2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Servers tab. On the page that appears, select one or more SInstance hosts to be removed.

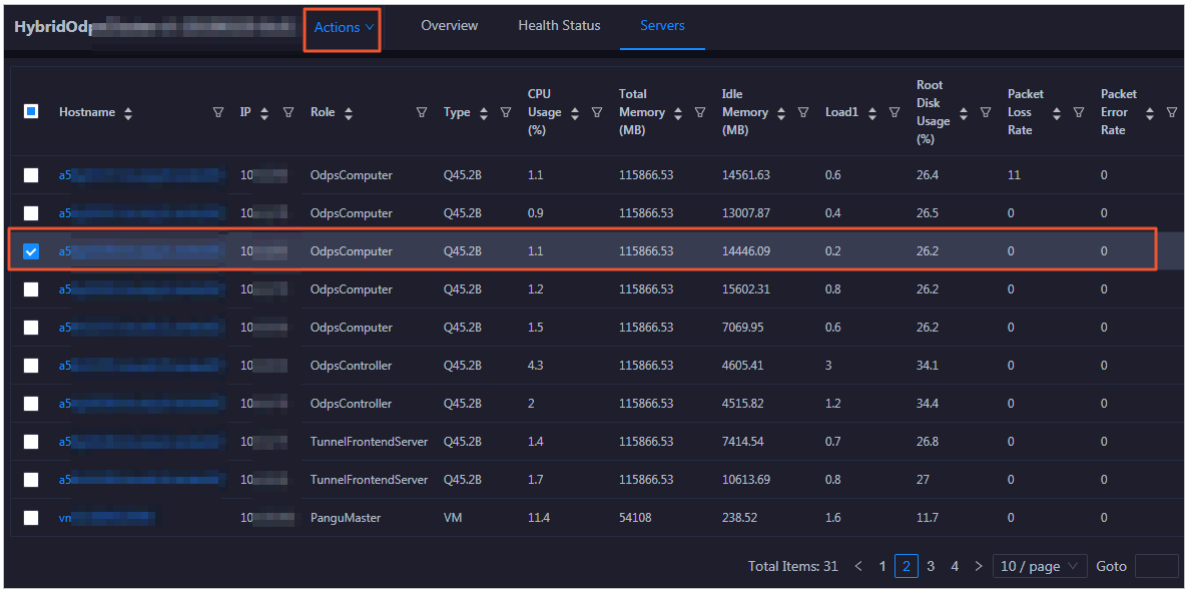

3. Choose Actions > Scale in Cluster in the upper-left corner. In the Scale in Cluster dialog box that appears, set the Hostname parameter.

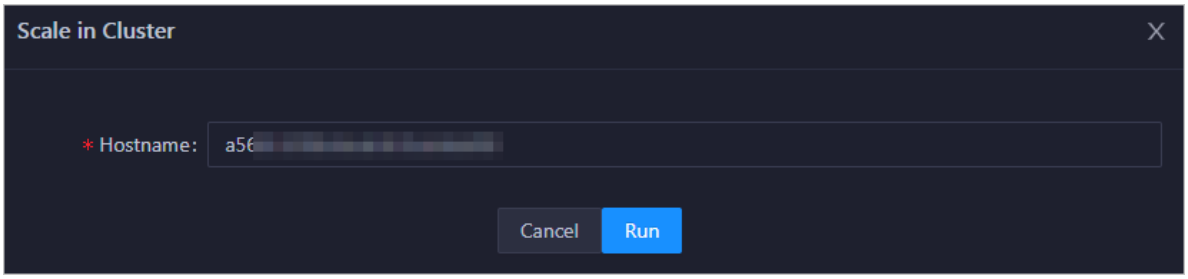

Hostname: the name of the host to be removed from the MaxCompute cluster. By default, the name of the selected host is used.

4. Click Run. A message appears, indicating that the action has been submitted.
5. View the scale-in status.

Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale in Cluster to view the scale-in history.

It may take some time for the cluster to be scaled in. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. If the status is RUNNING, click Details in the Details column to view the steps and progress of the scale-in.

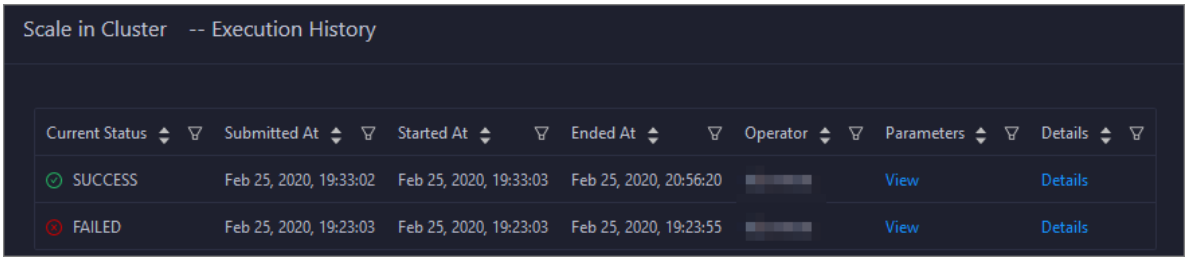

7. If the status is FAILED, click Details in the Details column to locate the failure cause. For more information, see *Locate the failure cause*.

Locate the failure cause

This section uses cluster scale-in as an example to describe how to locate the failure cause.

1. On the Clusters page, move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale in Cluster to view the scale-in history.
2. Click Details in the Details column of a failed execution to locate the failure cause.

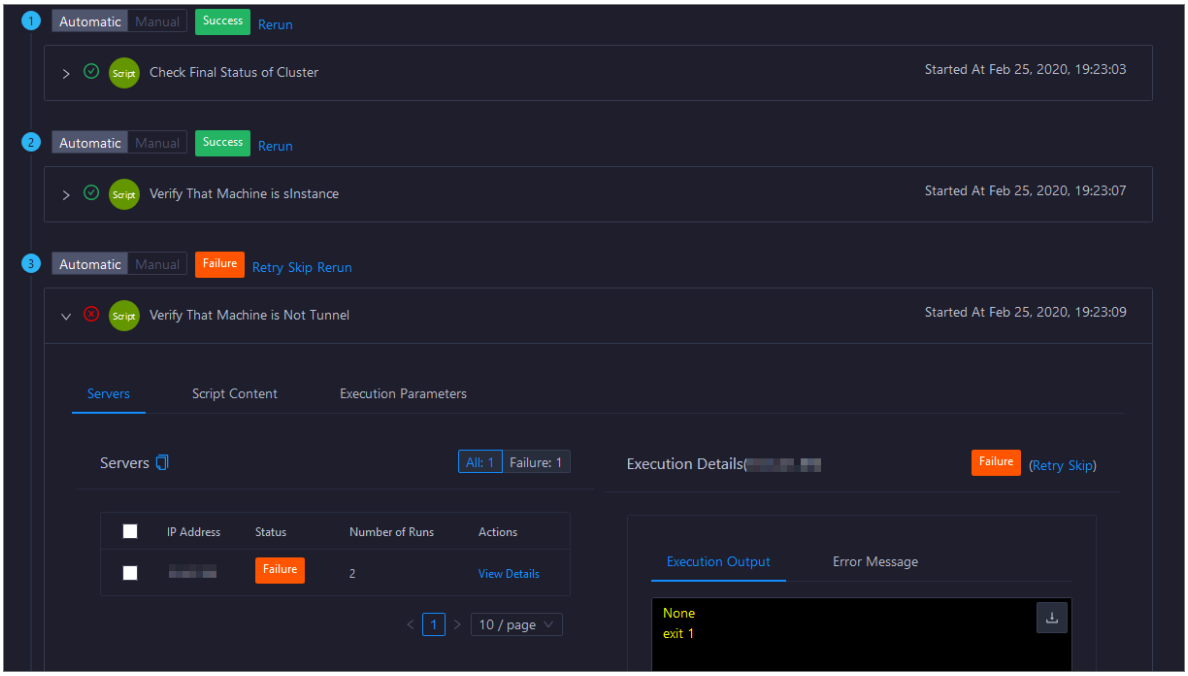

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.5.5 Host O&M

1.5.5.1 Host O&M overview

This topic describes the host O&M features of MaxCompute and how to access the MaxCompute host O&M page.

Host O&M features

- **Overview page: displays brief information about a host in a MaxCompute cluster . On this page, you can view the attributes, services, service roles, 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, and packet transmission for the host .**
- **Charts page: displays the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.**
- **Health Status page: displays the checkers of the selected host, including the checker details, check results for the host, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.**
- **Services page: displays the cluster, service instances, and service instance roles of a host.**

Entry

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

1.5.5.2 Host overview

The host overview page displays brief information about a host in a MaxCompute cluster. On this page, you can view the attributes, services, service roles, 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, and packet transmission for the host.

Entry

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

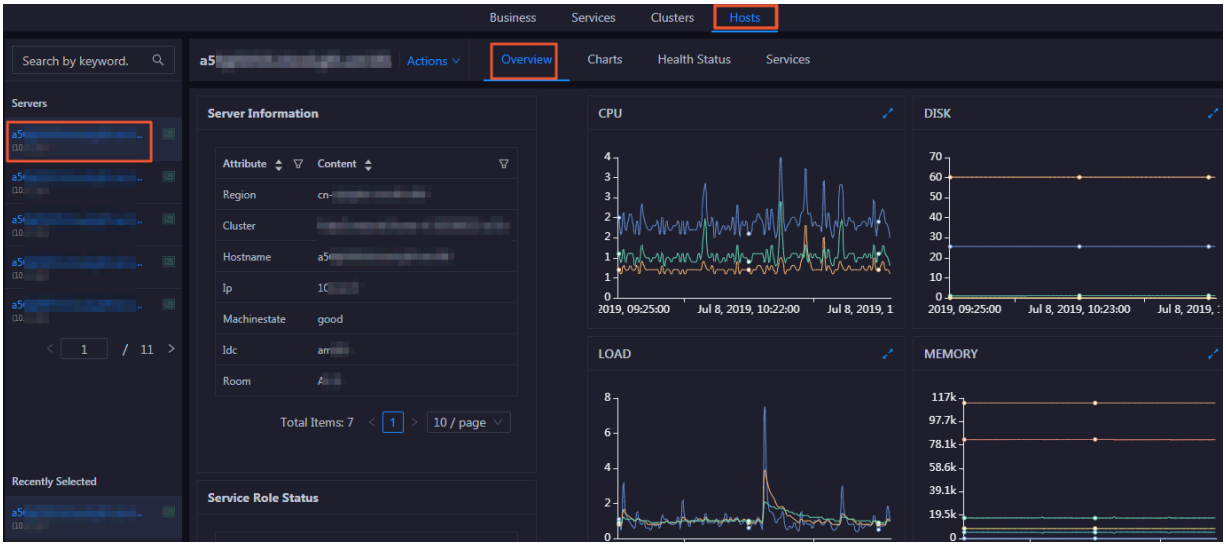

On the Overview page, you can view the attributes, services, service roles, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host.

Server Information

This section displays the information about the host, including the region, cluster, name, IP address, status, Internet data center (IDC), and server room of the host.

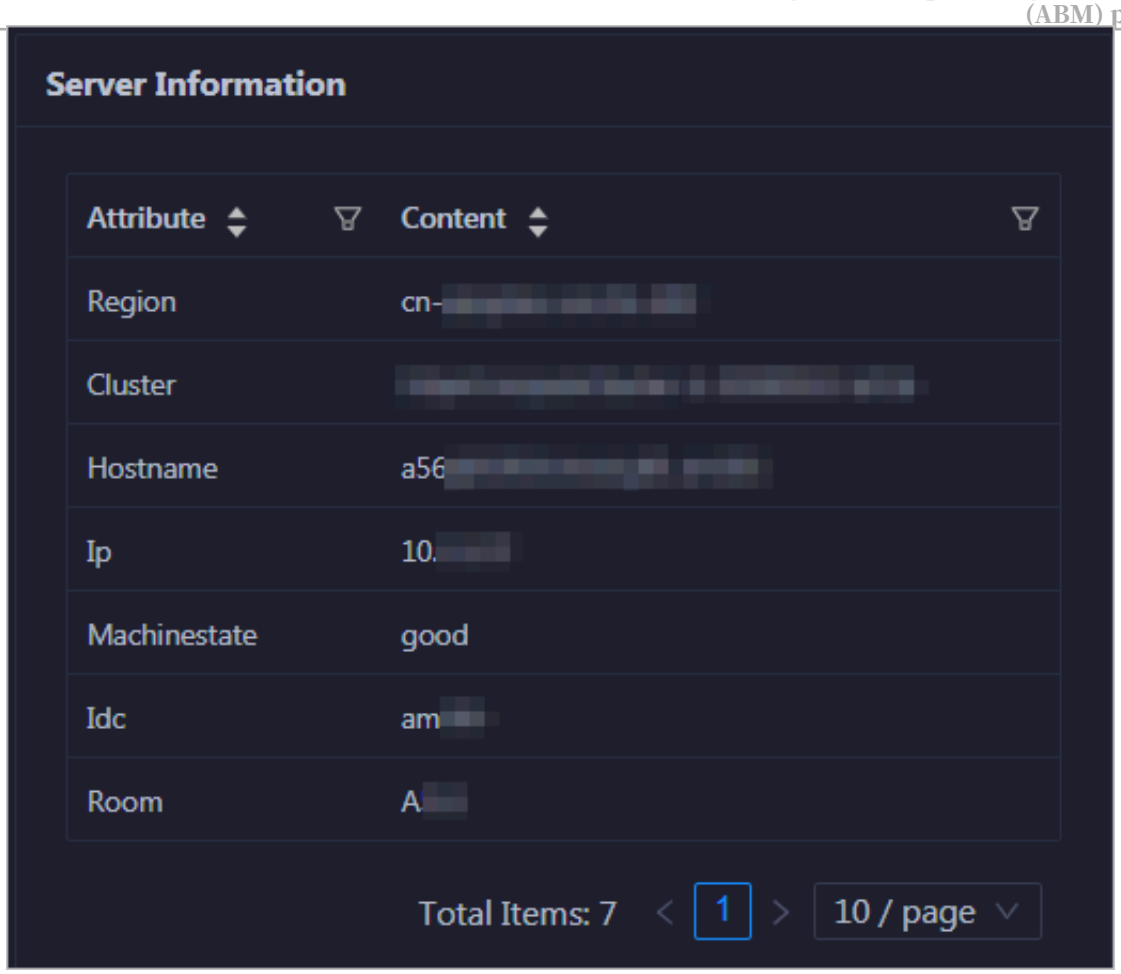

Service Role Status

This section displays the information about the services deployed on the host, including the roles, statuses, and number of services.

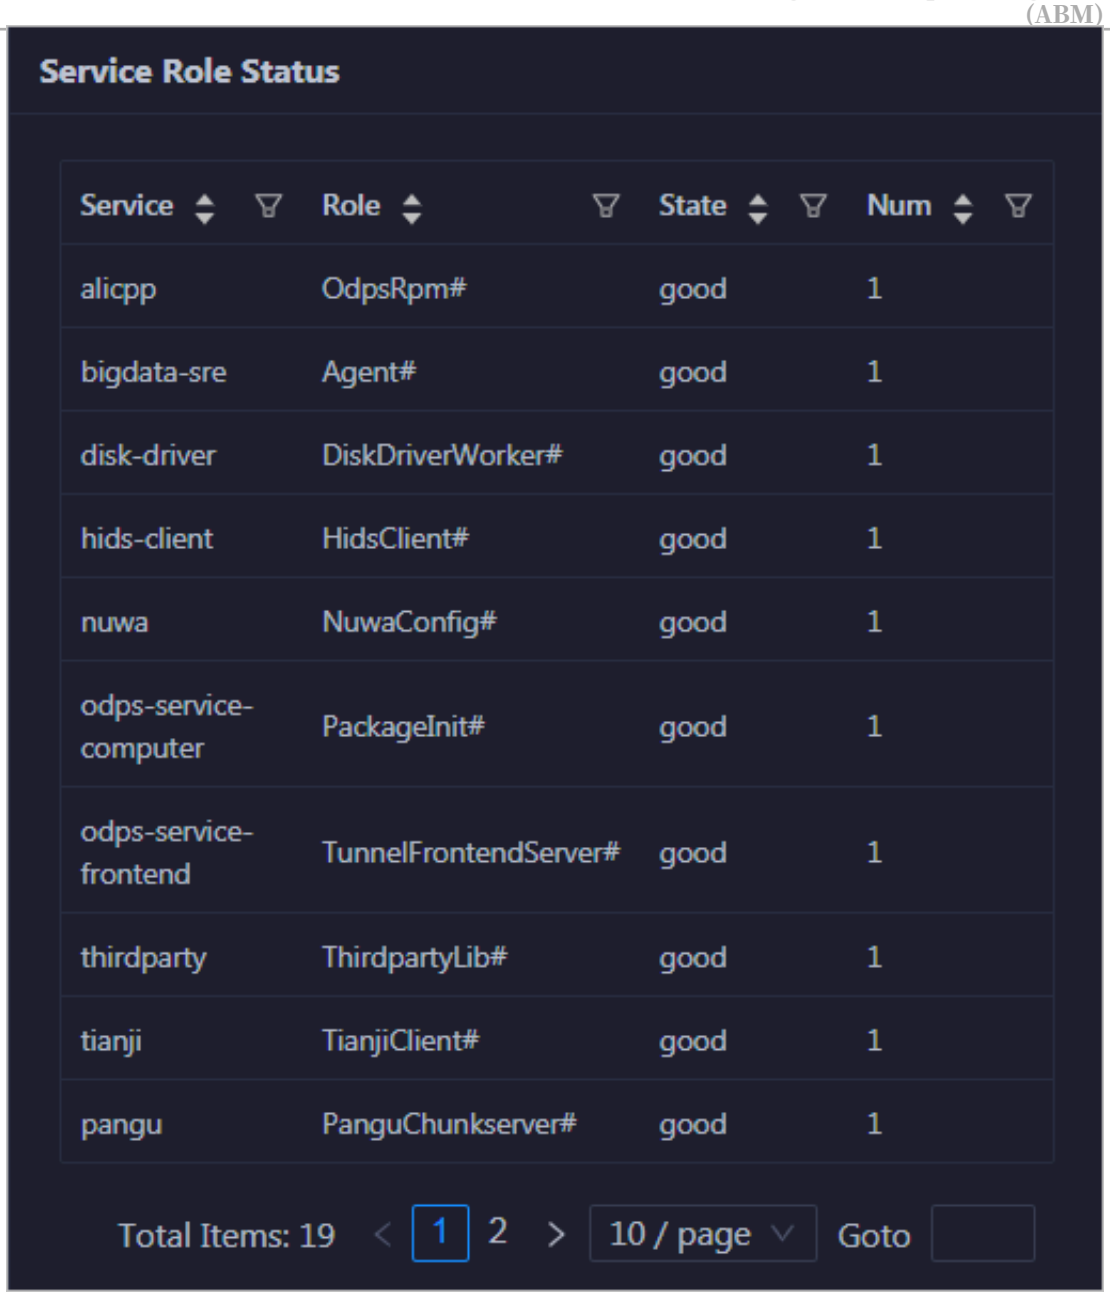

CPU

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 host over time in different colors.

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

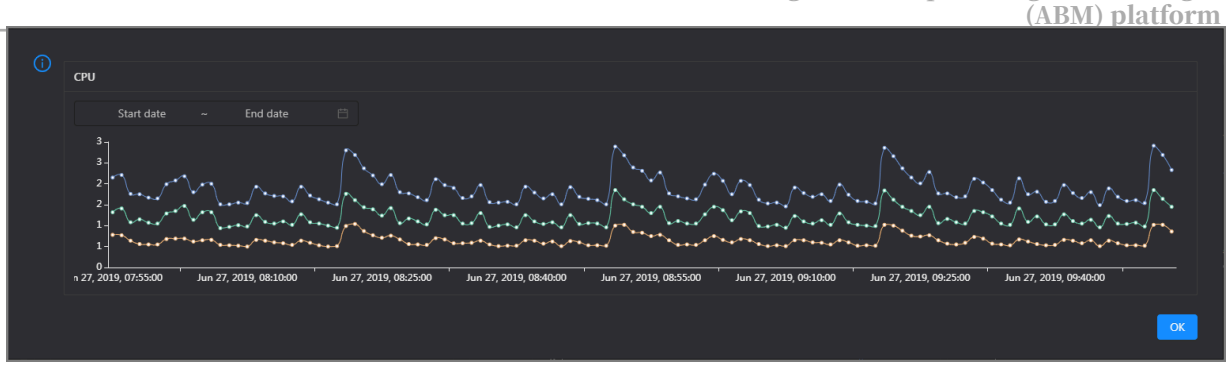

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

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

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

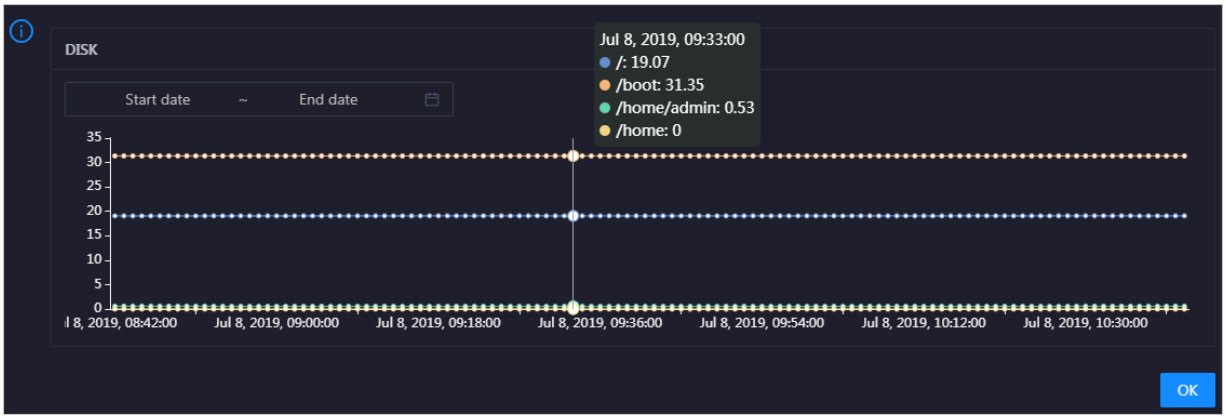

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

LOAD

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

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

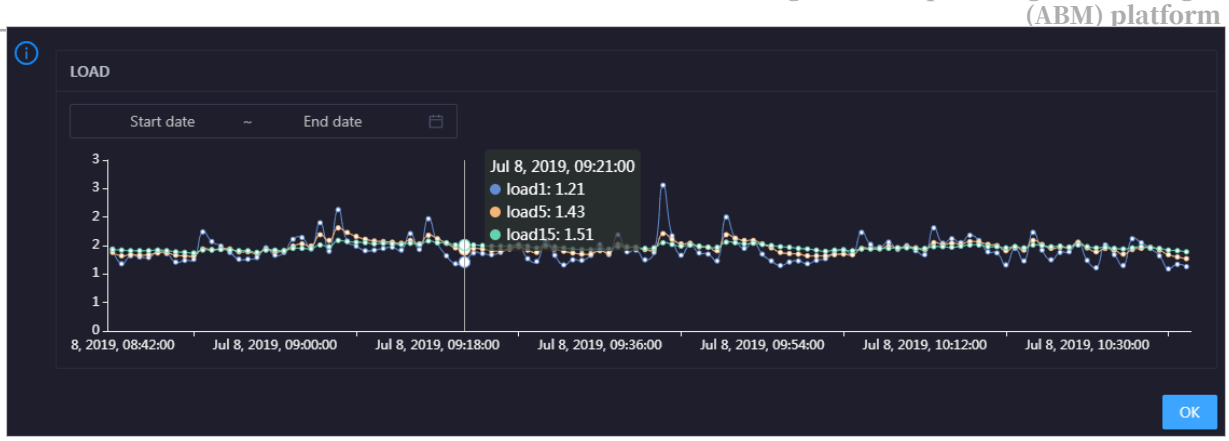

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 host over time in different colors.

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

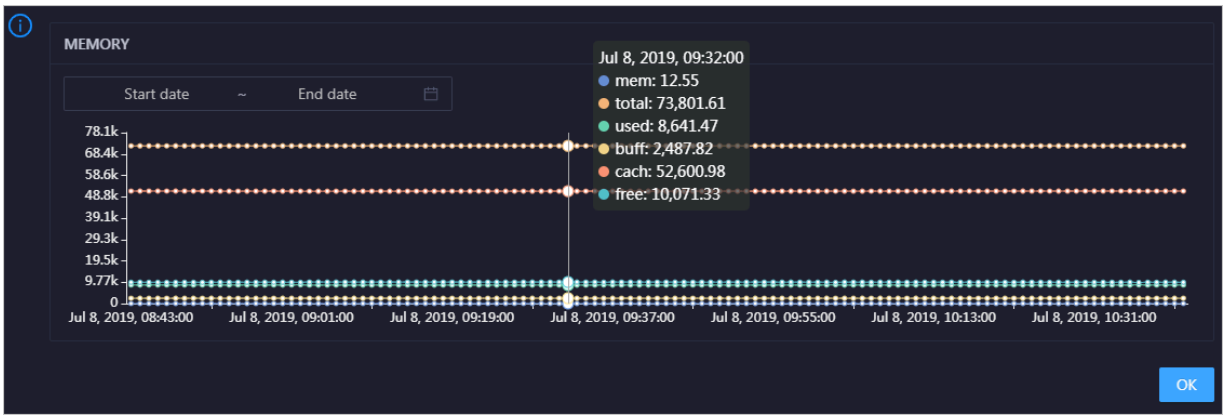

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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 host over time in different colors. These trend lines reflect the data transmissi on status of the host.

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

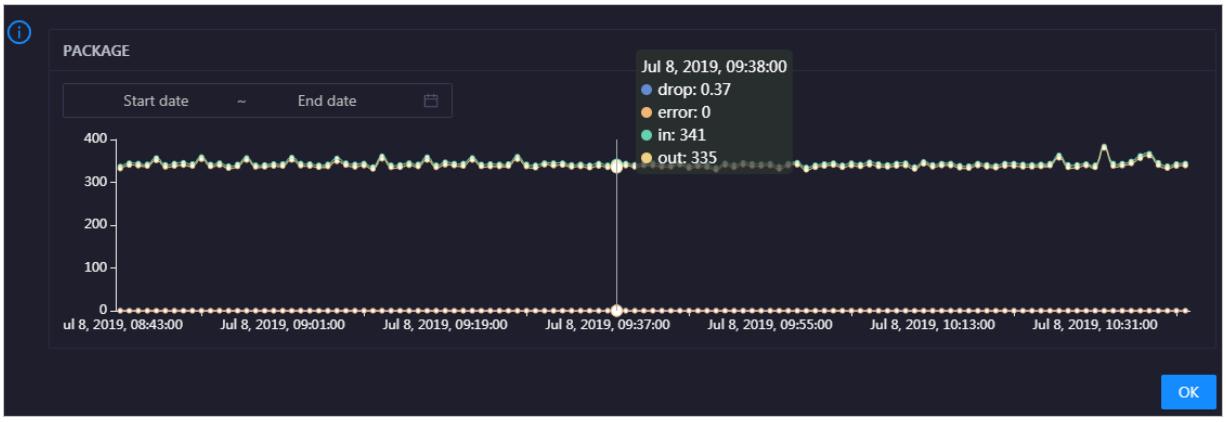

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

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

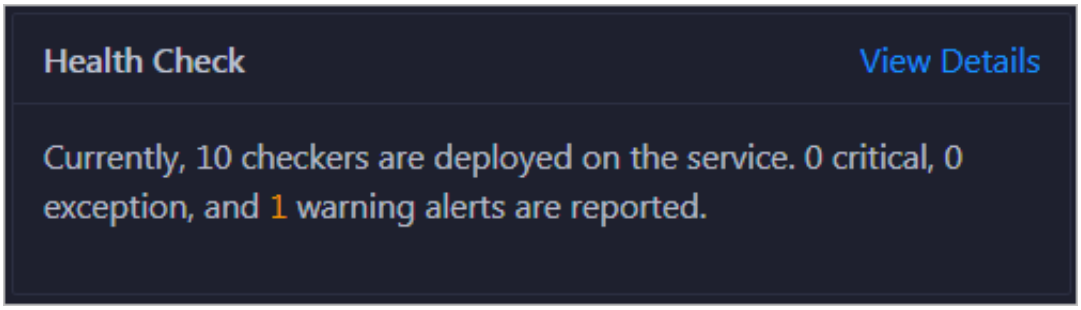

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

Health Check History

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

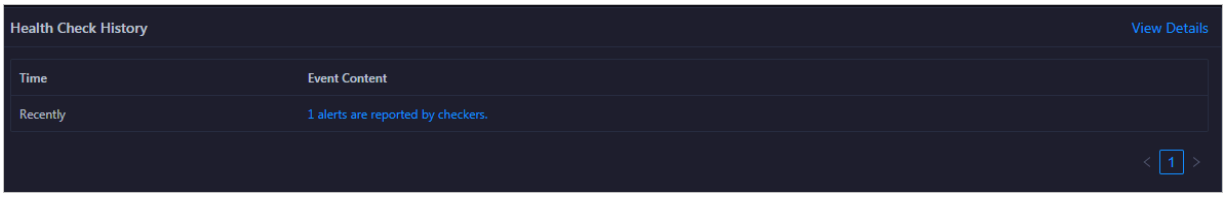

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

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

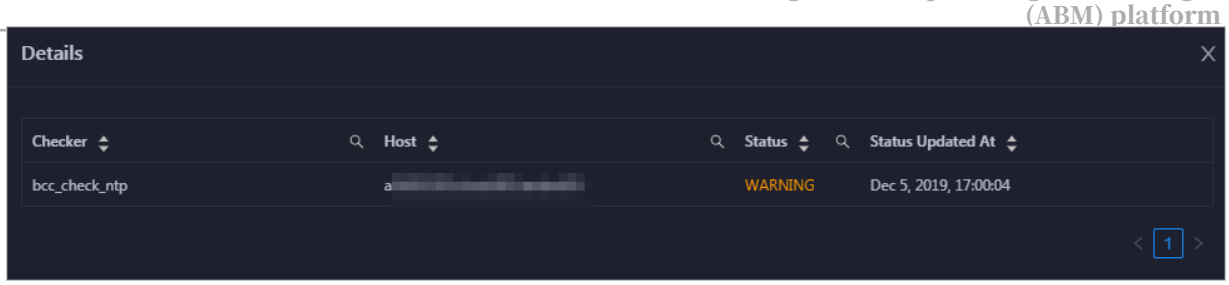

1.5.5.3 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.

The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see *Host overview*.

1.5.5.4 Host health

On the Health Status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

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

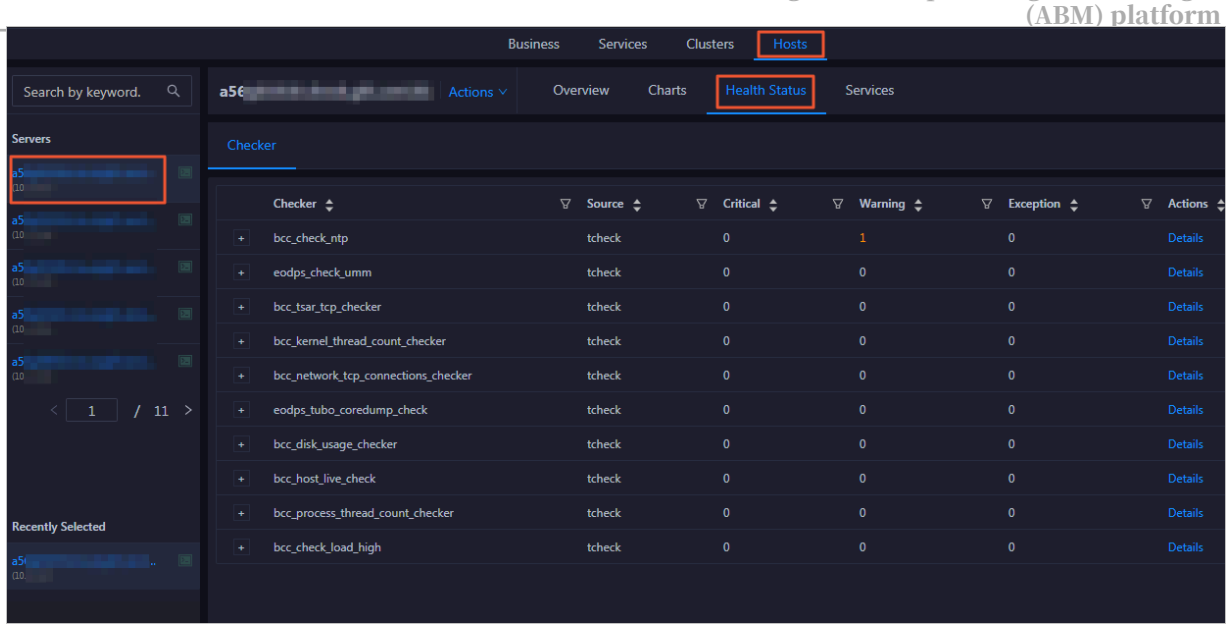

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. **On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.**

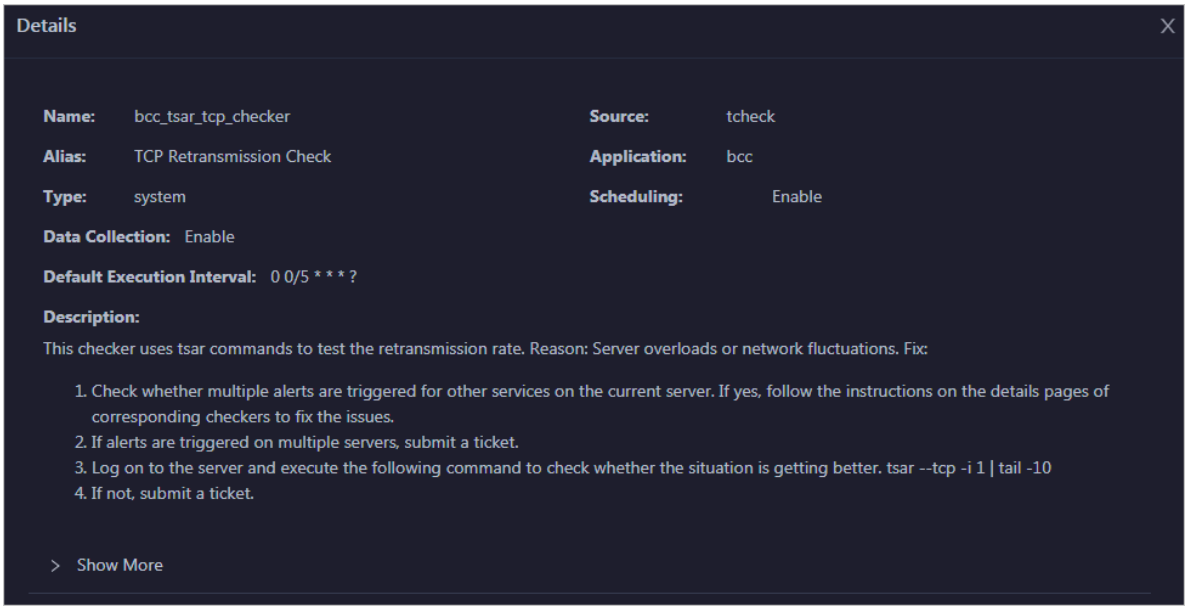

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

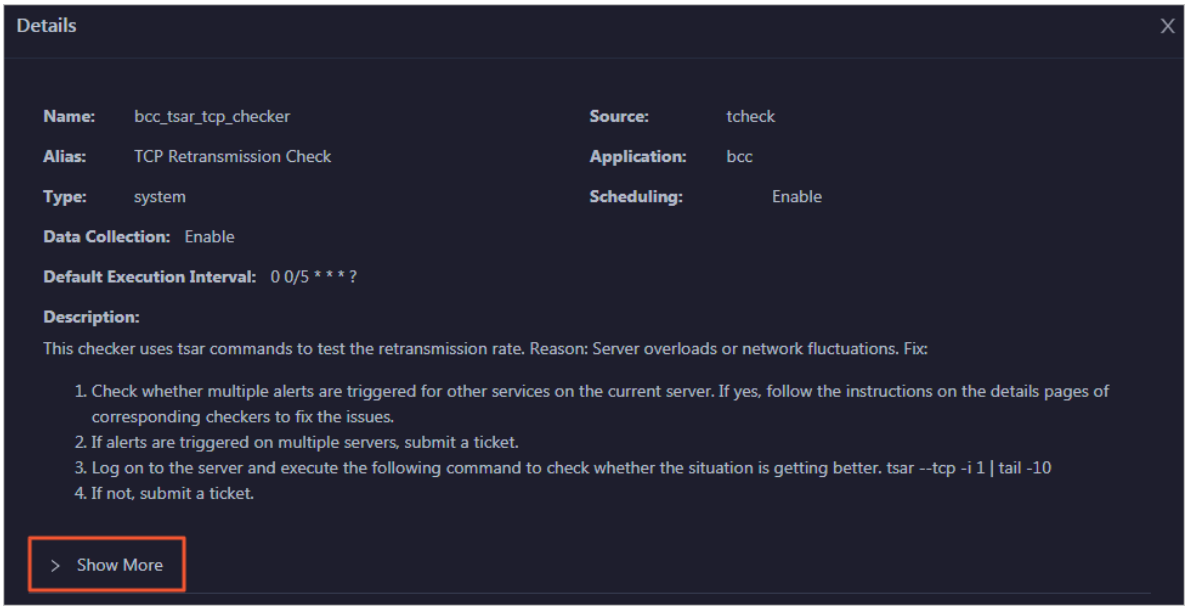

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

1. On the Health Status page, click + to expand a checker with alerts.

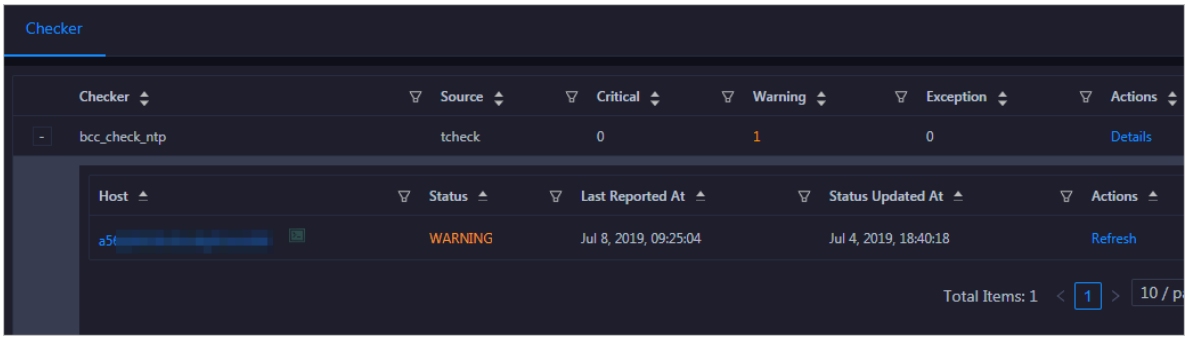

2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.

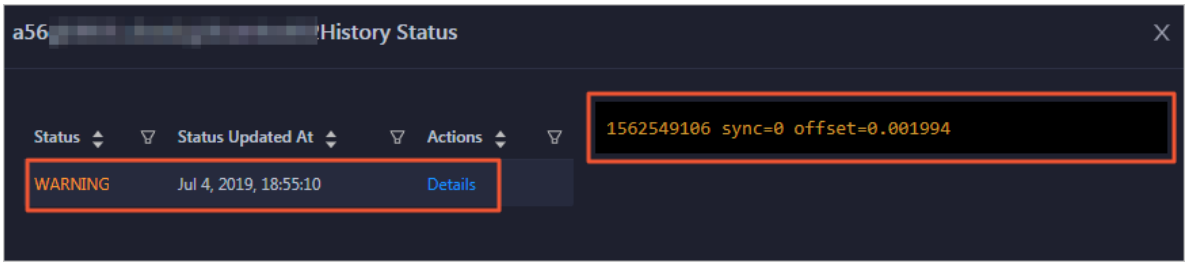

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

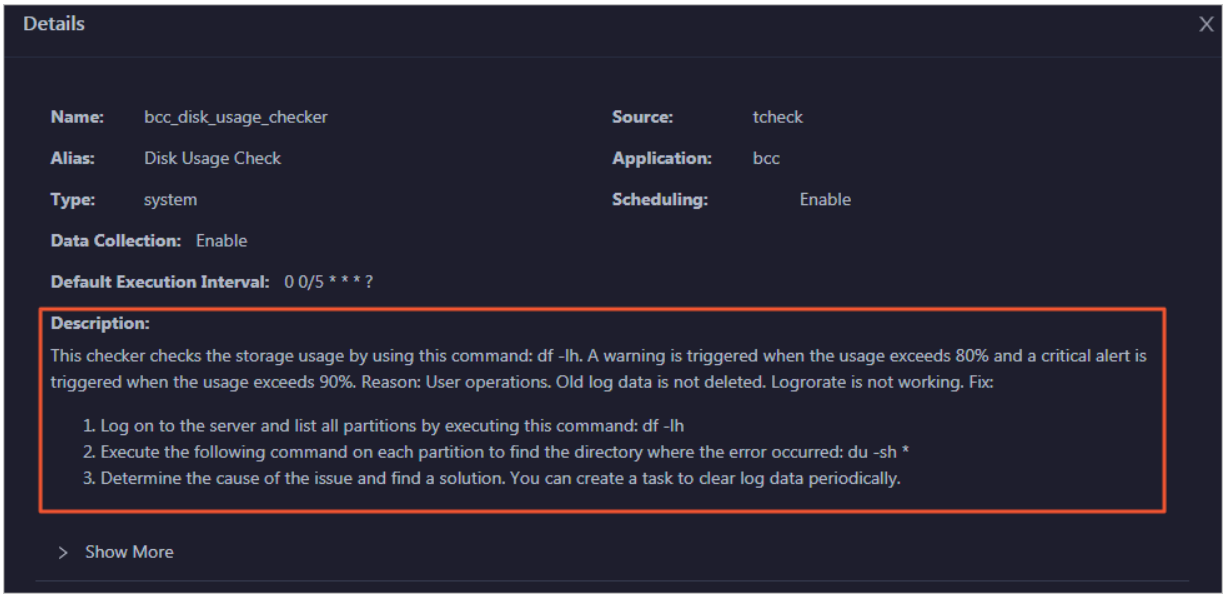

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

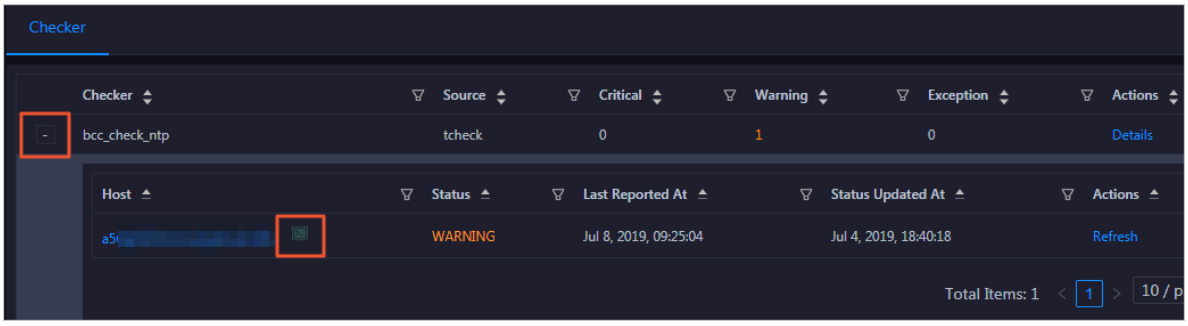

2. Click the Log On icon of a host. The TerminalService page appears.

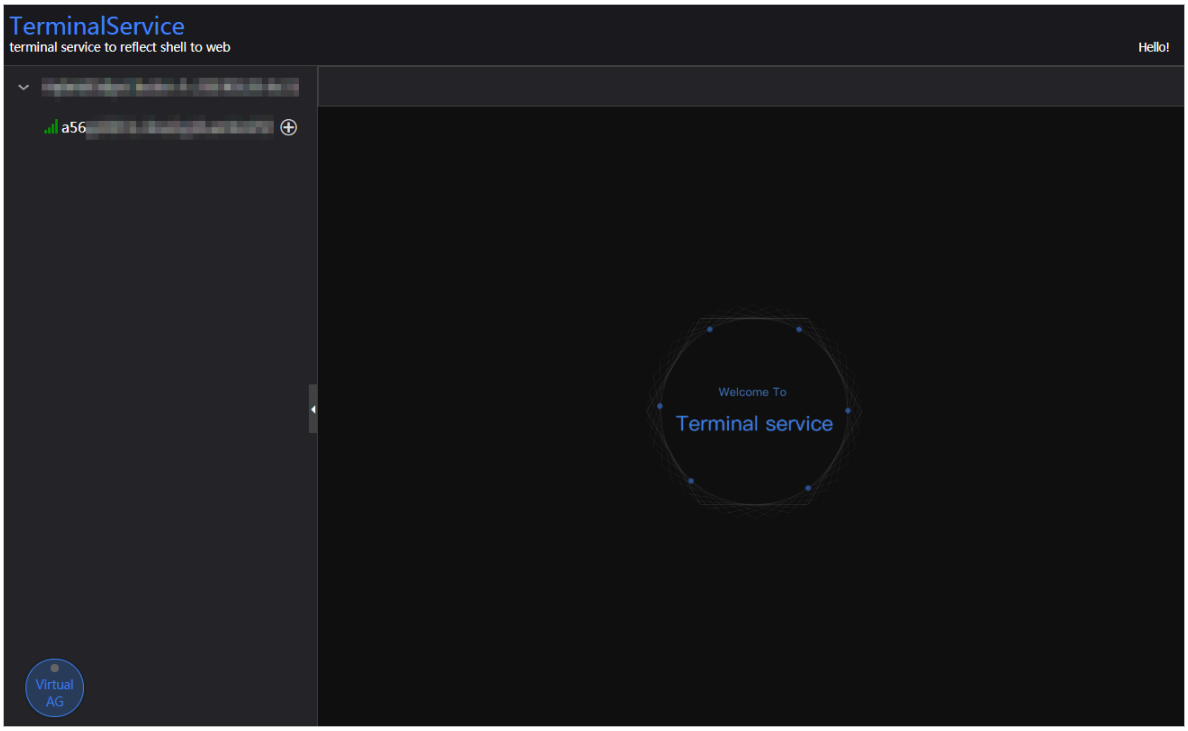

3. On the TerminalService page, click the hostname on the left to log on to the host.

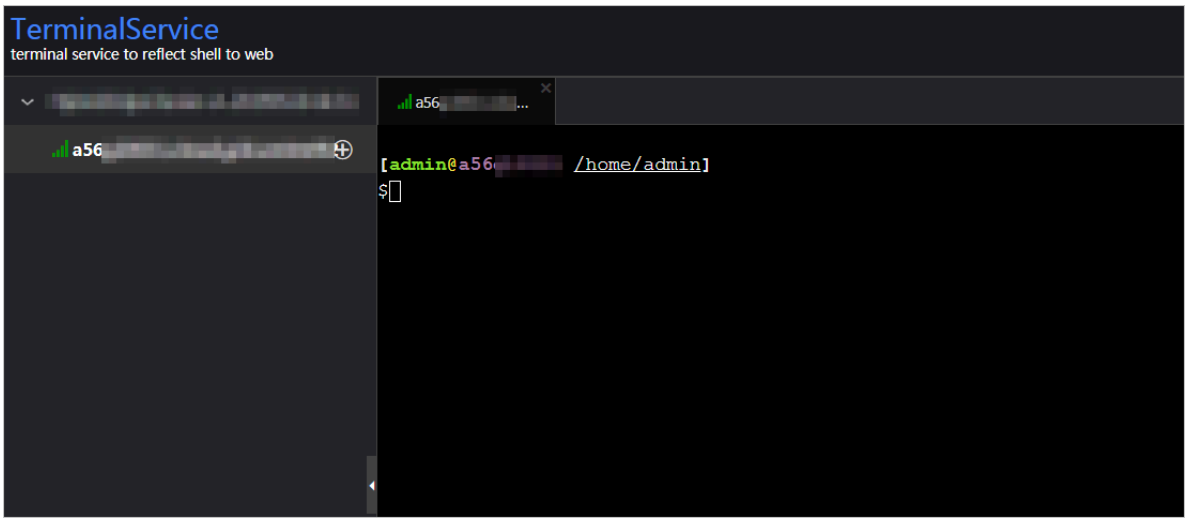

Run a checker again

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. In this way, you can check whether the alert is cleared.

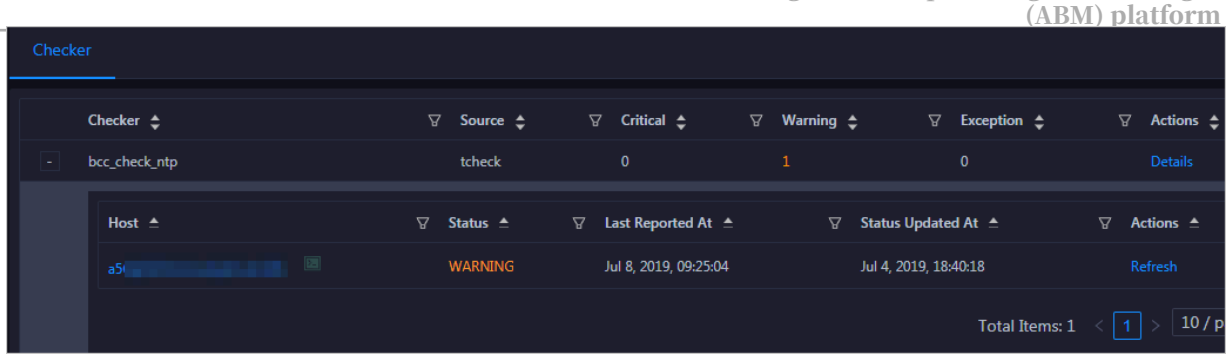

1.5.5.5 Host services

On the Services page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.

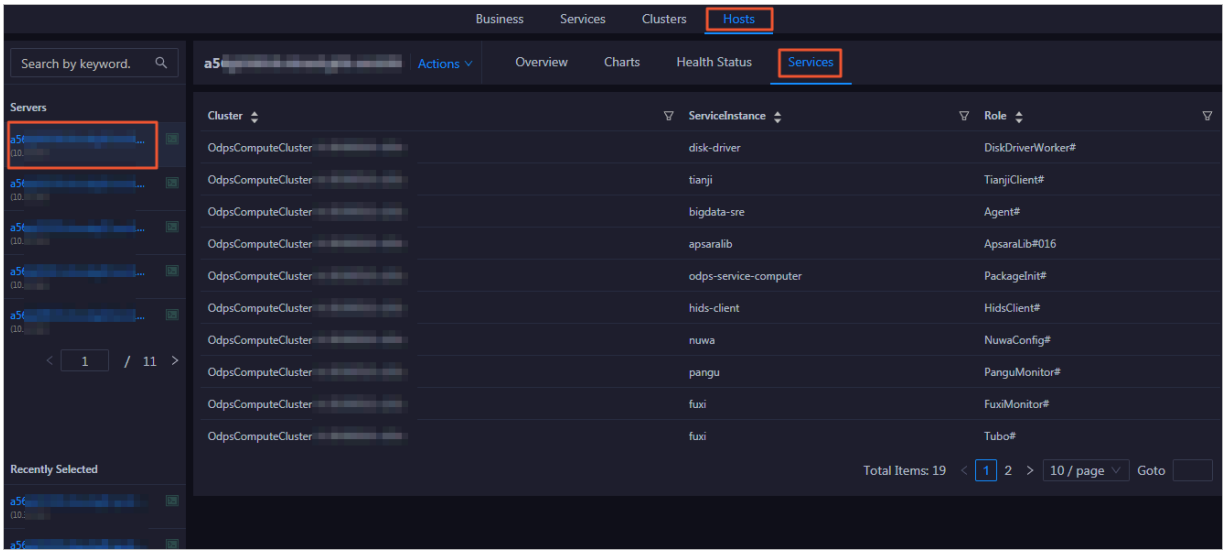

On the Services page, you can view the cluster, service instances, and service instance roles of the host.

1.6 DataWorks

1.6.1 DataWorks O&M overview

This topic describes the features of DataWorks O&M supported by Apsara Bigdata Manager (ABM) and how to access the DataWorks O&M page.

Modules

The modules provided by ABM for DataWorks O&M include the service, cluster, and host O&M modules. The following table describes them in detail.

Module	Feature	Description
Data Warehouse under Services	Overview	Displays the key operation metrics, including service overview, service status, instance scheduling information, and slot usage. On the page, you can also view the trend chart of the total number of finished tasks.
	Health Status	Displays 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.
	Instances	Displays the service roles of DataWorks.
	Slot	Displays the information about slot usage in DataWorks and allows you to change the number of slots in resource groups and hosts.
	Tasks	Displays the running status of DataWorks tasks.
	Settings	Allows you to change the values of configuration items for various service roles in DataWorks.
	Scale-up for Normal Hosts and Scale-down for Normal Hosts	Allows you to scale in or out a DataWorks cluster.
Clusters	Overview	Displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. 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 cluster.

Module	Feature	Description
	Health Status	Displays 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.
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.
	Health Status	Displays the checkers of the selected host, including the checker details, check results for the host, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.

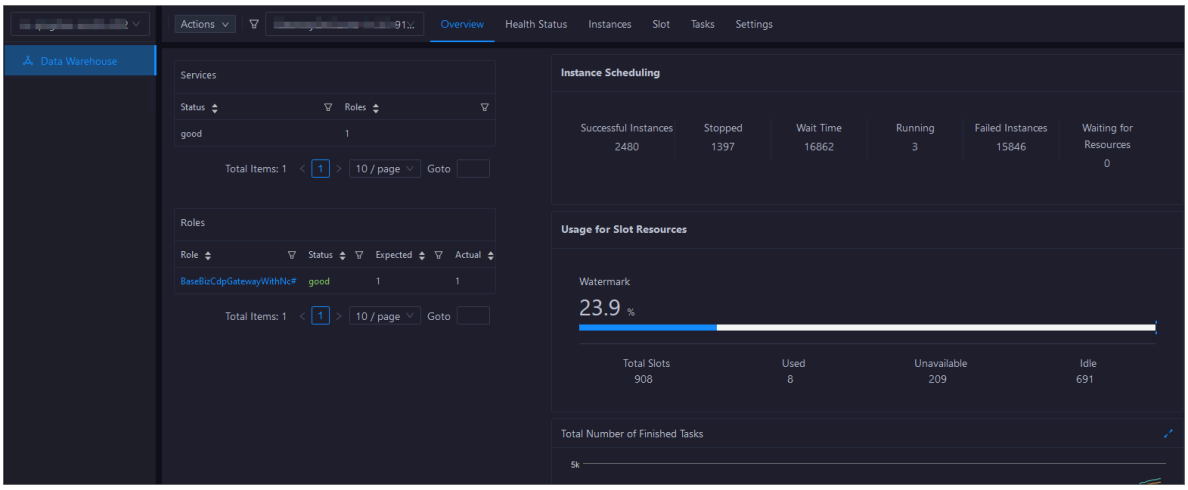

The O&M page includes three modules, namely, Services, Clusters, and Hosts.

1.6.2 Service O&M

1.6.2.1 Service overview

The DataWorks overview page displays the key operation metrics, including service overview, service status, instance scheduling information, and slot usage. On the page, you can also view the trend chart of the total number of finished tasks.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.

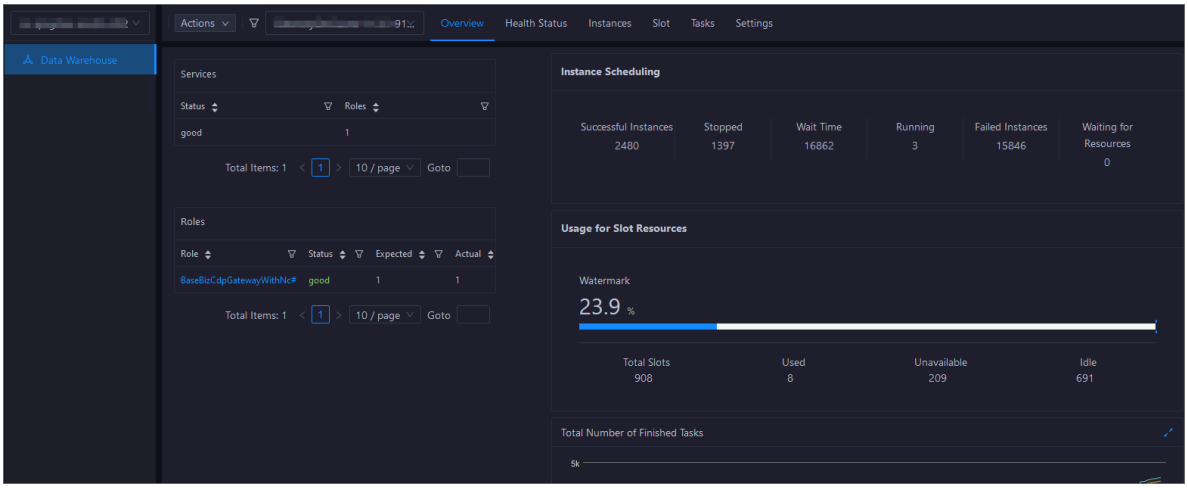

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

Roles

This section displays all DataWorks service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

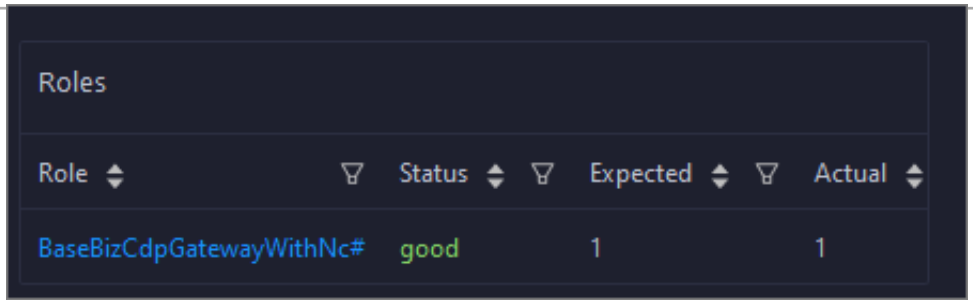

Instance Scheduling

This section displays the number of successful instances, number of instances not running, waiting duration, number of running instances, number of failed instances, and number of instances waiting for resources.

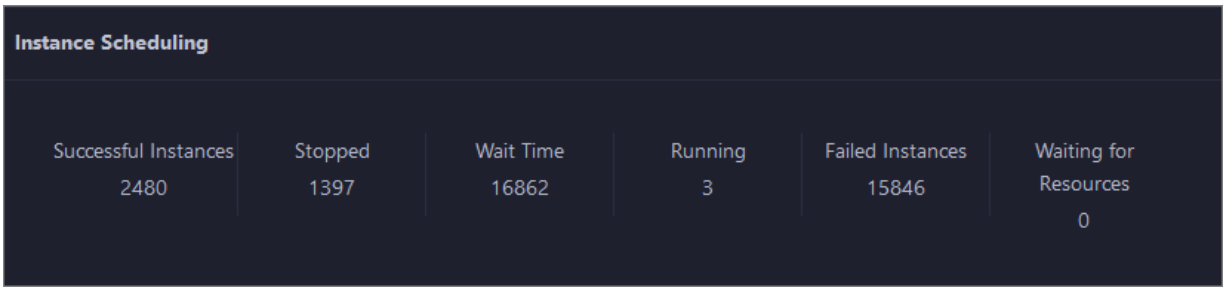

Usage for Slot Resources

This section displays the total number of slots, the number of used slots, the number of unavailable slots, and the number of idle slots for DataWorks.

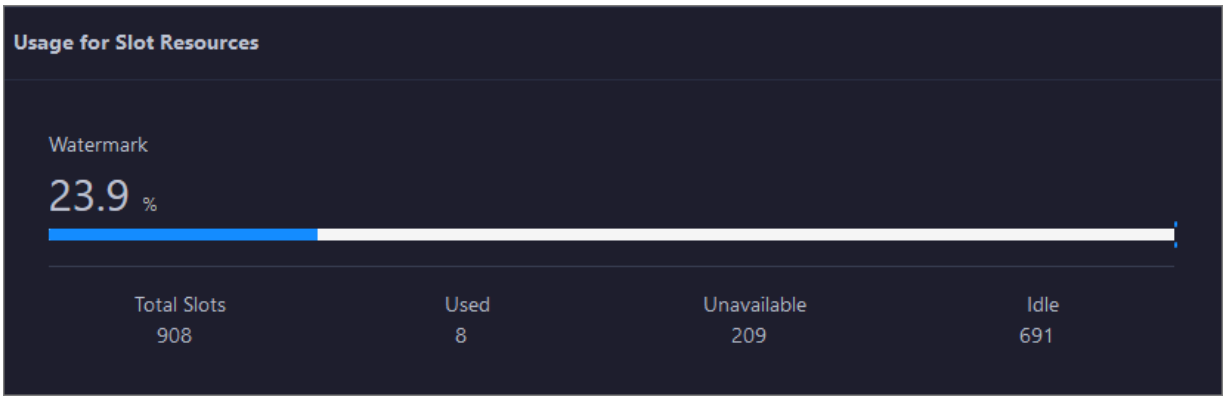

Note:
Slots are resources that can be used by DataWorks for instance scheduling.

Total Number of Finished Tasks

This section displays the trend chart of the total number of finished tasks. The trend chart displays the trend lines of the number of tasks finished yesterday, the

number of tasks finished today, and the average number of tasks finished each day over time in different colors.

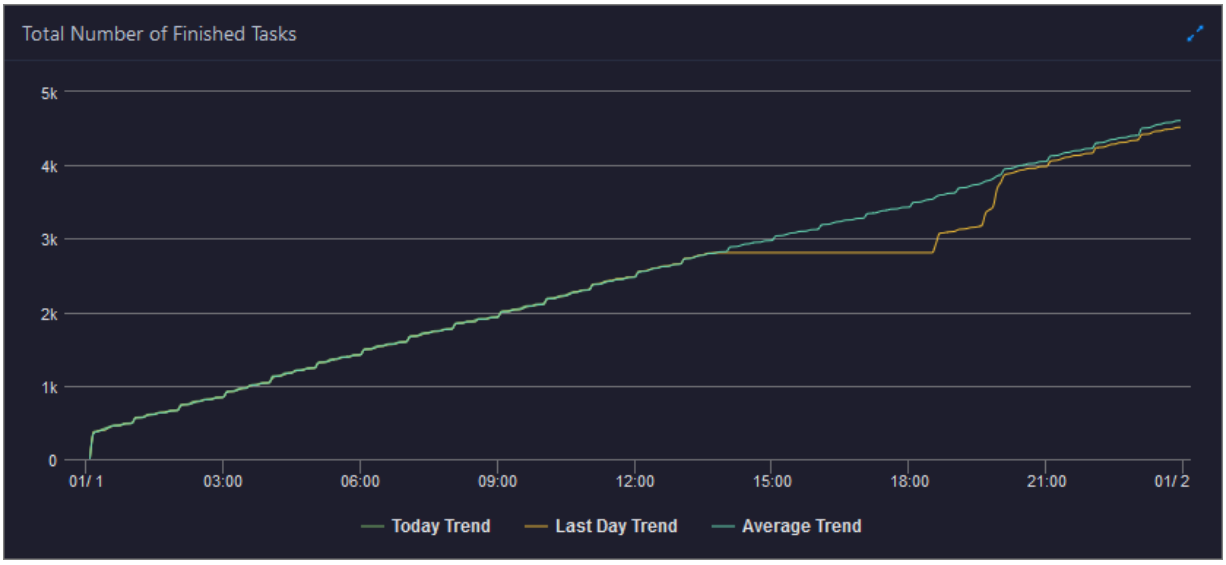

1.6.2.2 Service health

On the Health Status page for DataWorks, 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

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.

4. Select a cluster from the drop-down list, and then click the Health Status tab. The Health Status page appears.

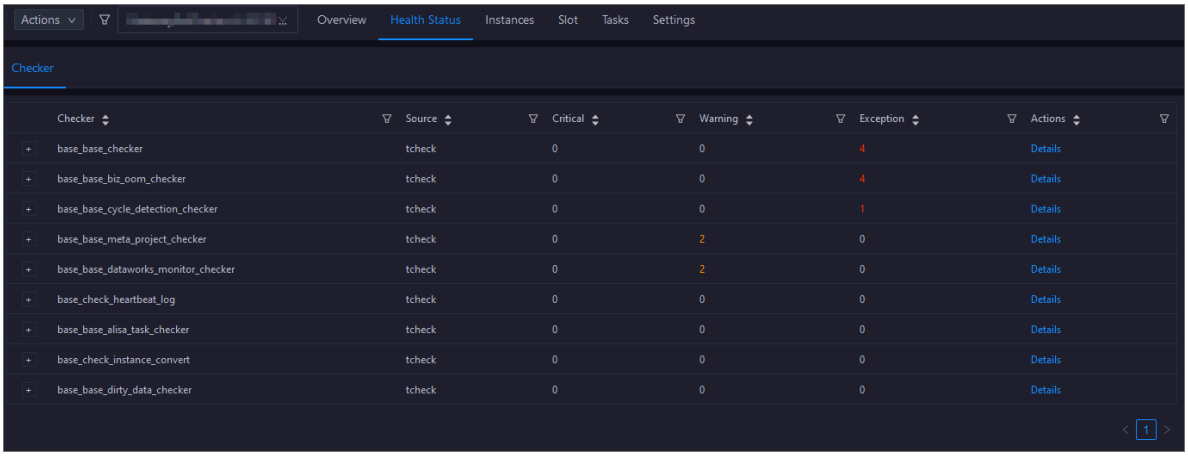

The Health Status page displays all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

Supported operations

On the Health Status page, you can view checker details, hosts with alerts, and alert causes. You can also log on to hosts with alerts, clear alerts, and run checkers again. For more information, see *Cluster health*.

1.6.2.3 Service instances

The Instances page displays information about all DataWorks service roles, including the name, status, and expected and actual numbers of machines in the final status.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.

4. Select a cluster from the drop-down list, and then click the Instances tab. The Instances page appears.

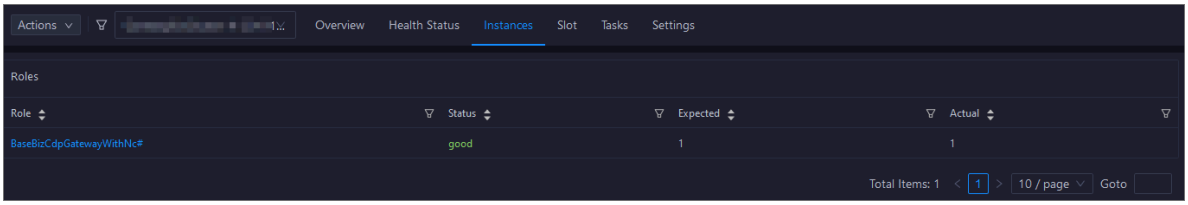

The Instances page displays information about all DataWorks service roles, including the status and the expected and actual numbers of machines in the final status. The statuses include good, bad, and upgrading.

Supported operations

You can filter or sort service roles based on a column to facilitate information retrieval on the Instances page. For more information, see *Common operations*.

1.6.2.4 Service slots

Slots are resources used to process tasks. Apsara Bigdata Manager (ABM) allows you to view the slot information of DataWorks clusters, resource groups, and hosts, including the maximum number of slots, the number of used slots, and the slot usage. You can also migrate resource groups, modify the number of slots for resource groups or hosts, and modify the host status.

Concepts

A data migration unit (DMU) represents the minimum operating capability required by a Data Integration task, that is, the data synchronization processing capability given limited CPU, memory, and network resources.

Resources measured by DMU are allocated by slot. Each DMU occupies two slots.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.

4. Select a cluster from the drop-down list, and then click the Slot tab. The Slot page appears.

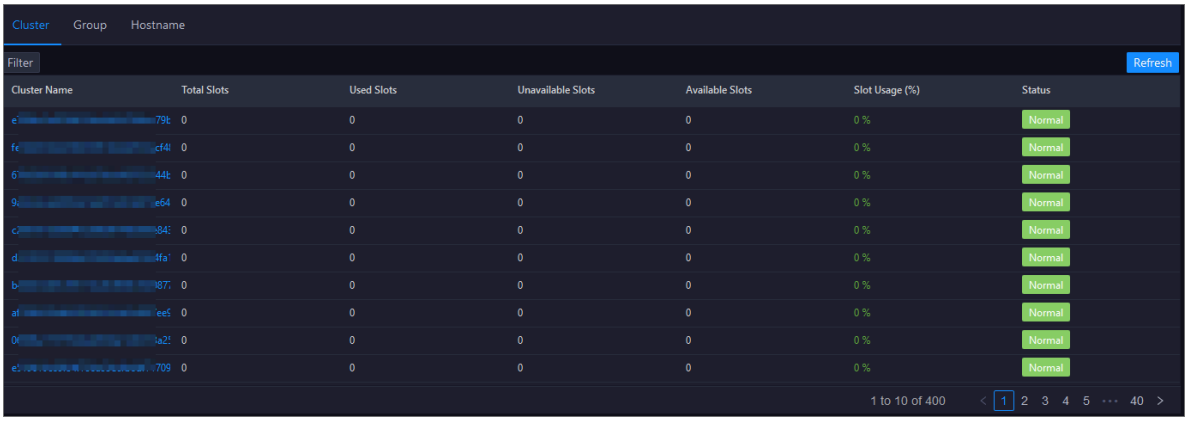

Cluster slots

Click the Cluster tab on the Slot page. The Cluster page appears.

The Cluster page displays the slot overview of all DataWorks clusters, including the total number of slots, the numbers of used slots and available slots, and the slot usage. The page also displays the cluster running status.

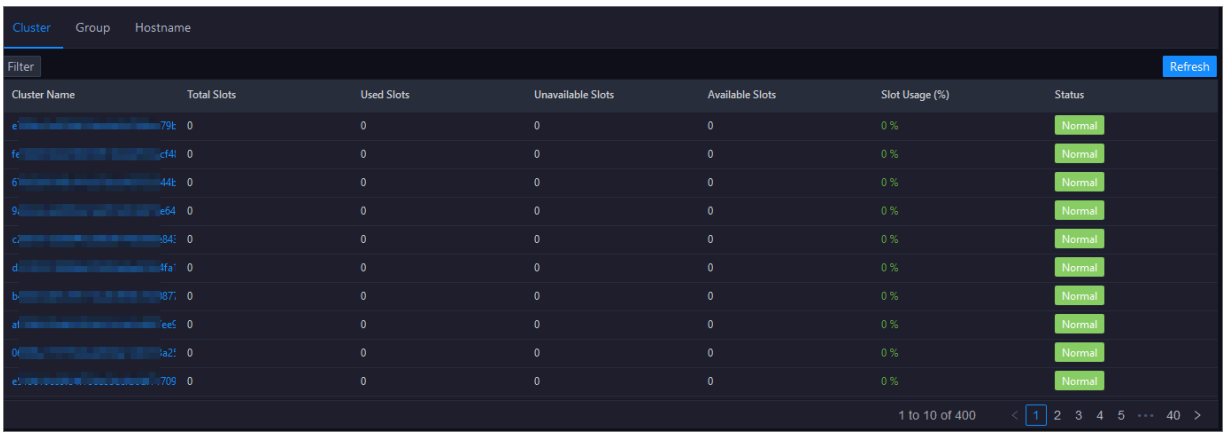

To view more information about slots of a specified cluster, click the name of the cluster.

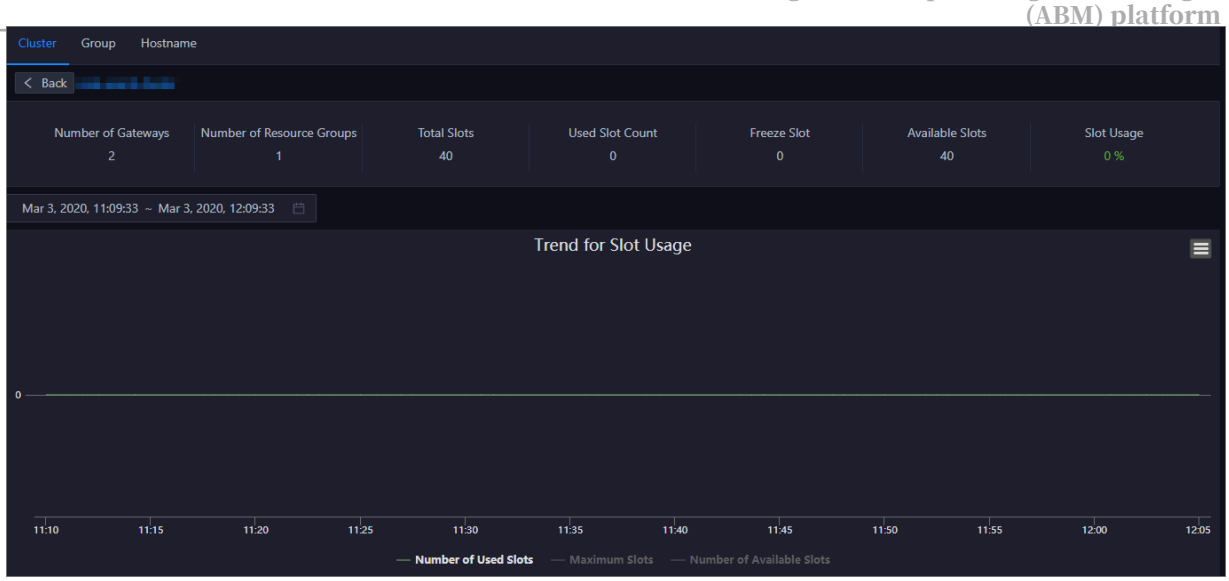

On the cluster details page, you can view the numbers of gateways, resource groups , slots, used slots, frozen slots, and available slots, and the slot usage of the cluster at the top. You can also view the trend chart of slot usage over time at the bottom. The trend chart displays the trend lines of the number of used slots, the maximum number of slots, and the number of available slots in different colors.

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 displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

Resource group slots

Click the Group tab on the Slot page. The Group page appears.

The Group page displays the slot overview of all DataWorks resource groups, including the maximum number of slots, the numbers of used slots and available slots, and the slot usage. The page also displays the name, cluster, project, and running status of each resource group.

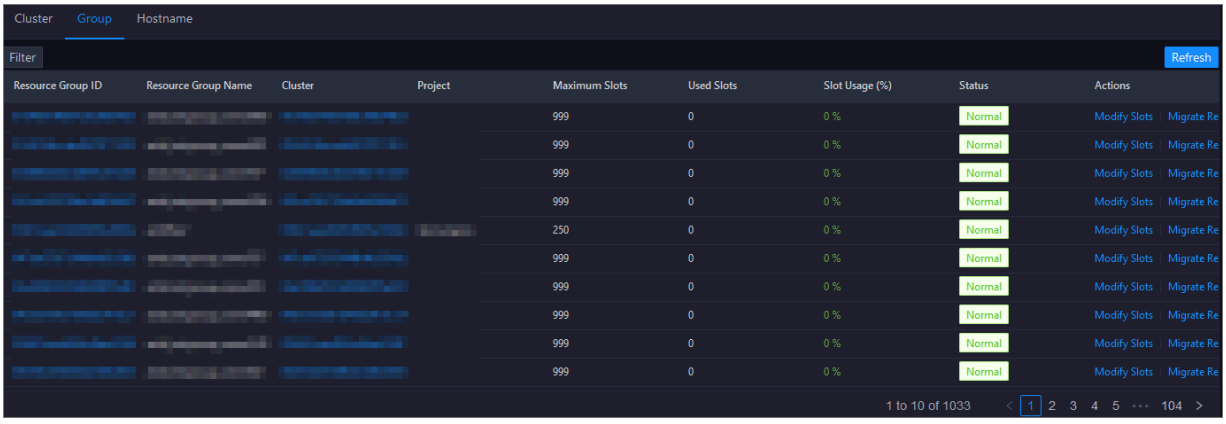

To view more information about slots of a specified resource group, click the ID of the resource group.

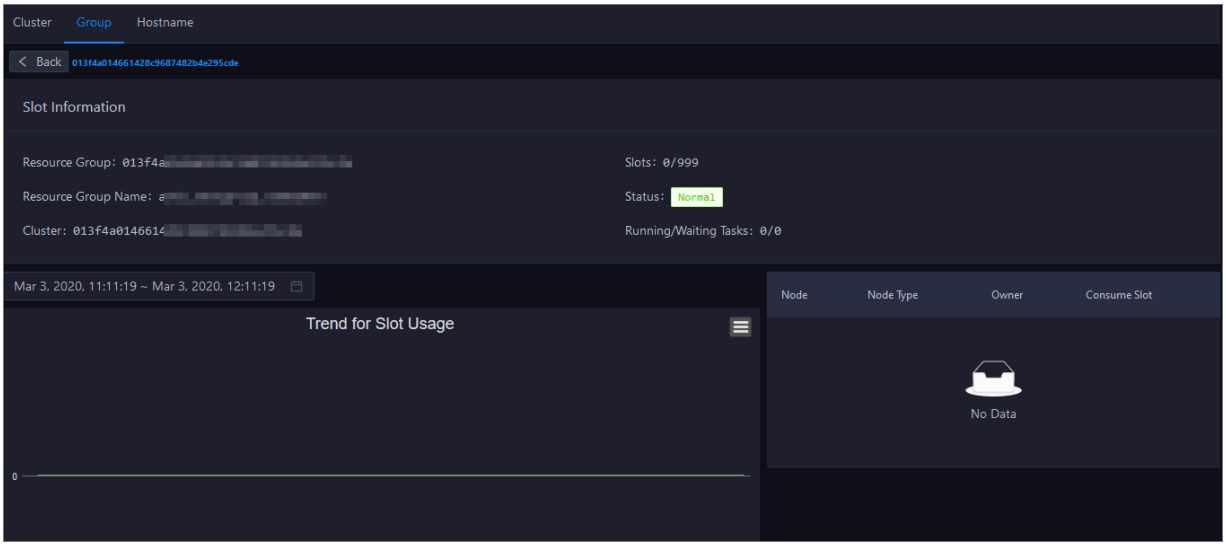

On the resource group details page, you can view the current slot information of the resource group, for example, the number of used slots and the maximum number of slots, at the top. You can also view the trend chart of slot usage over time, the nodes that occupy the slots, and the owners at the bottom. The trend chart displays the trend lines of the number of used slots, the maximum number of slots, and the number of available slots in different colors.

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 displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

Modify the number of resource group slots

If the number of slots in a resource group is insufficient or excessive, you can modify the number of slots to add or remove resources in the resource group.

1. On the Group page, find the target resource group and click Modify Slots in the Actions column.
2. In the dialog box that appears, set Maximum Slots.
3. Click Run. A message appears, indicating that the action has been submitted.

Migrate a resource group

If the slots in a cluster bound to a resource group are insufficient and cannot be increased, you can bind the resource group to another cluster.

1. On the Group page, find the target resource group and click Migrate Resource Group in the Actions column.
2. In the dialog box that appears, set Target Cluster.
3. Click Run. A message appears, indicating that the action has been submitted.

Host slots

Click the Hostname tab on the Slot page. The Hostname page appears.

The Hostname page displays the slot overview of all DataWorks hosts, including the maximum number of slots, the number of used slots, and the slot usage. The page also displays the IP address, cluster, running status, activeness, and monitoring status of each host.

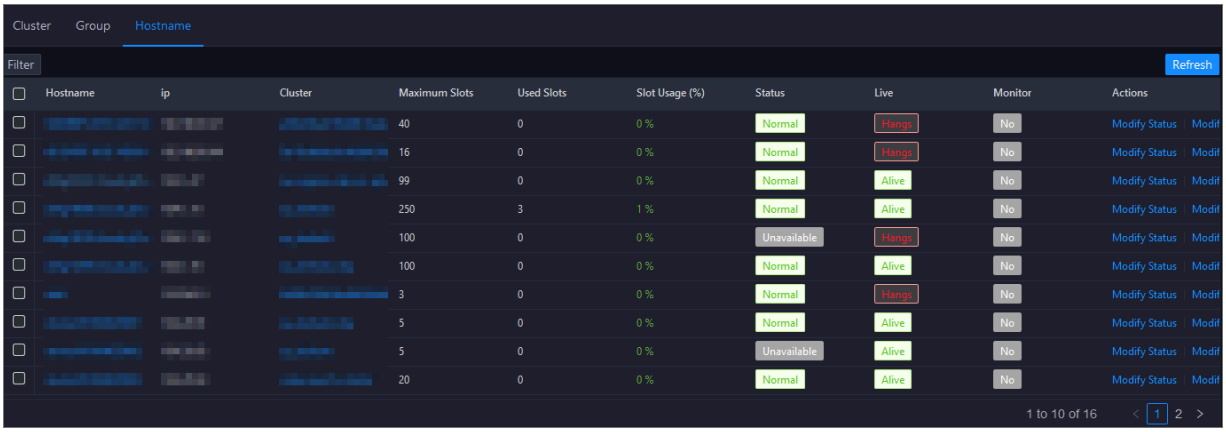

To view more information about slots of a specified host, click the name of the host.

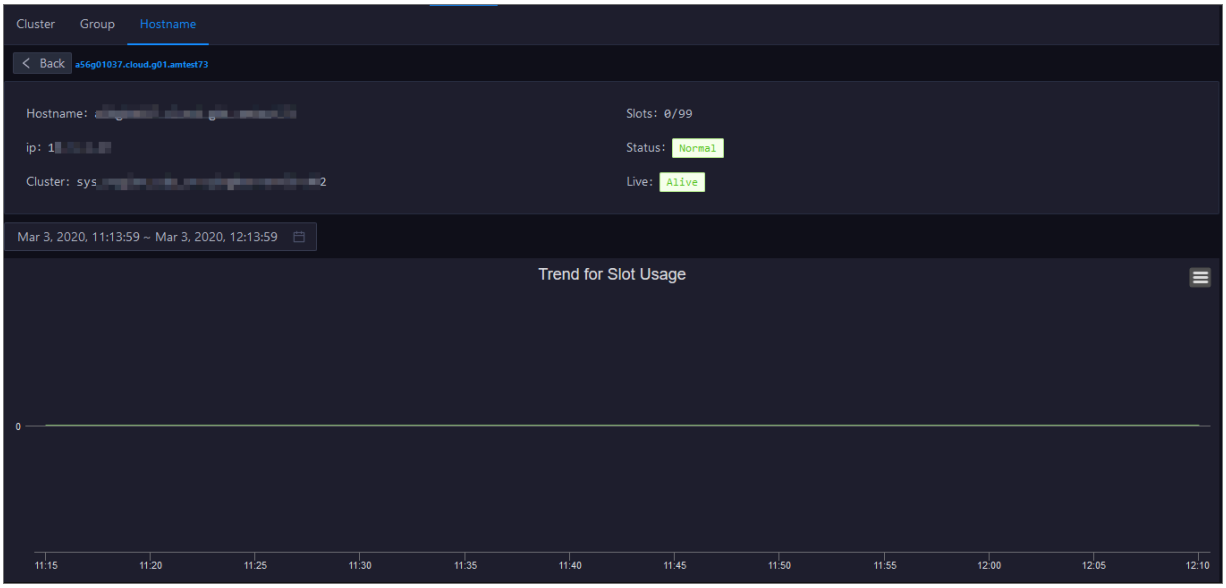

On the host details page, you can view the current slot information of the host, for example, the number of used slots and the maximum number of slots, at the top. You can also view the trend chart of slot usage over time at the bottom. The trend

chart displays the trend lines of the number of used slots, the maximum number of slots, and the number of available slots in different colors.

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 displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

Modify the host status

The host can be in the normal, unavailable, or suspended state. You can modify the host status as needed.

1. On the Hostname page, find the target host and click Modify Status in the Actions column.
2. In the dialog box that appears, set Status.
3. Click Run. A message appears, indicating that the action has been submitted.

Modify the number of host slots

If the number of slots in a host is insufficient or excessive, you can modify the number of slots to add or remove resources in the host.

1. On the Hostname page, find the target host and click Modify Slots in the Actions column.
2. In the dialog box that appears, set Maximum Slots.
3. Click Run. A message appears, indicating that the action has been submitted.

1.6.2.5 Service tasks

The Tasks page displays tasks created by a user in DataWorks. You can filter or sort tasks based on a column to facilitate information retrieval.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.

4. **Select a cluster from the drop-down list, and then click the Tasks tab. The Tasks page appears.**

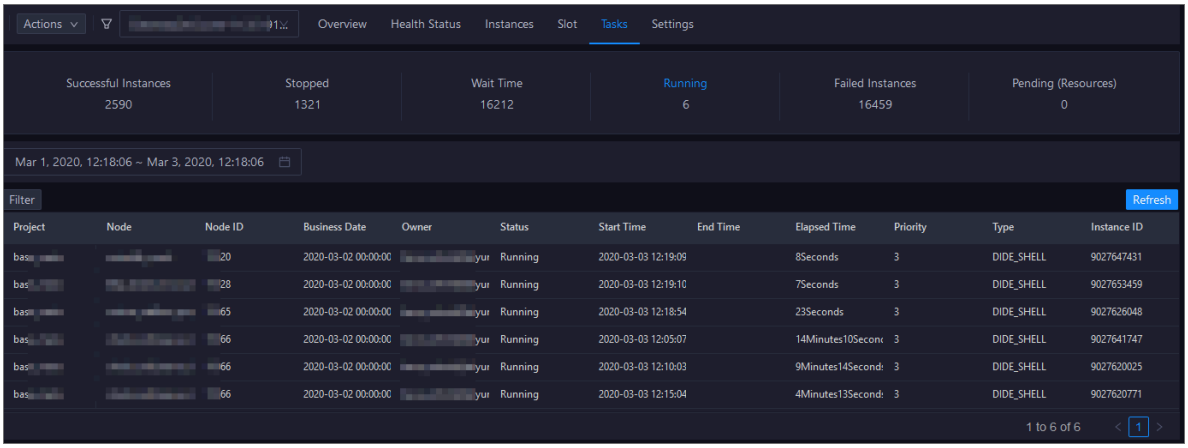

The Tasks page displays the task information of the current cluster, including the project name, node name, node ID, business date, owner, running status, start time, end time, running duration, priority, type, and instance ID.

Filter tasks by status

On the Tasks page, the respective number of tasks in all states is displayed at the top. Click a task state to display corresponding tasks in the list. By default, tasks in the Running state are displayed.

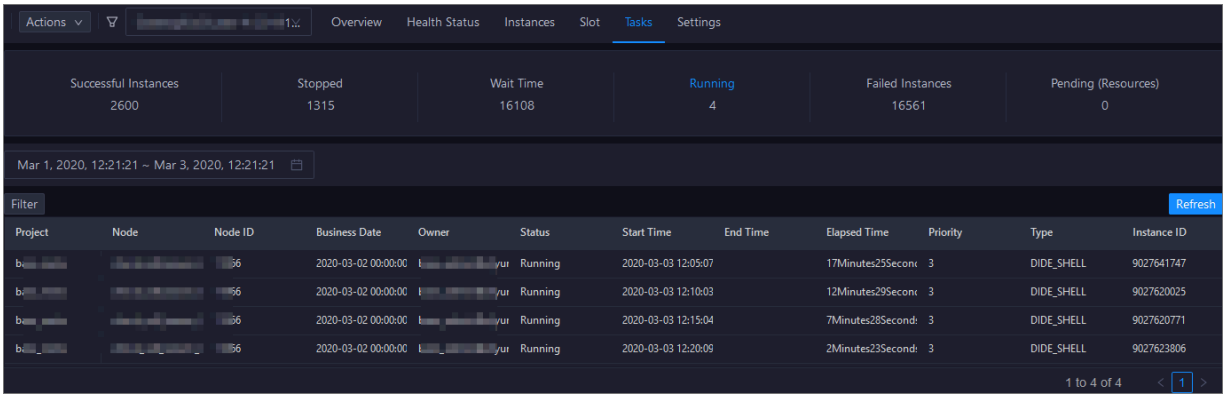

Filter tasks by time

Select a time period (both the date and time can be set) in the upper-left corner of the task list to view the tasks in the corresponding time period.

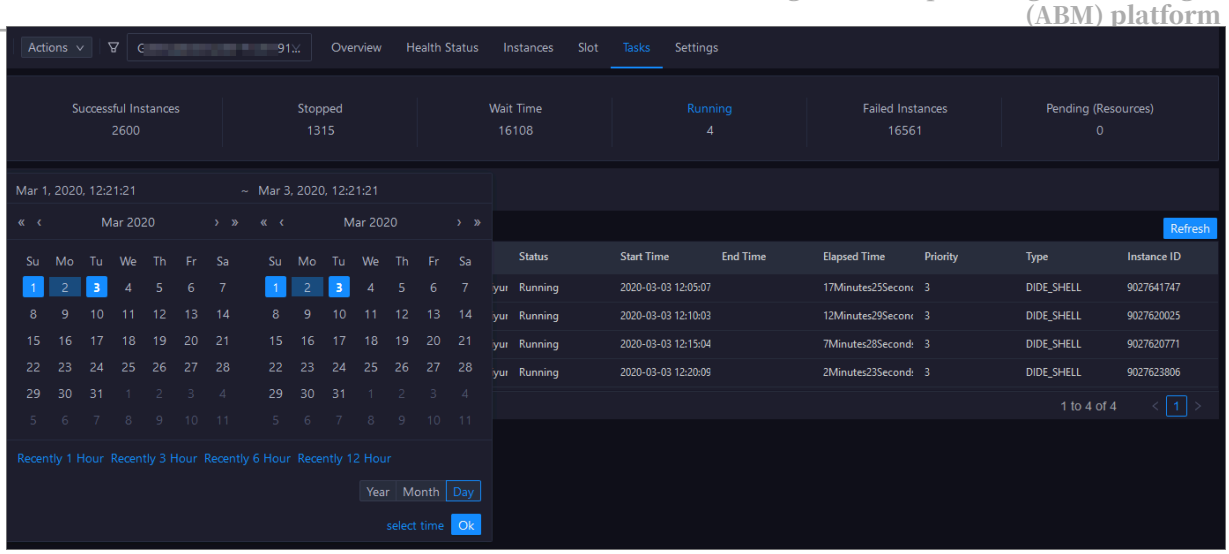

Other operations

You can filter tasks, sort tasks based on a column, and customize columns on the Tasks page. For more information, see *Common operations*.

1.6.2.6 Service settings

The Settings page allows you to change the values of configuration items for various service roles in DataWorks.

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.
4. Select a cluster from the drop-down list, and then click the Settings tab. The Settings page appears.

1.6.2.7 Cluster scaling

Apsara Bigdata Manager (ABM) supports DataWorks cluster scaling. To scale out a DataWorks cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the DataWorks cluster. To scale in a DataWorks cluster, remove physical hosts from the DataWorks cluster to the default cluster of Apsara Infrastructure Management Framework.

Background

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to

the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework can be considered as an available resource pool that provides resources for scaling out business clusters. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

When you scale out a DataWorks cluster, ABM adds physical hosts in the default cluster to the DataWorks cluster. When you scale in a DataWorks cluster, ABM removes physical hosts from the DataWorks cluster to the default cluster. The service roles of physical hosts in DataWorks include BaseBizCdpGatewayWithNc# and BaseBizGatewayWithNc#. DataWorks cluster scaling only supports these two service roles.

Prerequisites

- **Scale-out**
  - **The physical host to be added to a DataWorks cluster is in the default cluster of Apsara Infrastructure Management Framework.**
  - **If you use a host as a template host for scale-out, the service role of the host is BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#.**
- **Scale-in**

  If you use a host as a template host for scale-in, the service role of the host is BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#.

> **Note:**
> **You can go to the DataWorks page. Click O&M in the upper-right corner, and then click the Services tab. Click Data Warehouse in the left-side navigation pane, and then click the Instances tab. In the service role list, find the service role BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#, and then click the service role name to go to the Apsara Infrastructure Management Framework console to view the hosts with the service role BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#.**

Scale out a DataWorks cluster

You can add multiple hosts to a DataWorks cluster at a time to scale out the cluster. To achieve this, you need to specify an existing host as the template host. When you scale out the DataWorks cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.
4. Select a cluster from the drop-down list, and then click the Slot tab. The Slot page appears.
5. Click the Hostname tab on the Slot page, and then select a physical host whose service role is BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc# in the host list as the template host.

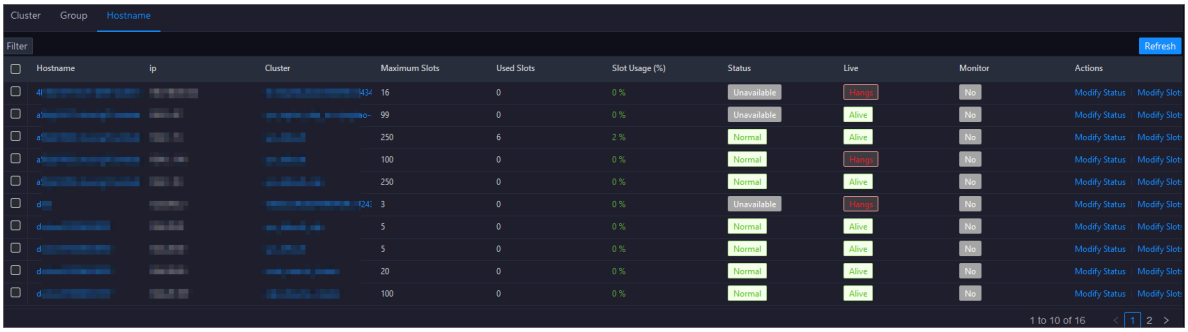

6. Choose Actions > Scale-up for Normal Hosts in the upper-left corner. In the Scale-up for Normal Hosts dialog box that appears, set relevant parameters.

   The parameters are described as follows:

   - Refer Hostname: the name of the template host. By default, the name of the selected host is used.
   - Hostname: the name of the host to be added to the DataWorks cluster. Enter the name of an available host in the default cluster for scale-out. To enter multiple hostnames, separate them with commas (,).
7. Click Run. A message appears, indicating that the action has been submitted.

8. View the scale-out status.

   Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale-up for Normal Hosts to view the scale-out history.

   It may take some time for the cluster to be scaled out. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

   If the status is RUNNING, click Details in the Details column to view the steps and progress of the scale-out.

   If the status is FAILED, click Details in the Details column to locate the failure cause. For more information, see *Locate the failure cause*.

Scale in a DataWorks cluster

You can remove physical hosts from a DataWorks cluster to the default cluster of Apsara Infrastructure Management Framework to scale in the DataWorks cluster.

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.
4. Select a cluster from the drop-down list, and then click the Slot tab. The Slot page appears.
5. Click the Hostname tab on the Slot page, and then select a physical host whose service role is BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc# in the host list as the template host.

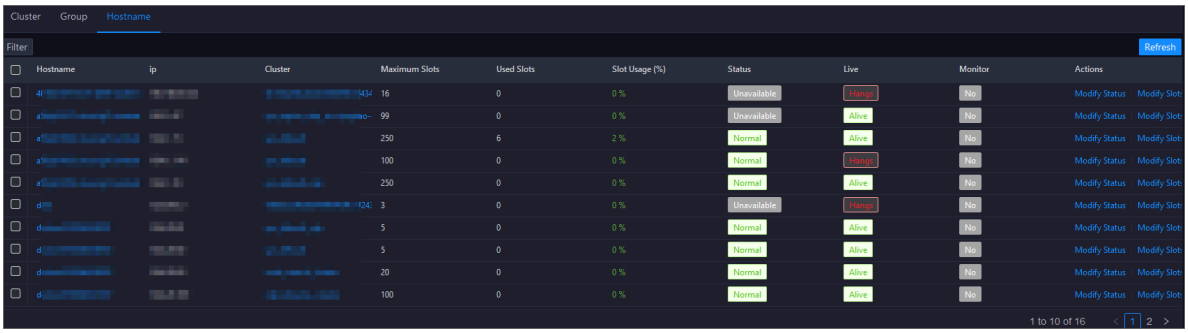

6. **Choose Actions > Scale-down for Normal Hosts in the upper-left corner. In the Scale-down for Normal Hosts dialog box that appears, set relevant parameters.**

   The parameters are described as follows:

   - **Hostname: the name of the host to be removed from the DataWorks cluster. By default, the name of the selected host is used.**
   - **Biz Name: the service role of the host to be removed from the DataWorks cluster. Select the actual service role from the drop-down list. Valid values: base-biz-cdpgatewaywithnc# and base-biz-gatewaywithnc#.**

7. **Click Run. A message appears, indicating that the action has been submitted.**
8. **View the scale-in status.**

   Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale-down for Normal Hosts to view the scale-in history.

   It may take some time for the cluster to be scaled in. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

   If the status is RUNNING, click Details in the Details column to view the steps and progress of the scale-in.

   If the status is FAILED, click Details in the Details column to locate the failure cause. For more information, see *Locate the failure cause*.

Locate the failure cause

This section uses cluster scale-out as an example to describe how to locate the failure cause.

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click DataWorks.**
3. **On the page that appears, click O&M in the upper-right corner. The Overview page under the Data Warehouse page appears.**
4. **Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale-up for Normal Hosts to view the scale-out history.**

5. In the scale-out history dialog box, click Details in the Details column of a failed execution to locate the failure cause.

   You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.6.3 Cluster O&M

1.6.3.1 Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. 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 cluster.

Entry

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

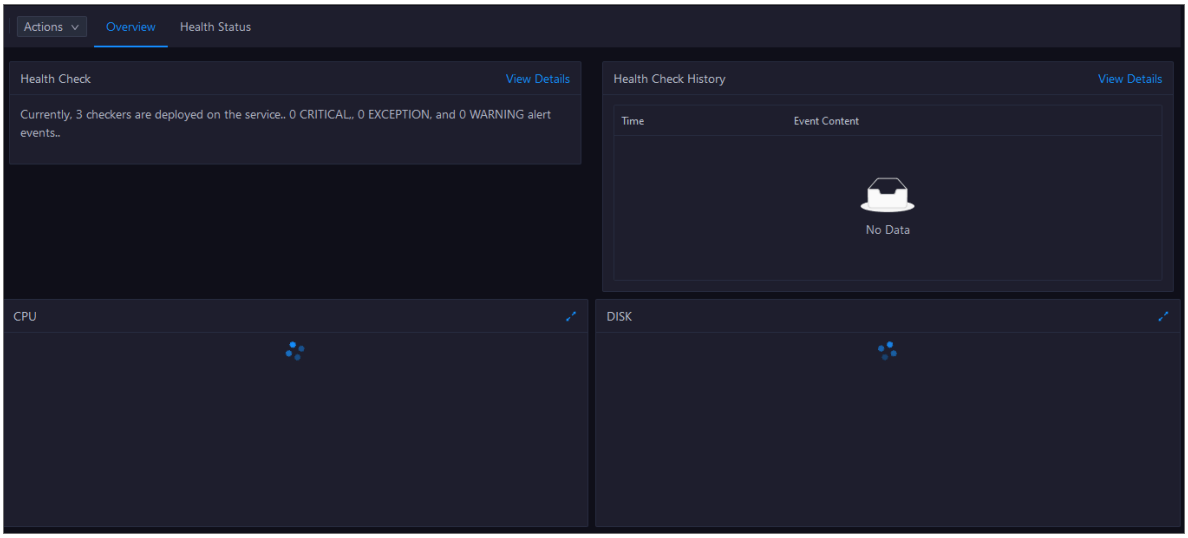

Health Check

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

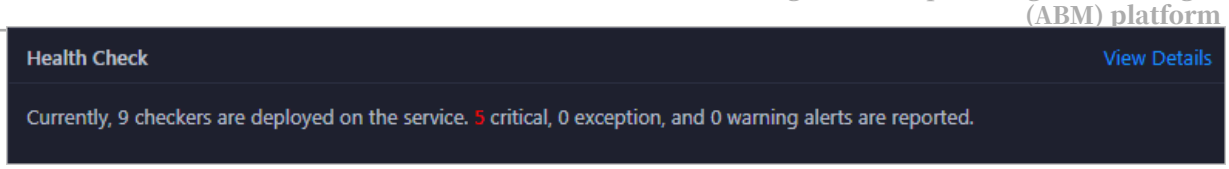

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Cluster health*.

Health Check History

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

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

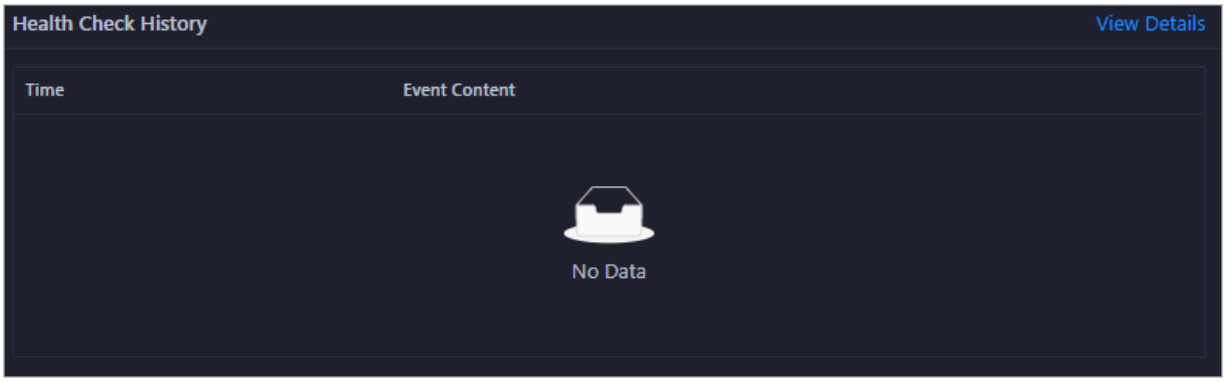

You can click the event content of a check to view the exception items.

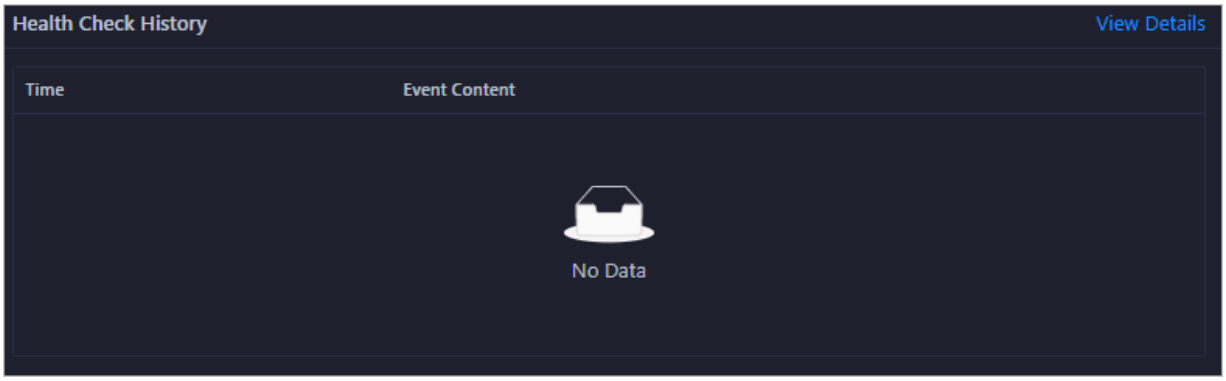

CPU

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 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.

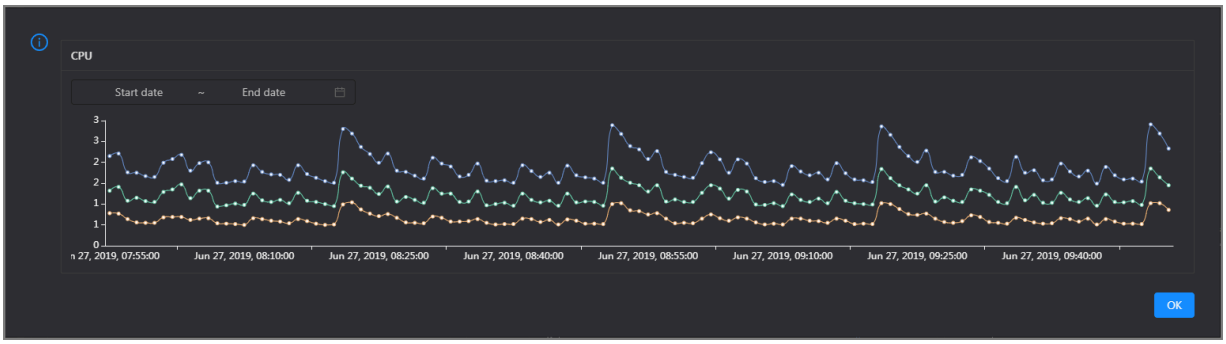

DISK

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

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

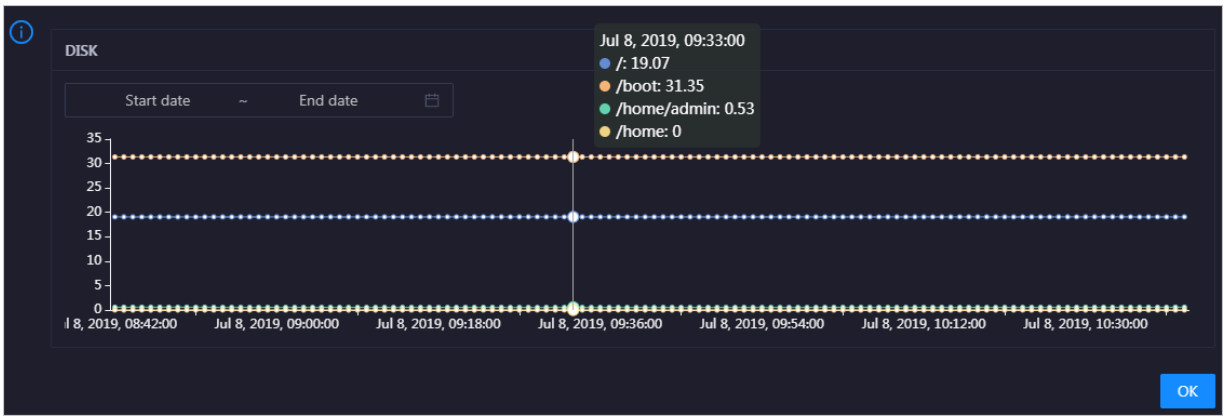

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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 cluster over time in different colors.

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

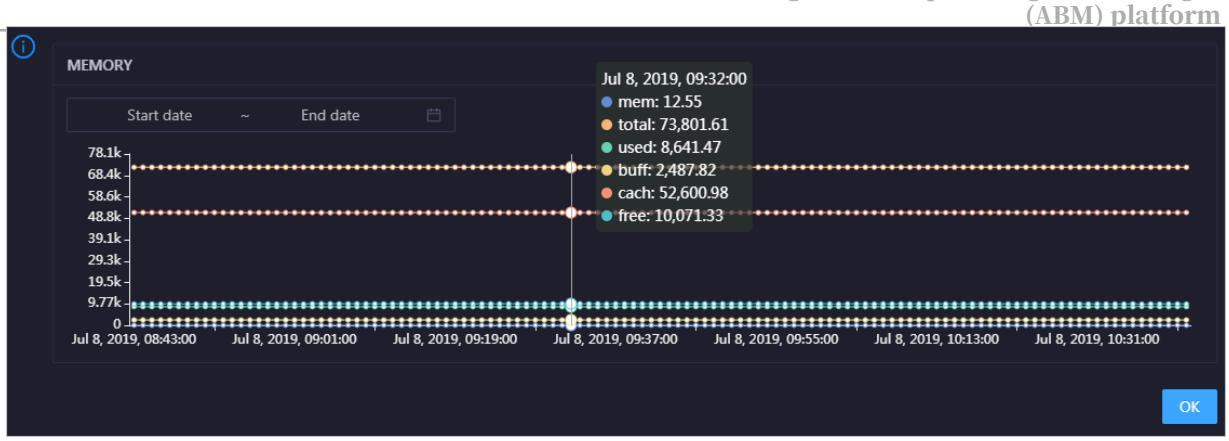

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

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

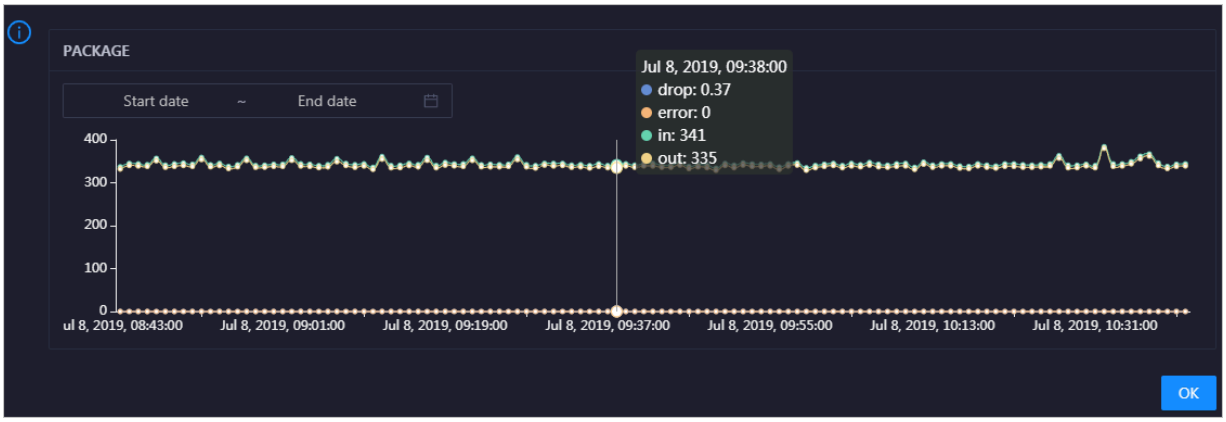

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

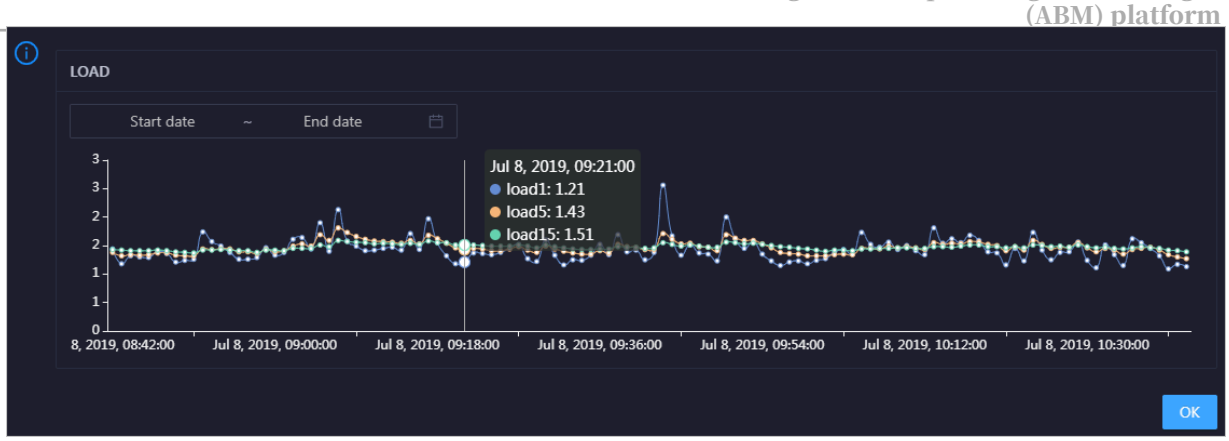

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

1.6.3.2 Cluster health

On the 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

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click DataWorks.**
3. **On the page that appears, click O&M in the upper-right corner.**
4. **Click the Clusters tab at the top of the O&M page.**
5. **On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page appears.**

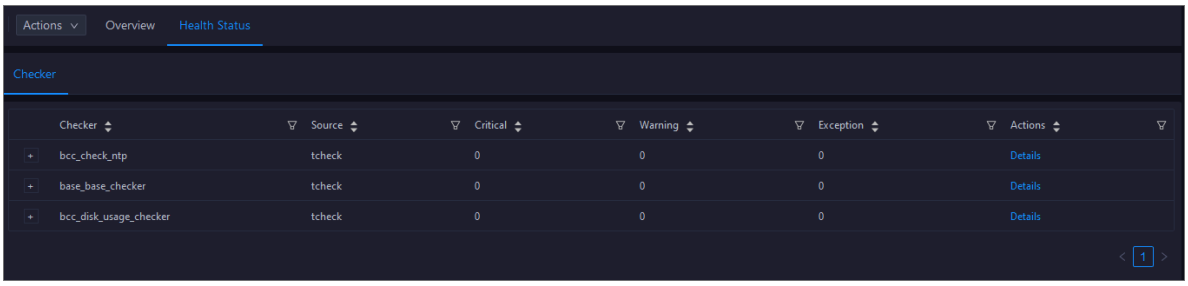

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to

the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. **On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.**

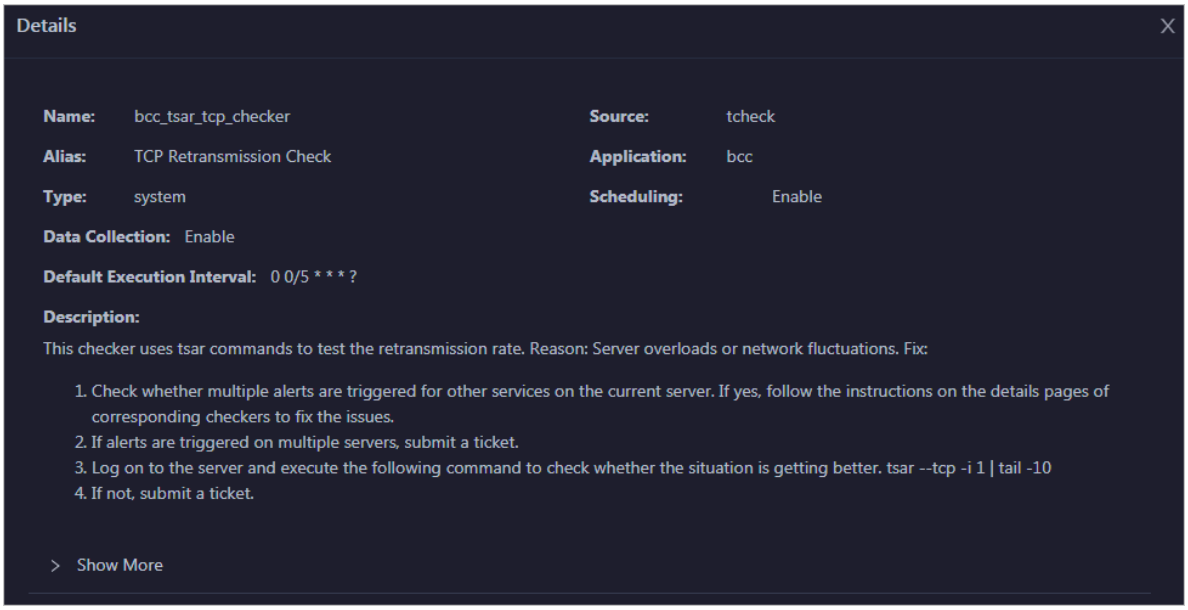

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

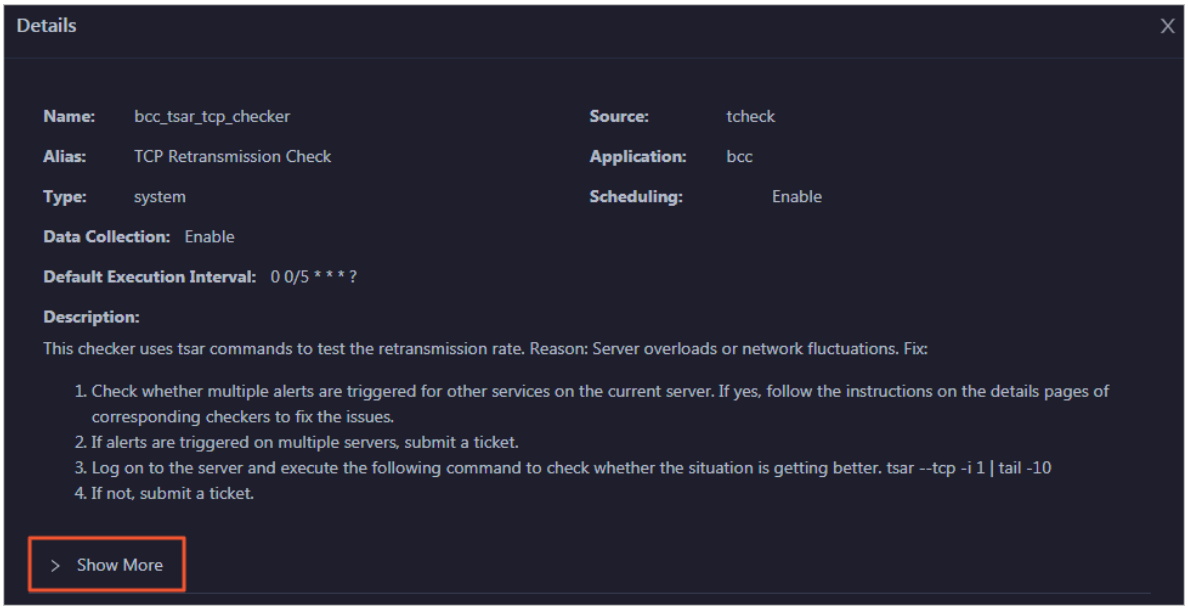

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View hosts with alerts and the alert causes

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

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

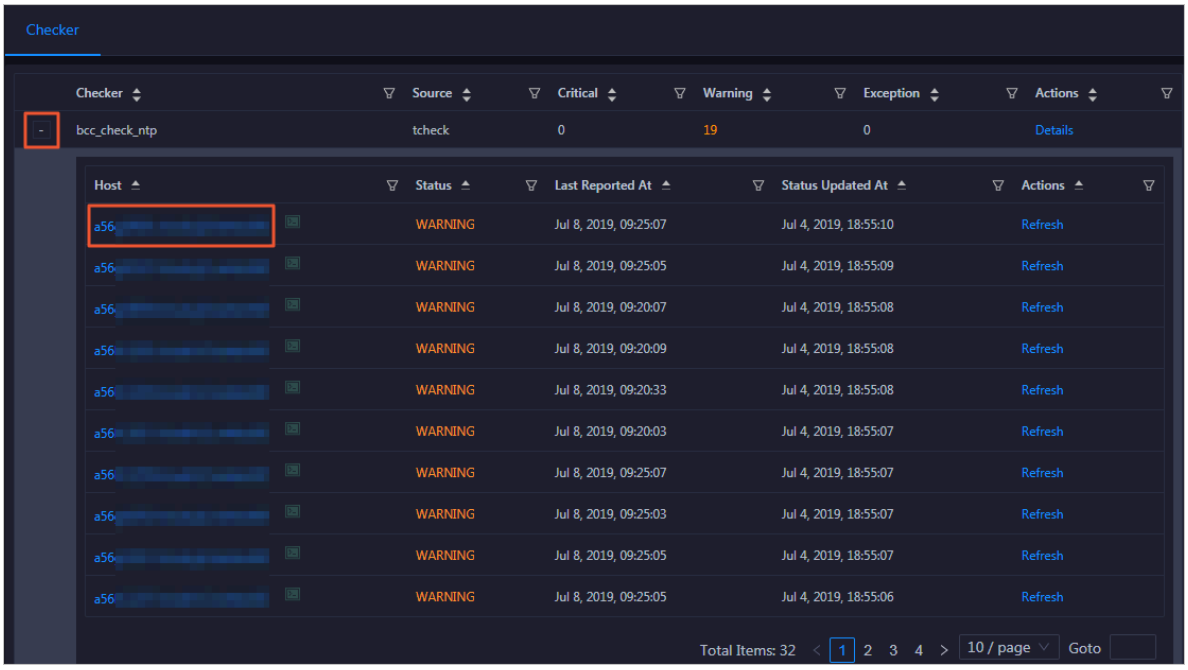

2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.

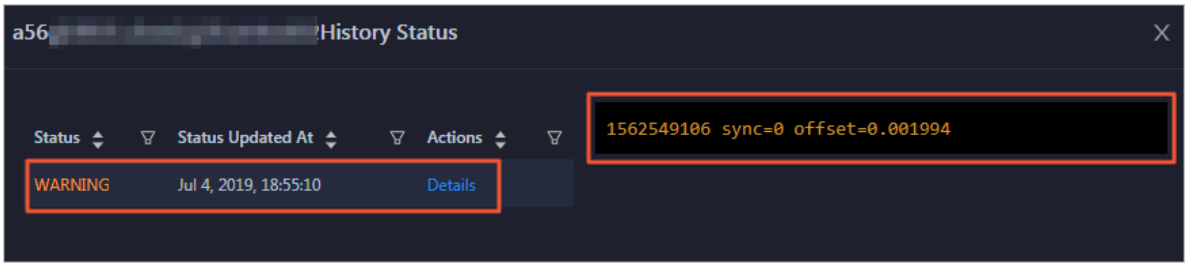

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

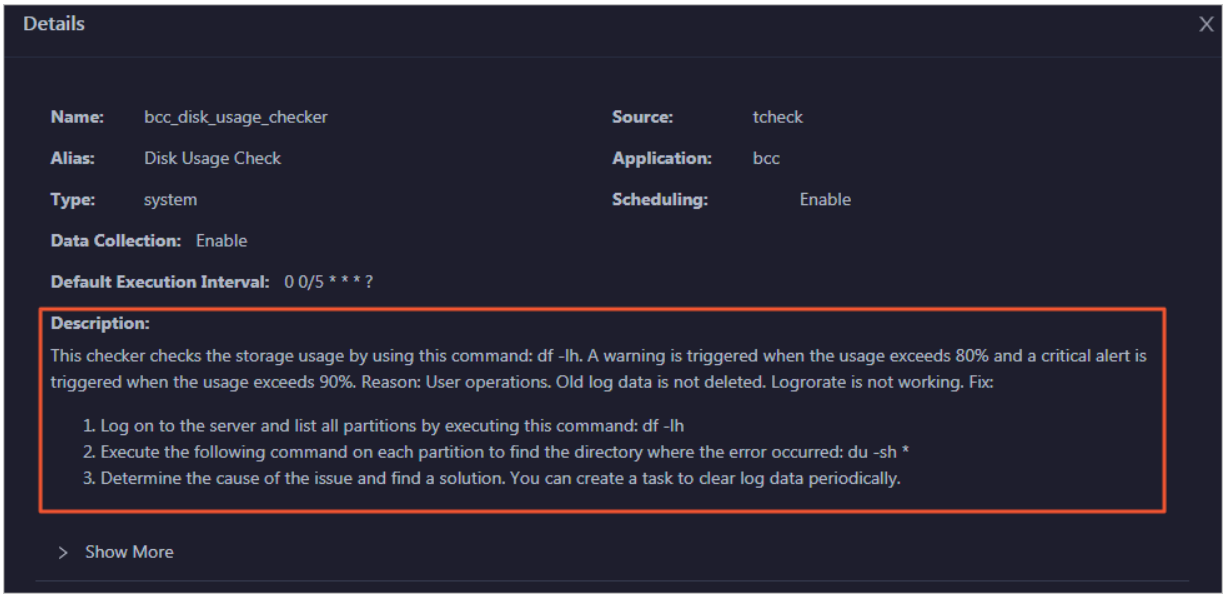

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. On the Health Status page, click + to expand a checker with alerts.

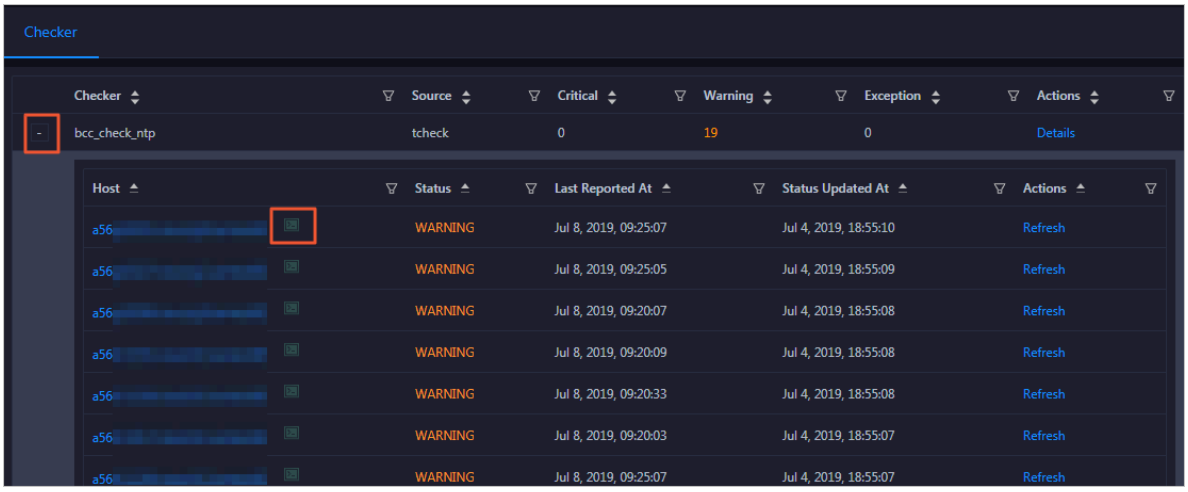

2. Click the Log On icon of a host. The TerminalService page appears.

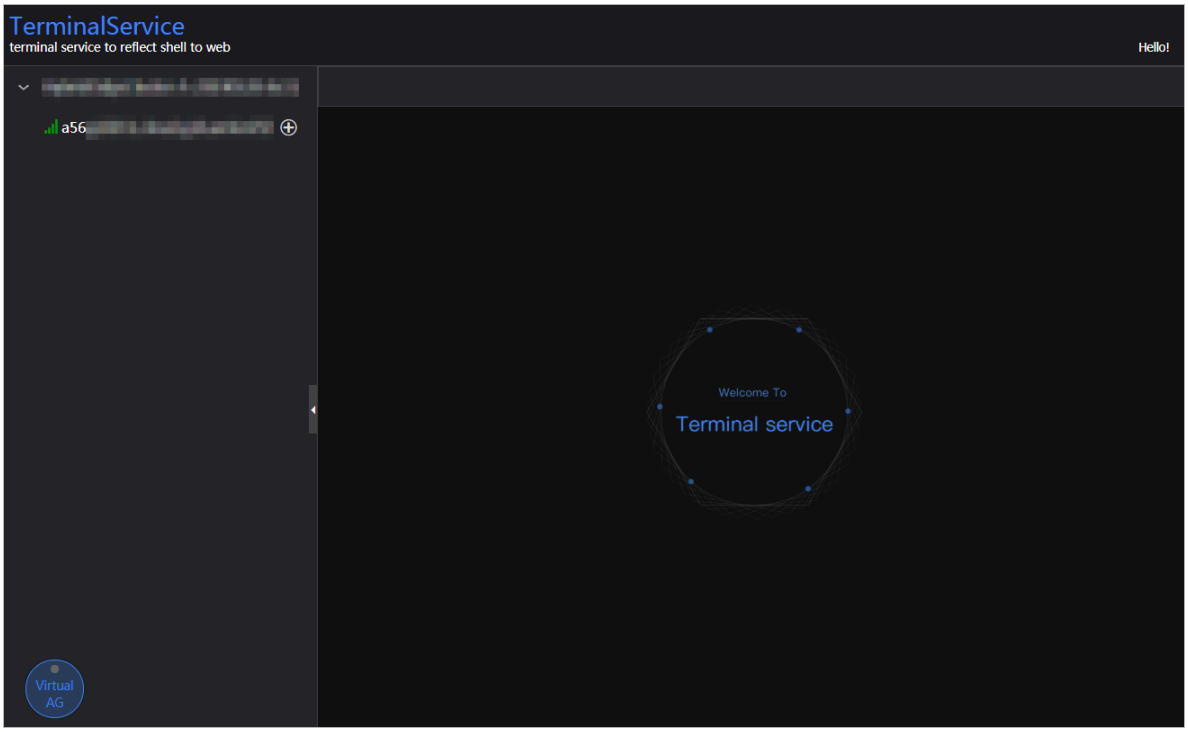

3. On the TerminalService page, click the hostname on the left to log on to the host.

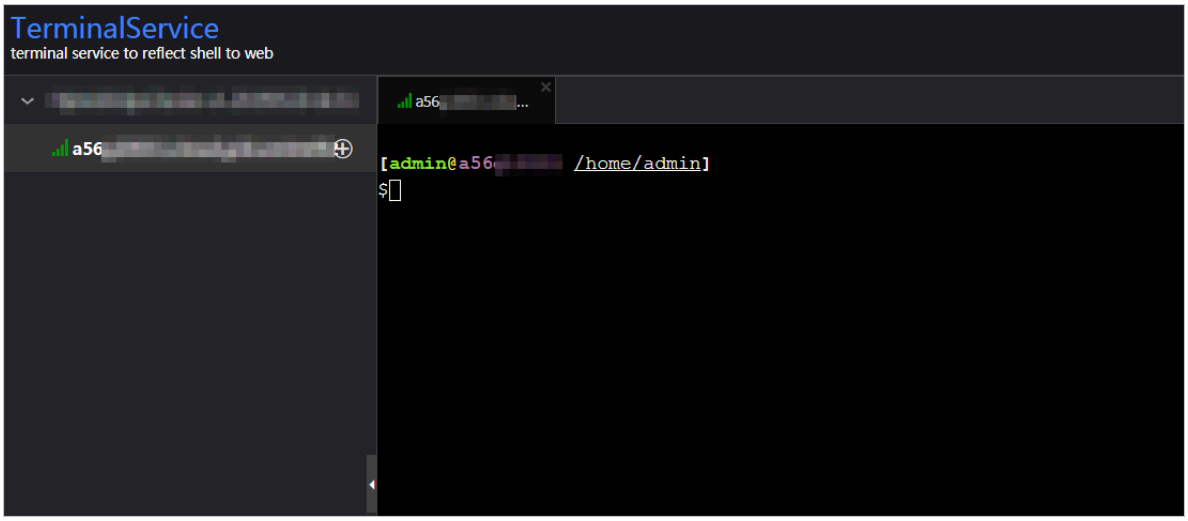

Run a checker again

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. In this way, you can check whether the alert is cleared.

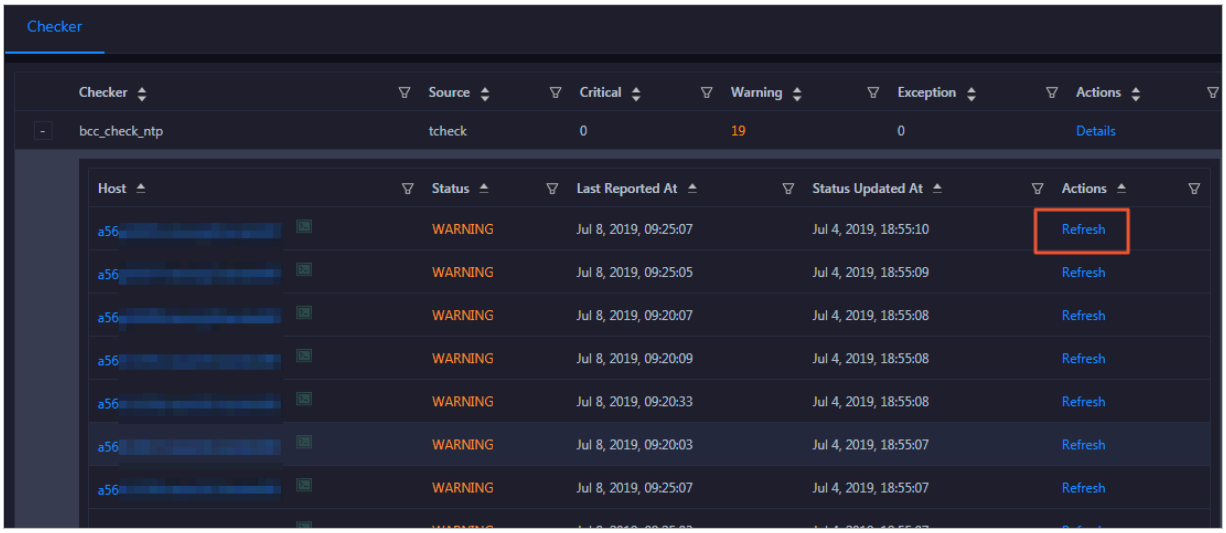

1.6.4 Host O&M

1.6.4.1 Host overview

The host overview page displays the overall running information about a host in a DataWorks 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

1. *Log on to the ABM console*.
2. **Click in the upper-left corner, and then click DataWorks.**
3. **On the page that appears, click O&M in the upper-right corner.**
4. **Click the Hosts tab at the top of the O&M page.**
5. **On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page 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.

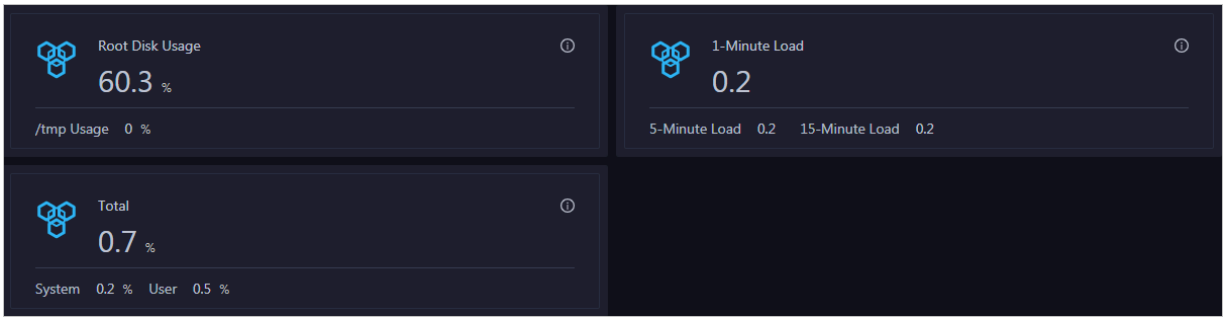

CPU

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 host over time in different colors.

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

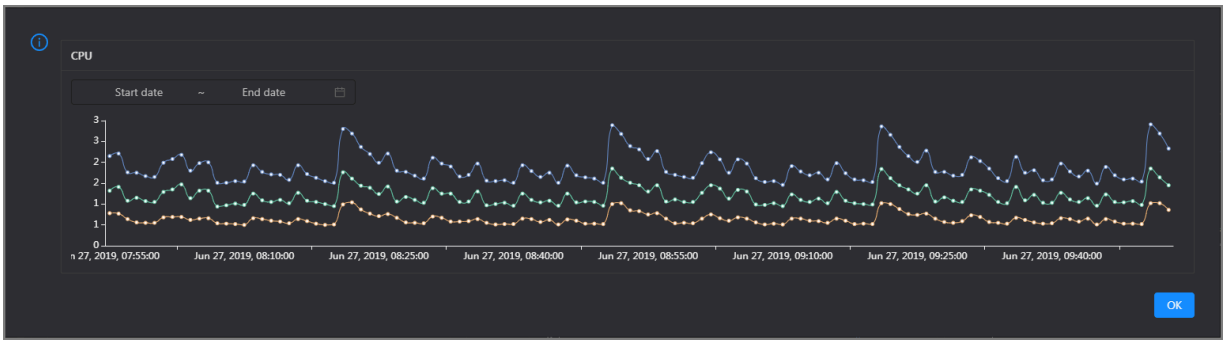

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

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

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

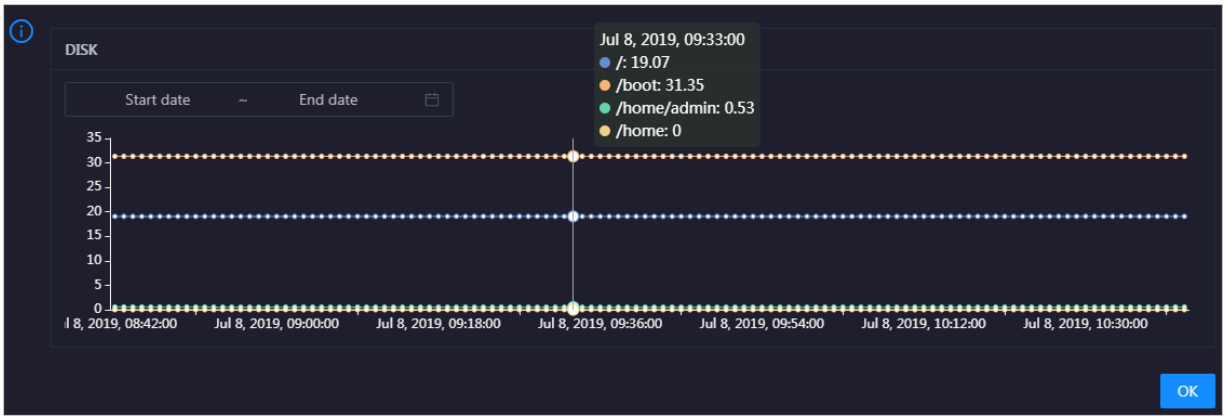

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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 host over time in different colors.

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

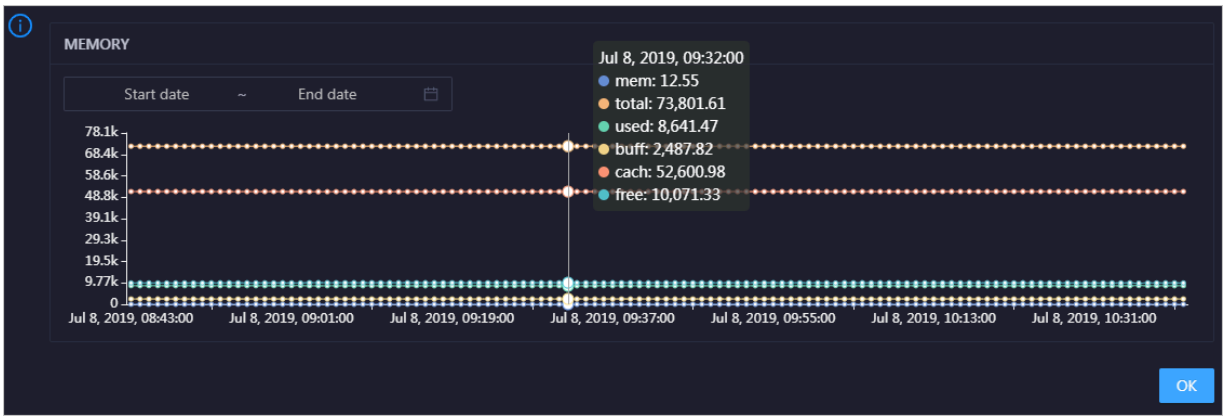

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

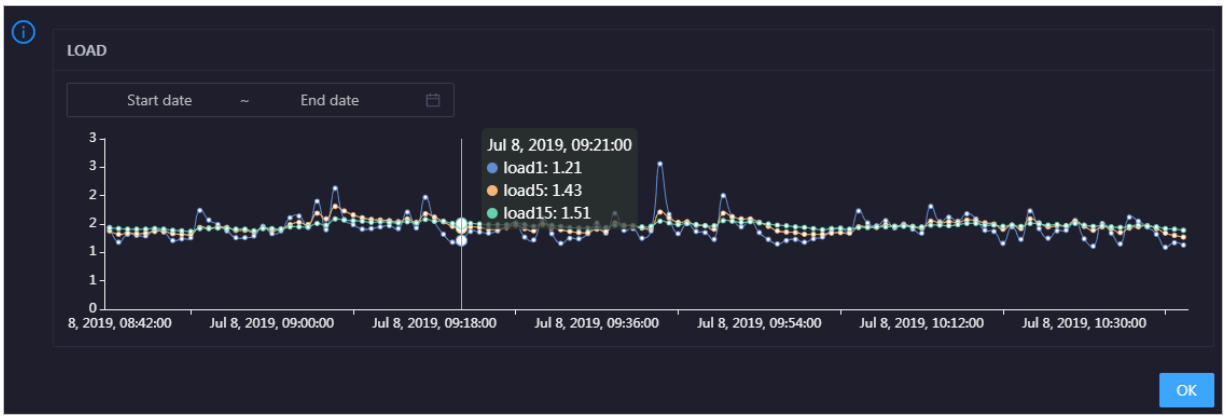

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 host over time in different colors. These trend lines reflect the data transmissi on status of the host.

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

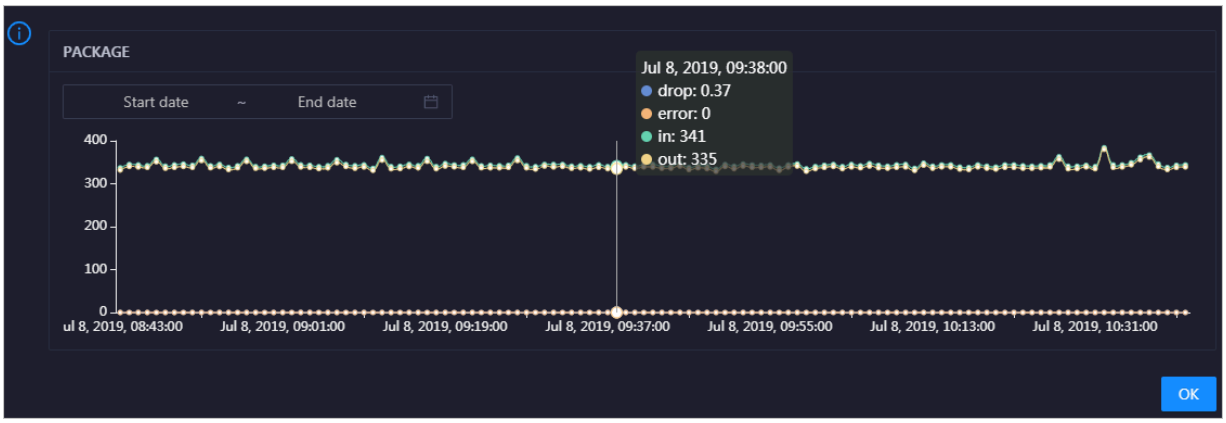

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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 ESTABLISHE D 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.

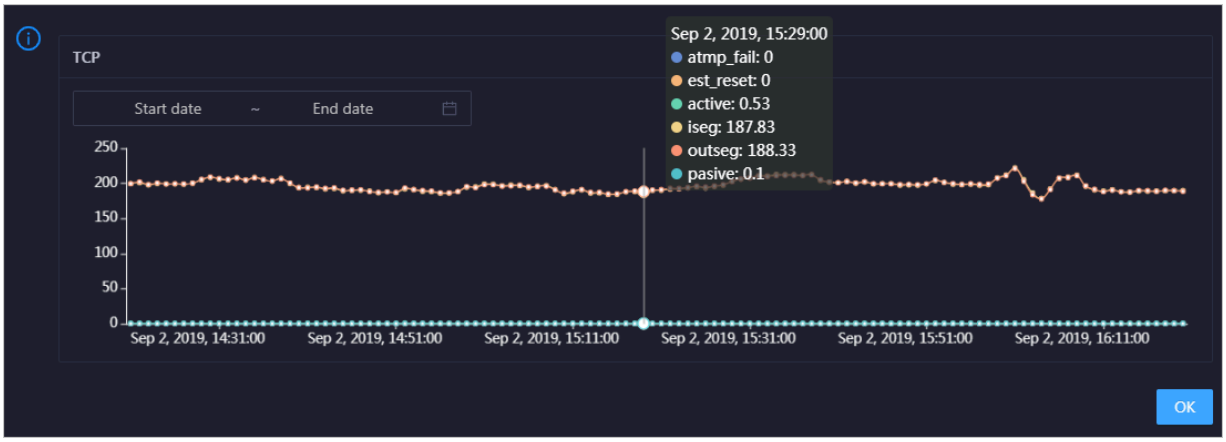

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

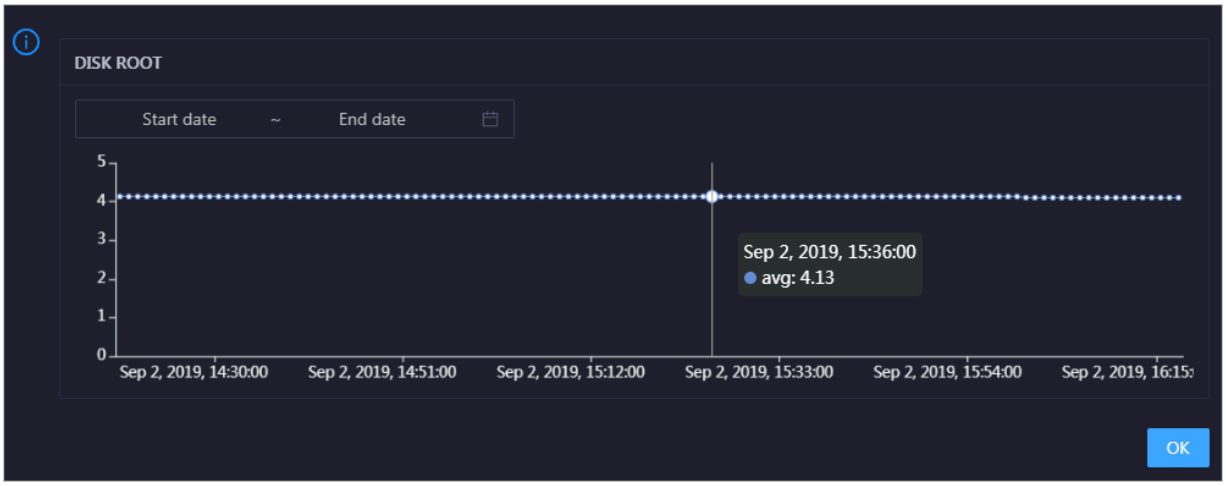

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

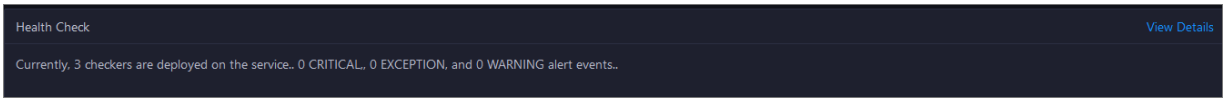

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

Health Check History

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

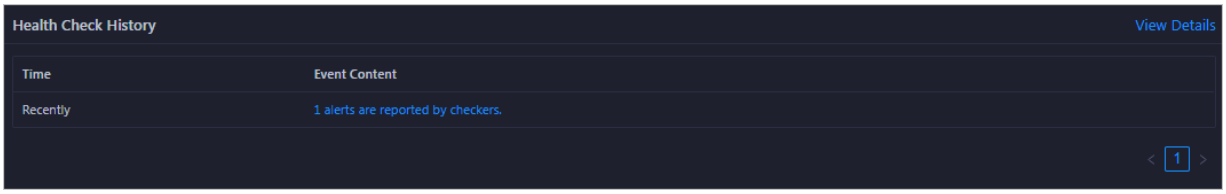

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

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

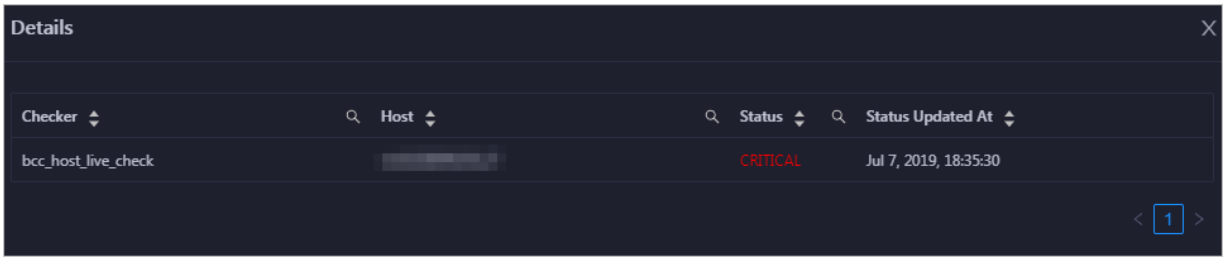

1.6.4.2 Host health

On the Health Status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click DataWorks.
3. On the page that appears, click O&M in the upper-right corner.
4. Click the Hosts tab at the top of the O&M page.
5. On the Hosts page, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page appears.

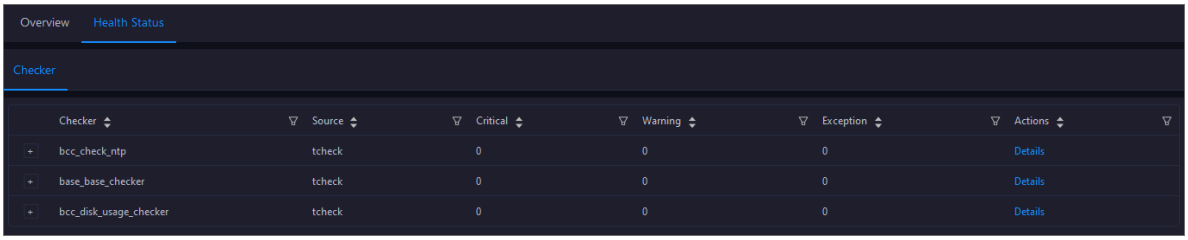

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

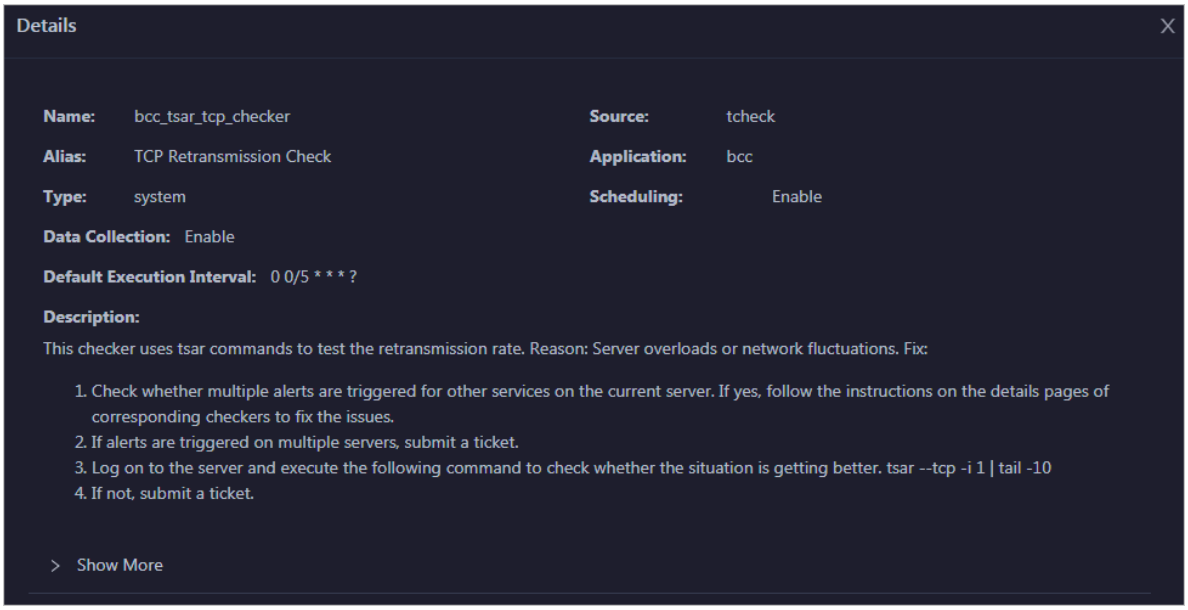

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

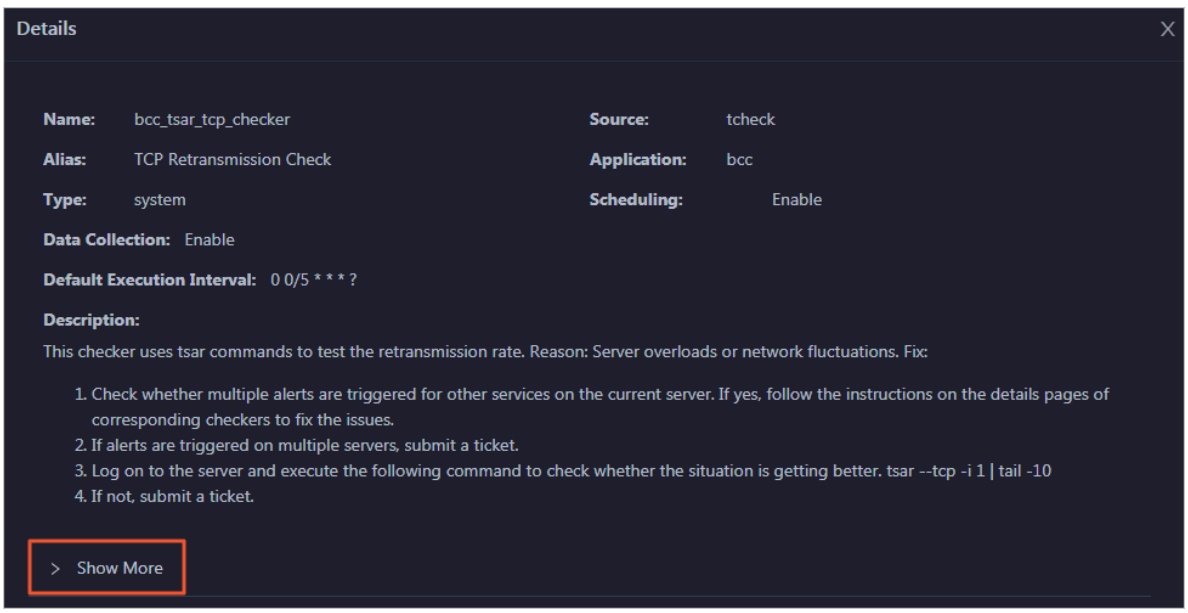

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

1. **On the Health Status page, click + to expand a checker with alerts.**

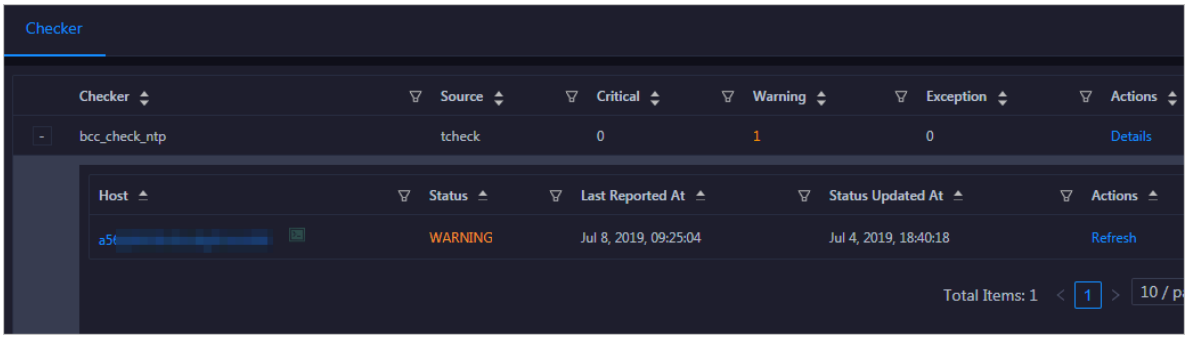

2. **Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

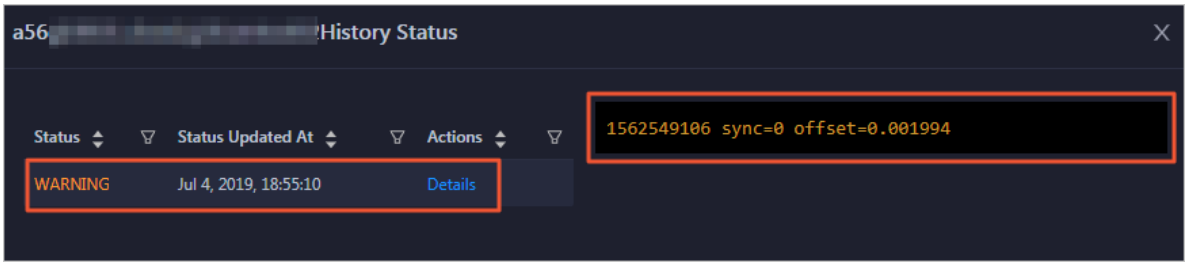

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

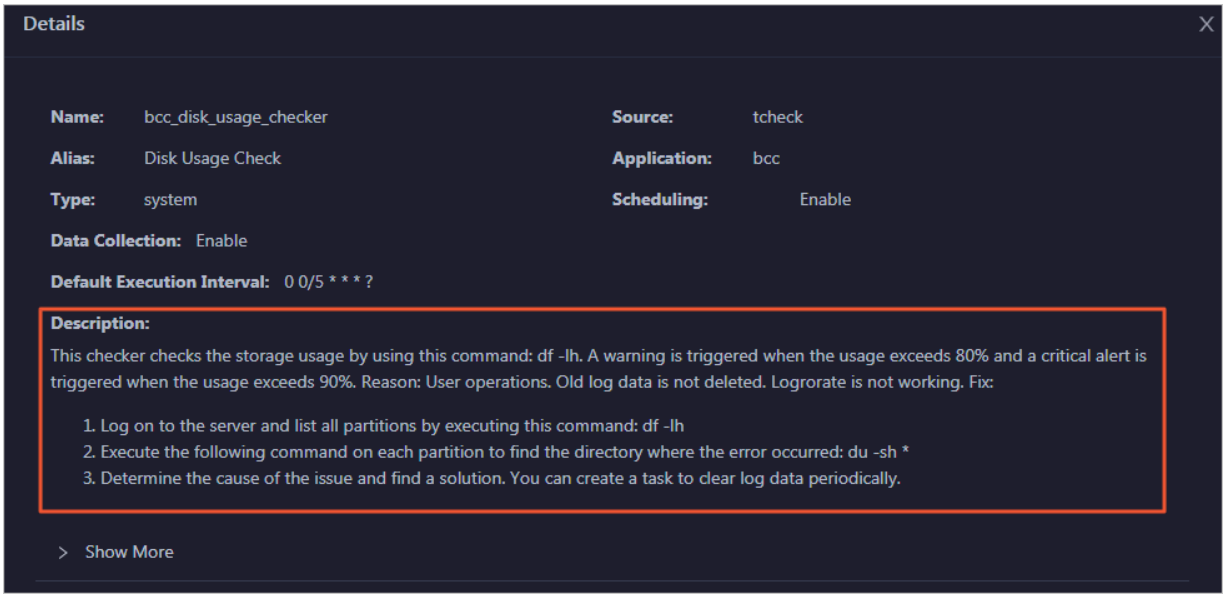

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

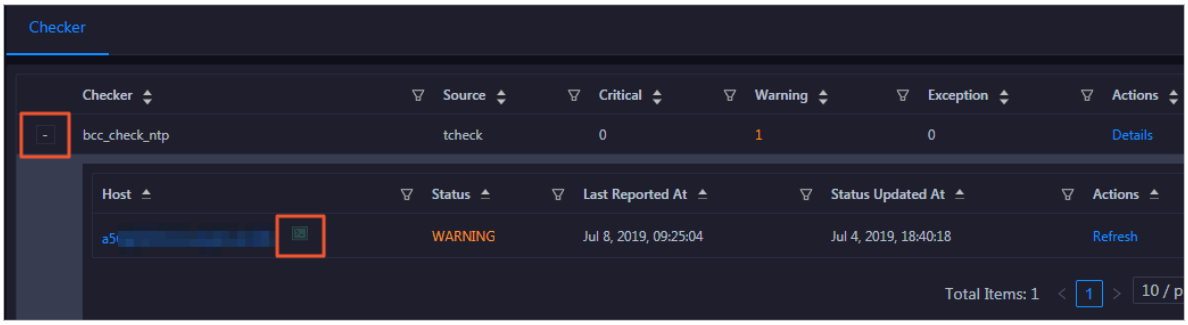

2. **Click the Log On icon of a host. The TerminalService page appears.**

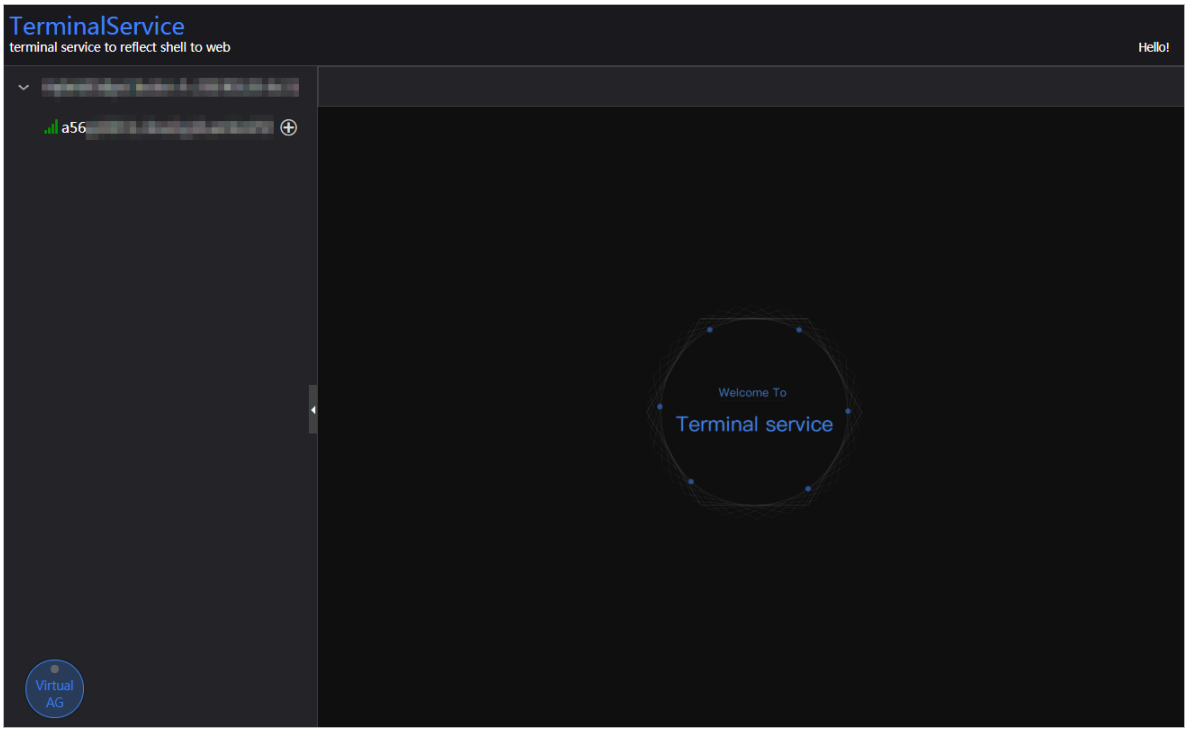

3. On the TerminalService page, click the hostname on the left to log on to the host.

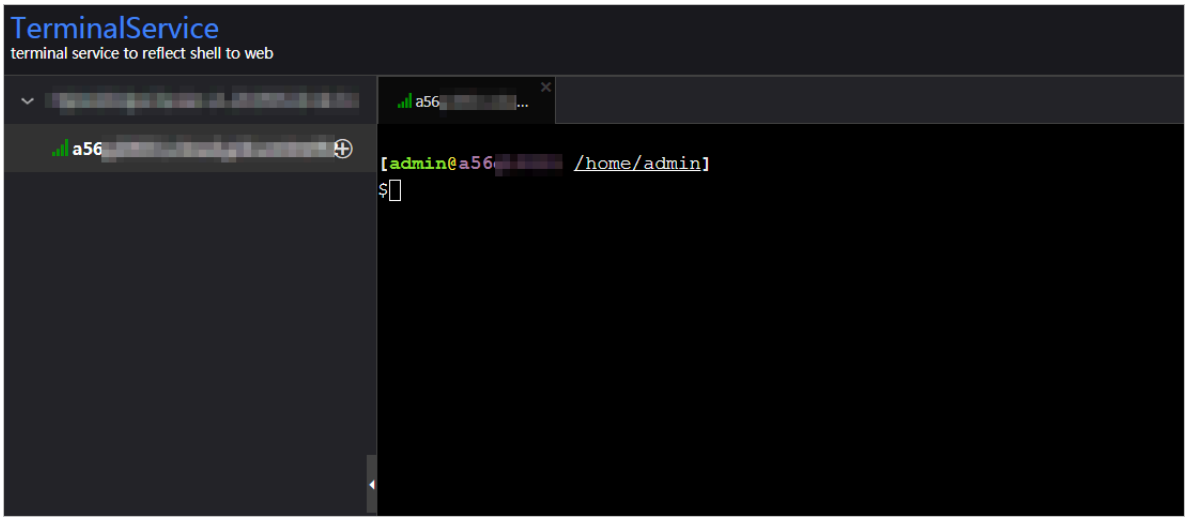

Run a checker again

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. In this way, you can check whether the alert is cleared.

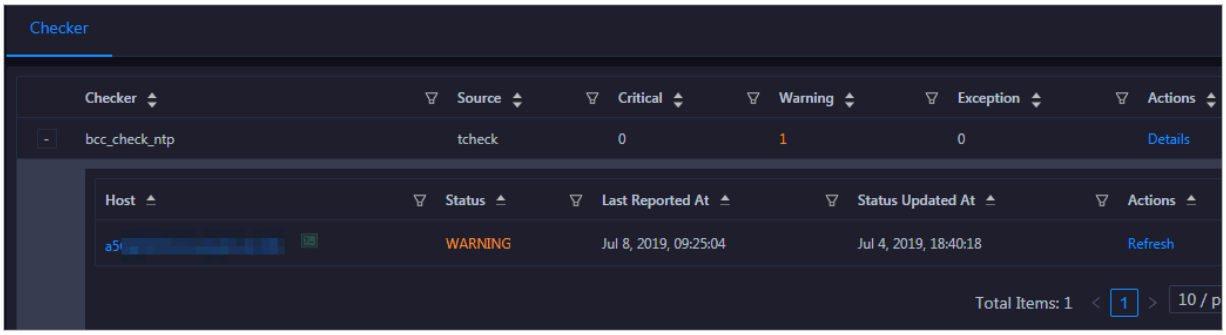

1.7 StreamCompute

1.7.1 StreamCompute O&M overview

This topic describes the features of Realtime Compute O&M supported by Apsara Bigdata Manager (ABM) and how to access the Realtime Compute O&M page.

Modules

Realtime Compute O&M includes business O&M, service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Business	Projects	Displays information about all projects in Realtime Compute.
	Jobs	Displays information about all jobs in Realtime Compute, and supports job diagnosis and analysis.
	Queues	Displays information about all queues in Realtime Compute.
Services	Blink	Displays the overview of the Blink service in Realtime Compute.
	Yarn	Displays the overview and health status of the YARN service in Realtime Compute.
	HDFS	Displays the overview and health status of the HDFS service in Realtime Compute.
Clusters	Overview	Displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster.
	Health Status	Displays 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.
	Hosts	Displays the information about hosts in a cluster , including the hostname, IP address, role, type , CPU usage, memory usage, root disk usage, packet loss rate, and packet error rate.
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.

Module	Feature	Description
	Health Status	Displays the checkers of the selected host, including the checker details, check results for the host, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click StreamCompute.
3. On the page that appears, click O&M in the upper-right corner. The Business page appears.

   The O&M page includes four modules, namely, Business, Services, Clusters, and Hosts.

1.7.2 Business O&M

1.7.2.1 Projects

This topic describes how to view information about the projects in Realtime Compute and how to go to the Queue Analysis page from the Projects page.

Projects

On the Business page, click Projects in the left-side navigation pane. The Projects page for Realtime Compute appears.

The Projects page displays information about the projects in Realtime Compute, including the name, BRS, queue, used compute units (CUs), total CUs, CU usage percentage, and number of jobs.

Go to the Queue Analysis page

The Queue Analysis page displays the status and resource usage of a queue, and information about jobs running in the queue, so that you can quickly know the running status of the queue.

On the Projects page, click a queue in the Queue column. The Queue Analysis page of the queue appears. For more information about the Queue Analysis page and operations that you can perform on this page, see *Queues*.

1.7.2.2 Jobs

This topic describes how to view information about the jobs in Realtime Compute and how to go to the Job Analysis page, Queue Analysis page, and Realtime Compute console from the Jobs page.

Jobs

On the Business page, click Jobs in the left-side navigation pane. The Jobs page for Realtime Compute appears.

The Jobs page displays information about the jobs in Realtime Compute, including the job names, users who created the jobs, projects to which the jobs belong, queues where the jobs are running, transactions per second (TPS) in the inbound direction, job latency, requested compute units (CUs), job statuses, and start time.

Go to the Realtime Compute console

On the Jobs page, click the content in the Failover column of a job to go to the Realtime Compute console.

Go to the Job Analysis page

The job analysis feature allows you to diagnose jobs to quickly troubleshoot job failures.

On the Jobs page, click a job in the Name column. The Job Analysis page of the job appears. For more information about the Job Analysis page and operations that you can perform on this page, see *Job analysis*.

Go to the Queue Analysis page

The Queue Analysis page displays the status and resource usage of a queue, and information about jobs running in the queue, so that you can quickly know the running status of the queue.

On the Jobs page, click a queue in the Queue column. The Queue Analysis page of the queue appears. For more information about the Queue Analysis page and operations that you can perform on this page, see *Queues*.

1.7.2.3 Queues

Apsara Bigdata Manager (ABM) allows you to view the information about the queues in Realtime Compute, including the queue names, queue statuses, minimum numbers of CPU cores and minimum memory capacity guaranteed for

the queues, maximum numbers of CPU cores and maximum memory capacity available for the queues, and numbers of jobs running in the queues.

Queues

On the Business page, click Queues in the left-side navigation pane. The Queues page for Realtime Compute appears.

The Queues page displays information about the queues in Realtime Compute, including the clusters to which the queues belong, queue names, queue statuses, requested compute units (CUs), minimum CUs guaranteed, maximum CUs available, and numbers of jobs running in the queues.

Go to the Queue Analysis page

The Queue Analysis page displays the status and resource usage of a queue, and information about jobs running in the queue, so that you can quickly know the running status of the queue.

On the Queues page, click a queue in the Queue column. The Queue Analysis page of the queue appears. For more information about the Queue Analysis page and operations that you can perform on this page, see *Queues*.

1.7.3 Service O&M

1.7.3.1 Blink

Apsara Bigdata Manager (ABM) allows you to view the overview of the Blink service in Realtime Compute.

On the Services page, click Blink in the left-side navigation pane. The Overview page for the Blink service appears.

The Overview page displays the overview, status, health check result, and health check history, as well as two core cluster metrics, transactions per second (TPS) and failover rate, of the Blink service.

1.7.3.2 Yarn

Apsara Bigdata Manager (ABM) allows you to view the overview and health status of the YARN service in Realtime Compute.

Overview

On the Services page, click Yarn in the left-side navigation pane. The Overview page for the YARN service appears.

The Overview page displays the health check result, health check history, application status, container status, node status, logical CPU usage, and logical memory usage for the YARN service.

Click View Details in the Health Check or Health Check History section. The Health Status page for the YARN service appears. On this page, you can view more details about the health check.

Heath status

On the Services page, click Yarn in the left-side navigation pane. Click the Health Status tab at the top of the Services page. The Health Status page for the YARN service appears.

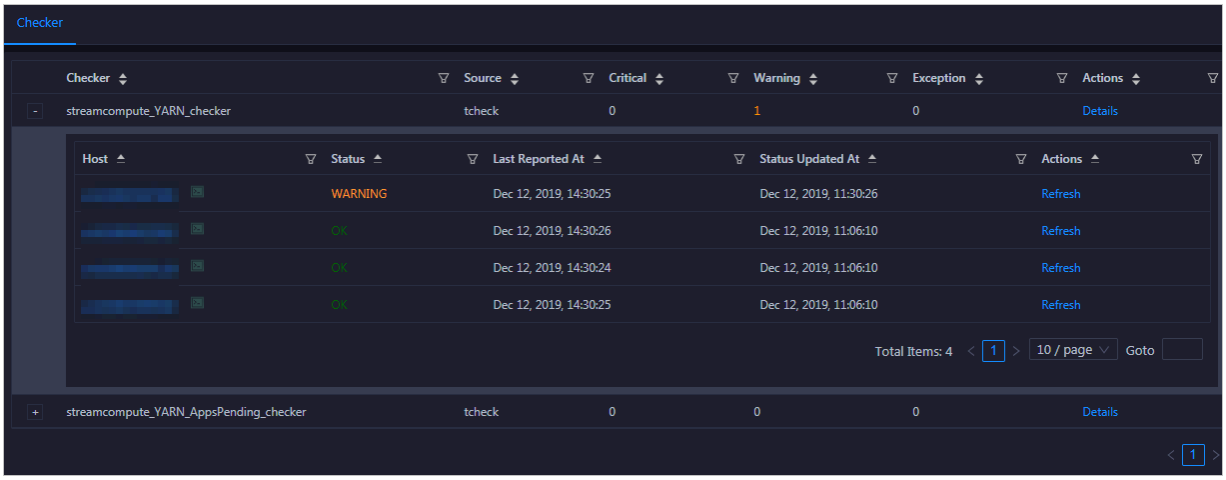

On the Health Status page, you can view all checkers of the YARN service and the check results for all hosts. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

The operations you can perform on the Health Status page for the YARN service are the same as those on the Health Status page for Realtime Compute clusters. For more information, see *Cluster health*.

1.7.3.3 HDFS

Apsara Bigdata Manager (ABM) allows you to view the overview and health status of the Hadoop Distributed File System (HDFS) service in Realtime Compute.

Overview

On the Services page, click HDFS in the left-side navigation pane. The Overview page for the HDFS service appears.

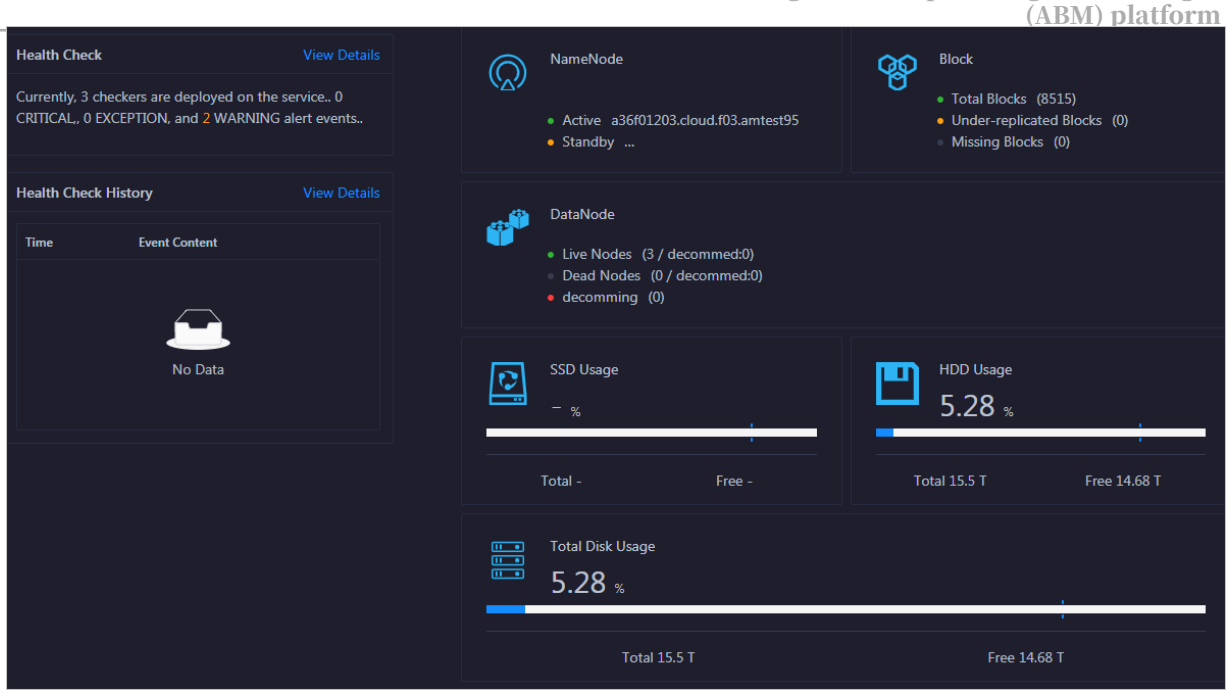

The Overview page displays the health check result, health check history, the information of NameNode, blocks, and DataNode, solid-state disk (SSD) usage, hard disk drive (HDD) usage, and total disk usage.

Click View Details in the Health Check or Health Check History section. The Health Status page for the HDFS service appears. On this page, you can view more details about the health check.

Health status

On the Services page, click HDFS in the left-side navigation pane. Click the Health Status tab at the top of the Services page. The Health Status page for the HDFS service appears.

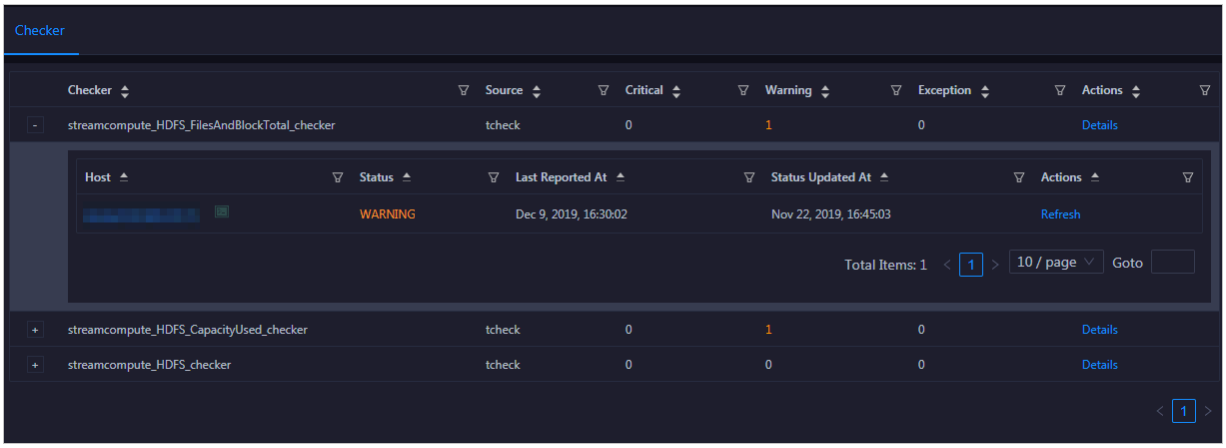

On the Health Status page, you can view all checkers of the HDFS service and the check results for all hosts in the cluster. The check results are divided into Critical,

Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

The operations you can perform on the Health Status page for the HDFS service are the same as those on the Health Status page for Realtime Compute clusters. For more information, see *Cluster health*.

1.7.4 Cluster O&M

1.7.4.1 Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the 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.

Hosts

This section displays all host statuses and the number of hosts in each status. The host statuses include good and bad.

Services

This section displays all services deployed in the cluster and the respective number of available and unavailable services.

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 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.

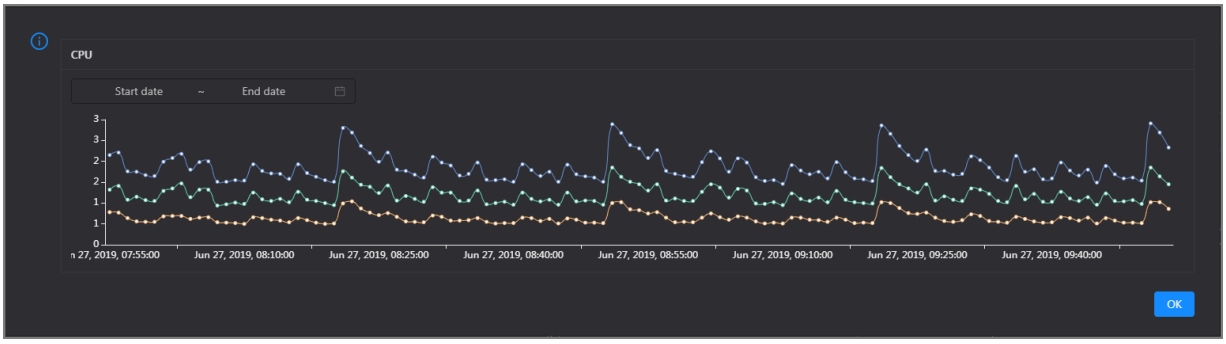

DISK

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

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

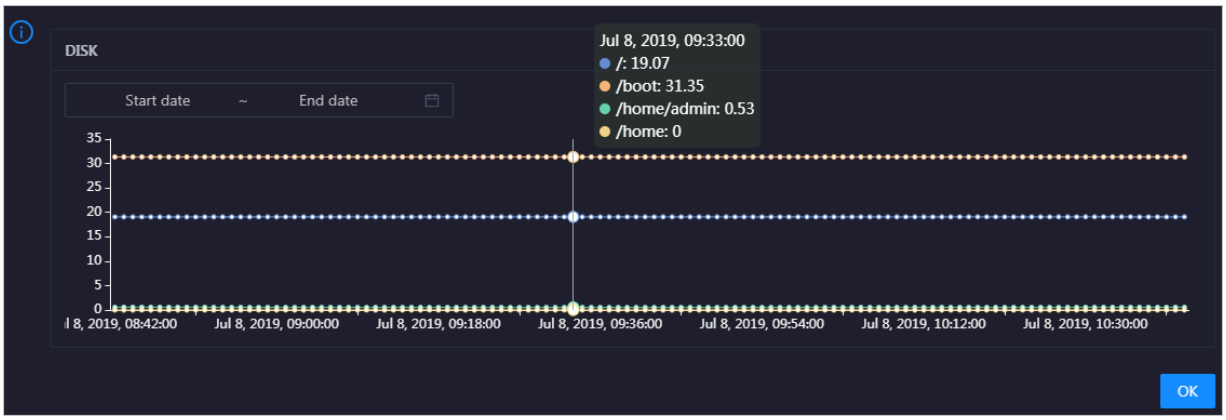

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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 cluster over time in different colors.

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

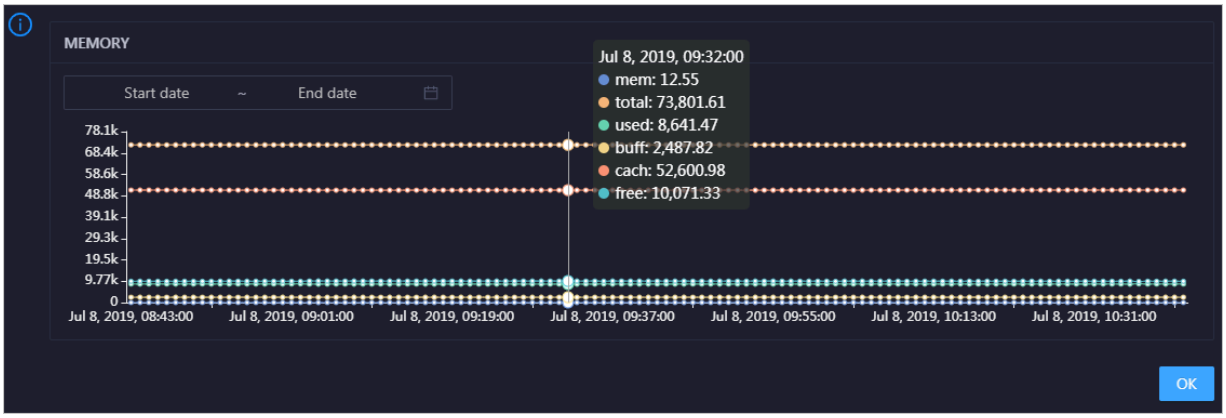

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

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

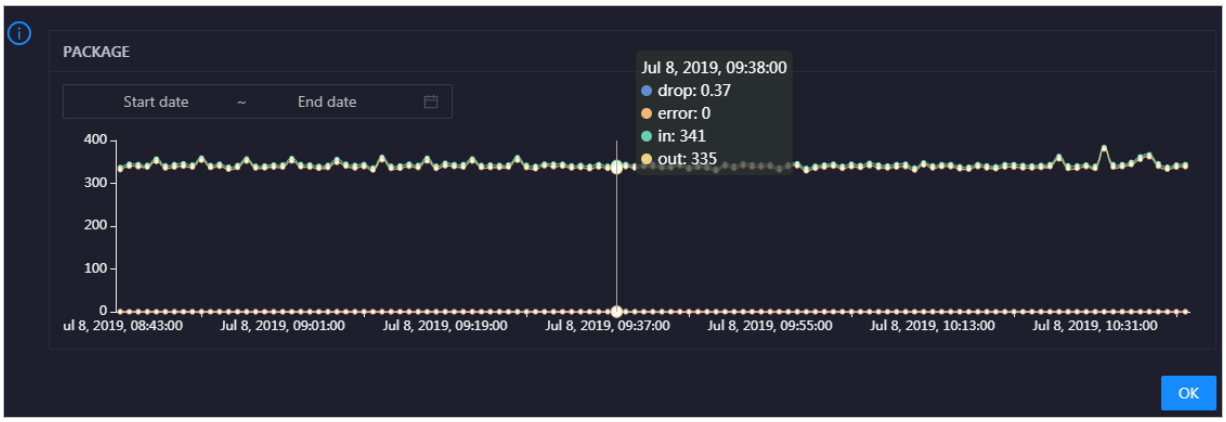

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

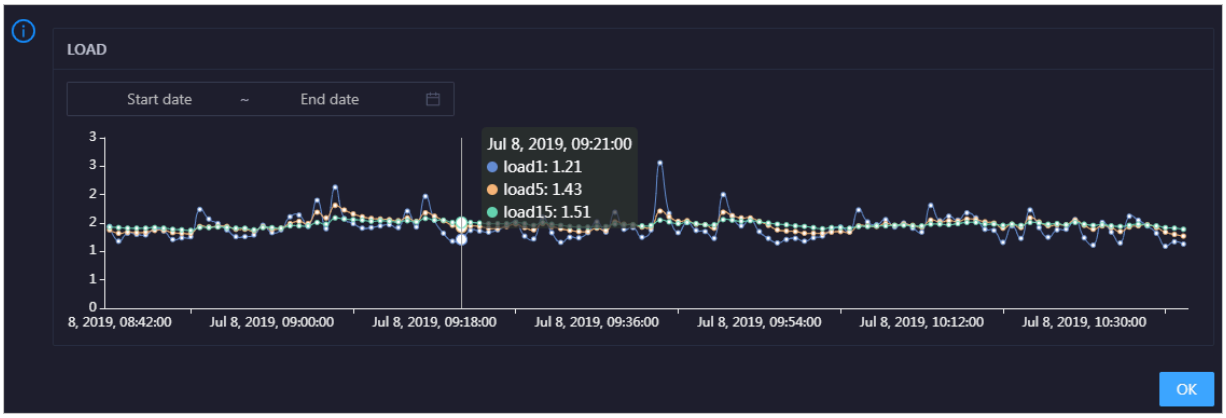

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

Health Check

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

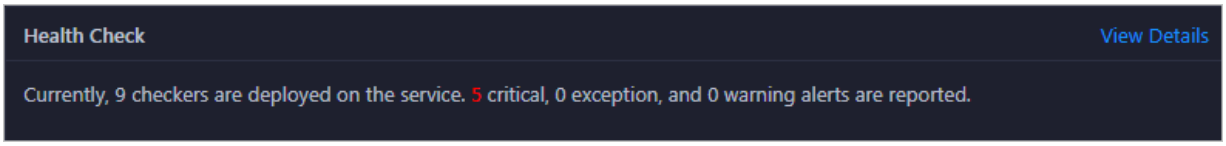

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Cluster health*.

Health Check History

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

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

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

1.7.4.2 Cluster health

On the 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.

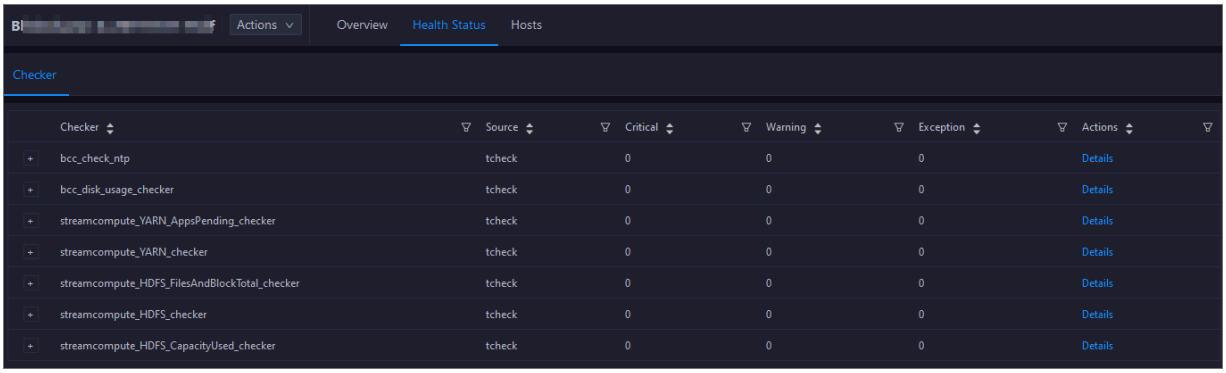

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

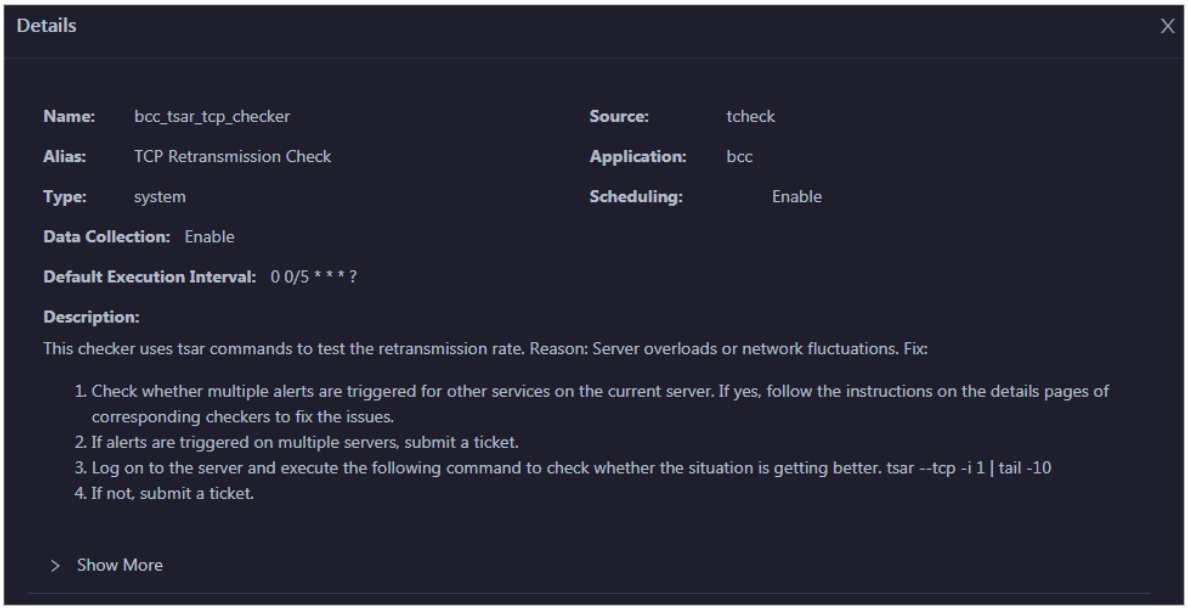

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

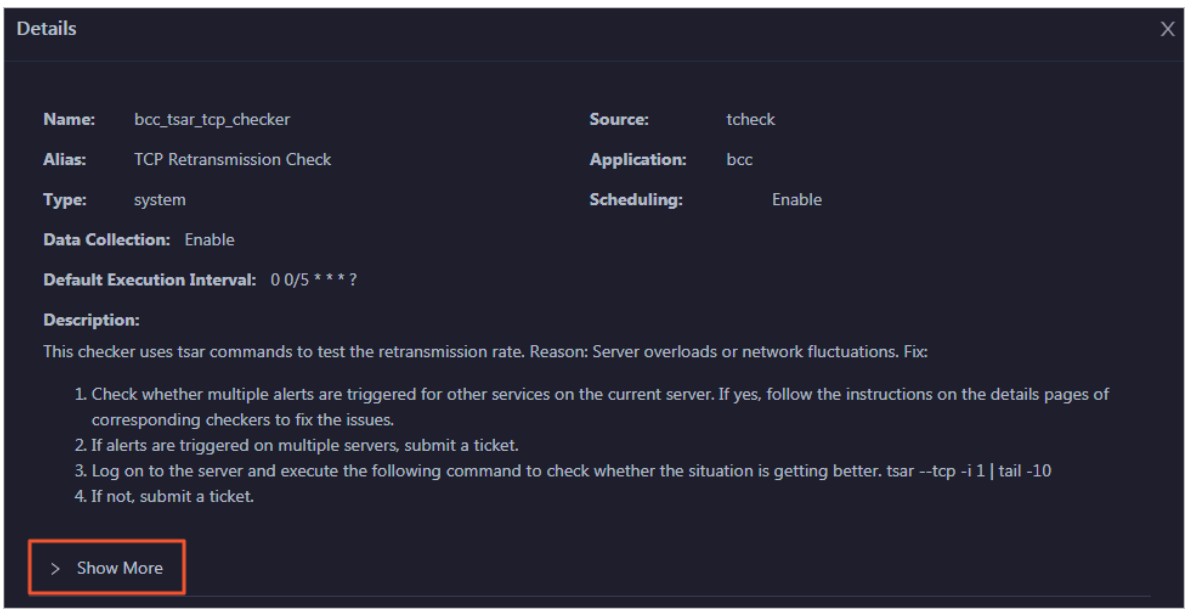

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View hosts with alerts and the alert causes

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

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

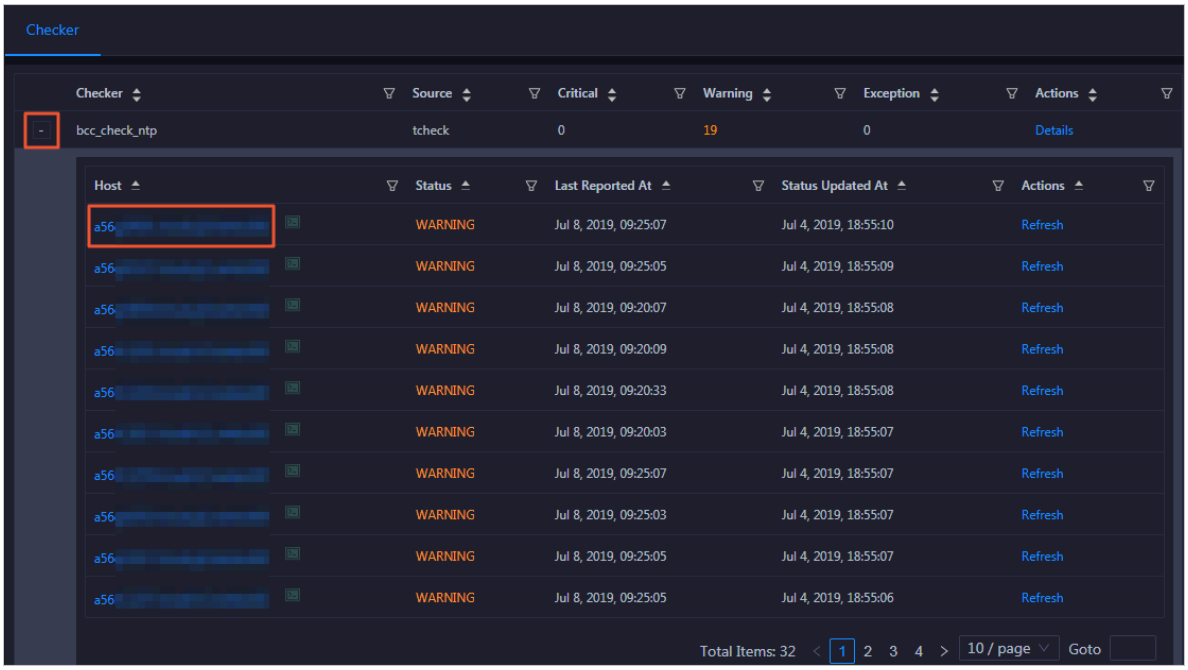

2. **Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

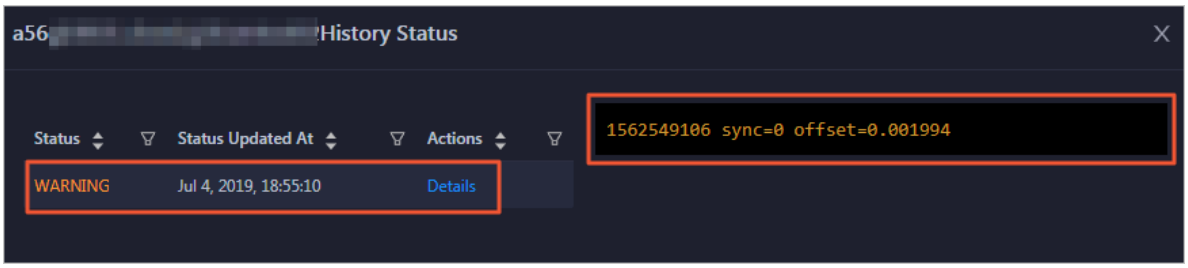

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

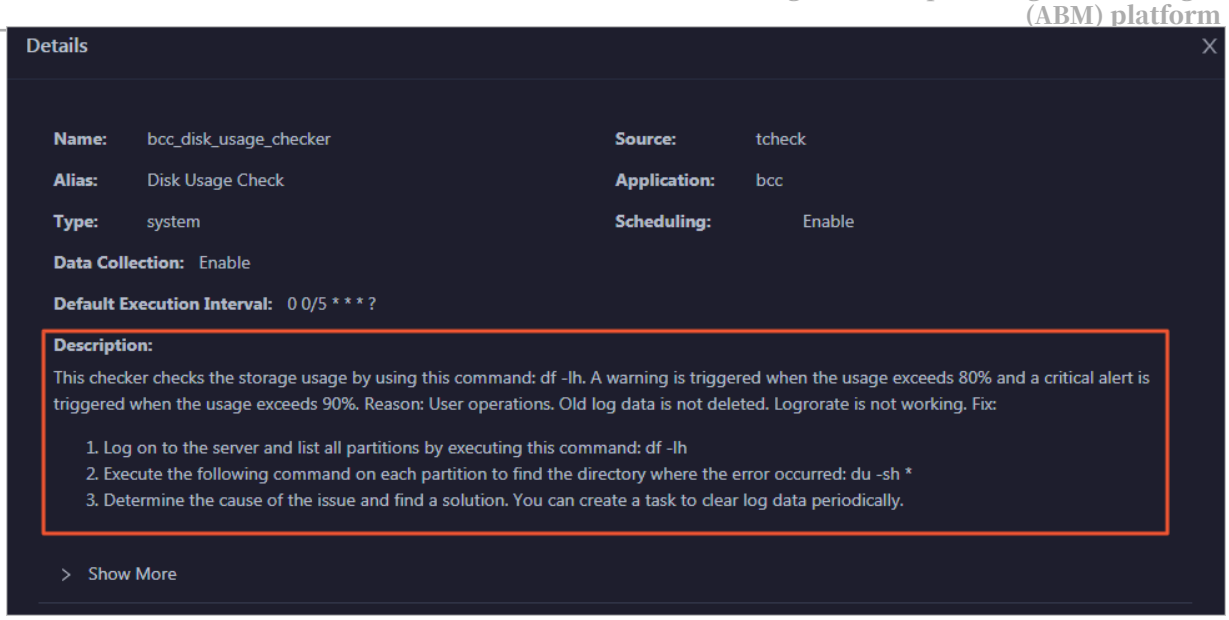

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

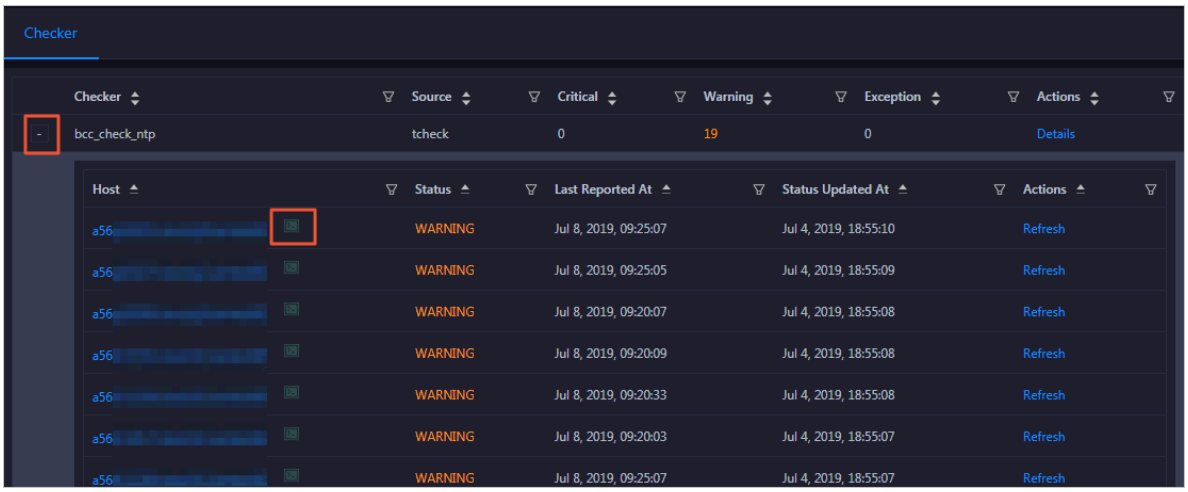

2. Click the Log On icon of a host. The TerminalService page appears.

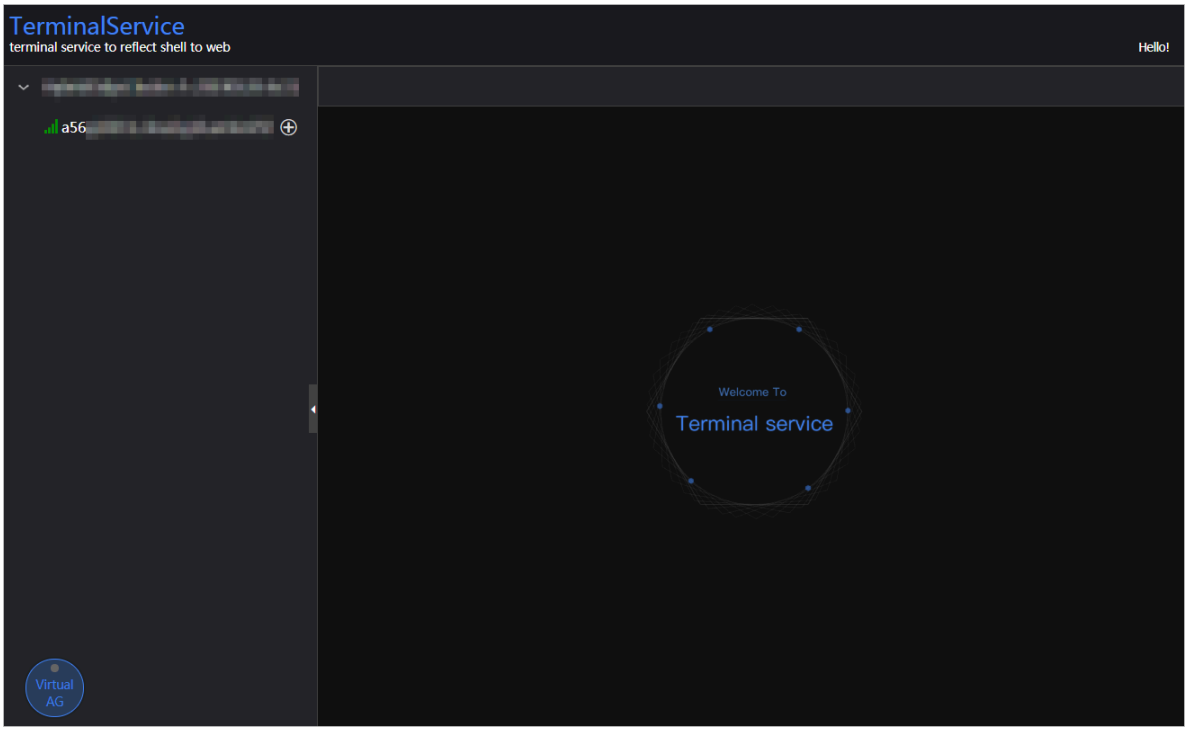

3. On the TerminalService page, click the hostname on the left to log on to the host.

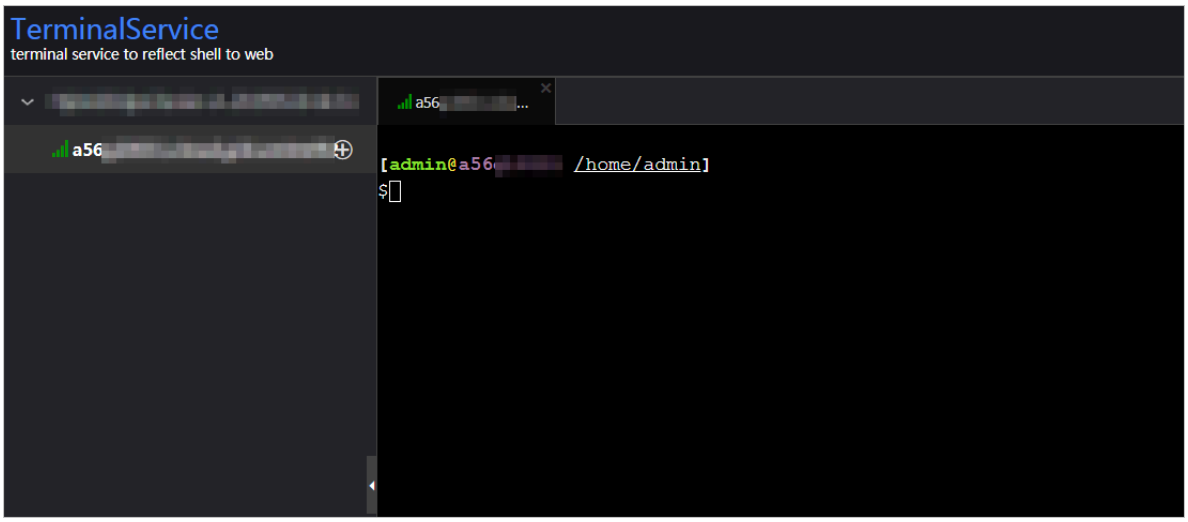

Run a checker again

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. In this way, you can check whether the alert is cleared.

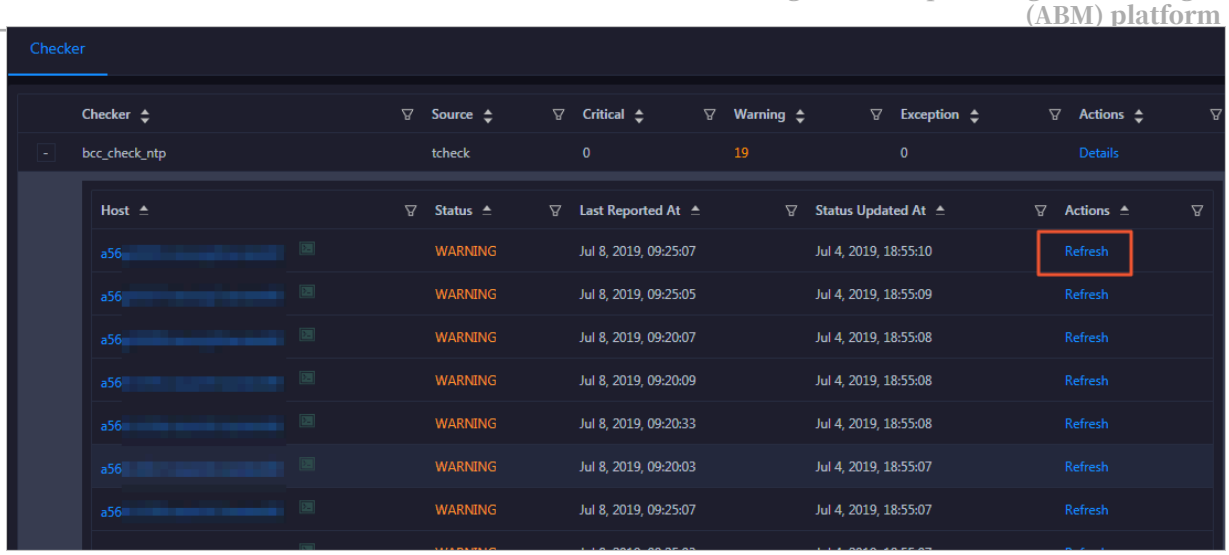

1.7.4.3 Hosts

The Hosts page displays information about hosts, including the hostname, IP address, role, type, CPU usage, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

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

To view more information about a host, click the name of the host. The Overview tab of the Hosts page appears. For more information, see *Host overview*.

1.7.4.4 Cluster scale-out

Apsara Bigdata Manager (ABM) allows you to scale out a Realtime Compute cluster by adding physical hosts. Cluster scale-out refers to the process of adding physical hosts in the default cluster of Apsara Infrastructure Management Framework to a Realtime Compute cluster. Currently, scale-out is only available for worker nodes in a Realtime Compute cluster.

Prerequisites

- **Your ABM account is granted the required permissions to perform O&M operations on Realtime Compute.**
- **Hosts whose service type is blink are deployed in the default cluster of Apsara Infrastructure Management Framework.**

Background

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to

the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework can be considered as an available resource pool that provides resources for scaling out business clusters. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

Step 1: Obtain the name of the host to be added to a Realtime Compute cluster

Before the scale-out, obtain the name of the host in the default cluster of Apsara Infrastructure Management Framework.

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click TIANJI to log on to the Apsara Infrastructure Management Framework console.
3. In the top navigation bar of the page that appears, choose Operations > Machine Operations.
4. On the Machine Operations page that appears, search for a host whose service type is blink in the default cluster. Copy the name of the host.

Step 2: Add the host to a Realtime Compute cluster

You can add multiple hosts to a Realtime Compute cluster at a time to scale out the cluster. To achieve this, you need to specify an existing host as the template host. When you scale out the Realtime Compute cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. On the O&M page of the ABM console, click the Clusters tab. On the page that appears, select a cluster in the left-side navigation pane. Click the Hosts tab, and then select a host whose role is Worker as the template host.

2. **Choose Actions > Scale out Cluster in the upper-left corner. In the Scale out Cluster dialog box that appears, set relevant parameters.**

   The parameters are described as follows:

   - **Refer Hostname: the name of the template host. By default, the name of the selected host is used.**
   - **Hostname: the name of the host to be added to the Realtime Compute cluster. The drop-down list displays all available hosts in the default cluster for scale-out. You can select one or more hosts from the drop-down list.**

3. **Click Run. A message appears, indicating that the action has been submitted.**
4. **View the scale-out status.**

   Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale out Cluster to view the scale-out history.

   It may take some time for the cluster to be scaled out. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

Step 3: View the scale-out progress

If the status is RUNNING, click Details in the Details column to view the steps and progress of the scale-out.

Step 4: Optional. Locate the cause of a scale-out failure

If the status is FAILED, click Details in the Details column to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.7.4.5 Cluster scale-in

Apsara Bigdata Manager (ABM) allows you to remove physical hosts to scale in a Realtime Compute cluster. Cluster scale-in refers to the process of removing physical hosts from a Realtime Compute cluster to the default cluster of Apsara Infrastructure Management Framework. Currently, scale-in is only available for the worker nodes in a Realtime Compute cluster.

Prerequisites

- **Your ABM account is granted the required permissions to perform O&M operations on Realtime Compute.**
- **More than three worker nodes are deployed in the current cluster. A Realtime Compute cluster creates three replicas for data by default. At least three worker nodes are required. Make sure that the cluster has at least three worker nodes after scale-in.**
- **Resources of the cluster, including the disk, CPU, and memory, are checked and still sufficient if the cluster is scaled in. For more information about how to check CPU and memory usage, see *Yarn*. You can run the `df` command to check disk usage.**

> **Notice:**
> **Scale-in triggers a job failover on hosts. If the cluster resources are insufficient after scale-in, the failover fails. This leads to negative effects on your business.**

Background

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework can be considered as an available resource pool that provides resources for scaling out business clusters. ABM allows you to scale in or out a cluster for your business. To scale out a cluster , add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

You can remove multiple hosts from a Realtime Compute cluster at a time to scale in the cluster.

Procedure

1. **On the O&M page of the ABM console, click the Clusters tab. On the page that appears, select a cluster in the left-side navigation pane. Click the Hosts tab, and then select one or more hosts whose role is Worker.**

2. **On the Clusters page, choose Actions > Scale in Cluster. The Scale in Cluster dialog box appears.**

   Hostname: the name of the host to be removed from the Realtime Compute cluster. By default, the name of the selected host is used.
3. **Click Run. A message appears, indicating that the action has been submitted.**
4. **View the scale-in status.**

   Move the pointer over Actions in the upper-left corner, and then click Execution History next to Scale in Cluster to view the scale-in history.

   It may take some time for the cluster to be scaled in. In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.
5. **View the scale-in progress.**

   If the status is RUNNING, click Details in the Details column to view the steps and progress of the scale-in.
6. **Locate the cause of a scale-in failure.**

   If the status is FAILED, click Details in the Details column to locate the failure cause.

   You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

1.7.5 Host O&M

1.7.5.1 Host overview

The host overview page displays the overall running information about a host in a Realtime Compute 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.

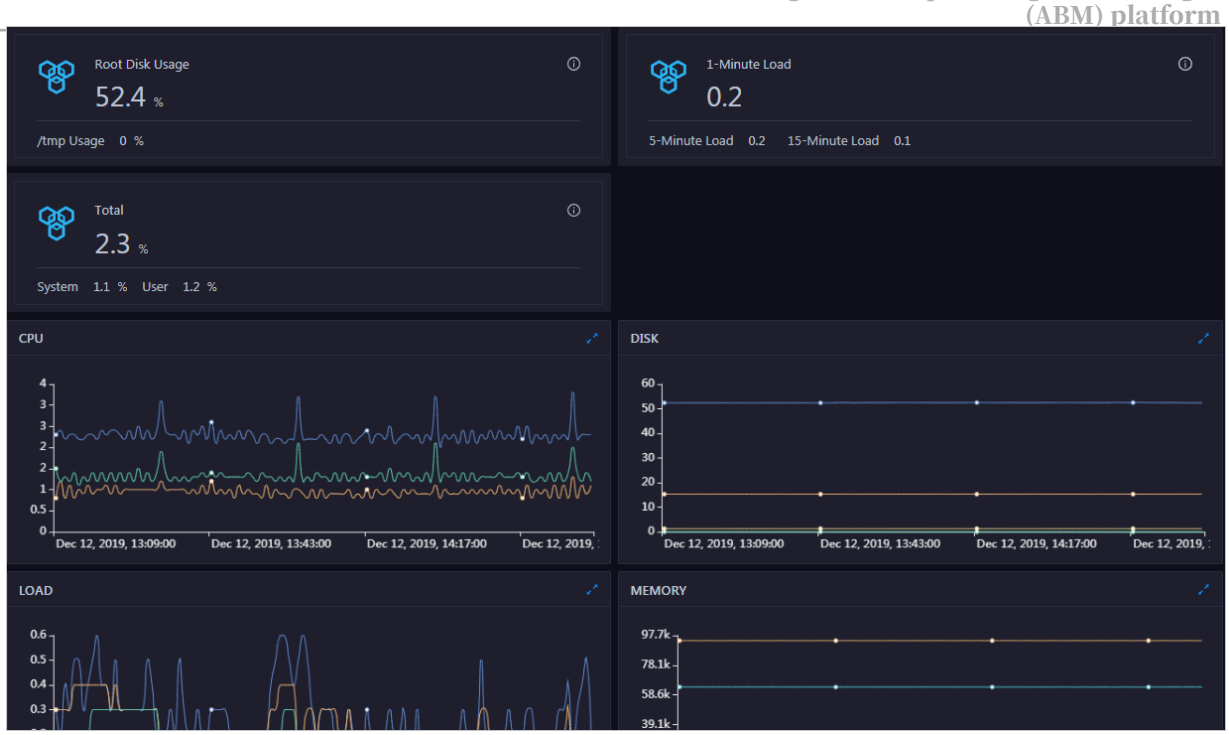

On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.

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.

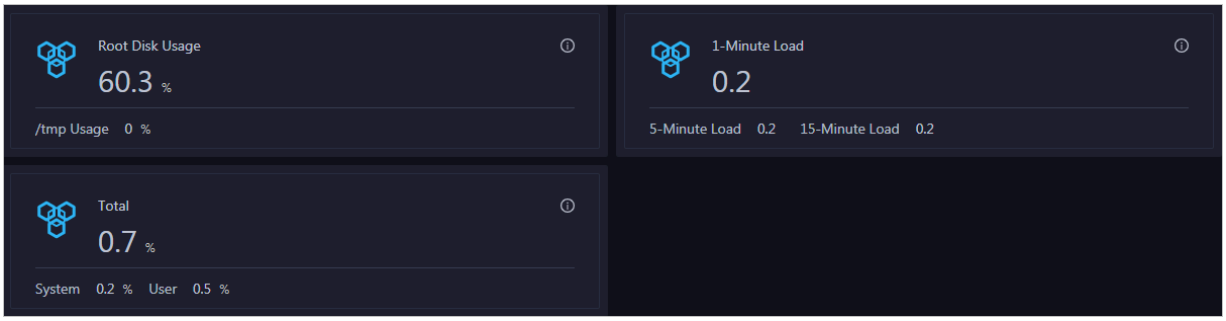

CPU

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 host over time in different colors.

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

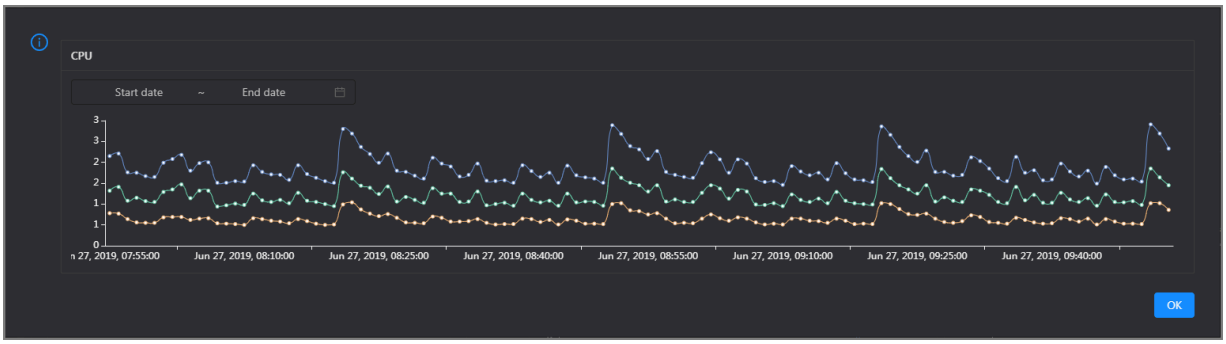

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

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

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

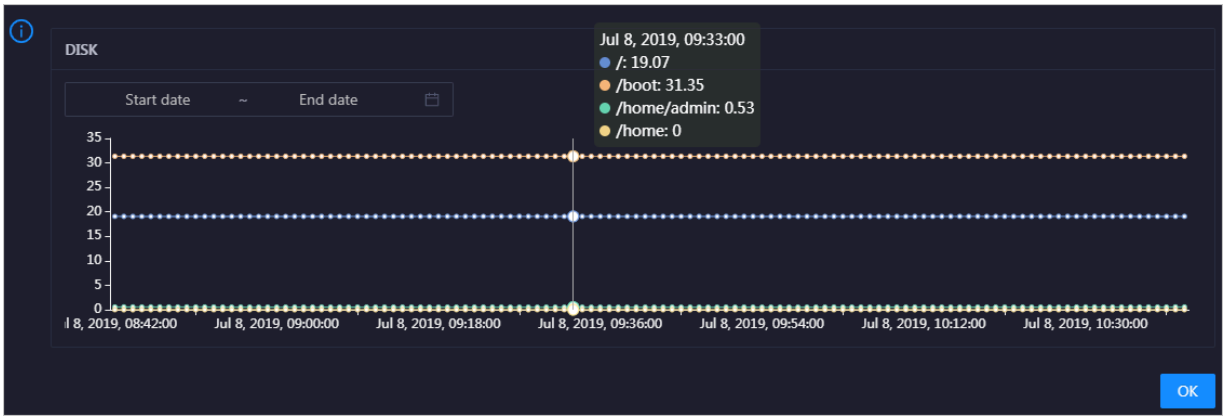

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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 host over time in different colors.

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

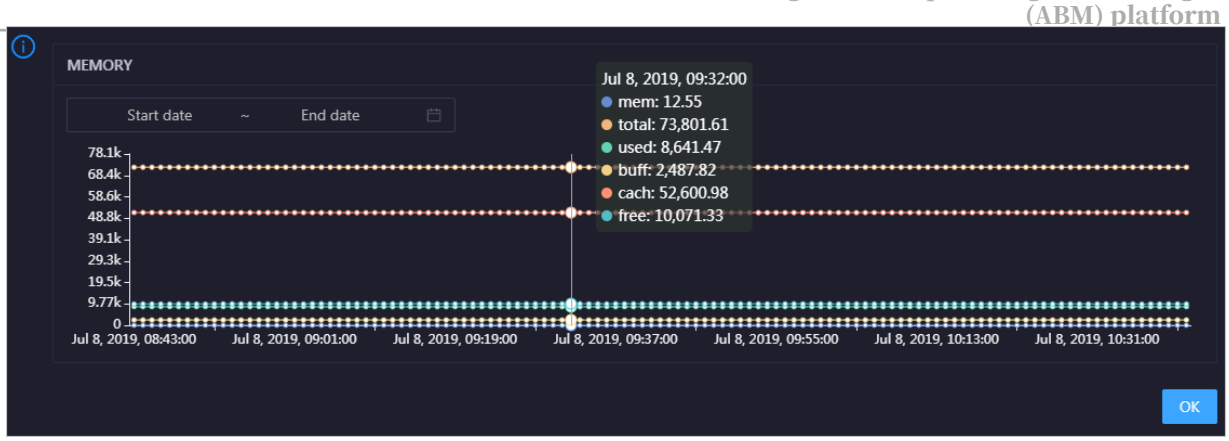

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

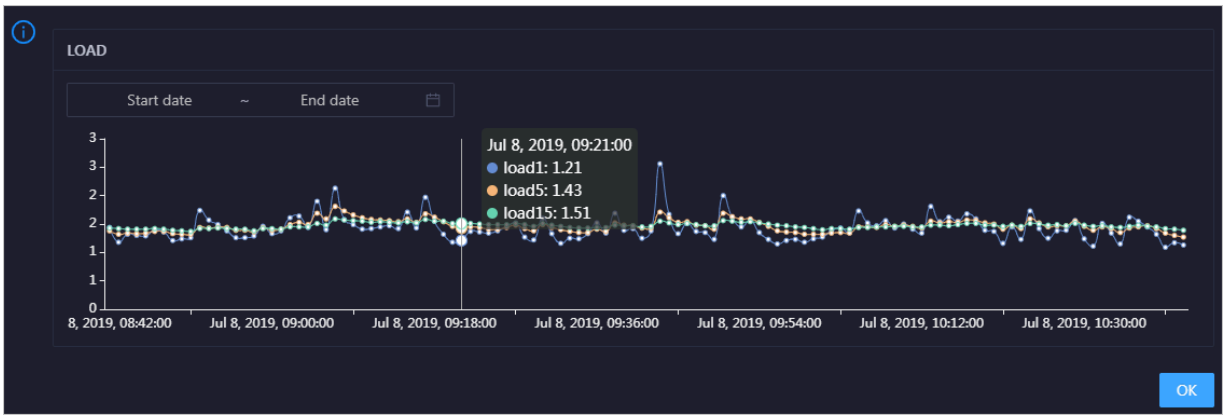

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 host over time in different colors. These trend lines reflect the data transmissi on status of the host.

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

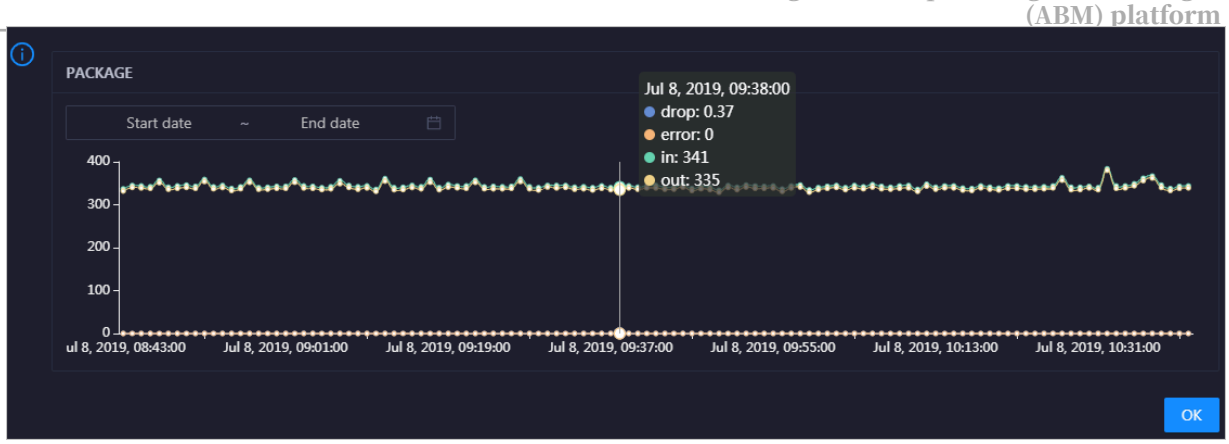

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

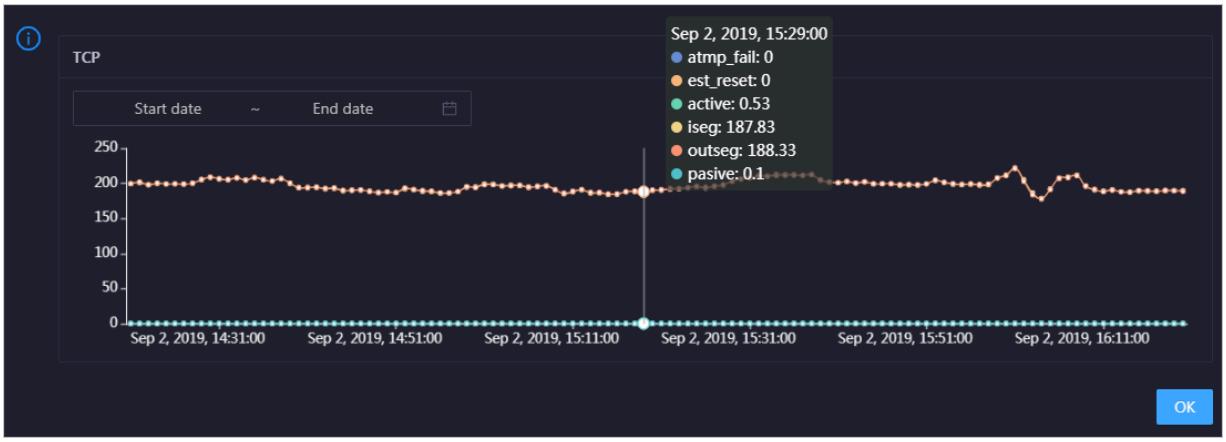

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

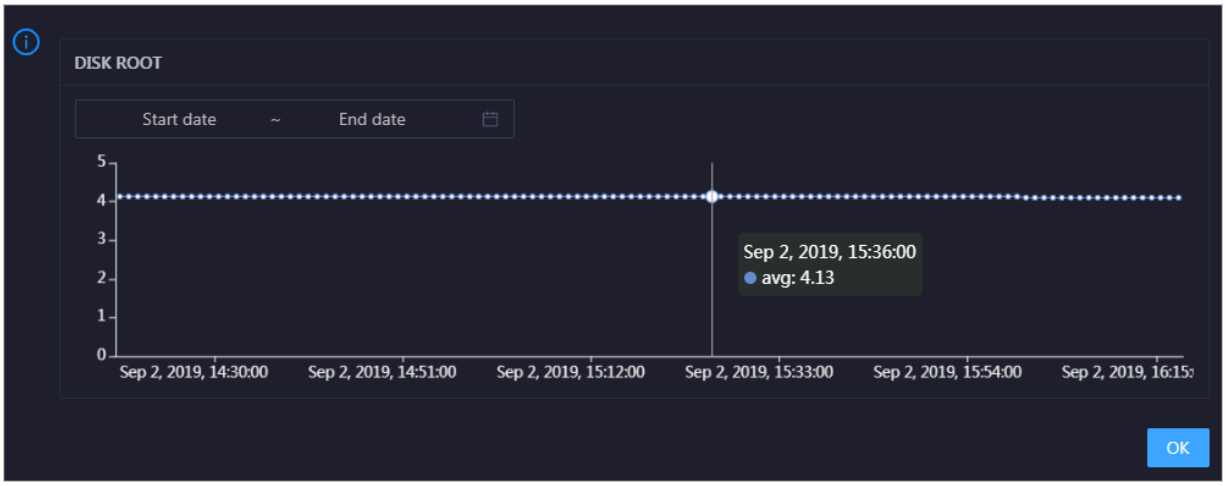

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

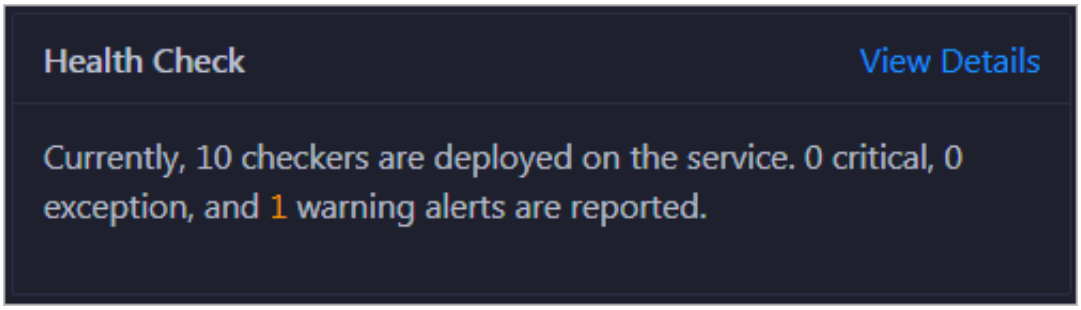

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

Health Check History

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

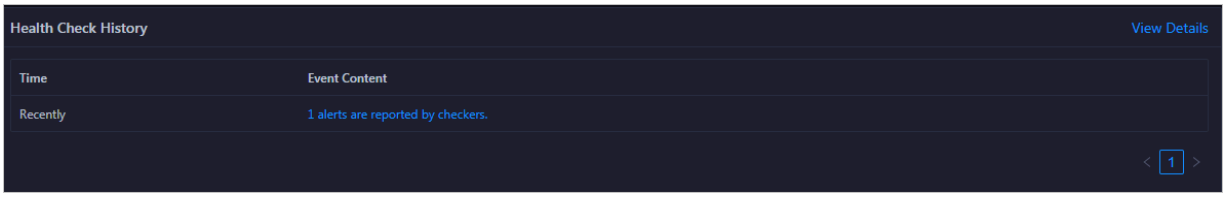

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

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

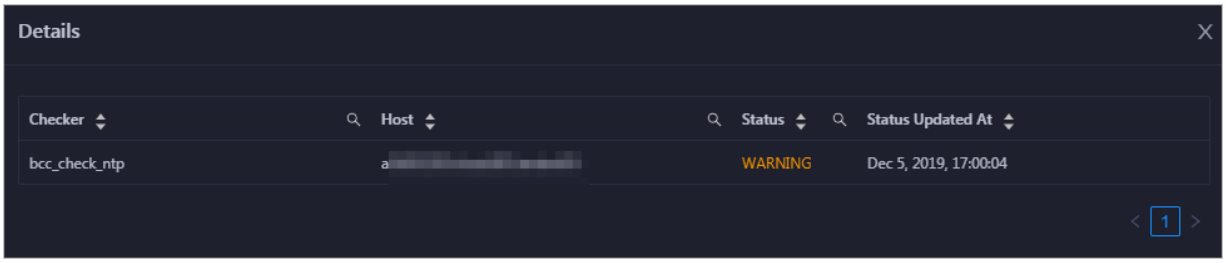

1.7.5.2 Host health

On the Health Status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

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

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. **On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.**

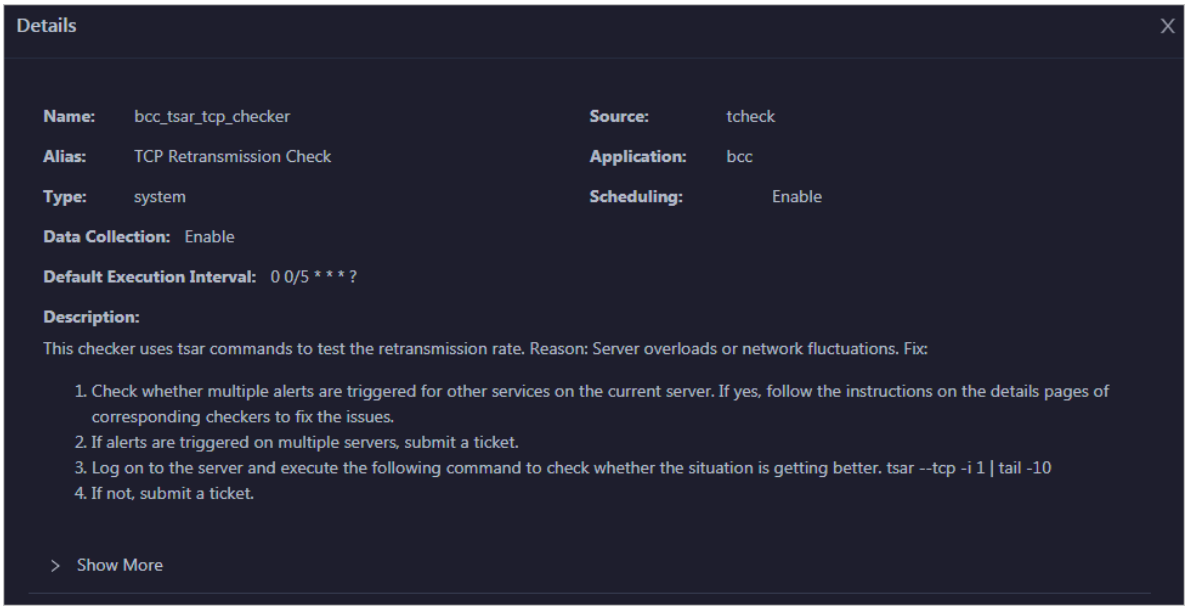

   The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

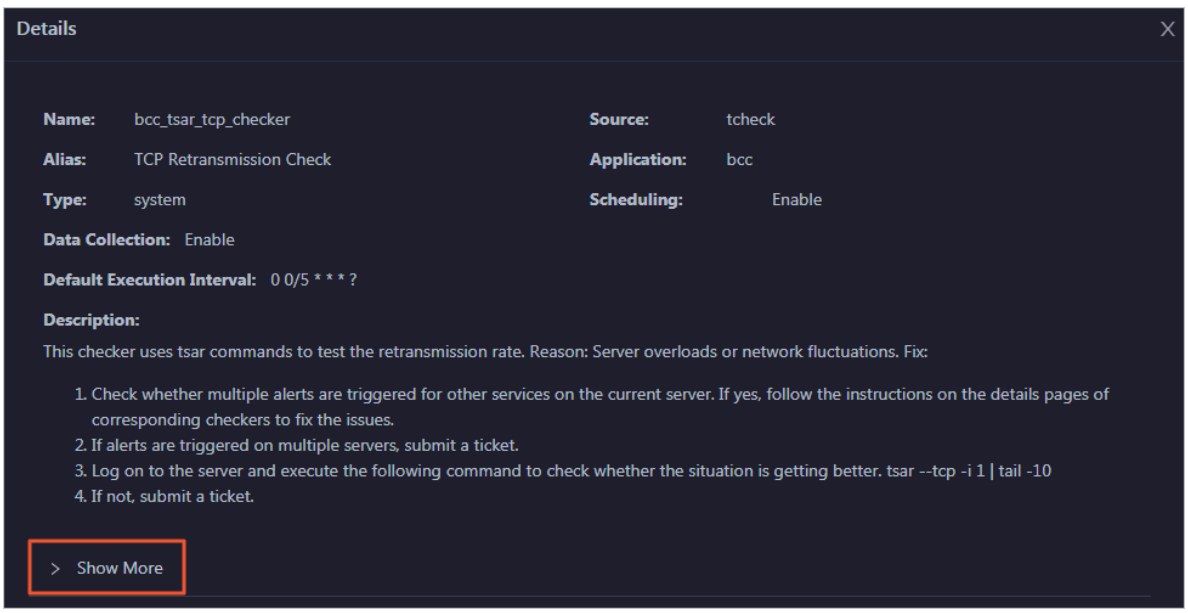

   You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

1. **On the Health Status page, click + to expand a checker with alerts.**

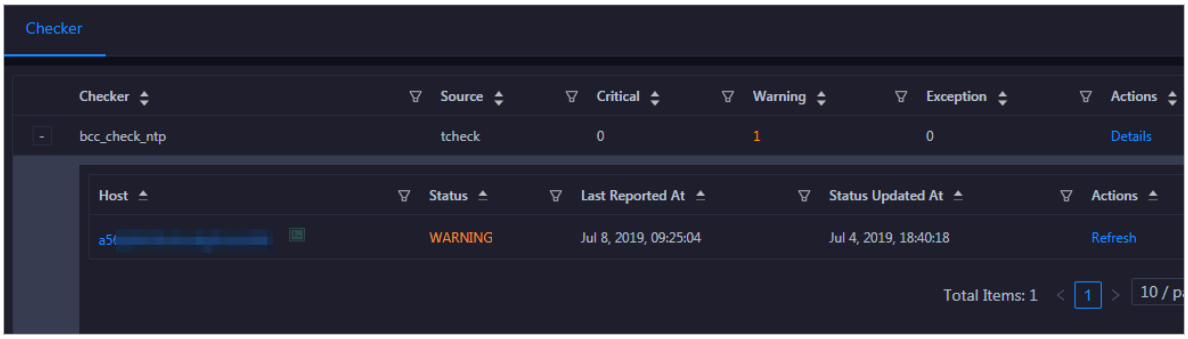

2. **Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

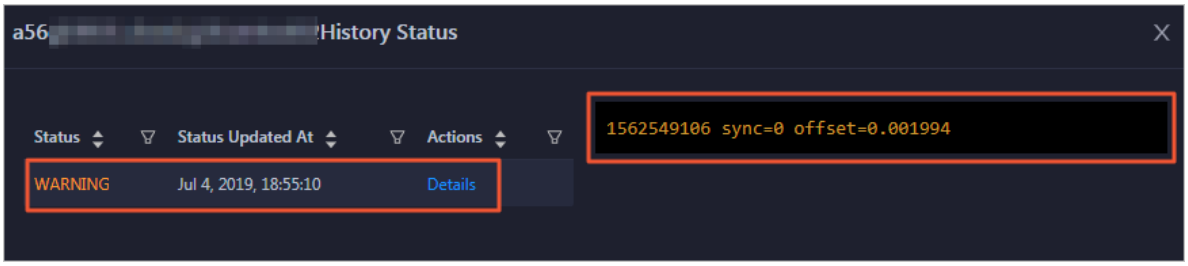

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

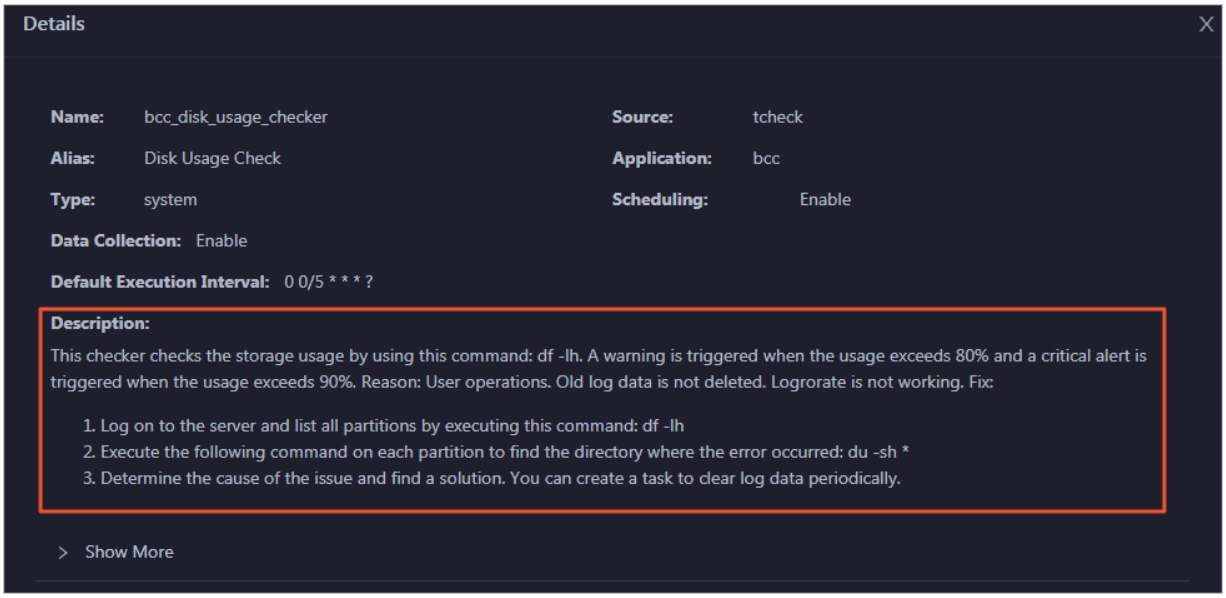

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

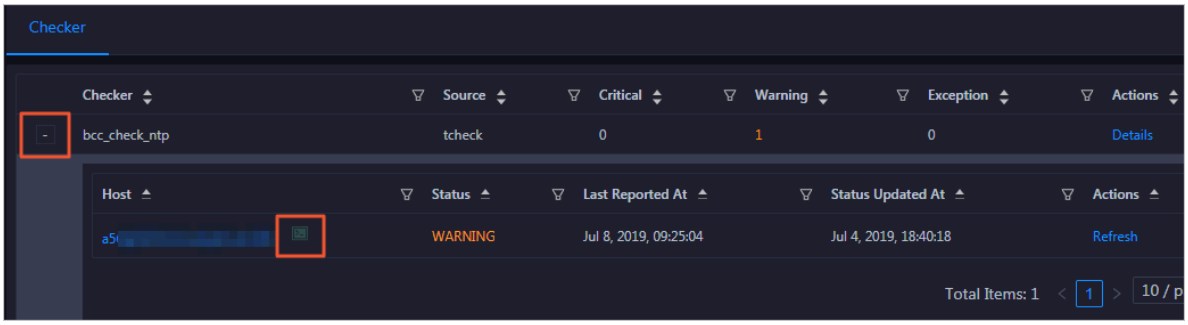

2. **Click the Log On icon of a host. The TerminalService page appears.**

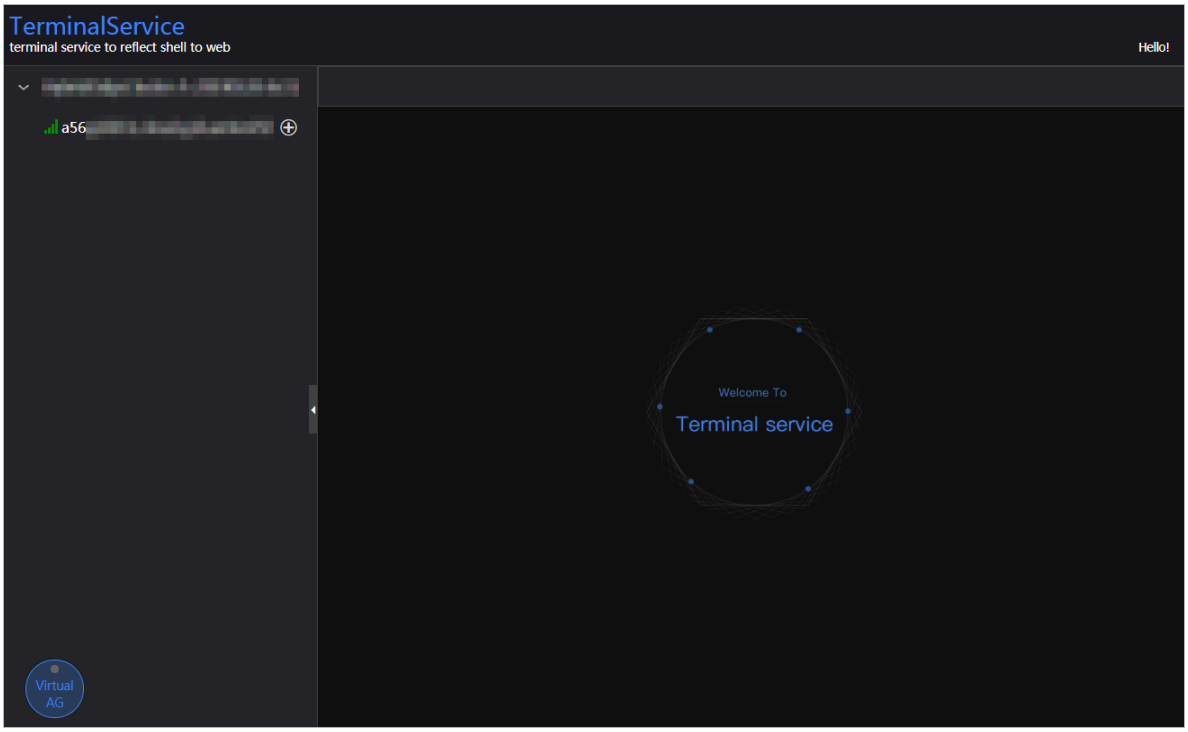

3. On the TerminalService page, click the hostname on the left to log on to the host.

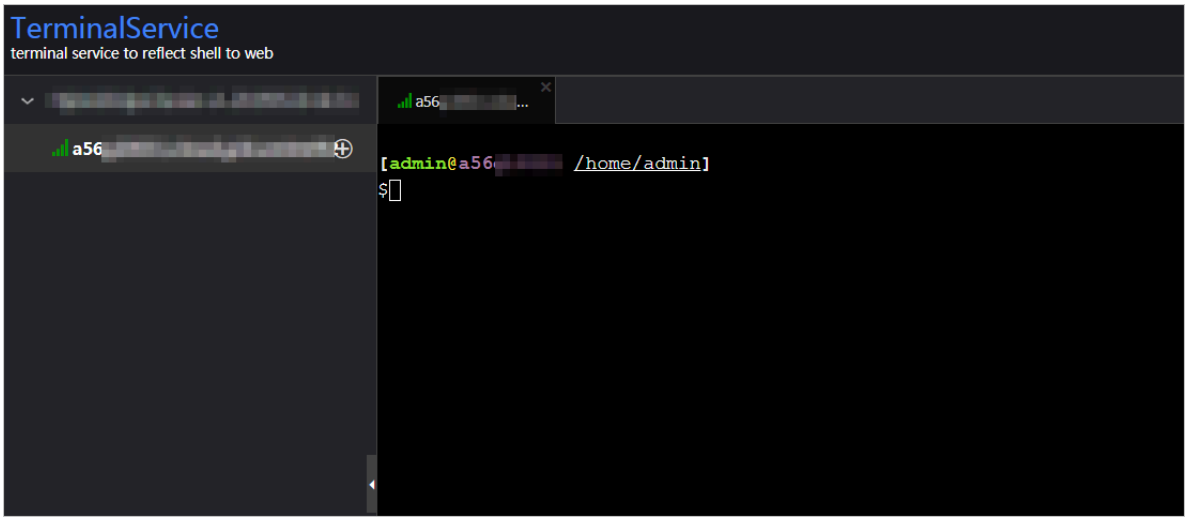

Run a checker again

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. In this way, you can check whether the alert is cleared.

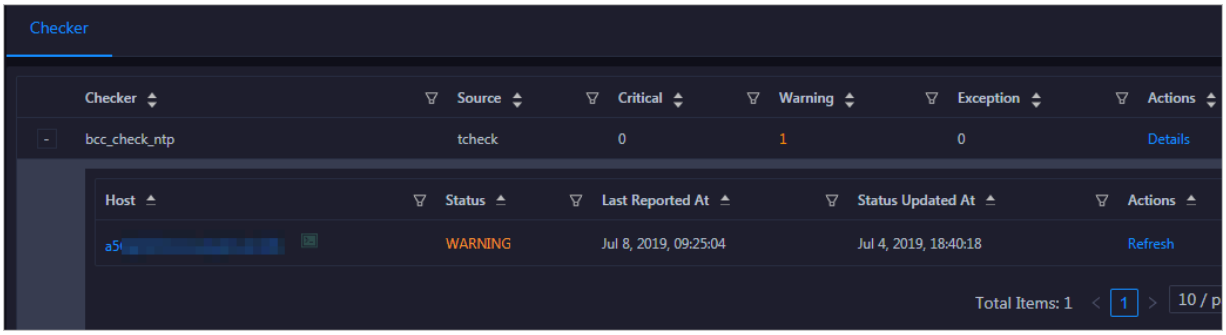

1.7.5.3 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.

The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see *Host overview*.

1.7.5.4 Host services

On the host service page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.

On the Services page, you can view the cluster, service instances, and service instance roles of the host.

1.7.6 Job and queue analysis

1.7.6.1 Job analysis

The job analysis feature allows you to diagnose jobs to quickly troubleshoot job failures.

Prerequisites

Jobs are in the running state.

Context

Job analysis has two steps, namely, Failover and Blink Metric. In the Blink Metric step, the system checks the latency, garbage collection (GC) time, transactions per second (TPS), the number of times of GC, data skew, and back pressure nodes of a job.

Procedure

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click StreamCompute.
3. On the page that appears, click Analyze in the upper-right corner. The Job Analysis page appears.

   You can also click Business on the O&M page, click Jobs in the left-side navigation pane, and then click a job name in the Name column to go to the Job Analysis page.
4. Select the job to be diagnosed and analyzed from the Select Job drop-down list.
5. In the Diagnosis section, click Start Diagnosis.

   After the diagnosis starts, the system automatically evaluates the time required for the diagnosis. Wait until the diagnosis is completed.
6. After the diagnosis is completed, click View Log to view the log details if the diagnosis result appears in red.

   The following table lists the metrics for job diagnosis.

Metric	Sub-metric	Description
Failover	N/A	Checks whether a failover is triggered for a job in a specified period and displays the information about the failover.
Blink Metric	Job Latency	Checks whether the latency of a subtask exceeds 10 minutes.
	Job GC	Checks whether the GC time of a Concurrent Low Pause Collector (CMS) exceeds 100 ms. This metric applies to all containers.
	Job TPS	Checks whether the TPS of a subtask is 0.

Metric	Sub-metric	Description
	Number of GC Times	Checks whether the number of the GC times exceeds 15 per minute. This metric applies to all containers.
	Data Skew	Checks whether the deviation of the input data size of each subtask in a task to the average input data size of all subtasks in the task exceeds 30%.
	Back Pressure Nodes	Checks whether each task has back pressure and finds the nodes that cause back pressure.

1.7.6.2 Queue analysis

The queue analysis page displays the basic information, resource information, and job list of a queue, so that you can quickly know the resource usage of the queue and locate job exceptions.

Procedure

1. *Log on to the ABM console*.
2. Click  in the upper-left corner, and then click StreamCompute.
3. On the page that appears, click Analyze in the upper-right corner. Then click Queue Analysis in the left-side navigation pane.

   You can also click Business on the O&M page, click Queues or Jobs in the left-side navigation pane, and then click a queue in the Queue column to go to the Queue Analysis page.

   The Queue Analysis page displays the following queue information:

   - Basic information: the status and name of the queue, the cluster and partition to which the queue belongs, and the number of jobs running in the queue.
   - Resource information: the minimum number of CPU cores and minimum memory capacity guaranteed as well as the maximum number of CPU cores and maximum memory capacity available for the queue.
   - Job list: information about all jobs in the queue, including the job names, users who created the jobs, projects to which the jobs belong, transactions per second (TPS) in the inbound direction, job latency, requested compute units (CUs), failover frequency, and start time.

4. **On the Queue Analysis page, select a cluster and queue respectively from the Select Cluster and Select Queue drop-down lists at the top to view the details of the specified queue.**

1.8 Quick BI

1.8.1 QuickBI O&M overview

This topic describes the features of Quick BI O&M supported by Apsara Bigdata Manager (ABM) and how to access the Quick BI O&M page.

Modules

Quick BI O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Services	Overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service in a cluster.
	Server	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Clusters	Overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmissi on, TCP connection, and root disk usage for a cluster.
	Health Status	Displays 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.

Module	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.
	Health Status	Displays the checkers of the selected host, including the checker details, check results, check history, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click Quick BI.
3. On the page that appears, click O&M in the upper-right corner. The Services page appears.

The O&M page includes three modules, namely, Services, Clusters, and Hosts.

1.8.2 Service O&M

1.8.2.1 Service overview

The service overview page lists all Quick BI 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

1. **At the top of the O&M page, click Services.**
2. **On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.**
3. **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.

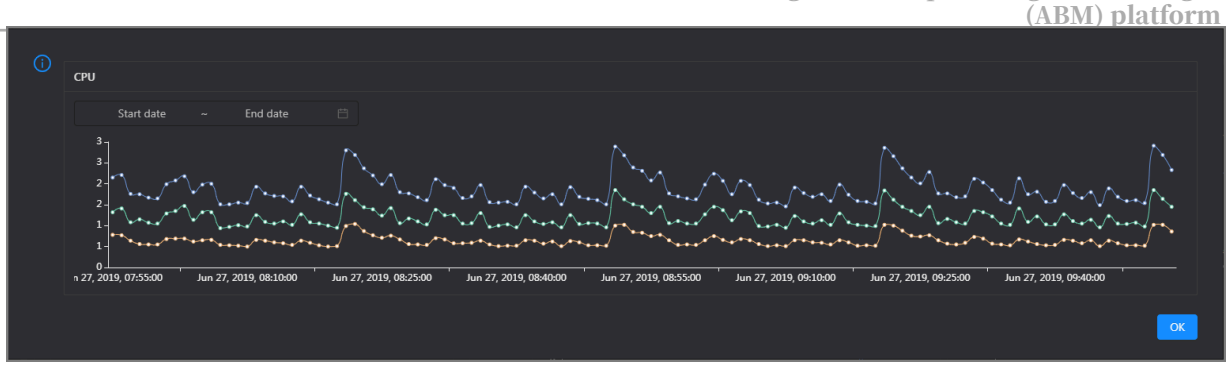

DISK

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.

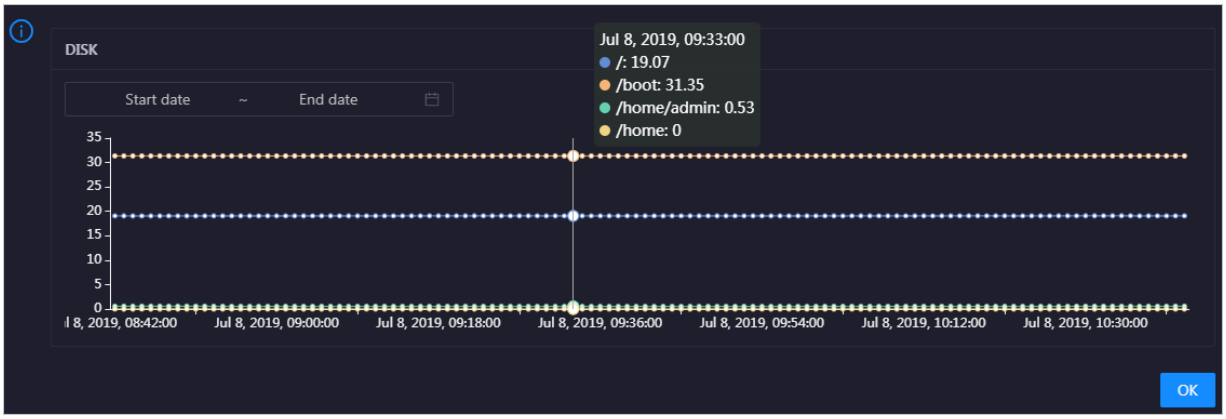

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

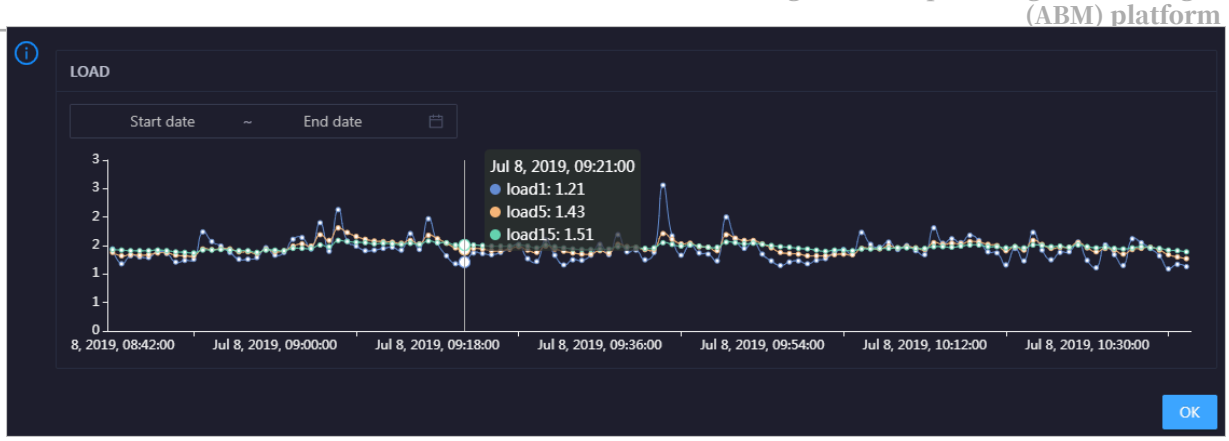

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

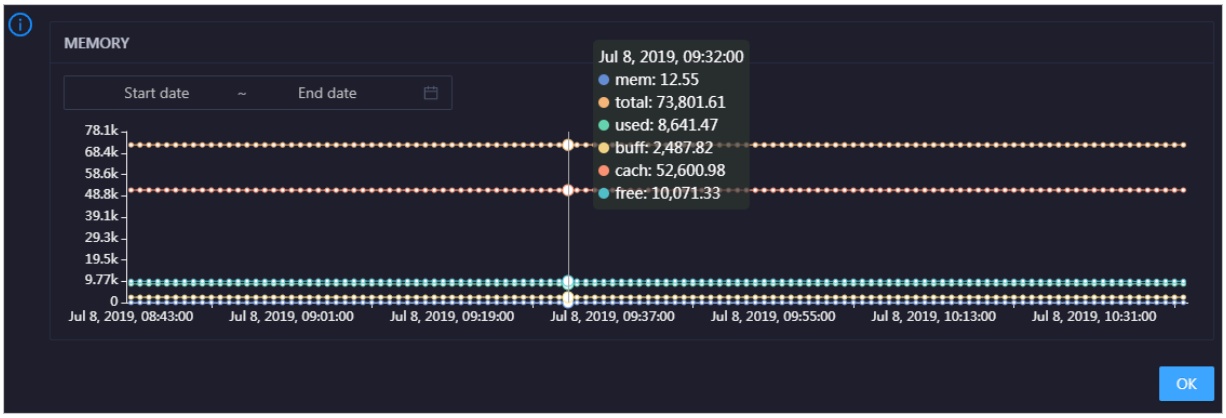

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

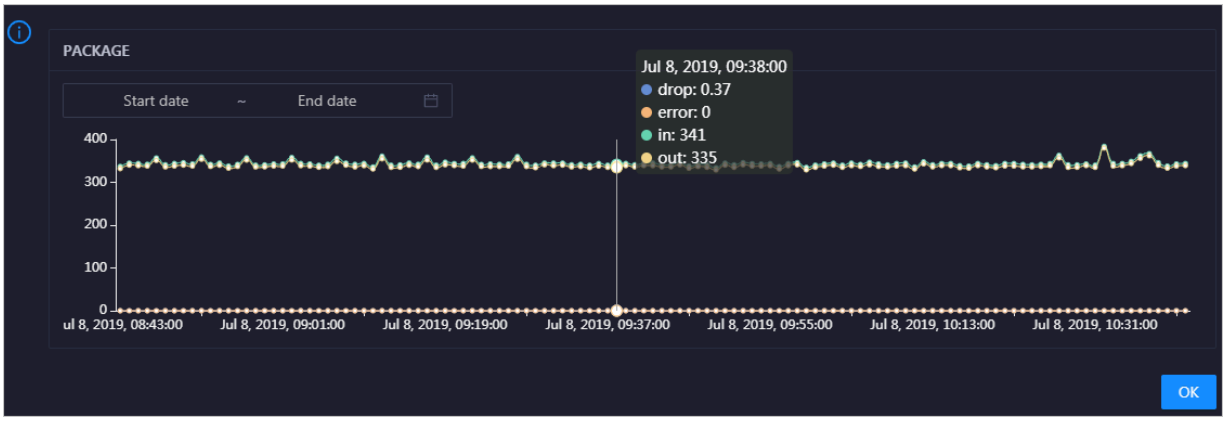

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

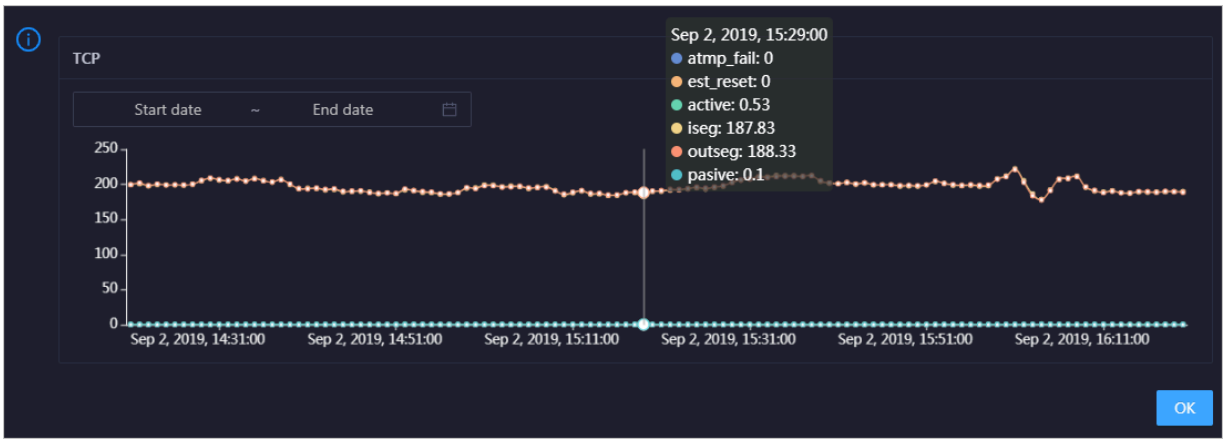

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

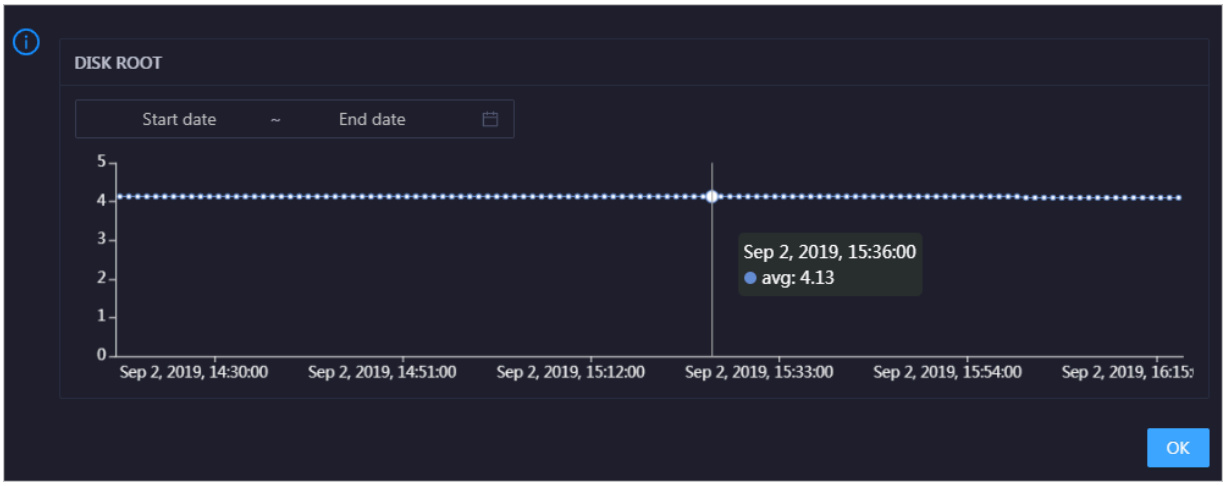

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

1.8.2.2 Service hosts

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

1. **At the top of the O&M page, click Services.**
2. **On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.**
3. **Click the Server tab. The Server page for the service appears.**

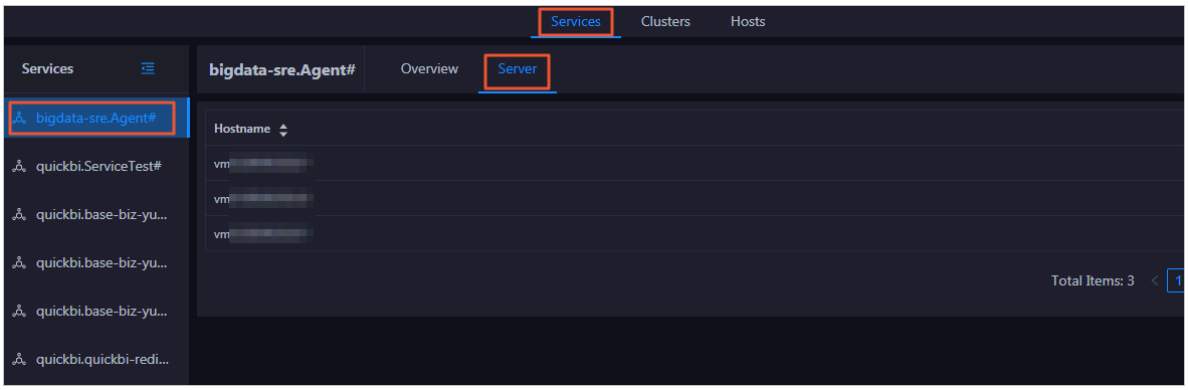

On the Server page, you can view the hosts where the selected service is run.

1.8.3 Cluster O&M

1.8.3.1 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

1. **At the top of the O&M page, click Clusters.**
2. **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.**

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 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.

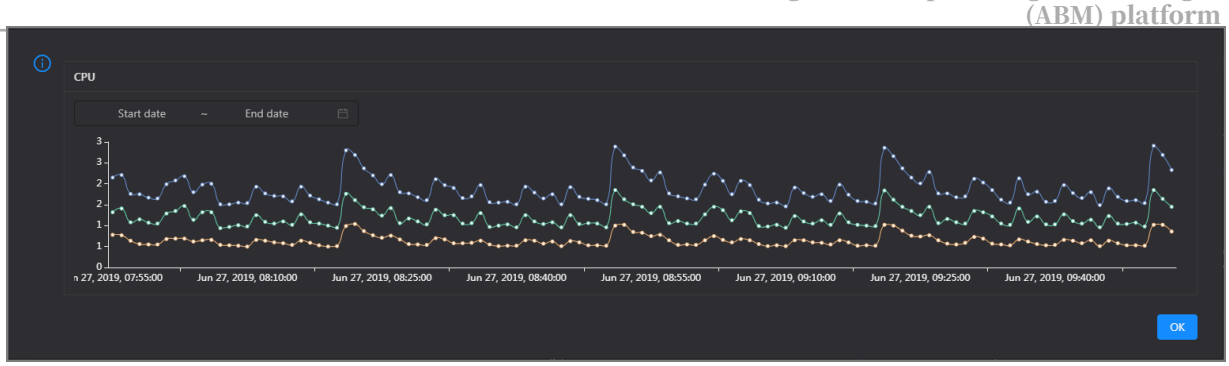

DISK

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

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

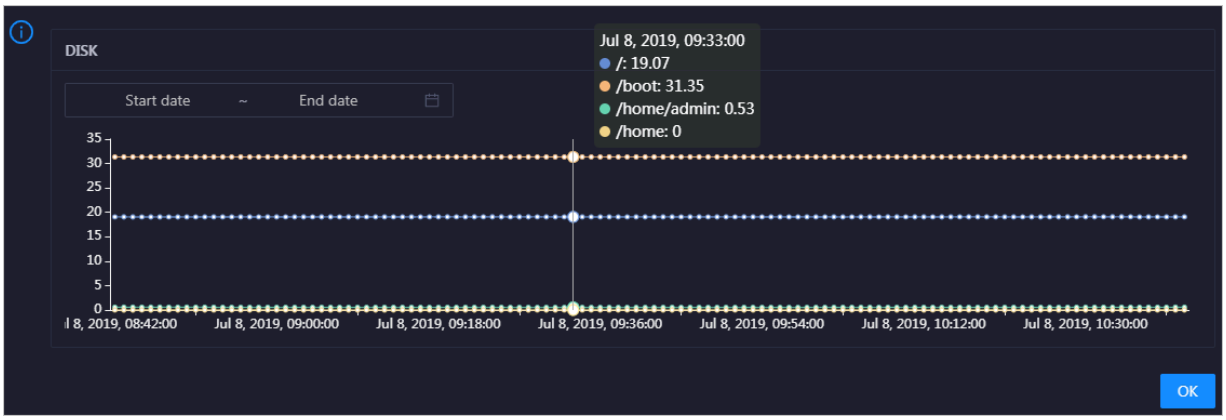

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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 cluster over time in different colors.

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

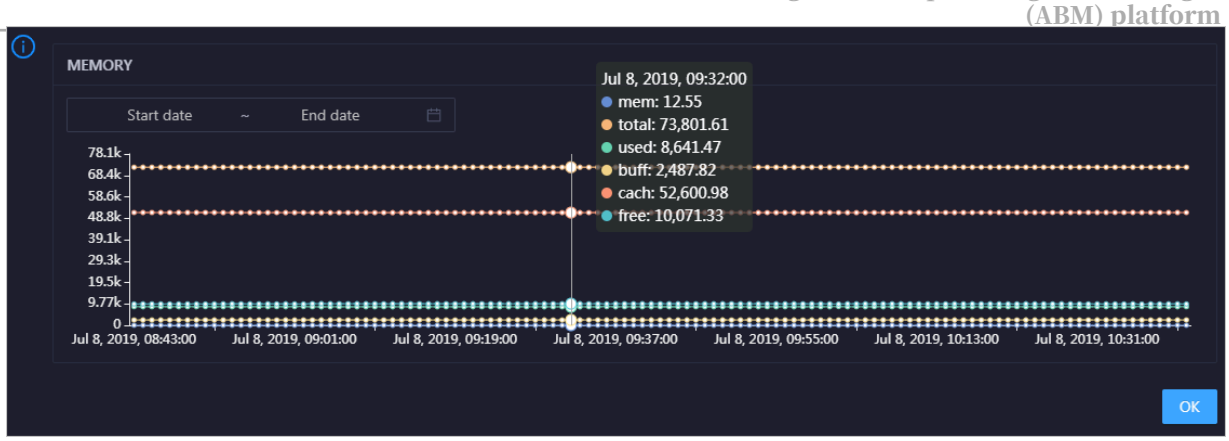

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

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

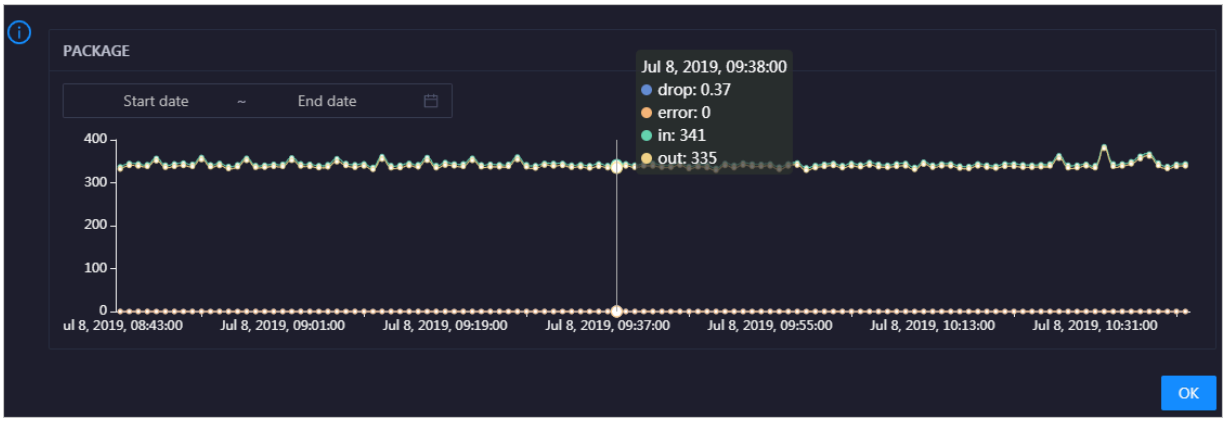

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

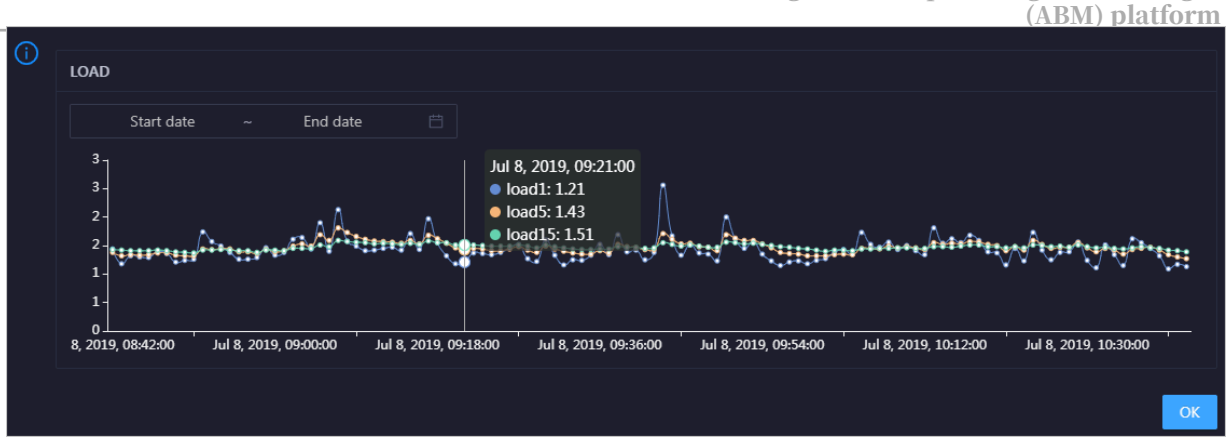

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

1.8.3.2 Cluster health

On the 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

1. At the top of the O&M page, click Clusters.
2. 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 page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

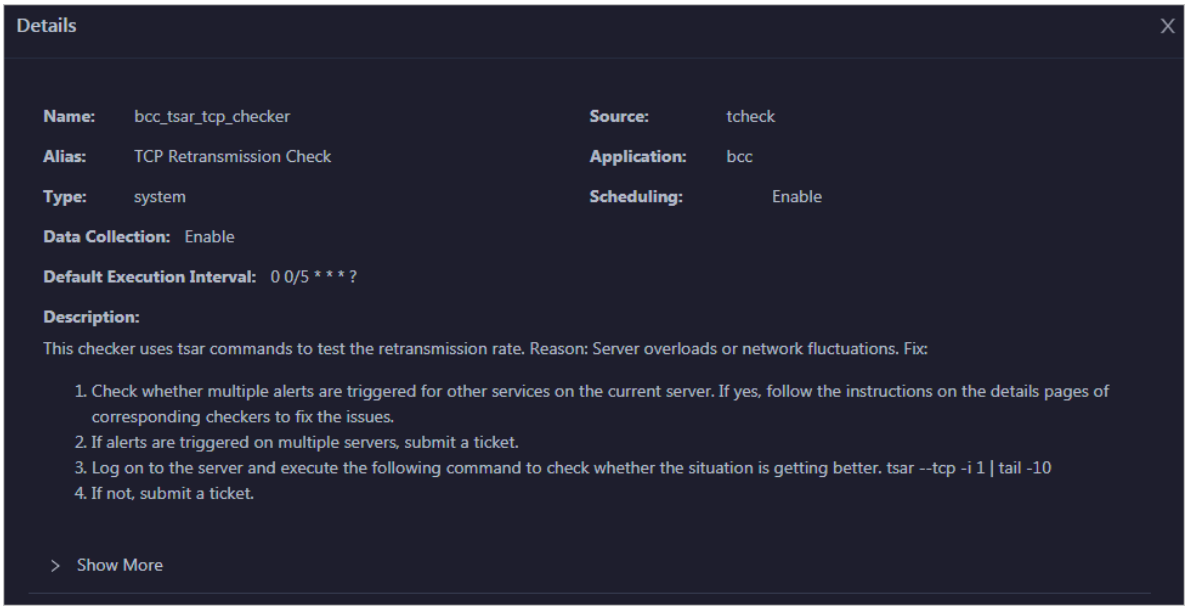

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

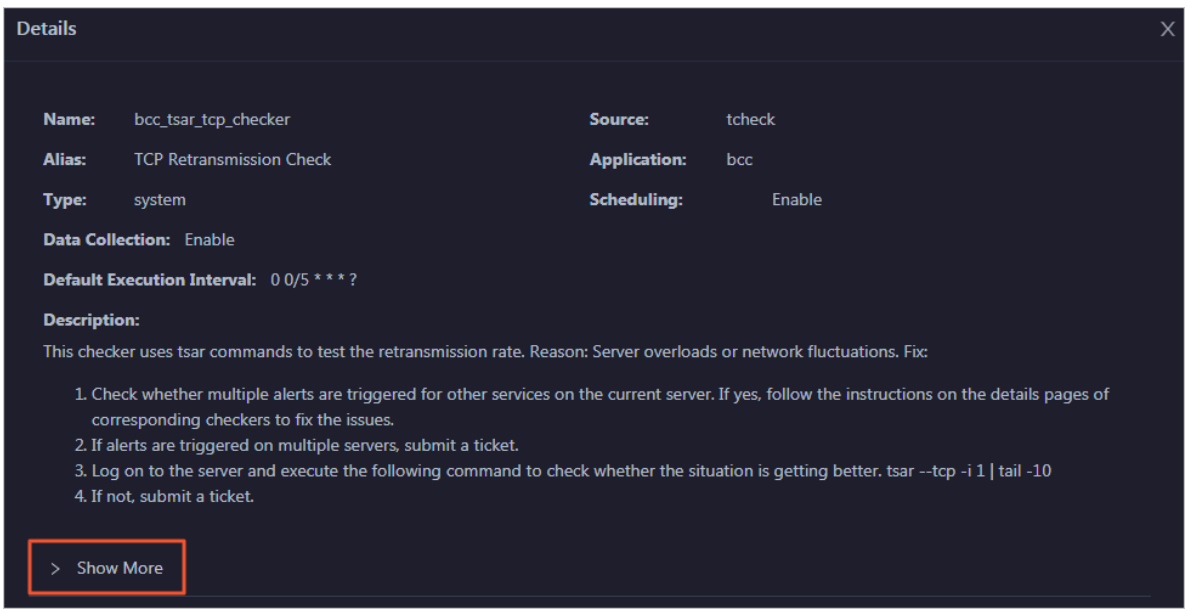

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View hosts with alerts and the alert causes

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

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

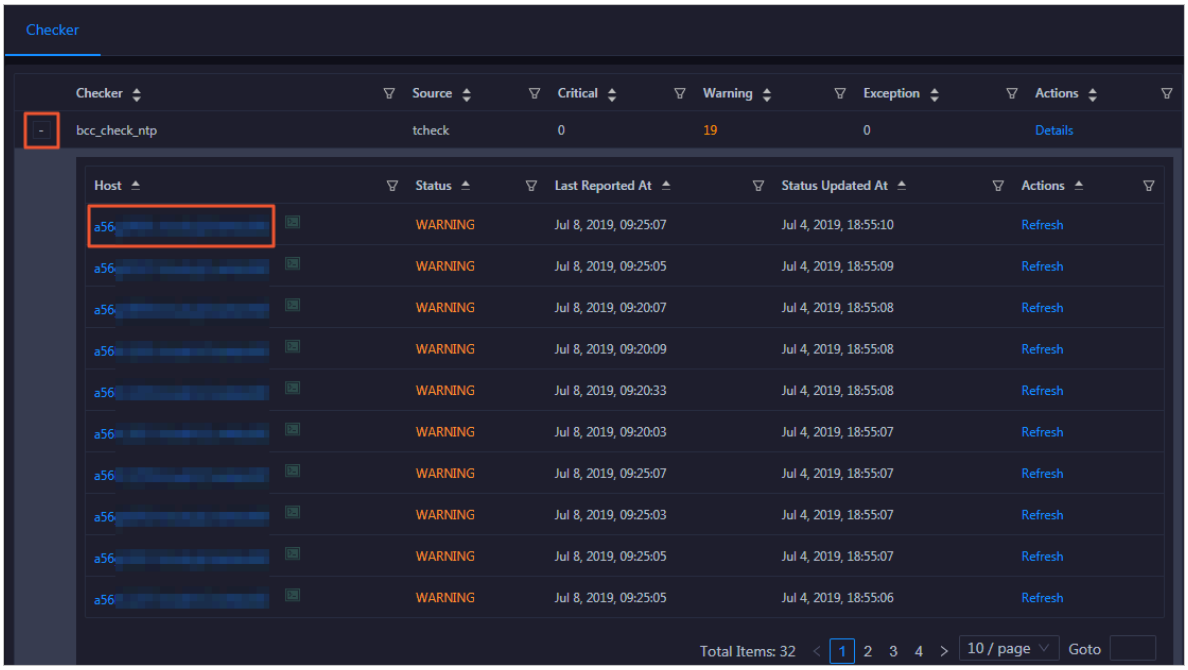

2. **Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

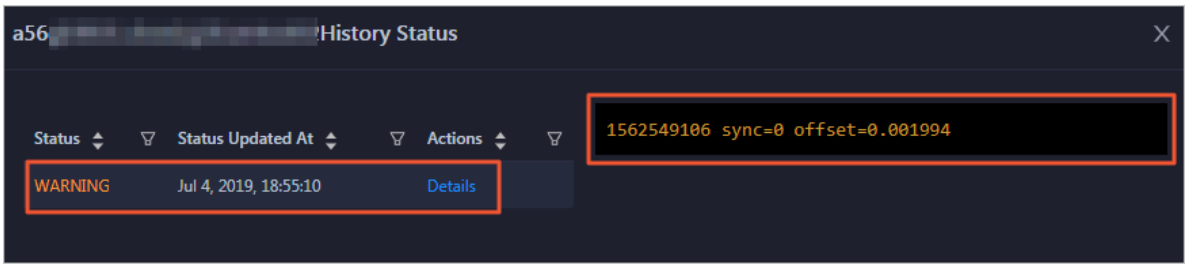

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

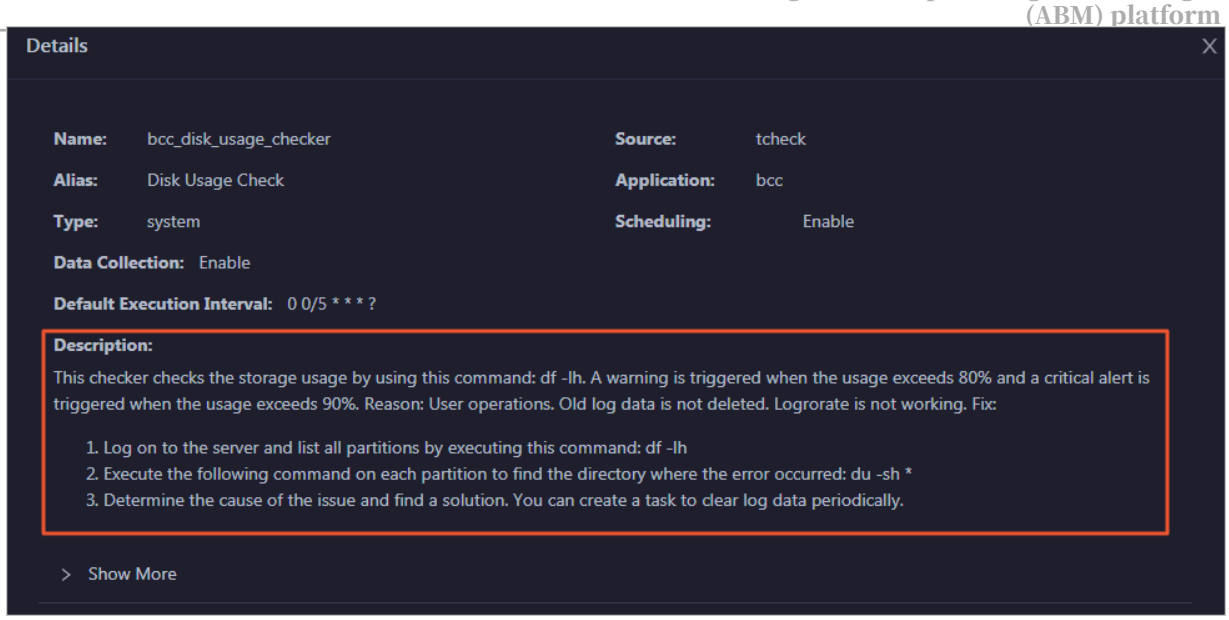

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

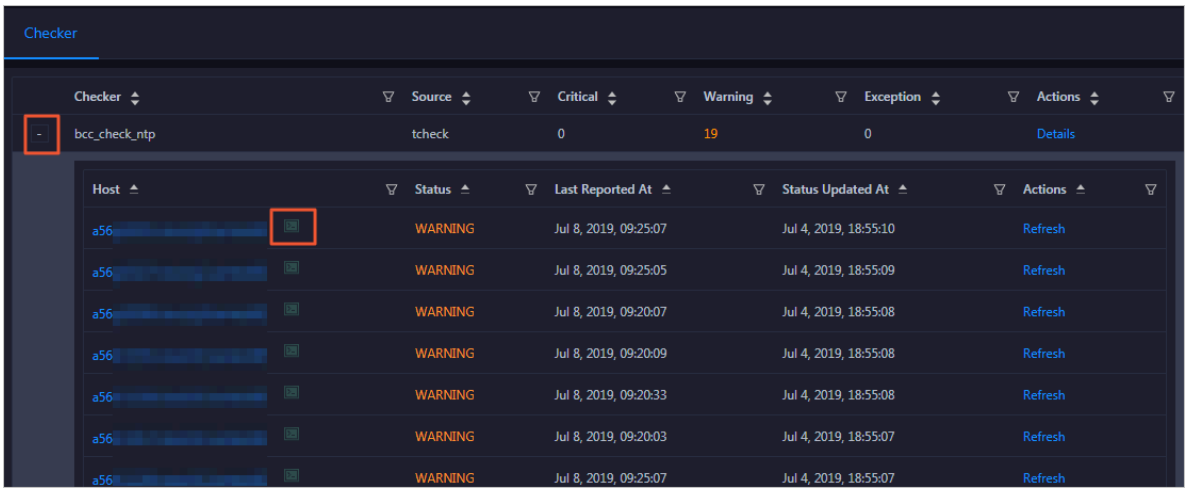

2. Click the Log On icon of a host. The TerminalService page appears.

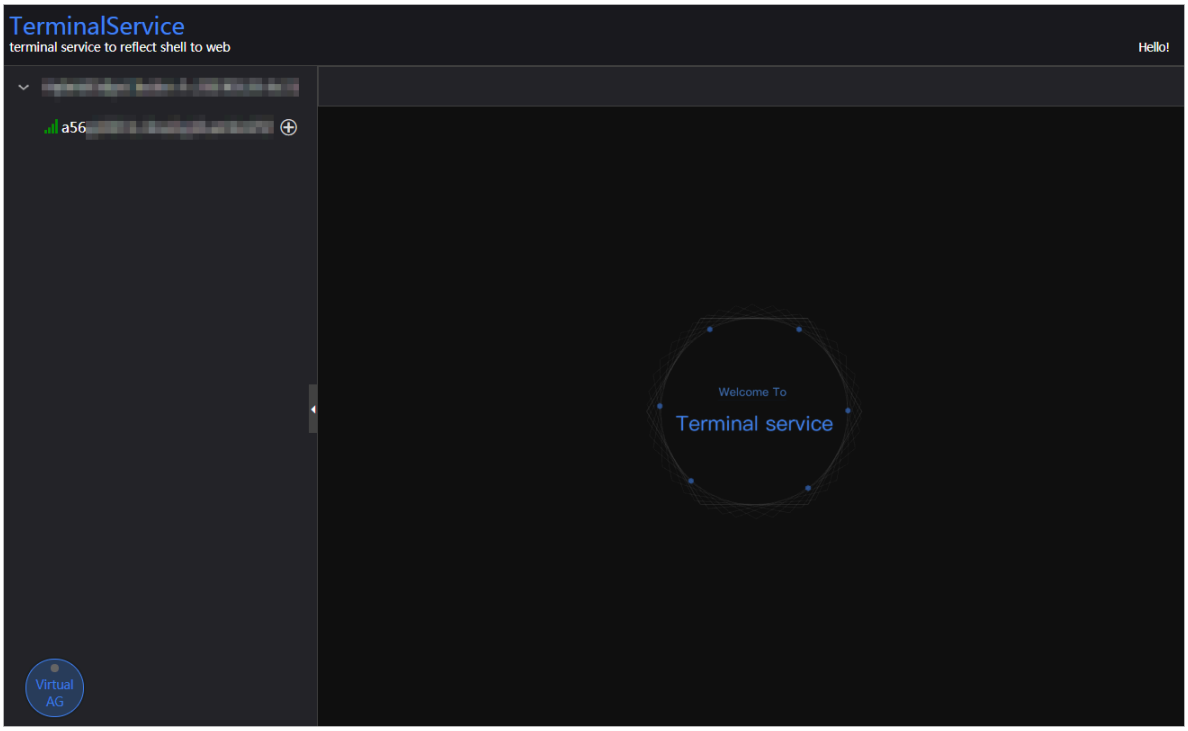

3. On the TerminalService page, click the hostname on the left to log on to the host.

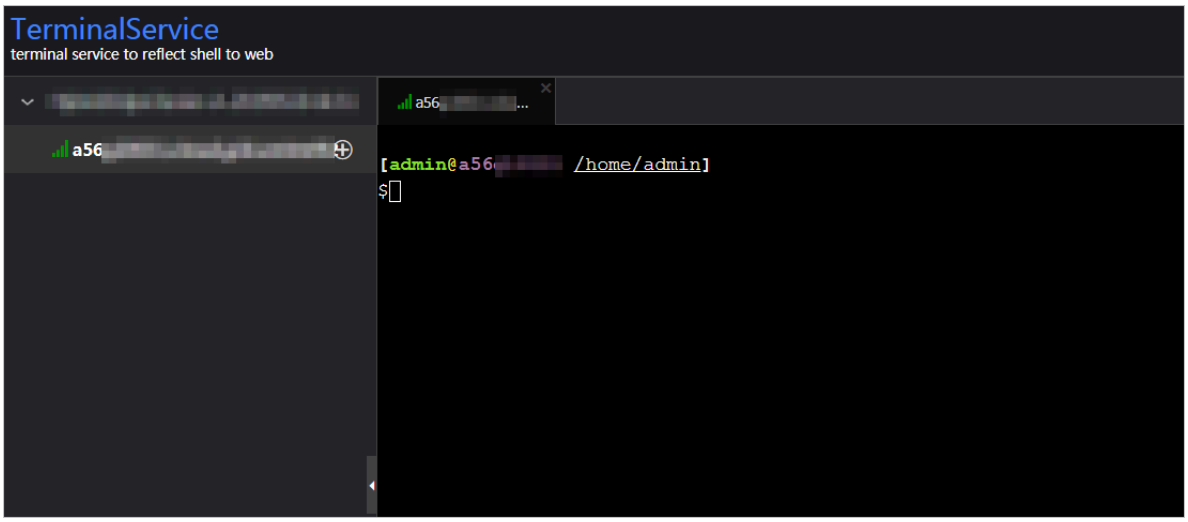

Run a checker again

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. In this way, you can check whether the alert is cleared.

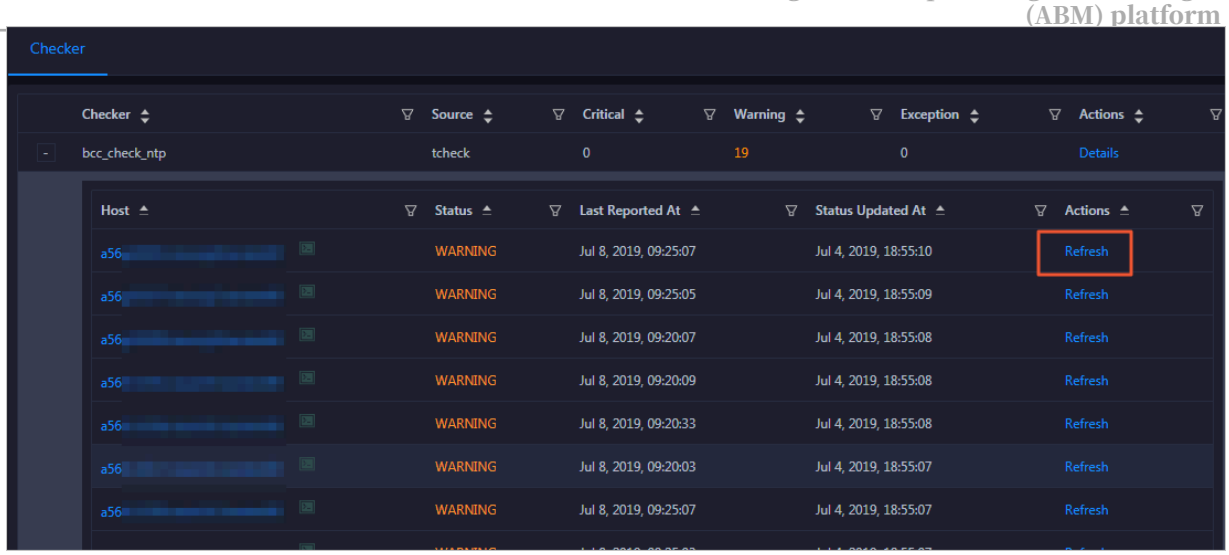

1.8.4 Host O&M

1.8.4.1 Host overview

The host overview page displays the overall running information about a host in an ElasticSearch 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, 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.

Entry

1. **At the top of the O&M page, click Hosts.**

2. On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.

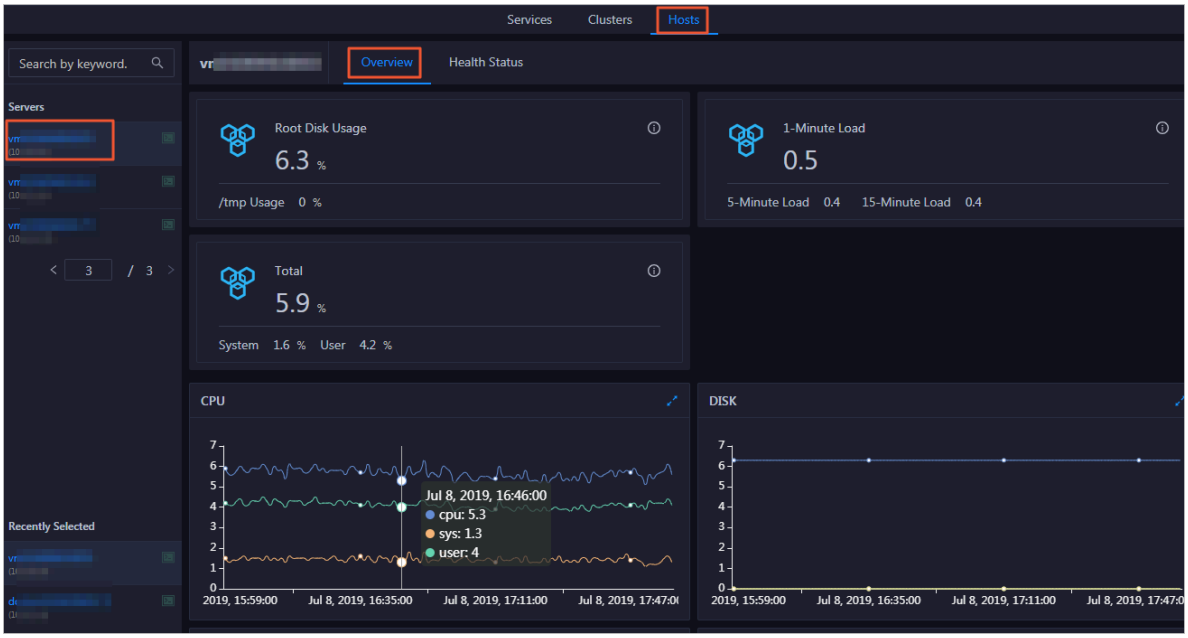

On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

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.

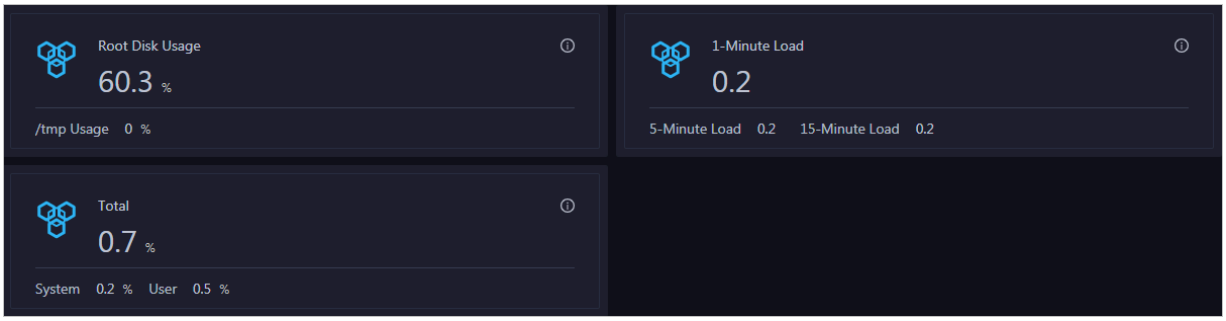

CPU

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 host over time in different colors.

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

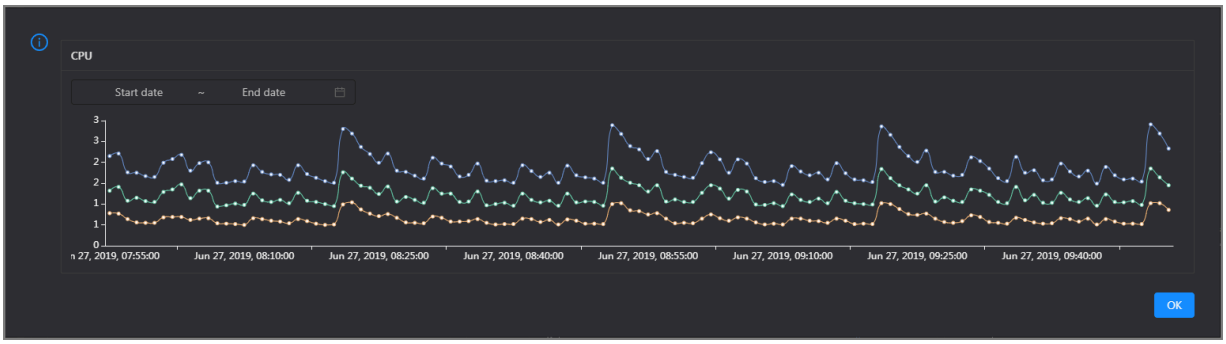

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

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

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

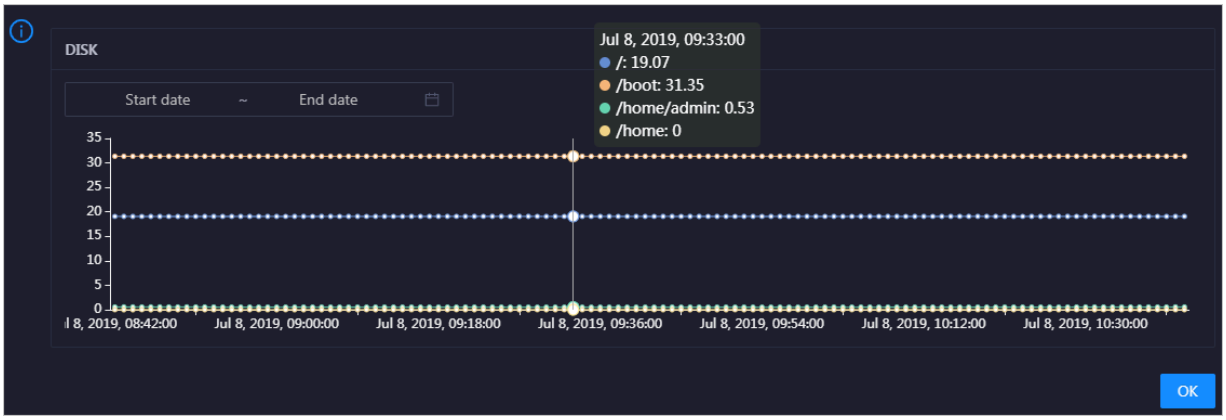

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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 host over time in different colors.

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

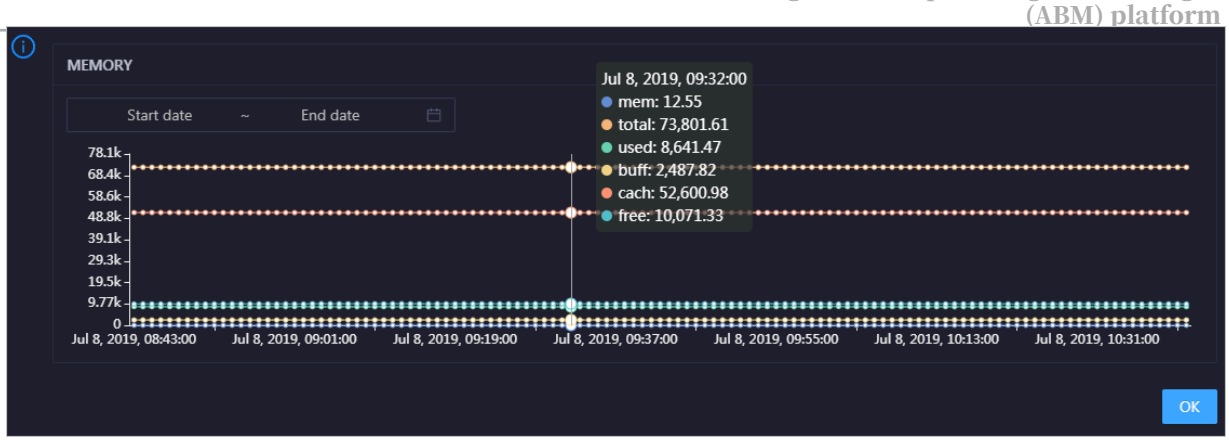

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

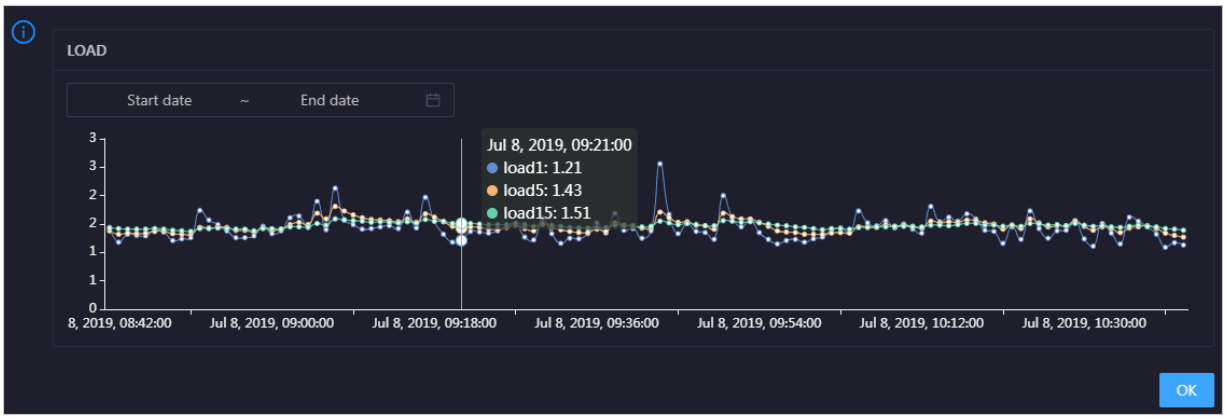

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 host over time in different colors. These trend lines reflect the data transmissi on status of the host.

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

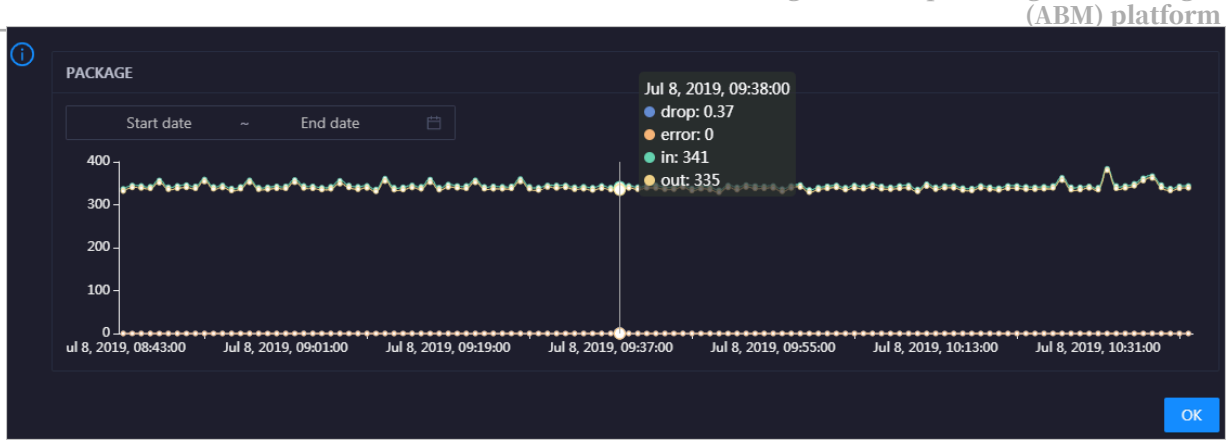

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

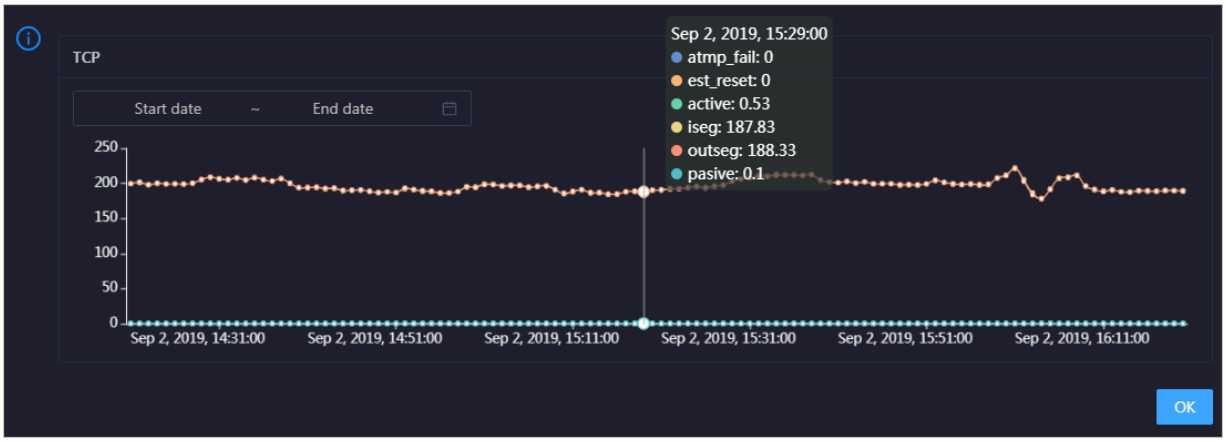

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

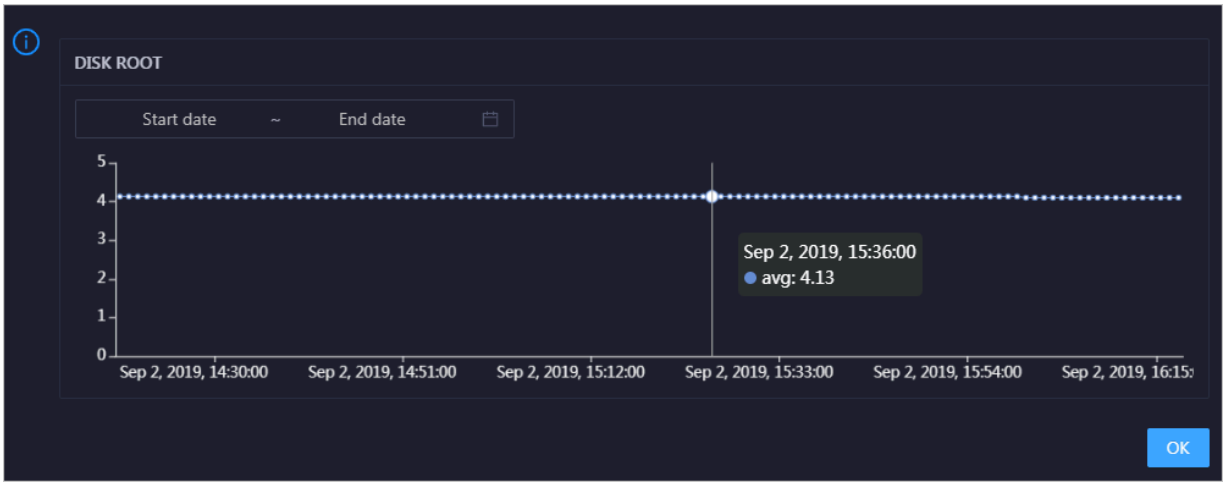

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

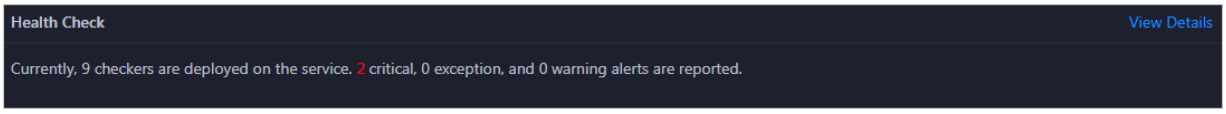

Click View Details to go to the *Host health* page. On this page, you can view the health check details.

Health Check History

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

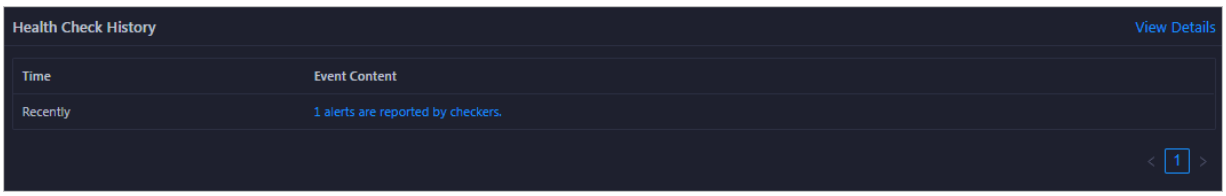

Click View Details to go to the *Host health* page. On this page, you can view the health check details.

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

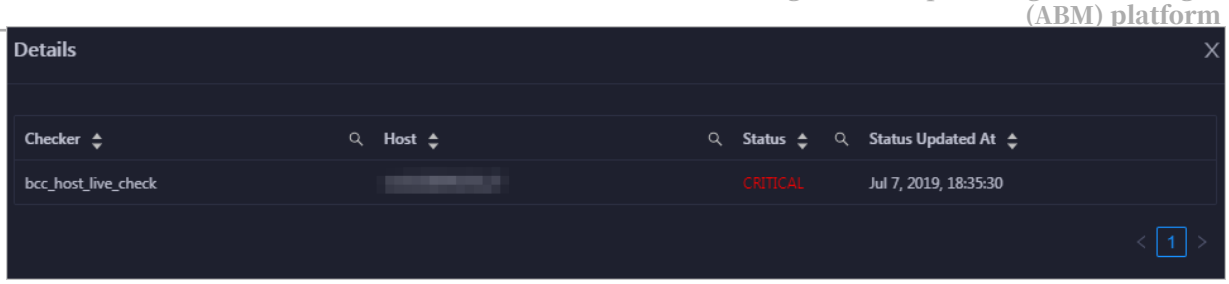

1.8.4.2 Host health

On the Health Status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. **At the top of the O&M page, click Hosts.**
2. **On the page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.**

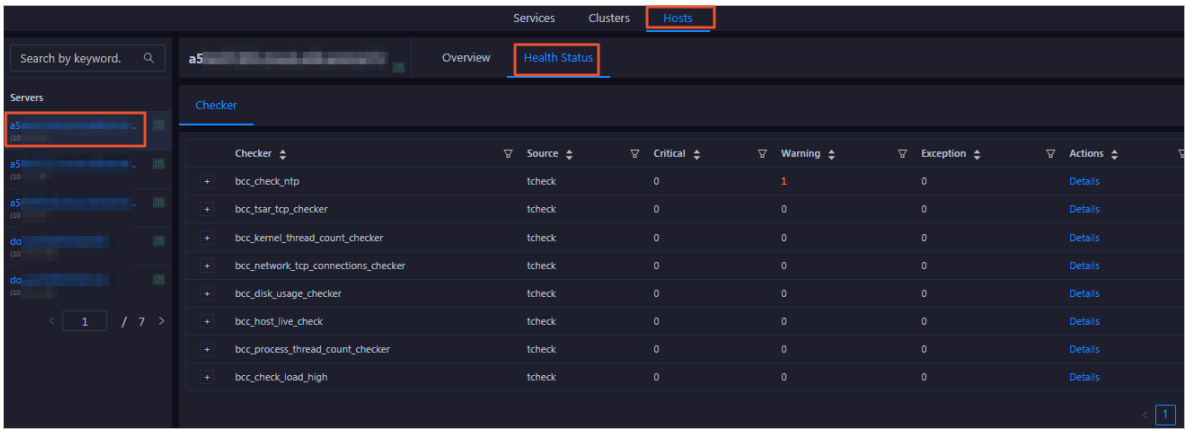

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

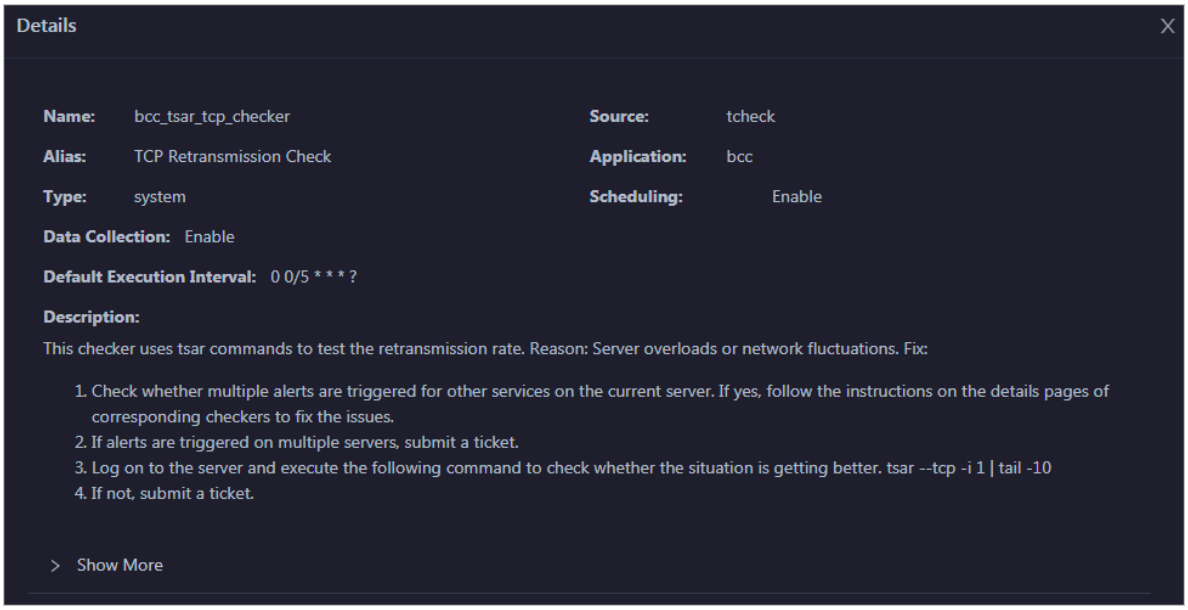

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

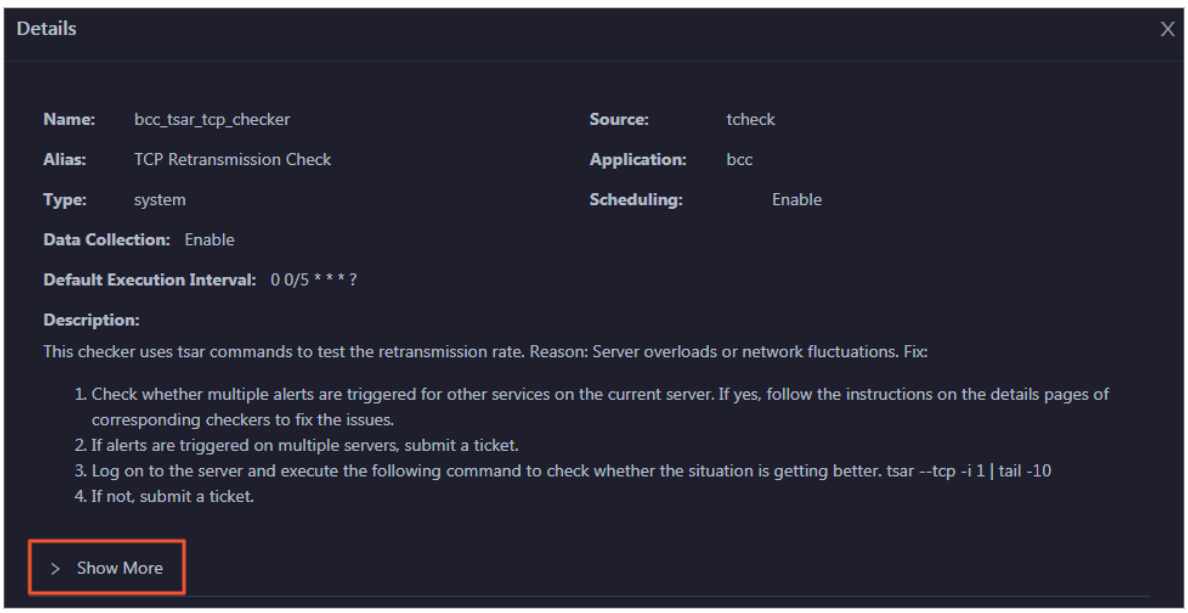

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

1. **On the Health Status page, click + to expand a checker with alerts.**

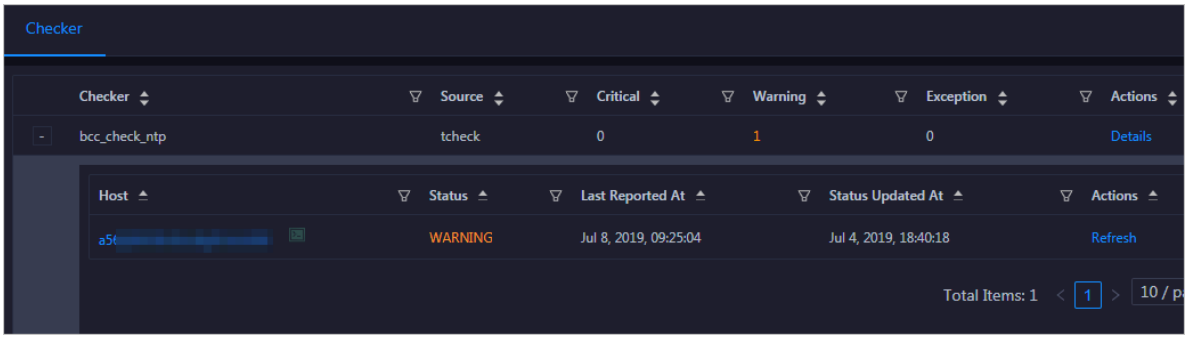

2. **Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

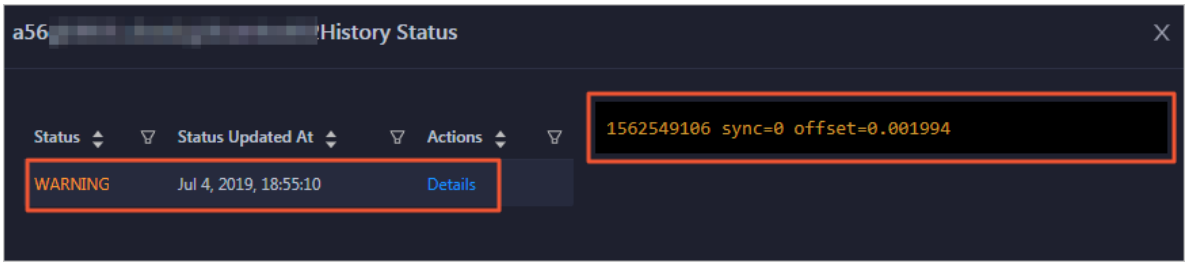

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

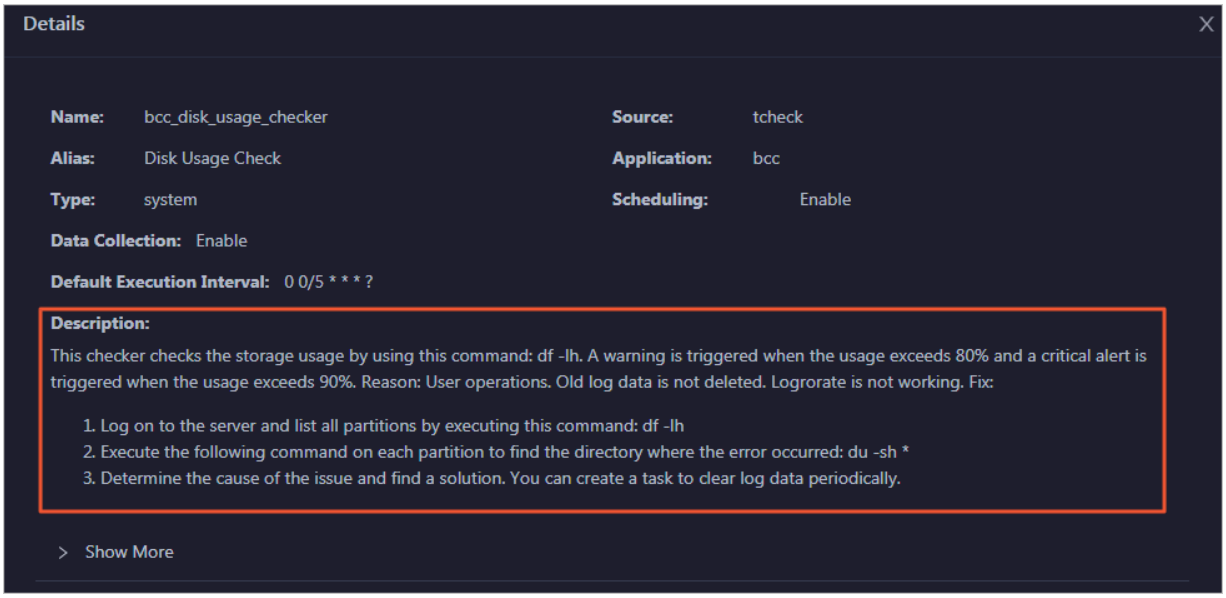

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

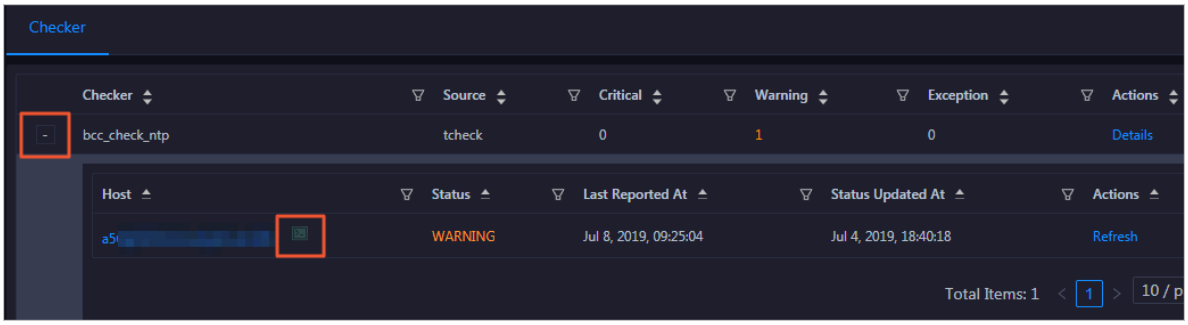

2. **Click the Log On icon of a host. The TerminalService page appears.**

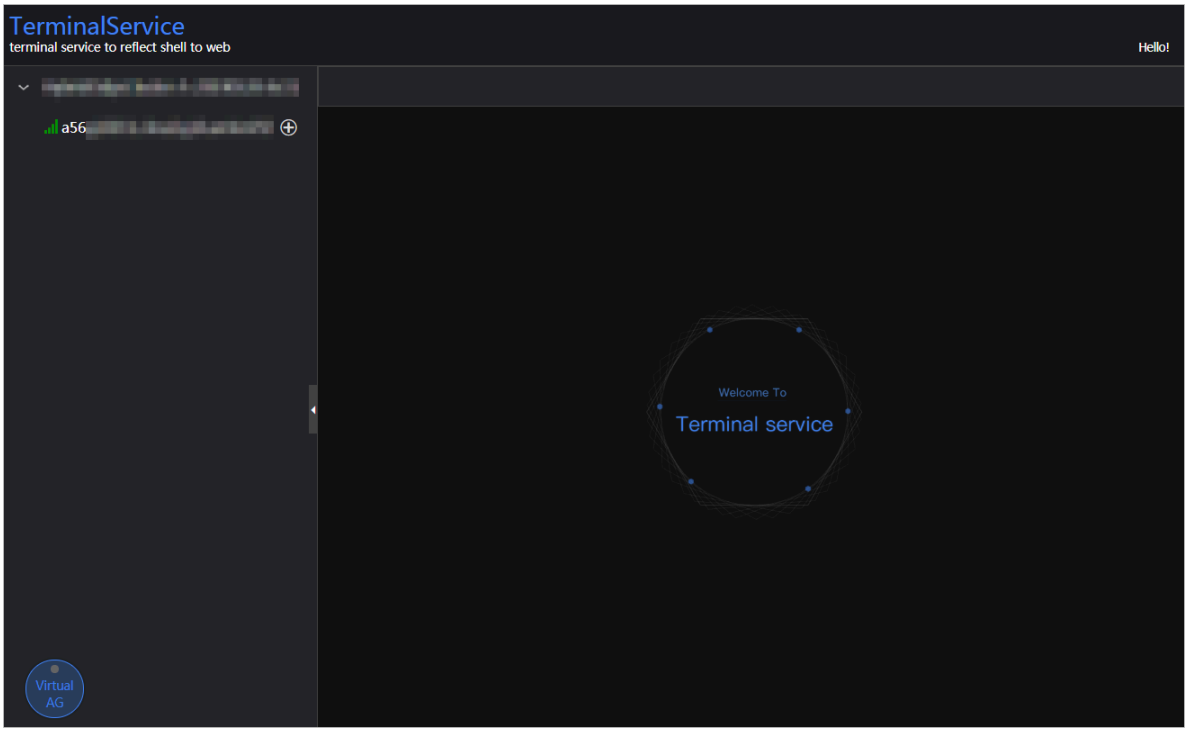

3. On the TerminalService page, click the hostname on the left to log on to the host.

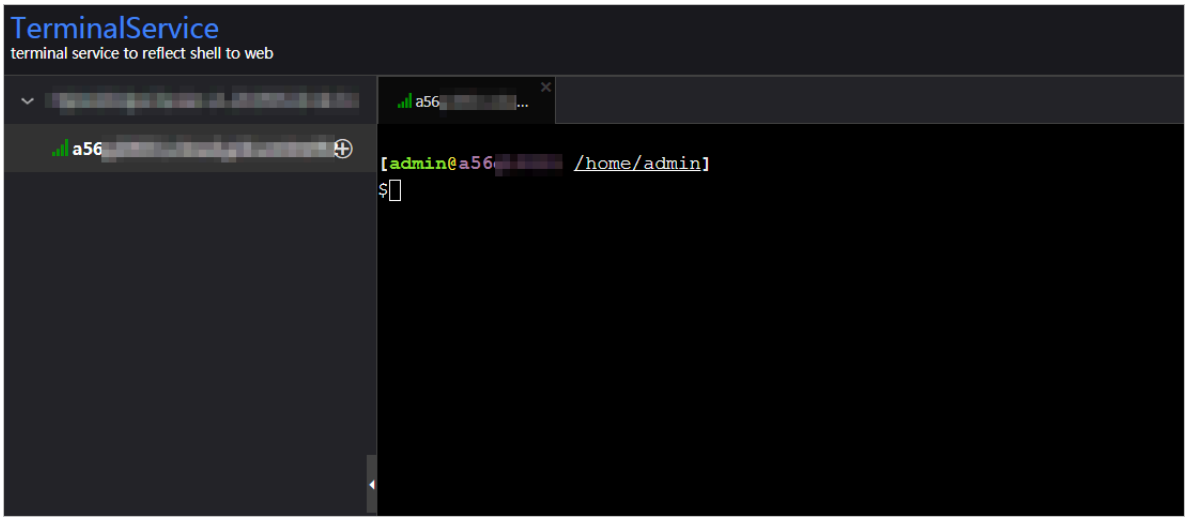

Run a checker again

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. In this way, you can check whether the alert is cleared.

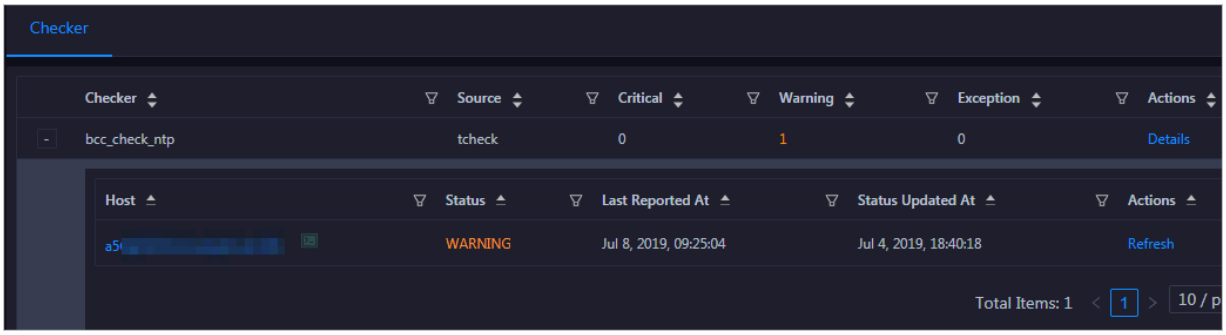

1.9 PAI

1.9.1 PAI O&M overview

This topic describes the features of Machine Learning Platform for Artificial Intelligence (PAI) O&M supported by Apsara Bigdata Manager (ABM) and how to access the PAI O&M page.

Modules

PAI O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Services	Overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service in a cluster.
	Server	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Clusters	Overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmissi on, TCP connection, and root disk usage for a cluster.
	Health Status	Displays 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.
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.
	Health Status	Displays the checkers of the selected host, including the checker details, check results, check history, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. *Log on to the ABM console*.
2. Click in the upper-left corner, and then click PAI.

3. On the page that appears, click O&M in the upper-right corner. The Services page appears.

The O&M page includes three modules, namely, Services, Clusters, and Hosts.

1.9.2 Service O&M

1.9.2.1 Service overview

The service overview page lists all PAI 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

1. At the top of the O&M page, click Services.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.

3. Click the Overview tab. The Overview page for the service appears.

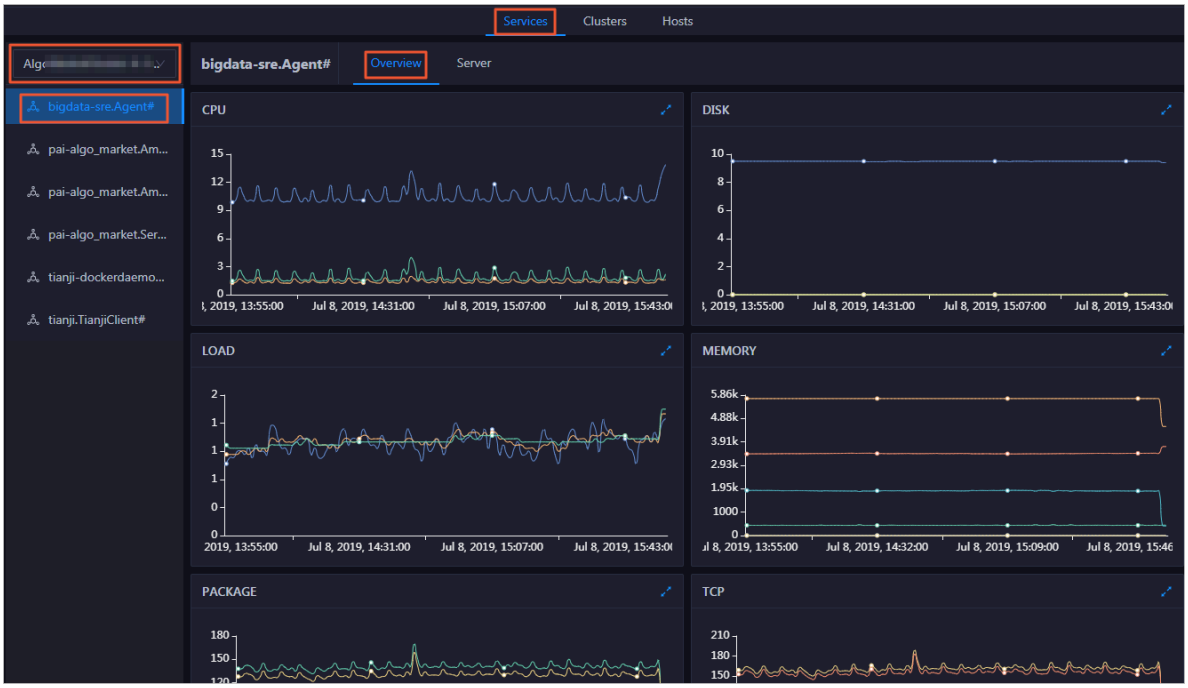

On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

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.

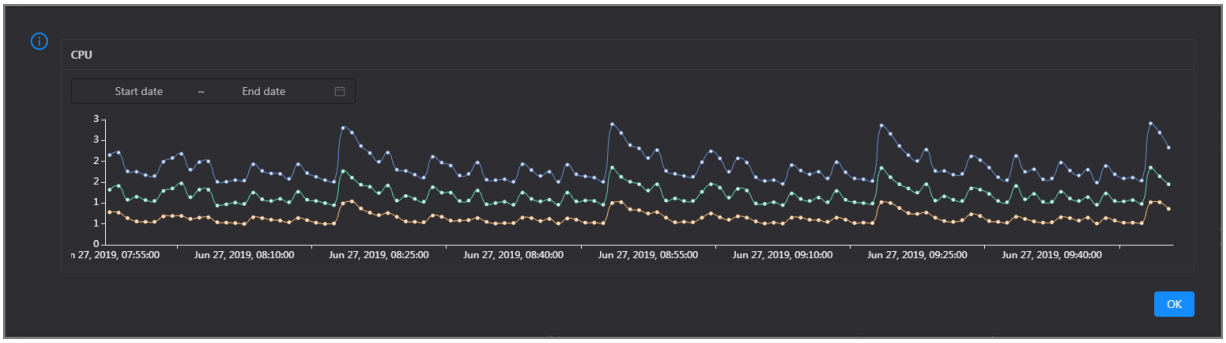

DISK

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.

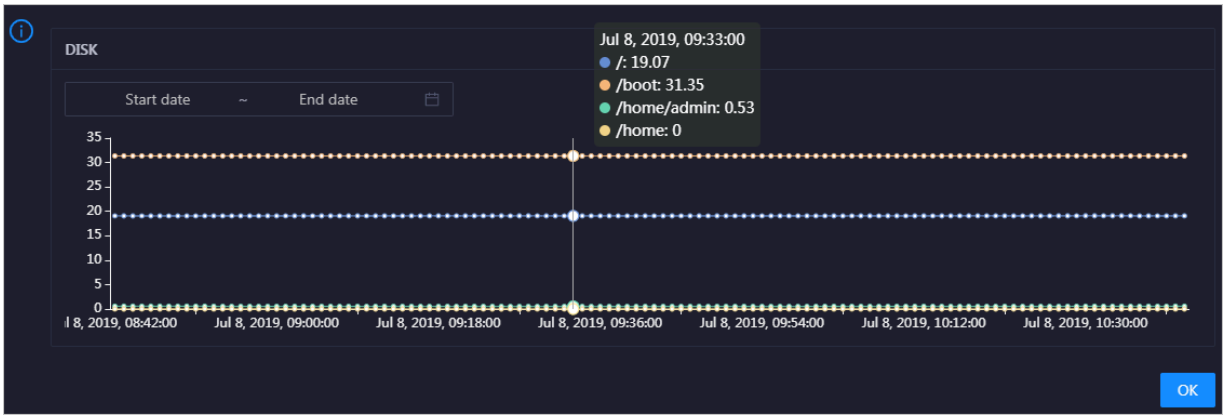

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

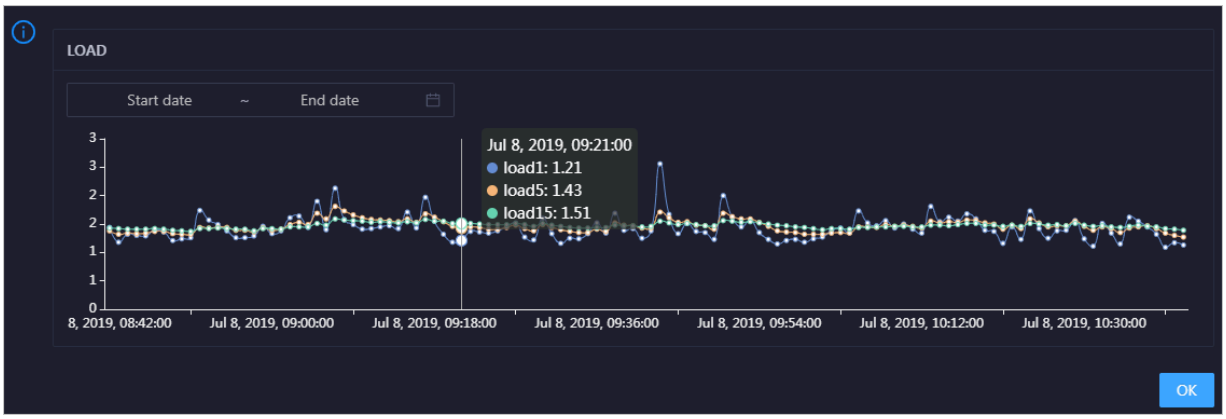

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

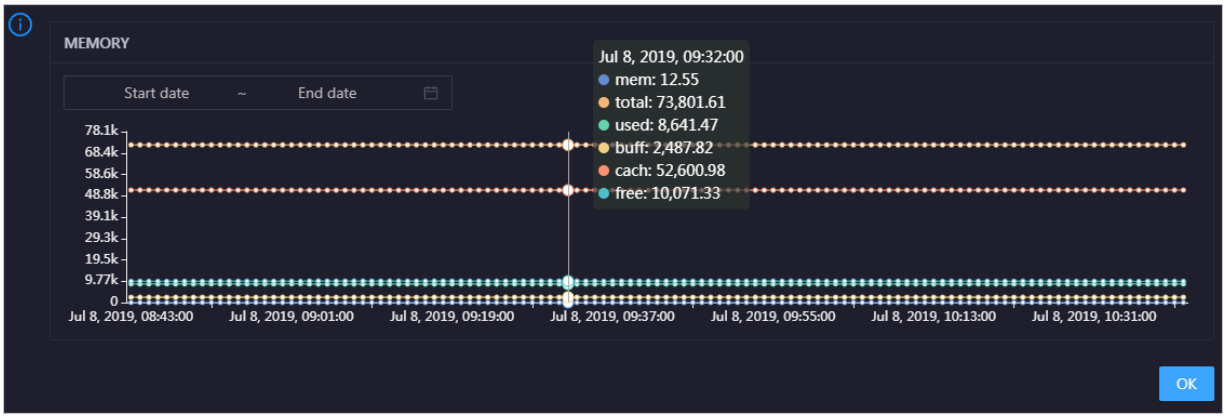

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

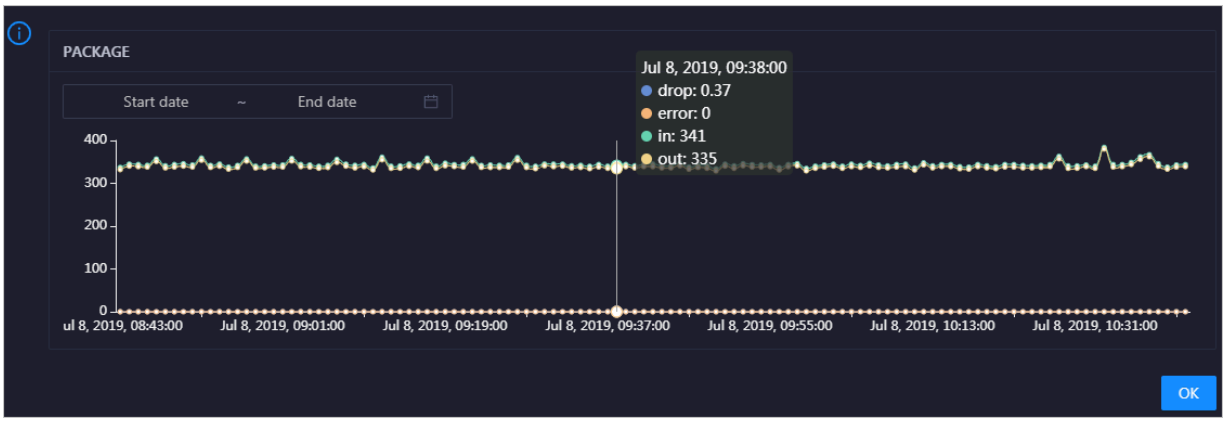

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

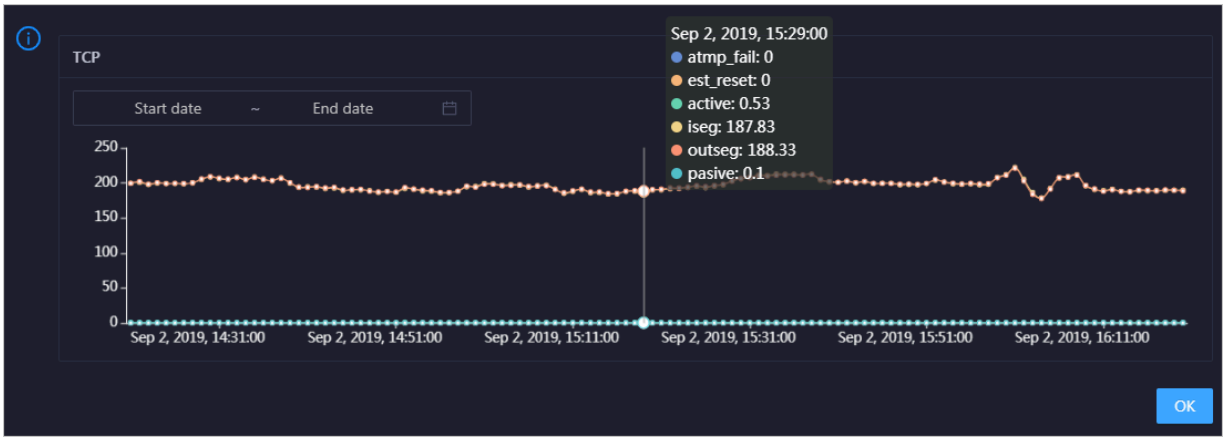

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

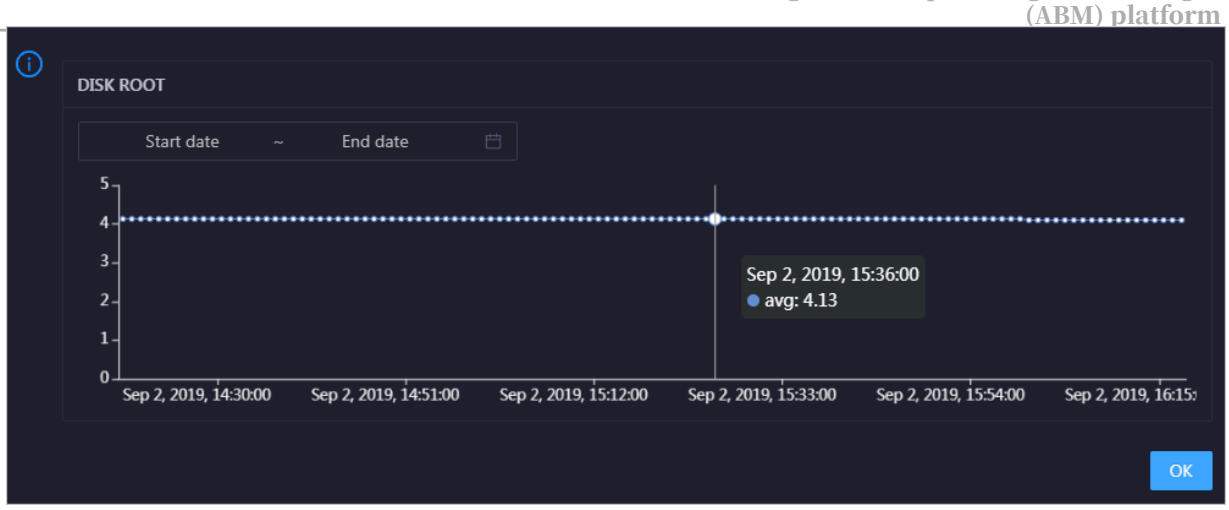

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

1.9.2.2 Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each Machine Learning Platform for Artificial Intelligence (PAI) service role so that you can understand the service deployment on hosts.

1. At the top of the O&M page, click Services.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.
3. Click the Server tab. The Server page for the service appears.

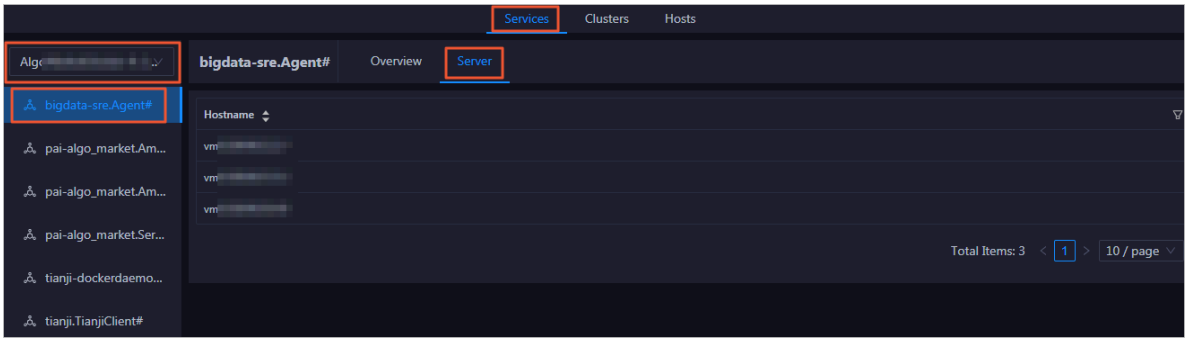

On the Server page, you can view the hosts where the selected service is run.

1.9.3 Cluster O&M

1.9.3.1 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

1. **At the top of the O&M page, click Clusters.**
2. **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.**

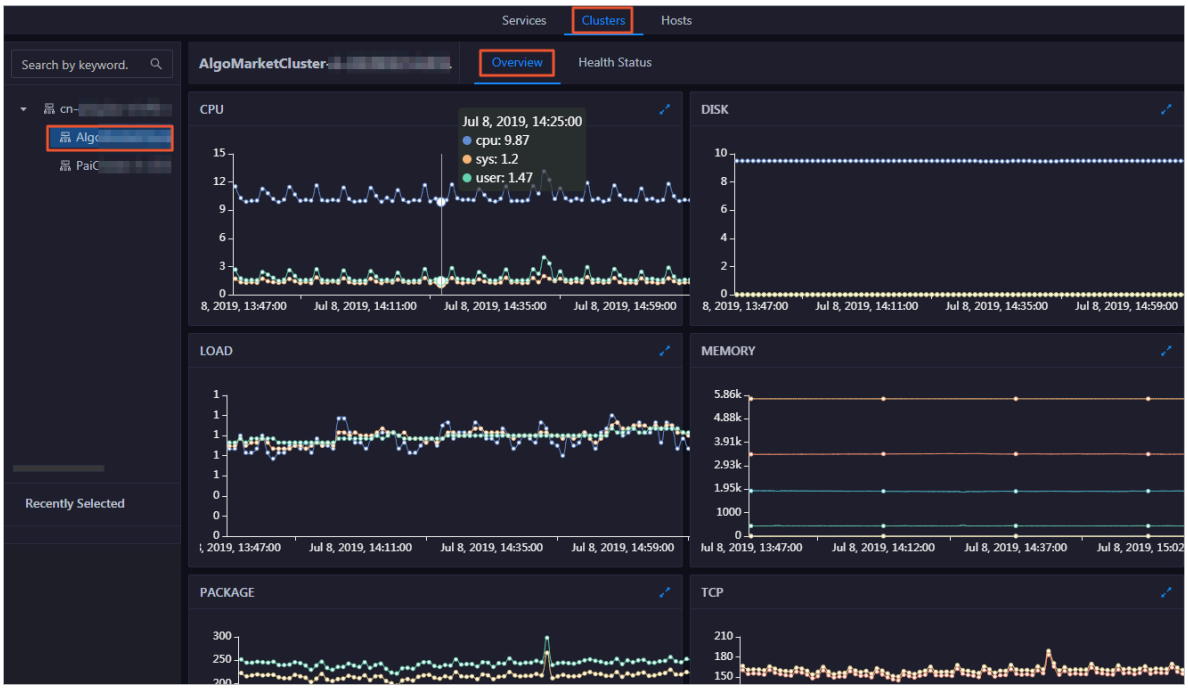

CPU

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 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.

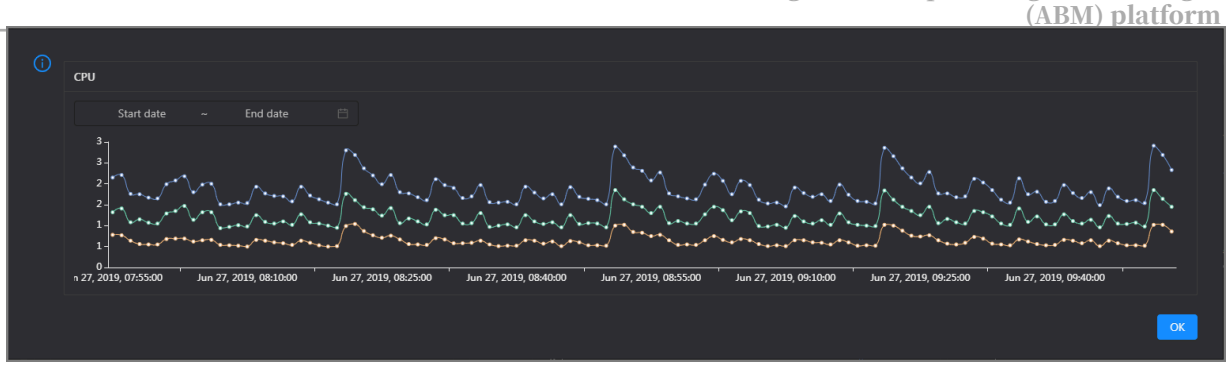

DISK

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

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

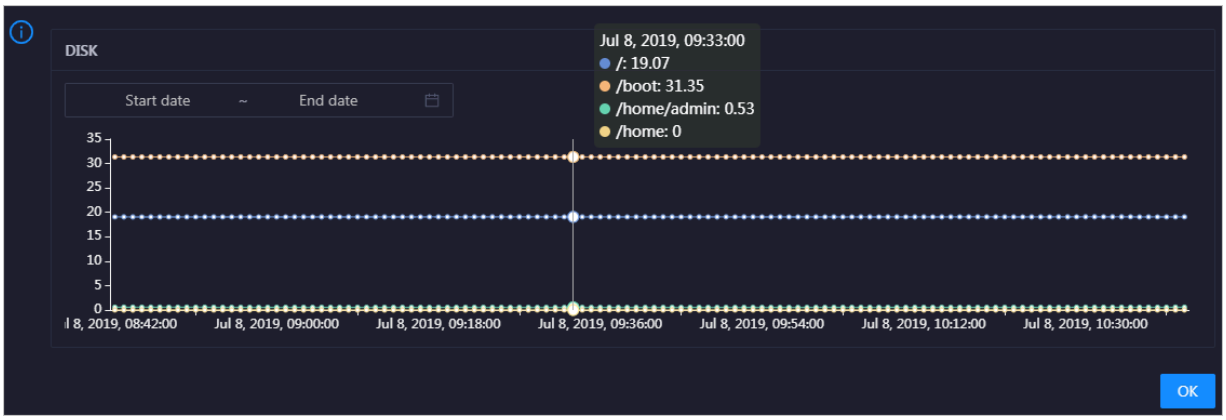

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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 cluster over time in different colors.

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

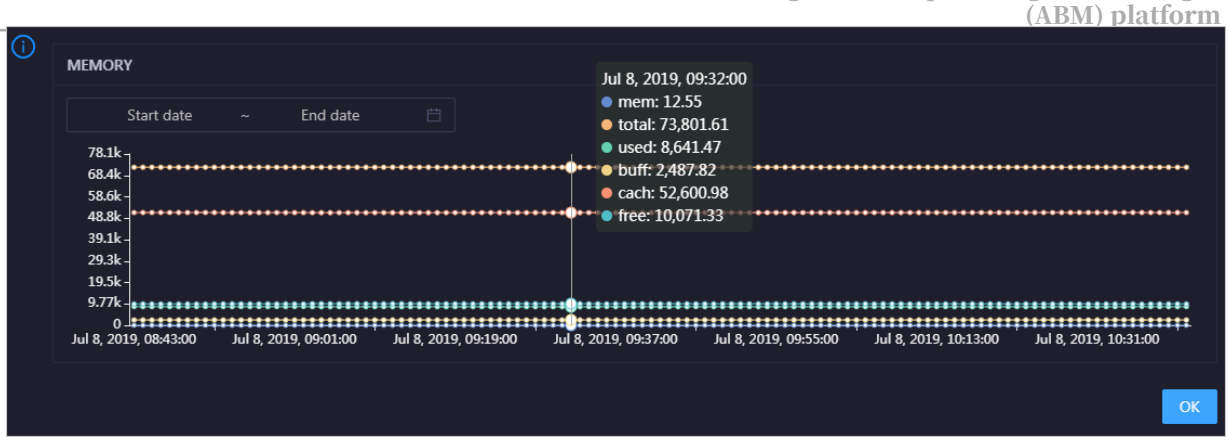

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

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

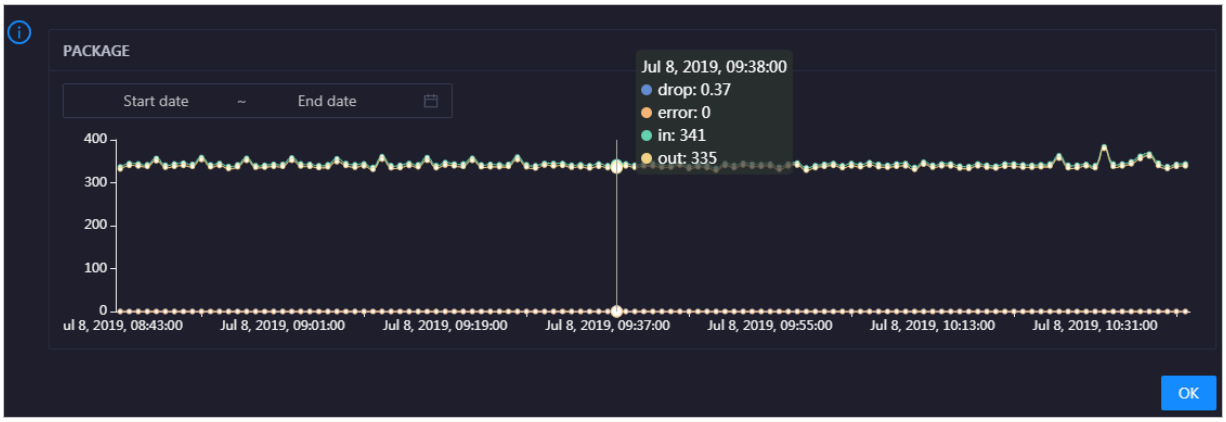

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

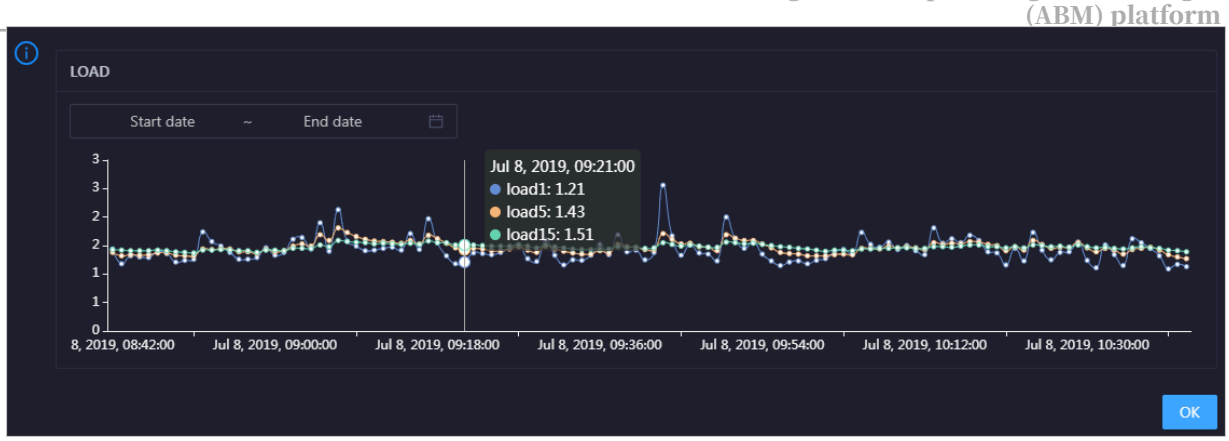

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

1.9.3.2 Cluster health

On the 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

1. **At the top of the O&M page, click Clusters.**
2. **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 page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

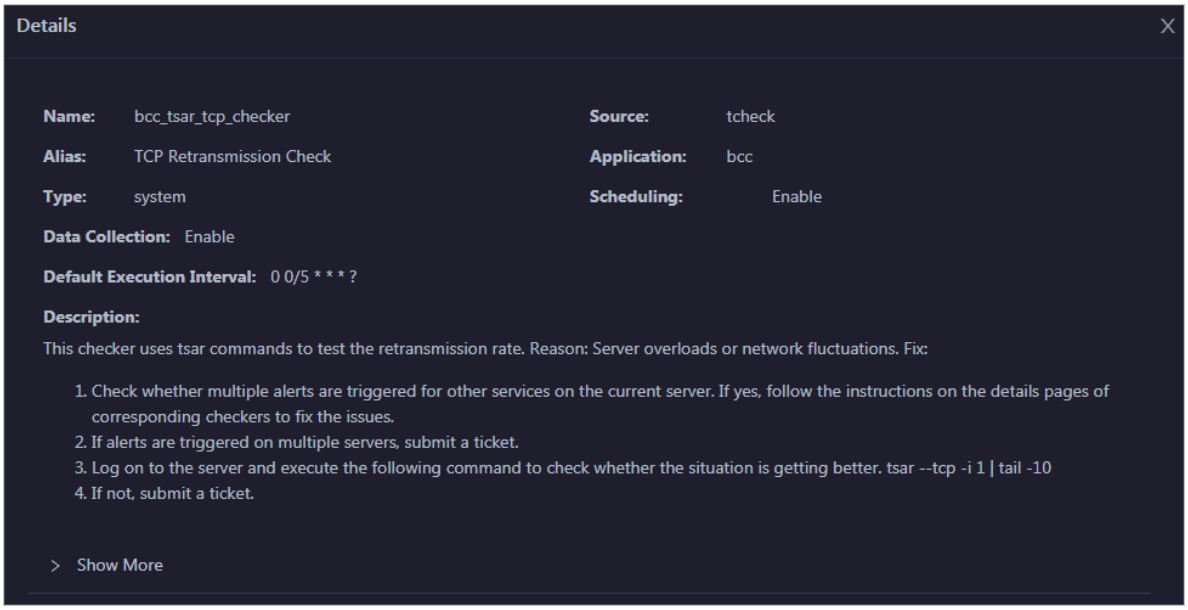

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

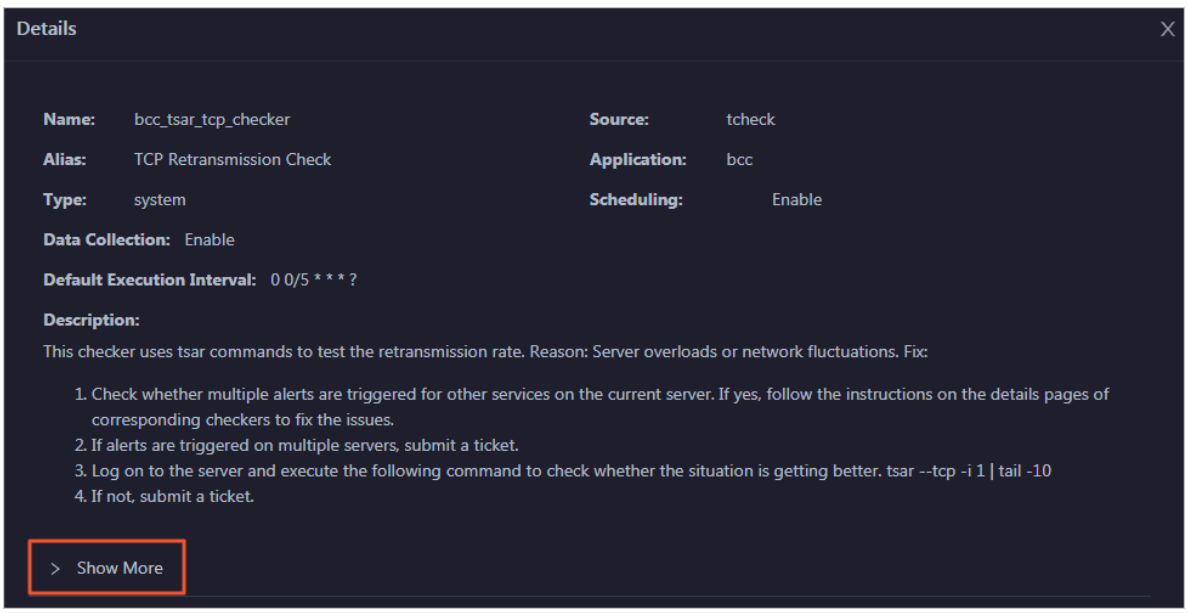

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View hosts with alerts and the alert causes

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

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

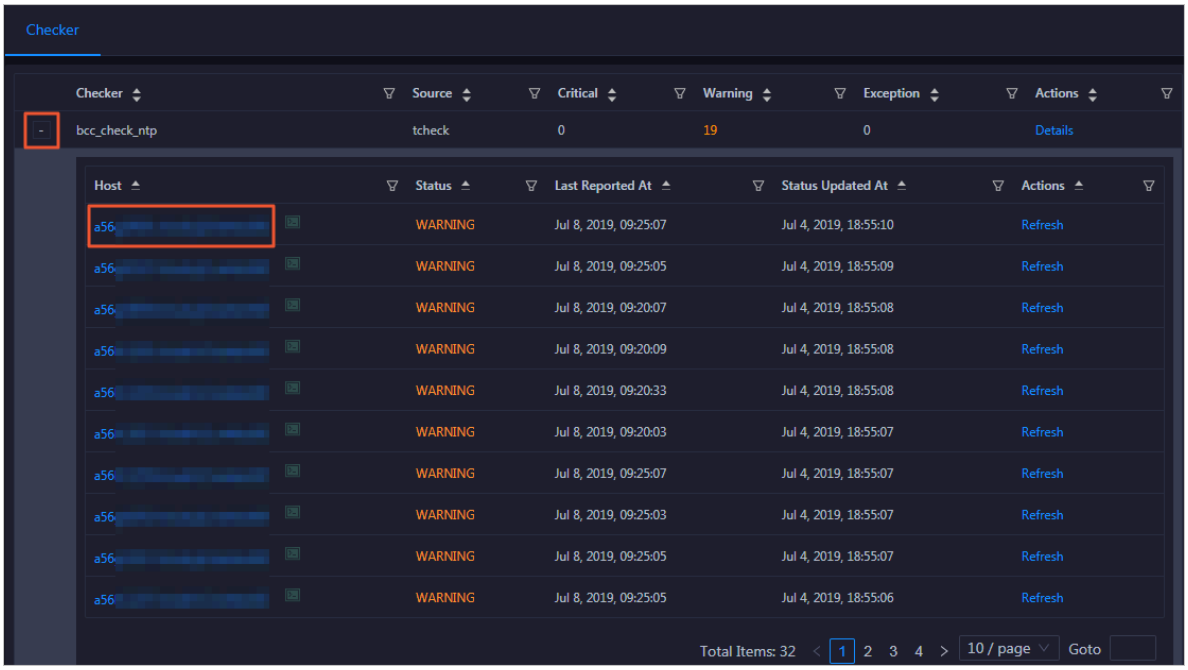

2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.

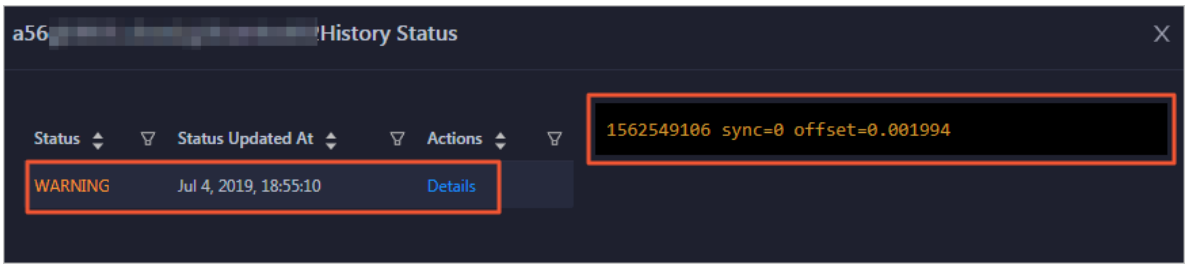

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

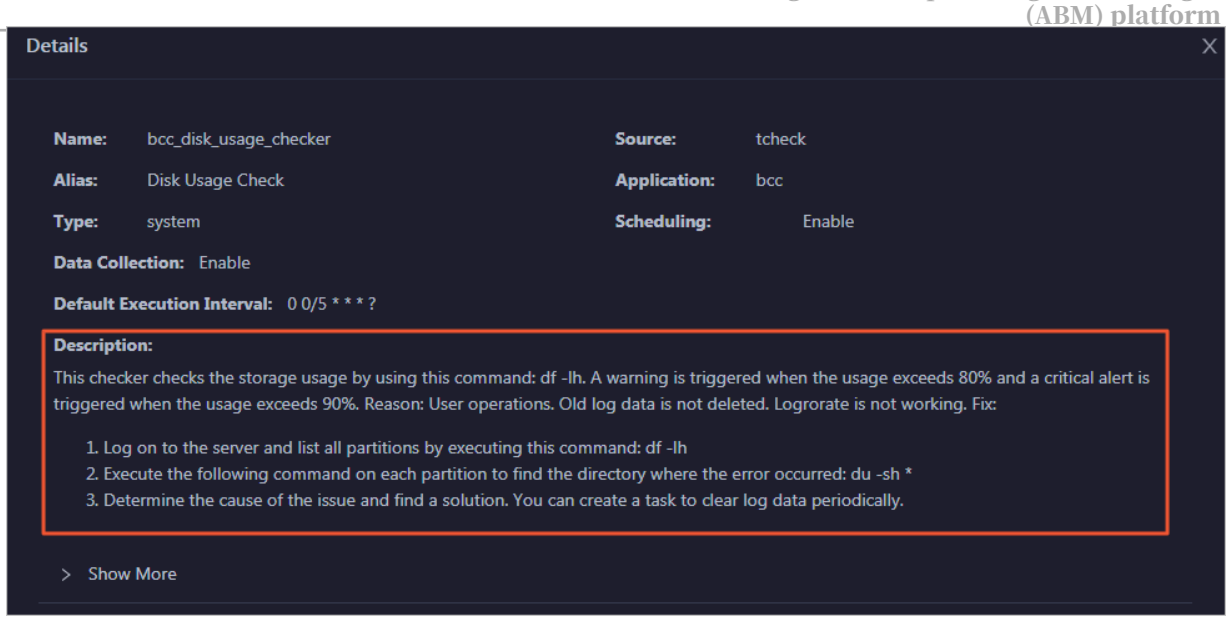

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

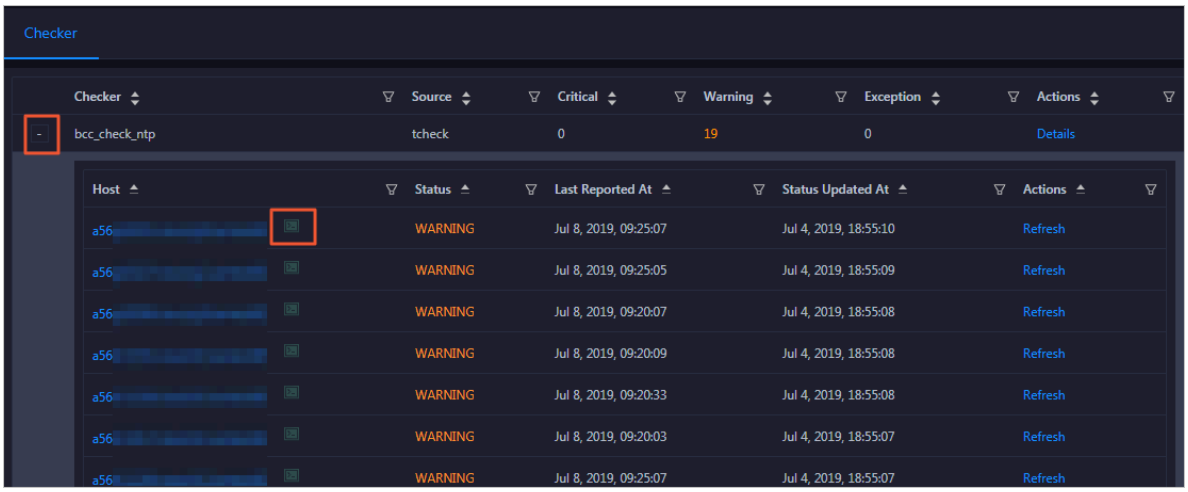

2. Click the Log On icon of a host. The TerminalService page appears.

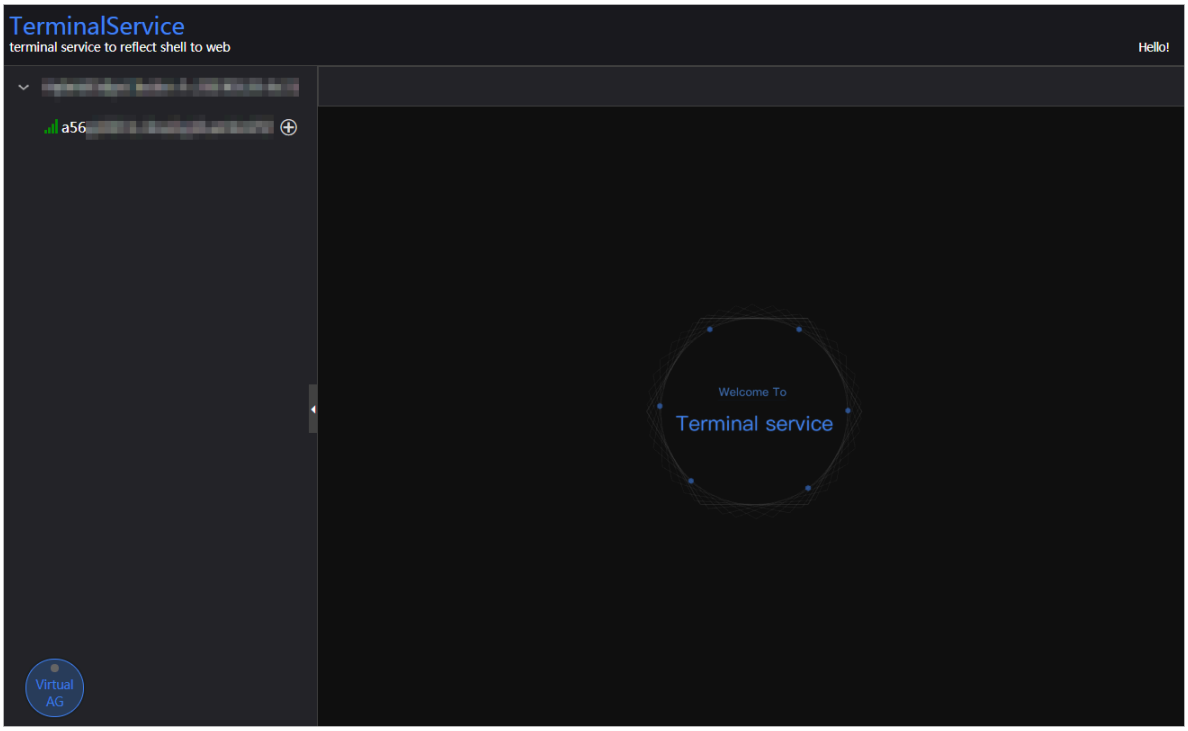

3. On the TerminalService page, click the hostname on the left to log on to the host.

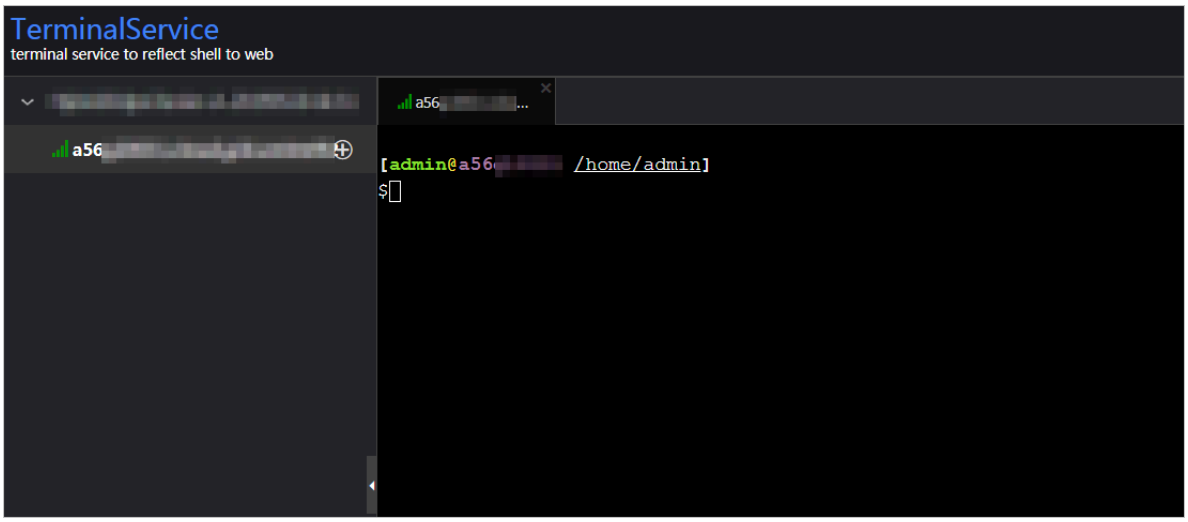

Run a checker again

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. In this way, you can check whether the alert is cleared.

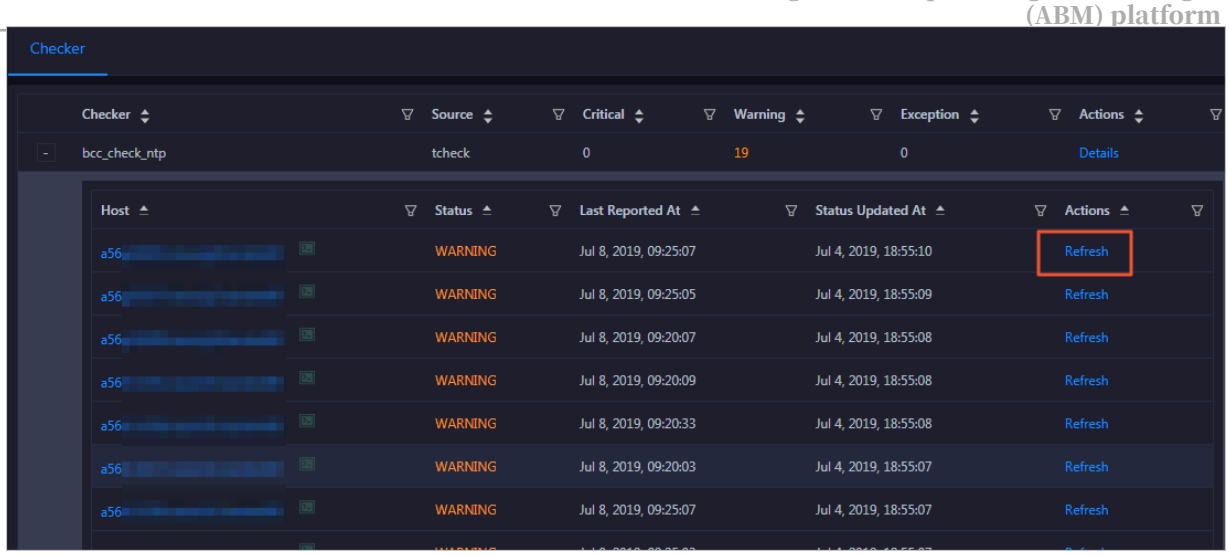

1.9.4 Host O&M

1.9.4.1 Host overview

The host overview page displays the overall running information about a host in a Machine Learning Platform for Artificial Intelligence (PAI) 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.

Entry

1. **At the top of the O&M page, click the Hosts tab.**

2. On the page that appears, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.

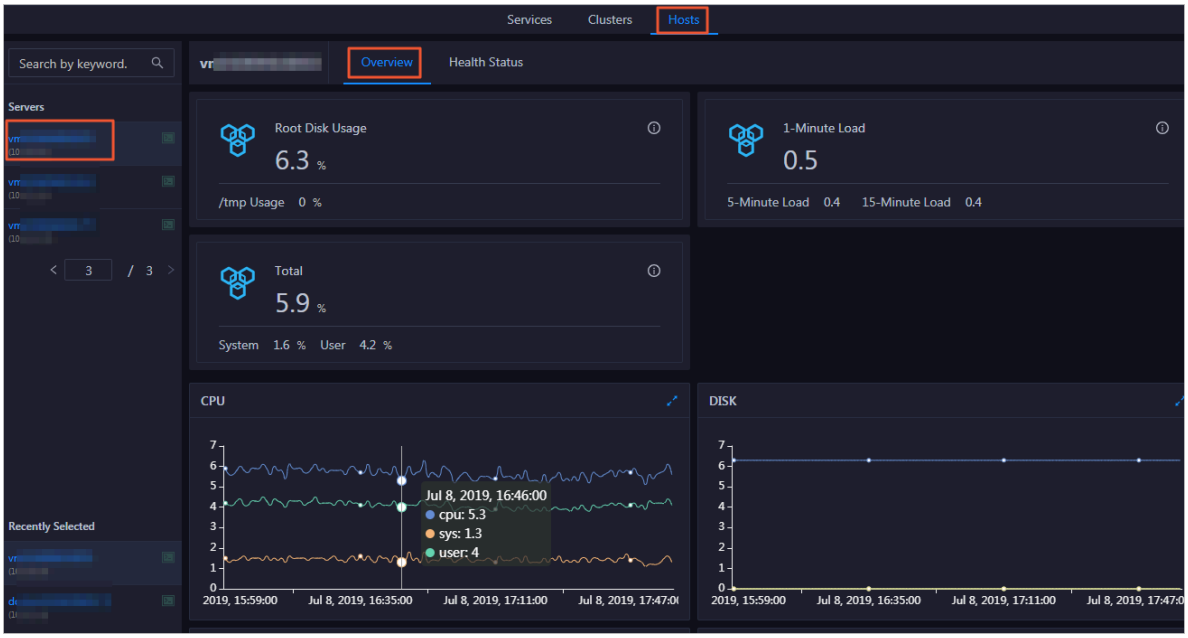

On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

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.

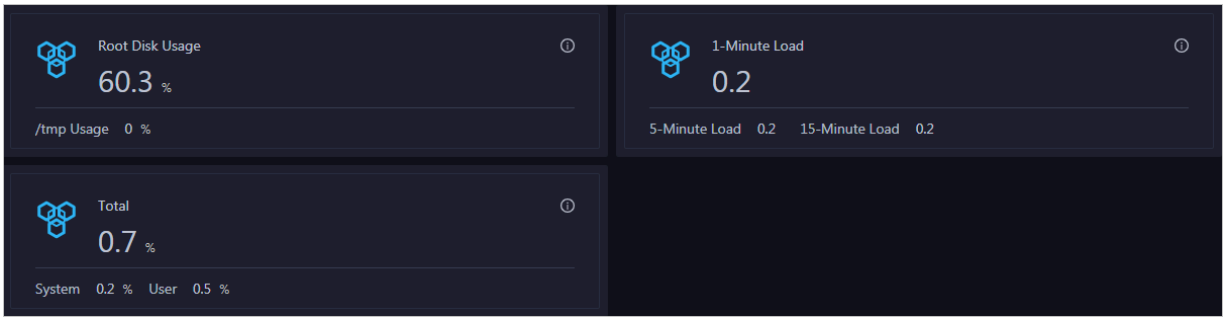

CPU

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 host over time in different colors.

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

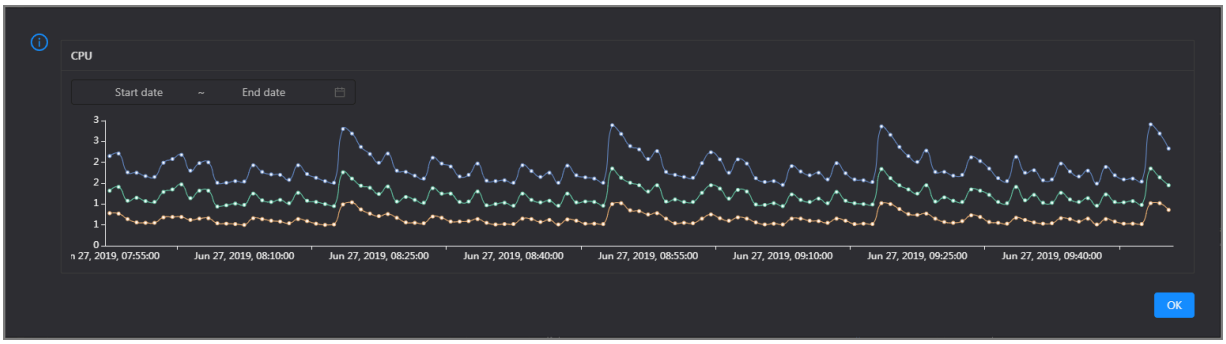

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

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

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

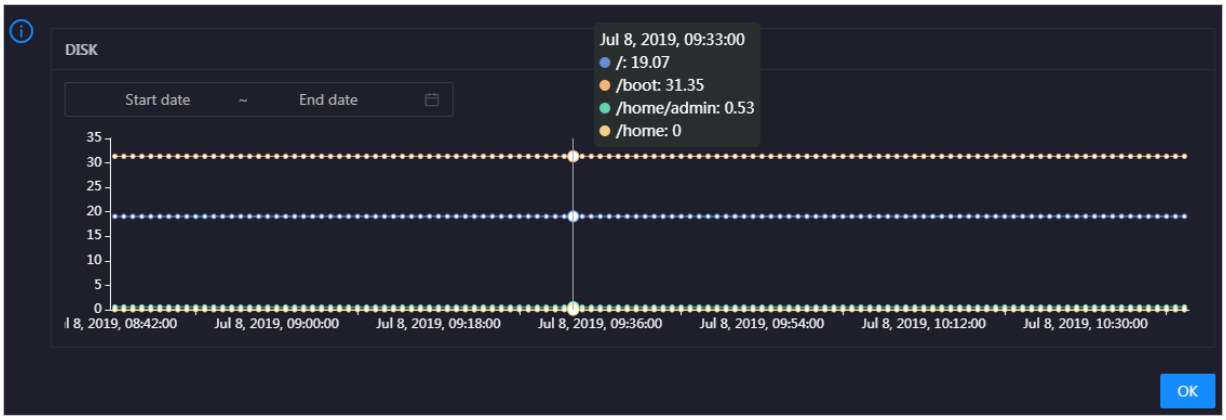

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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 host over time in different colors.

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

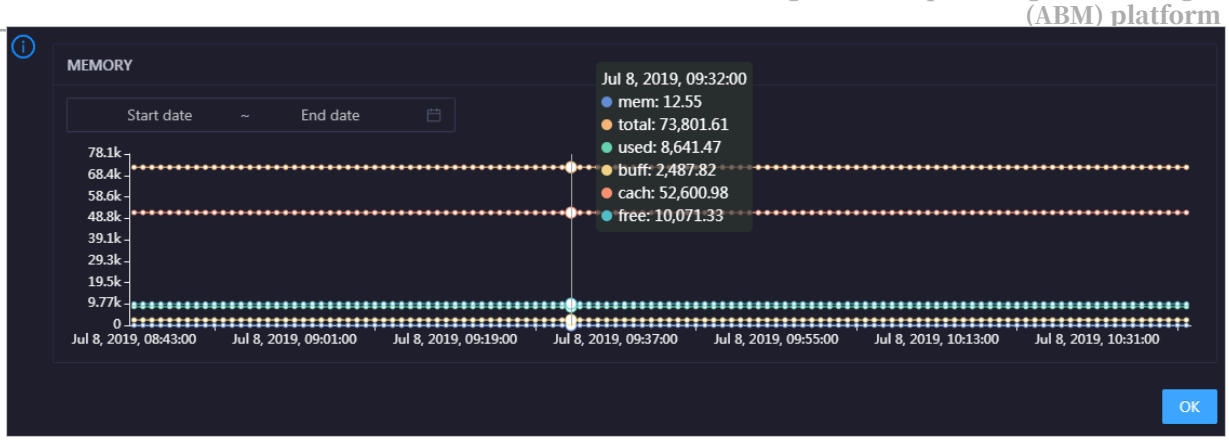

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

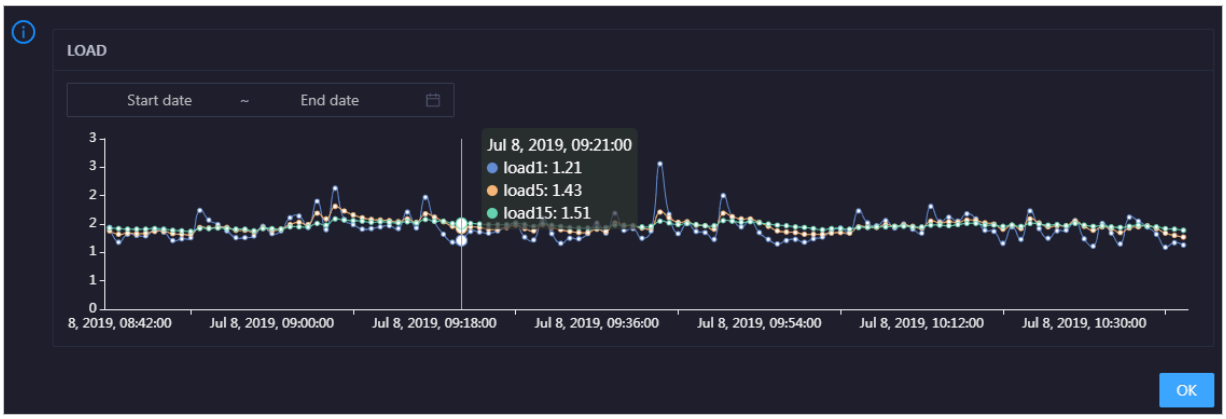

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 host over time in different colors. These trend lines reflect the data transmissi on status of the host.

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

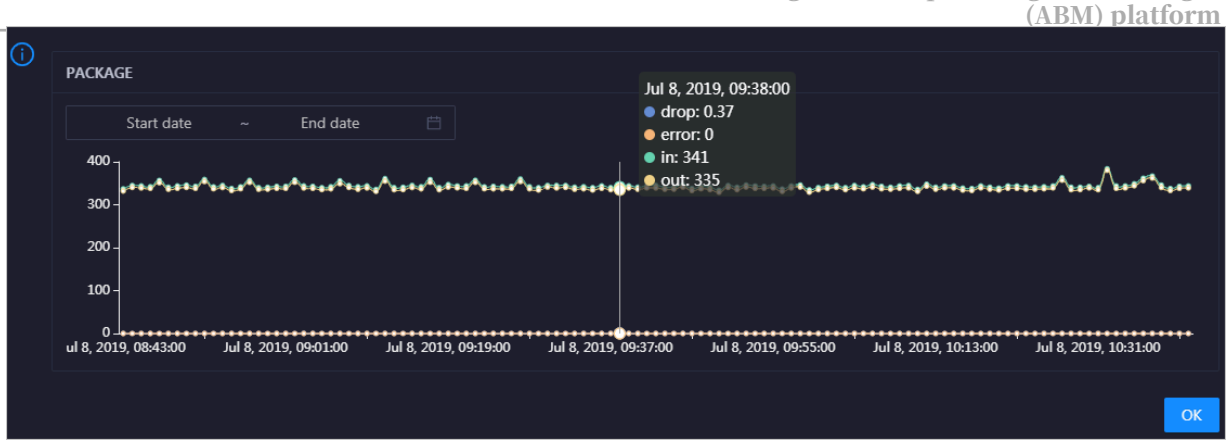

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

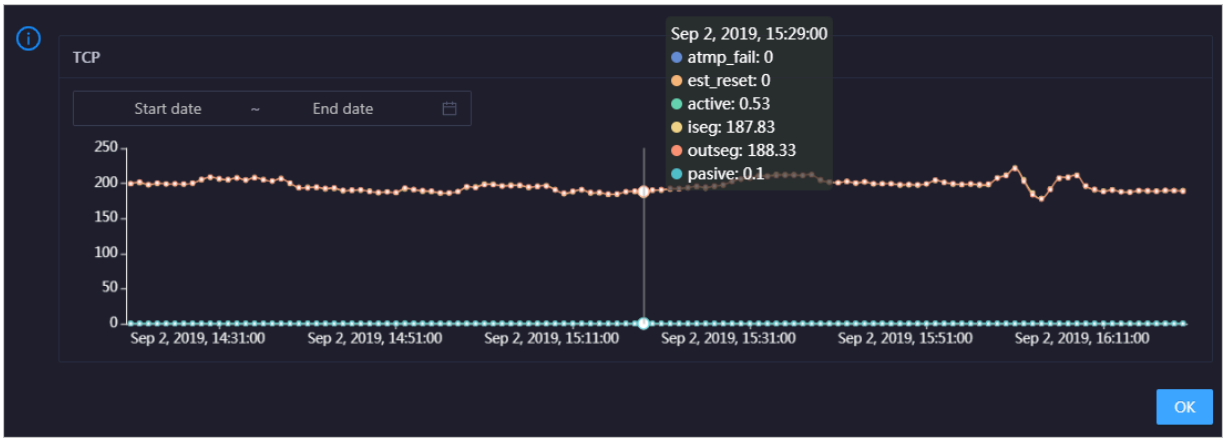

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

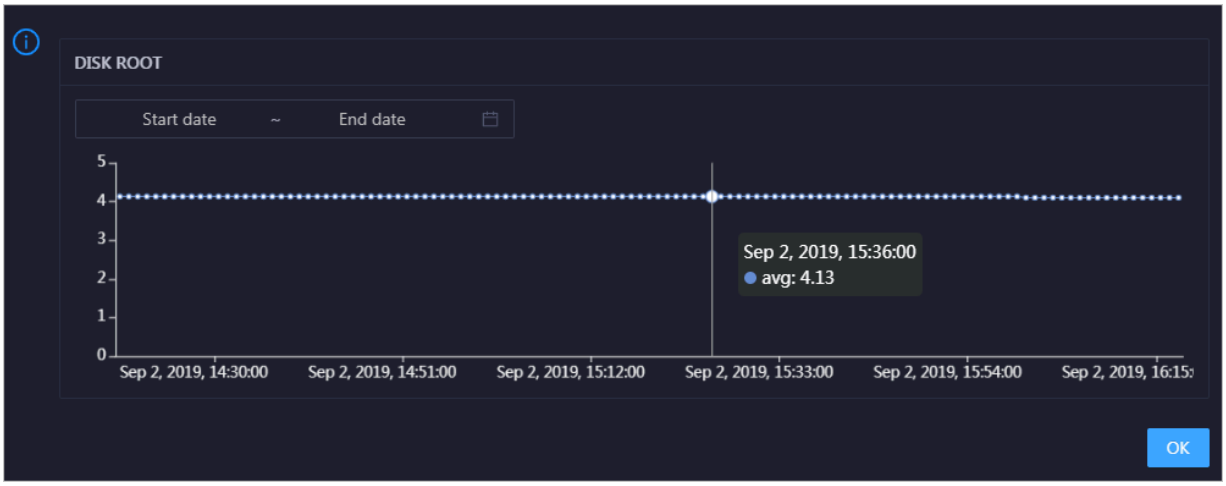

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

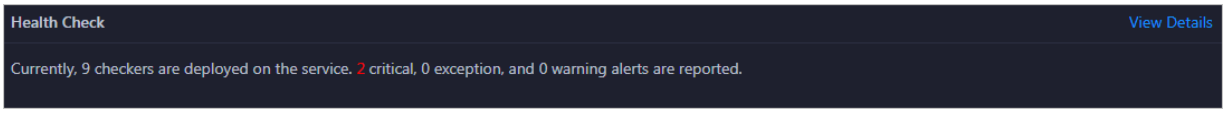

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

Health Check History

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

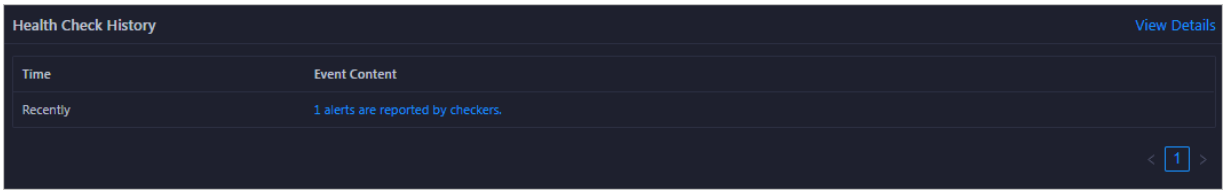

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see *Host health*.

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

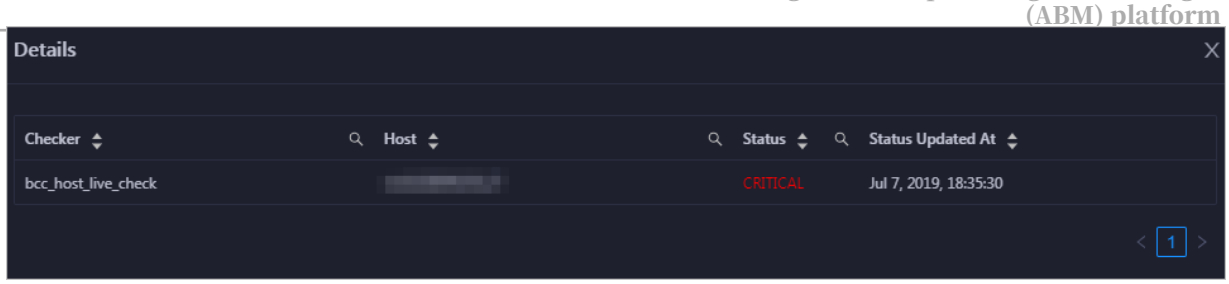

1.9.4.2 Host health

On the Health Status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to the host and perform manual checks on the host.

Entry

1. **At the top of the O&M page, click Hosts.**
2. **On the Hosts page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.**

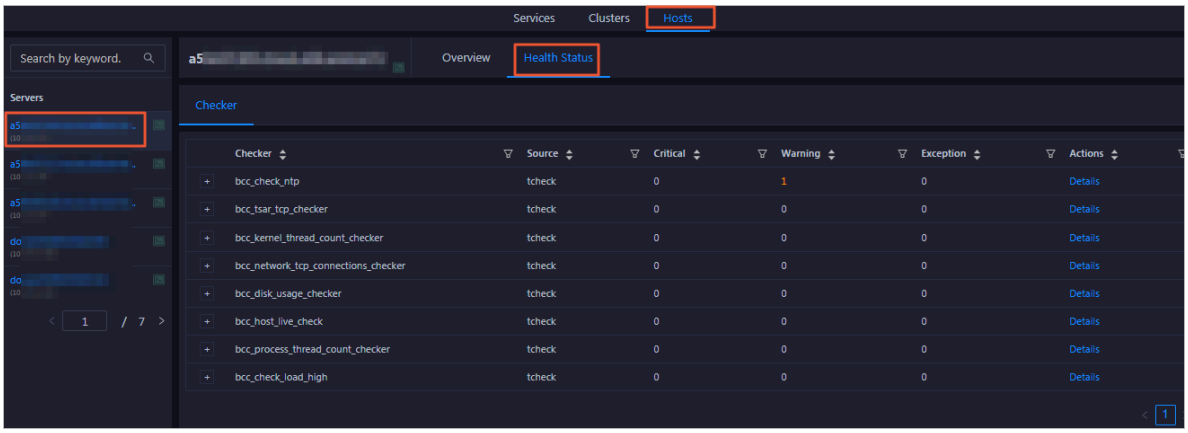

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, and Exception. They are displayed in different colors. Pay attention to the check results, especially the Critical and Warning results, and handle them in a timely manner.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

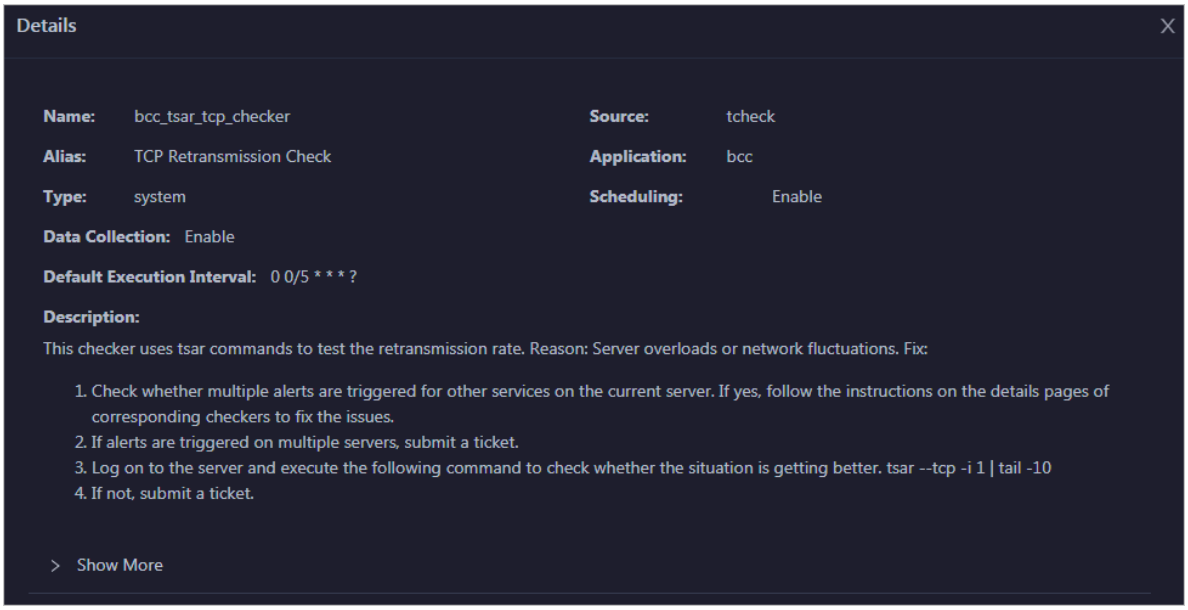

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

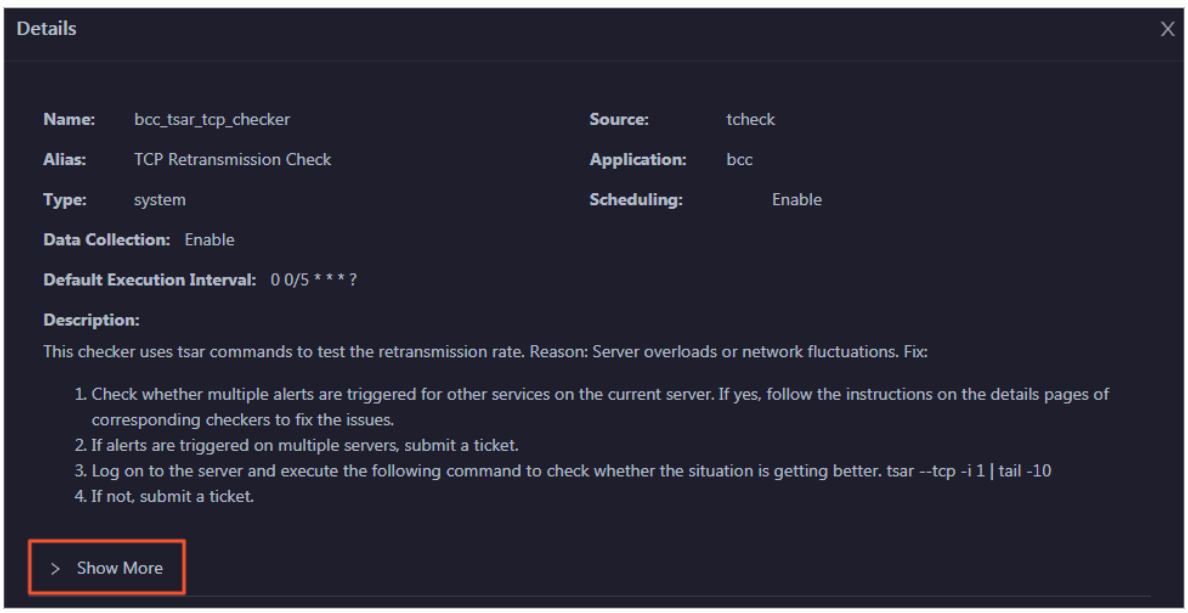

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

1. **On the Health Status page, click + to expand a checker with alerts.**

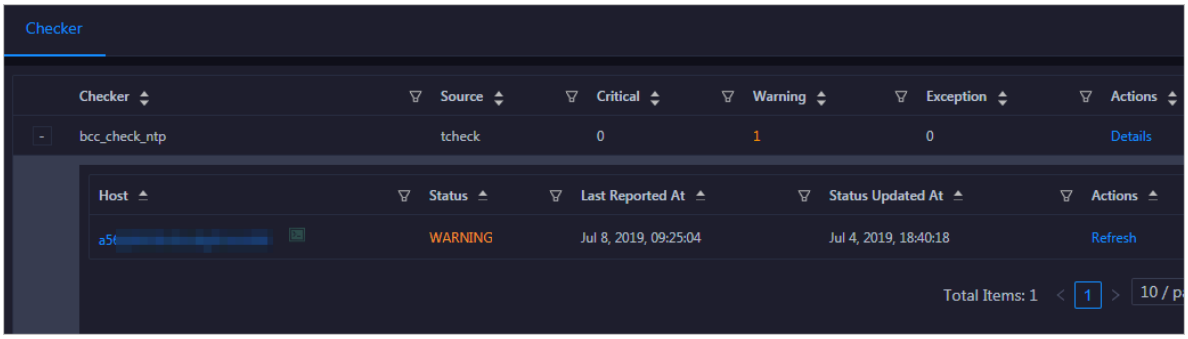

2. **Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**

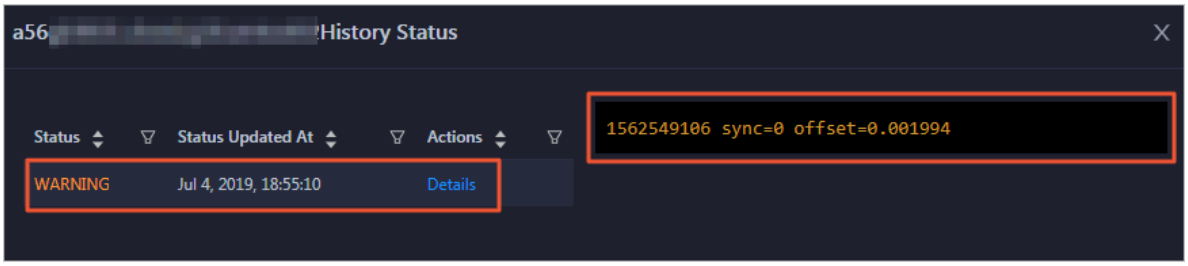

Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

1. **On the Health Status page, click + to expand a checker with alerts.**

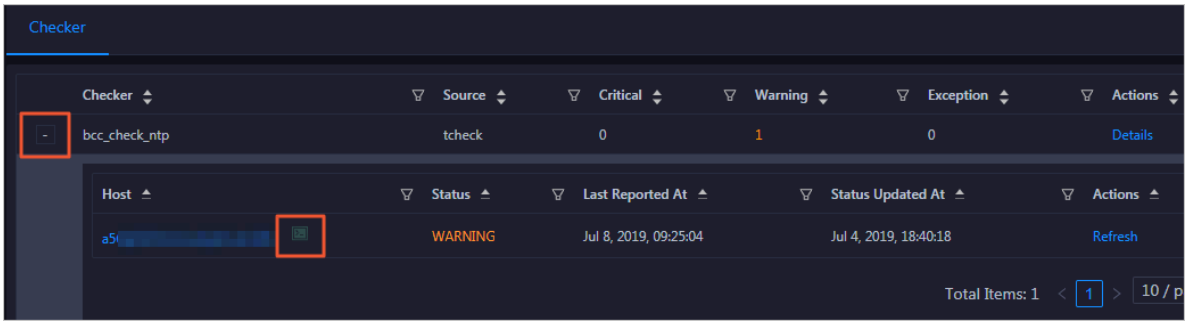

2. **Click the Log On icon of a host. The TerminalService page appears.**

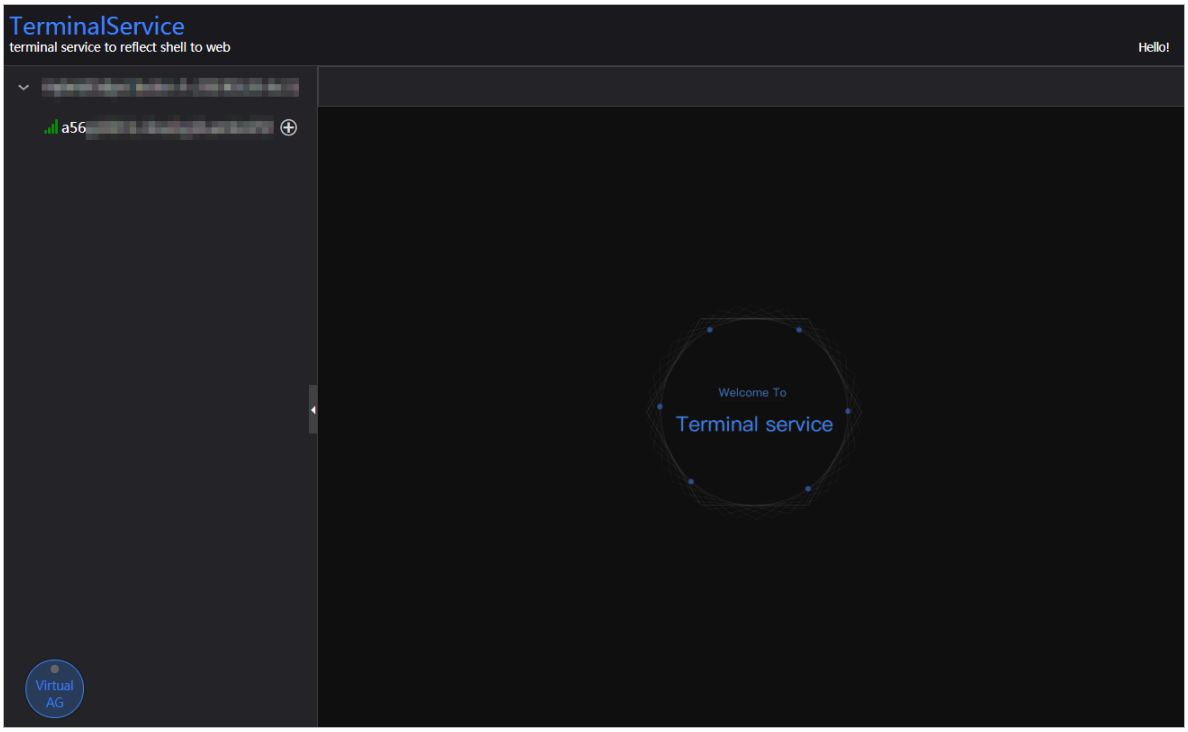

3. On the TerminalService page, click the hostname on the left to log on to the host.

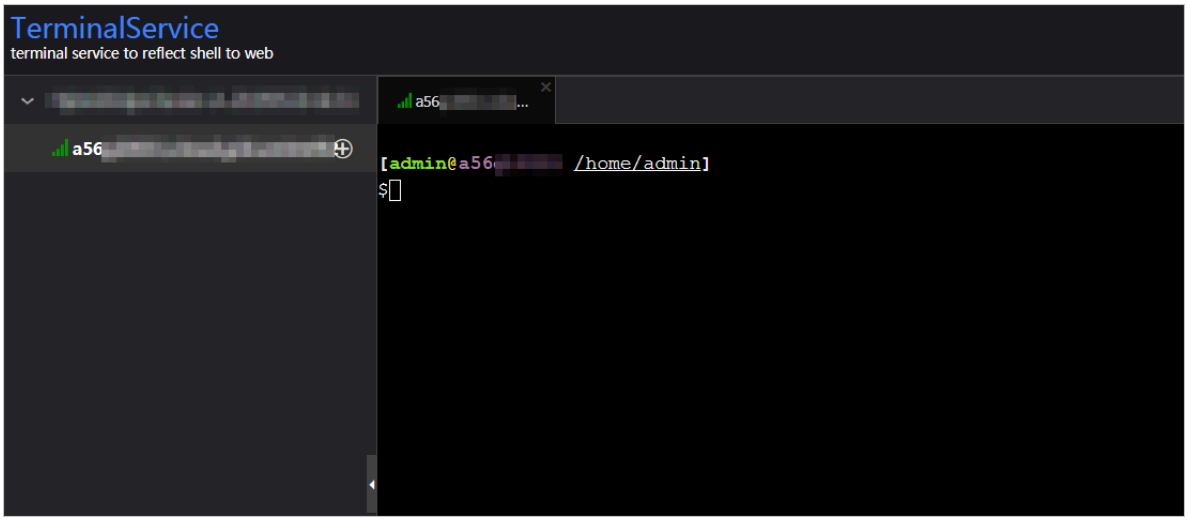

Run a checker again

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. In this way, you can check whether the alert is cleared.

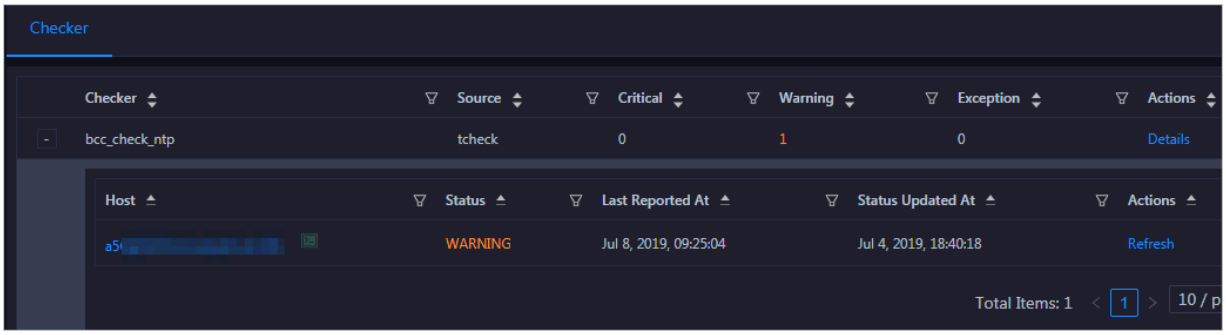

1.10 Management

1.10.1 Overview

The management module is the configuration and software management center of Apsara Bigdata 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.**

1.10.2 Jobs

1.10.2.1 Overview

This topic describes the Jobs page and concepts related to jobs.

Apsara Bigdata Manager (ABM) runs jobs to perform O&M operations on big data services. Jobs, also known as service O&M tasks, are O&M operations performed on physical devices in the cluster. The Jobs page consists of the Job Execution and Job Management tabs.

Concepts

Concepts related to jobs include:

- **Ordinary job: a job that can only be manually run.**
- **Cron job: a job that is automatically run based on timer settings.**
- **Scheme: a job template provided by ABM. You can use schemes to generate jobs.**
- **Atom: a step template provided by ABM. You can use atoms as steps when generating jobs.**
- **Ordinary step: a step that you need to create when using schemes to generate jobs. Step types include the following: command execution, script execution, file push, API call, and manual step.**
- **Atomic step: a step that you can directly use when using schemes to generate jobs .**

ABM provides common schemes and atoms that support most O&M scenarios.

Job Execution page

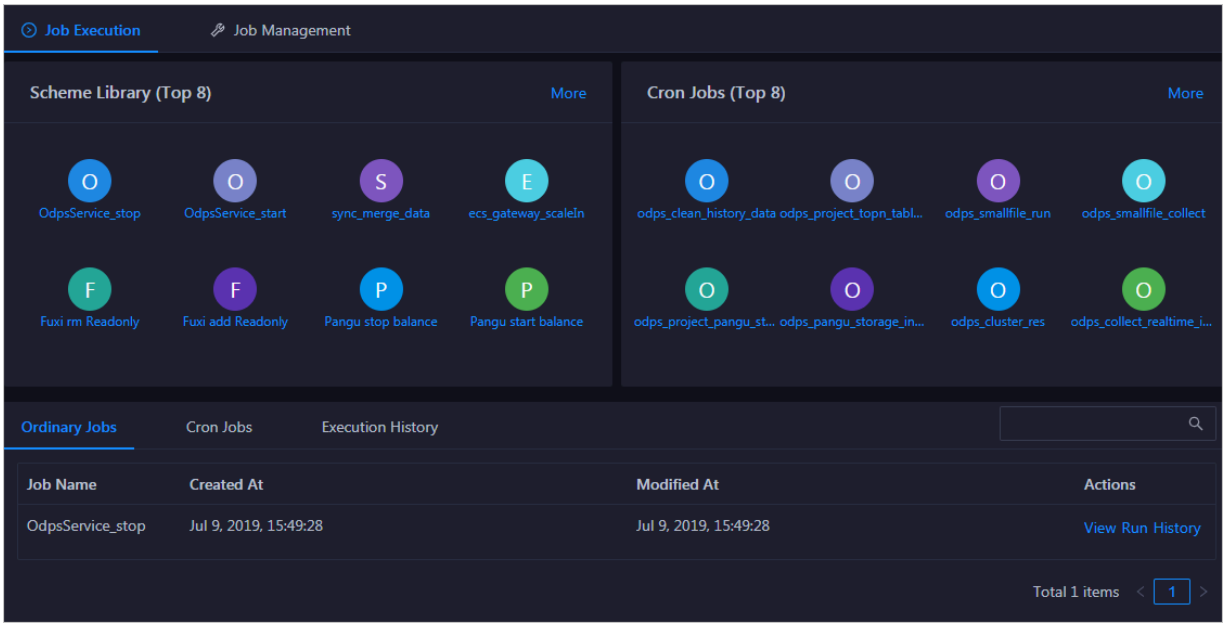

The Job Execution page provides the following features:

- **Ordinary 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.
- **Scheme Library (Top 8): dynamically displays the top 8 most used schemes.**
- **Cron Jobs (Top 8): dynamically displays the top 8 most used cron jobs.**
- **Execution History:**

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

  You can search for the execution record of a specific job by multiple conditions.

Job Management page

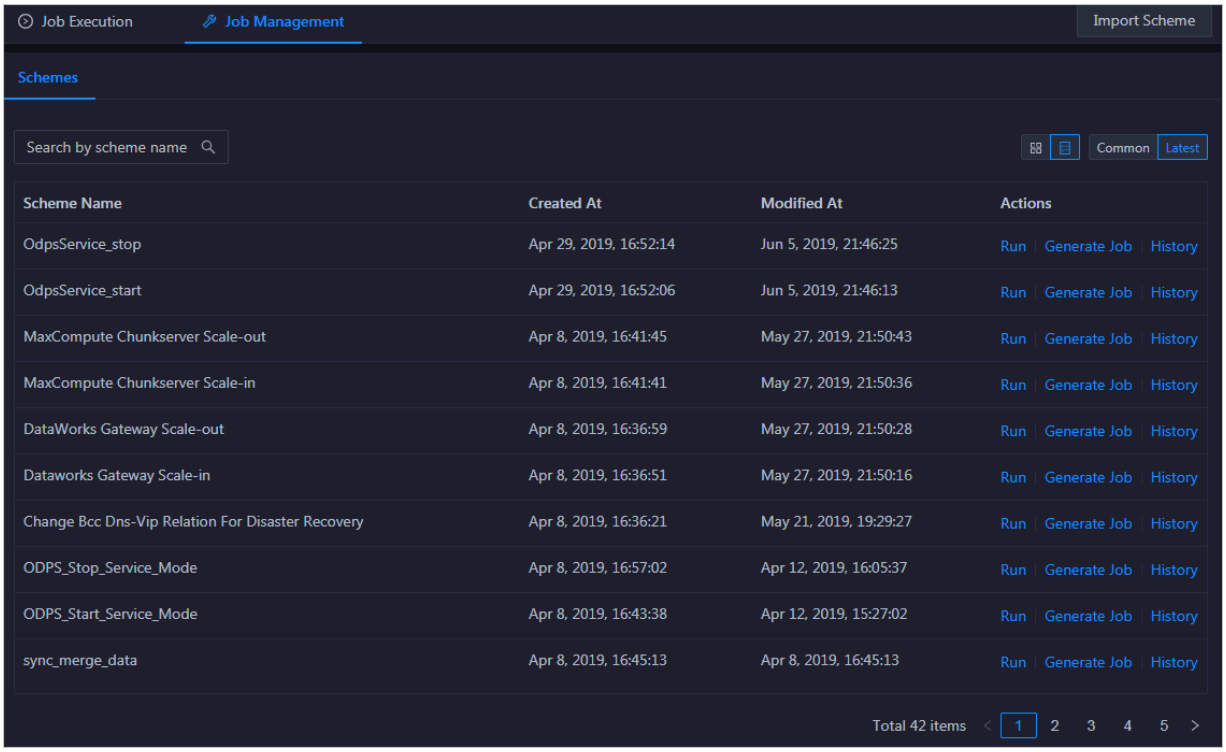

The Job Management page provides the following features:

- **You can generate and run jobs based on schemes and view the execution history of schemes.**
- **You can search for a specific scheme.**
- **You can view schemes in grid or list mode.**

1.10.2.2 Jobs

1.10.2.2.1 Run a job from a scheme

When you perform O&M operations, you can directly run jobs from schemes that meet your requirements. This enables you to quickly perform product O&M jobs.

Prerequisites

You must have an ABM administrator account.

Context

When you run a job from a scheme, you need to specify the Target Group and Global Variable parameters. The other parameters cannot be modified. If you want to modify the parameters, see *Create a job from a scheme*.

Running a job from a scheme is a one-time operation and does not generate a job on the Ordinary Jobs tab. You can view the history operations on the Execution History tab. For more information, see *View the execution history*.

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. **Run a job by using one of the following methods:**
   - **In the Scheme Library (Top 8) section, select a scheme.**

     

     Note:
     This method only allows you to choose a scheme from the top 8 most frequently used schemes.
   - **On the Jobs page, click the Job Management tab, and then click Run in the Actions column of a scheme in the Schemes list.**

4. **On the Run from Scheme page, you need to set Target Group and Variable Name as needed.**

   The instructions for setting Target Group and Variable Name are shown in *Table 1-1: Job parameters*.

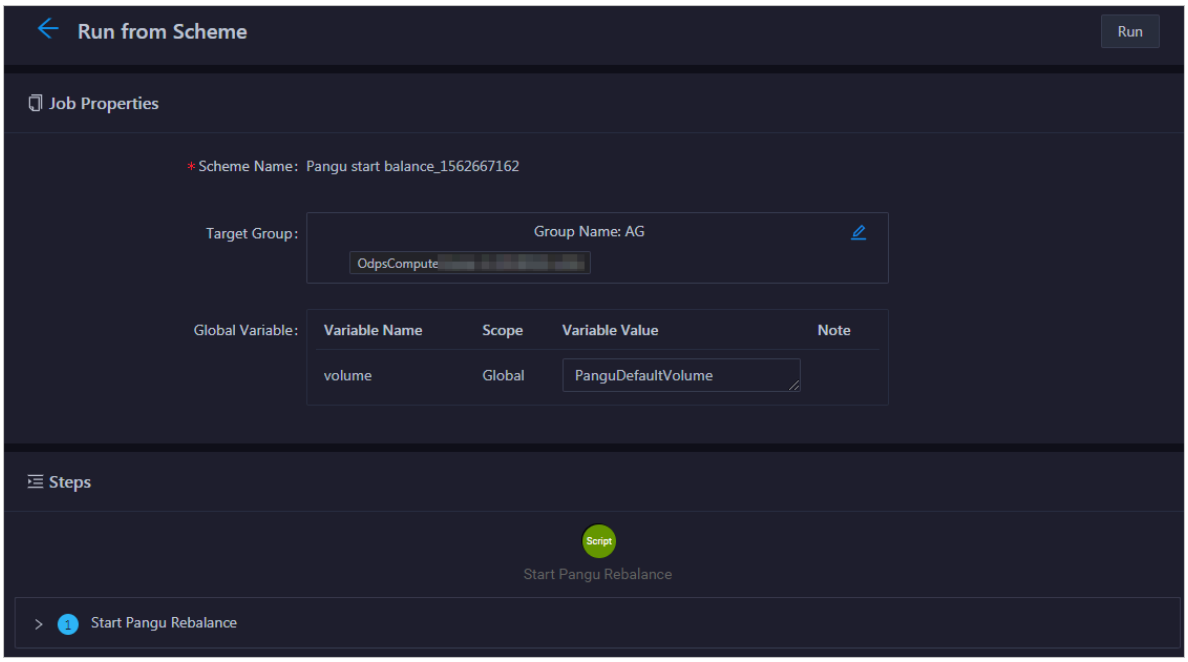

Table 1-1: Job parameters

Parameter	Description
Target Group	A collection of target nodes on which the operations are performed. After you have added nodes to target groups, you can select a value for Target Group based on your needs when you configure the steps. Click next to the target group, and set the nodes to be included in the target group as needed. When you add a node, you can either select the name of the node in Apsara Infrastructure Management Framework or enter the IP address of the node under Servers.
Global Variable	If global variables are set in the scheme, you need to enter the variable value.

5. After you have configured the preceding parameters, click Run in the upper-right corner.
6. Confirm the job risks in the displayed dialog box, and click Confirm Execution.

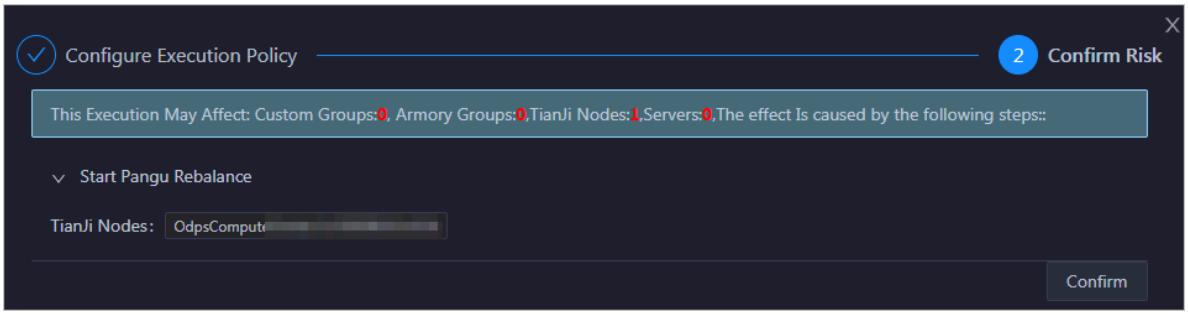

After you have confirmed, a record is automatically generated on the Execution History page. For more information, see *View the execution history*.

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

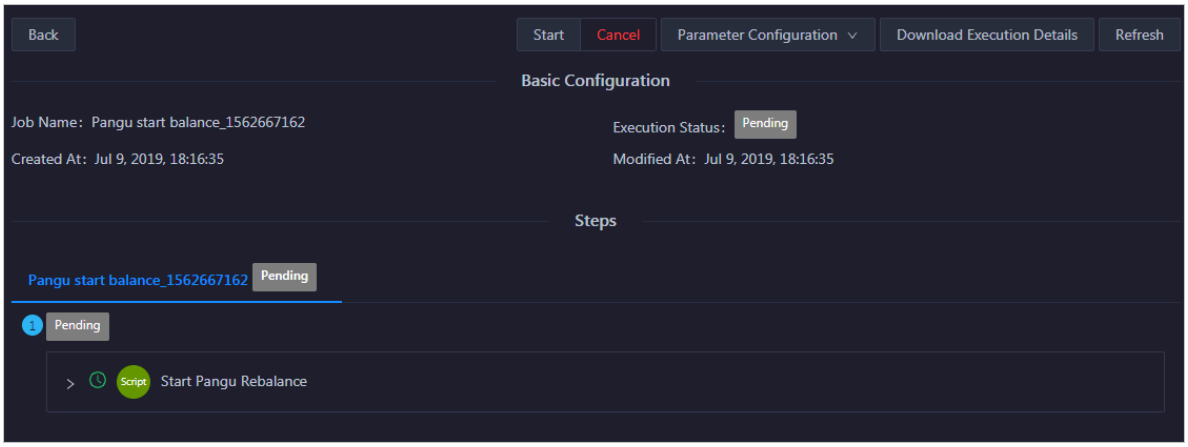

If you do not perform any operation and exit the job execution page, you can find a job record on the Execution History page. Click View to go to the job execution page again.

1.10.2.2.2 Create a job from a scheme

This topic describes how to generate a job from a scheme. You can generate both ordinary and cron jobs from schemes.

Prerequisites

An Apsara Bigdata Manager (ABM) administrator account is obtained.

Context

ABM allows you to create both ordinary and cron jobs from schemes. Settings for creating an ordinary job and a cron job are similar, but a schedule must be created for 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. On the Jobs page, click the Job Management tab, and click Generate Job in the Actions column of a scheme in the Schemes list.
4. On the Create Job page, set the parameters in the Job Properties and Steps sections as needed.

   For more information about the parameter configuration of Job Properties, see *Table 1-2: Job properties*.

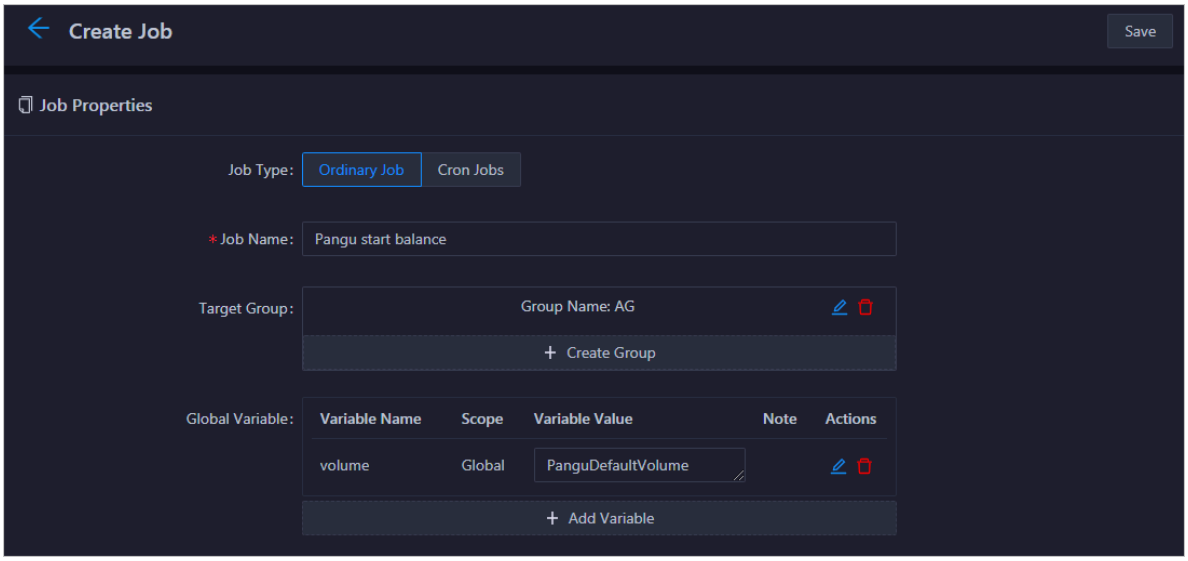

Table 1-2: Job properties

Parameter	Description
Job Type	The type of the job. · Ordinary Job: jobs that must be manually run. · Cron Jobs: jobs that automatically run based on a schedule. You can enter a cron expression or click Configure Cron Job to create a schedule. Cron expressions are based on crontab commands. If you are new to crontab commands, click Configure Cron Job to quickly set up a schedule.

Parameter	Description
Job Name	The name of the job. Set the job name based on the functionality of the job to be created so that the user understands what it is and can search for it.
Target Group	A collection of target nodes on which the operations are performed. After you have added nodes to the target groups, you can select the target group based on your needs when you configure the steps. After you have created a group, click to add nodes to the group. When you add a node, you can either select the name of the node in Tianji or enter the IP address of the node under Servers.
Global Variable	Click Add Variable and set the parameters in the dialog box that appears. The Scope parameter is used to set the scope of the variable. If it is set Global, it is valid for the entire job. If you select a certain step, it is only valid for this step.

5. On the Create Job page, add steps as needed.

The steps include ordinary steps and atomic steps.

- +Atomic Step: a range of built-in steps provided by the system.
- +Ordinary Step: Ordinary steps are classified into multiple types. Choose the required type and set the parameters accordingly. *Table 1-3: Parameters of command execution steps*, *Table 1-4: Parameters of script execution steps*, *Table 1-5: Parameters of file*

push steps, *Table 1-6: Parameters of API call steps*, and *Table 1-7: Parameters of manual steps* describe the parameters for different types of ordinary steps.

Table 1-3: Parameters of command execution steps

Section	Parameter	Description
N/A	Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Basic configuration	Target Node Group	The group of nodes on which the step is performed.
	Commands	The commands to be executed in this step.
	User Identity	The user who executes this step on the nodes, with a default setting of admin.
	Description	The description of the step.
Advanced Configuration	Input Context	Enable this option if you need to obtain the output of the previous step. When enabled, this step reads the file specified by the `$contextInput` variable to obtain the context.
	Output Context	Enable this option if you need to export the context to the next step. When enabled, this step writes the context to the file specified by the `$contextOutput` variable to export the context.
	Timeout Period	The maximum time period allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out. The default value is 60 seconds.

Section	Parameter	Description
	Retries	The number of times to retry the execution after a failure or timeout error occurs. The default value is 0.
	Retry Interval	The interval between two executions. The default value is 300 seconds. The retry interval is the period of time between the last timeout (or failure) and the next try.

Table 1-4: Parameters of script execution steps

Section	Parameter	Description
N/A	Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Basic configuration	Target Node Group	The group of nodes on which the step is performed.
	Script Content	Write the script based on the actual O&M requirements. Currently, Shell and Python are supported. You can write new scripts or upload local scripts to configure the script content.
	User Identity	The user who executes this step on the nodes, with a default setting of admin.
	Description	The description of the step.
Advanced Configuration	Input Context	Enable this option if you need to obtain the output of the previous step. When enabled, this step reads the file specified by the `$contextInput` variable to obtain the context.

Section	Parameter	Description
	Output Context	Enable this option if you need to export the context to the next step. When enabled, this step writes the context to the file specified by the `$contextOutput` variable to export the context.
	Timeout Period	The maximum time period allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out. The default value is 60 seconds.
	Retries	The number of times to retry the execution after a failure or timeout error occurs. The default value is 0.
	Retry Interval	The interval between two executions. The default value is 300 seconds. The retry interval is the period of time between the last timeout (or failure) and the next try.

Table 1-5: Parameters of file push steps

Parameter	Description
Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Target Node Group	The group of nodes to which the file is pushed.
Target Path	The directory to which the file is pushed.
File Permission	The permission of the file.
File Owner	The owner of the file.

Parameter	Description
File Content	Enter the file content in the code editor or upload a local file. After you enter or upload the content, specify the file name in the code editor.

Table 1-6: Parameters of API call steps

Parameter	Description
Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Target URL	The URL of the API.
HTTP Method	The type of request that you want to send. · GET: Query. · POST: Create. · PUT: Modify. · DELETE: Delete.
Content Format	The Content-Type field of the header in the HTTP packet. Select a value from the drop-down list.
APP NAME	APP NAME and APP KEY are included in the request to call APIs for authenticating permissions.
APP KEY	
BODY	The body of the HTTP request.
Timeout Period	The maximum time period allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out. The default value is 60 seconds.
Retries	The number of times to retry the execution after a failure or timeout error occurs. The default value is 0.

Parameter	Description
Retry Interval	The interval between two executions. The default value is 300 seconds. The retry interval is the period of time between the last timeout (or failure) and the next try.

Table 1-7: Parameters of manual steps

Parameter	Description
Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Document Content	The instructions to help relevant engineers complete this step.

6. To change the order of steps, click Sort in the upper-right corner of the Steps section and drag the steps to put them into the correct order.
7. After you have set the preceding parameters, click Save in the upper-right corner.

Result

If you created an ordinary job, it appears on the Ordinary Jobs tab. If you created a cron job, it appears on the Cron Jobs tab.

What's next

- If you created an ordinary job, you need to run it manually. For more information, see *Manually run a job*.
- If you created a cron job, you need to enable it. For more information, see *Enable or disable a cron job*. You can also manually run a cron job. For more information, see *Manually run a job*.

1.10.2.2.3 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.

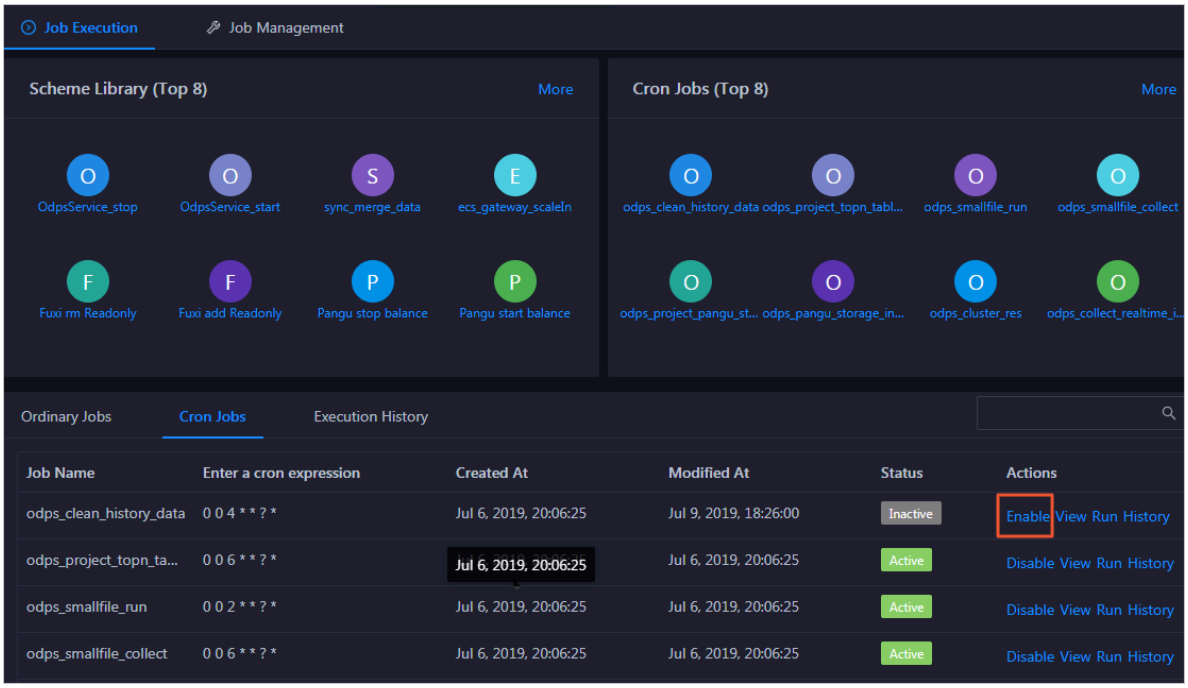

4. On the Cron Jobs page, you can enable or disable a cron job.
   - To enable a cron job in the inactive status, click Enable in the Actions column of the cron job.

     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.

1.10.2.2.4 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.

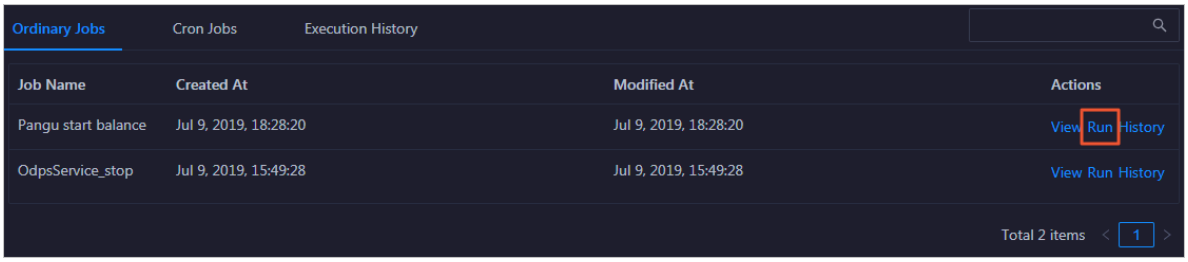

4. In the Ordinary Jobs list, click Run in the Actions column of a job.
5. Confirm the job risks in the dialog box that appears, and click Confirm.

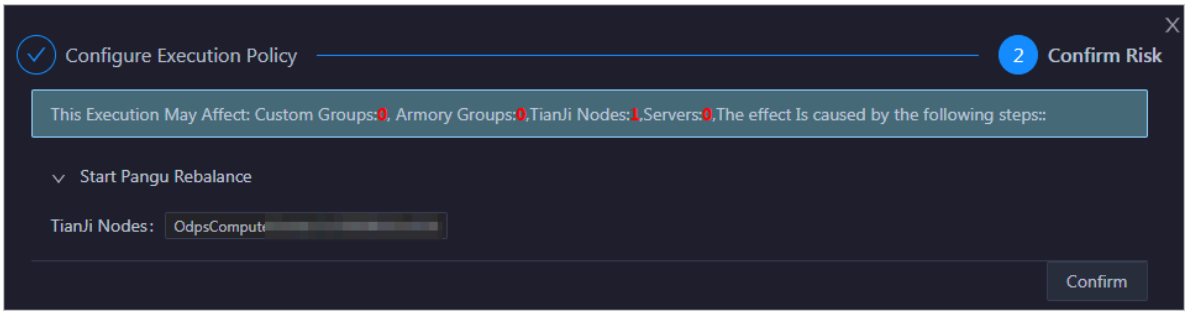

   After you have confirmed, a record is automatically generated on the Execution History page. For more information, see *View the execution history*.

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

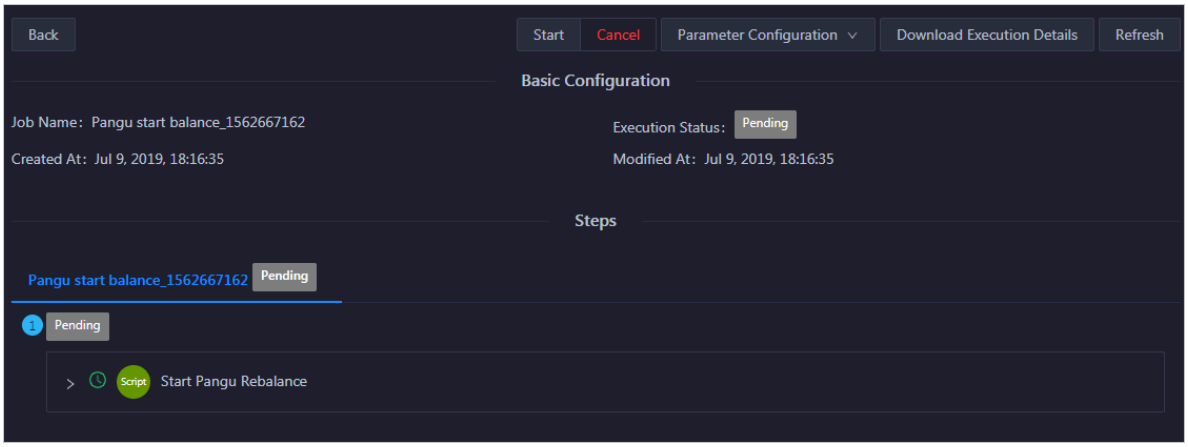

You can find the record about a job on the Execution History page, and click View to go to the detailed execution page.

1.10.2.2.5 View jobs

After you have created an ordinary job or a cron job, you can view job details, save the job as a scheme, and run the job in the jobs list.

Prerequisites

You must have an ABM administrator account.

Context

The topic describes how to view ordinary jobs. You can follow the same procedure to view cron jobs.

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.

4. Click View in the Actions column of an ordinary job to view its job details.

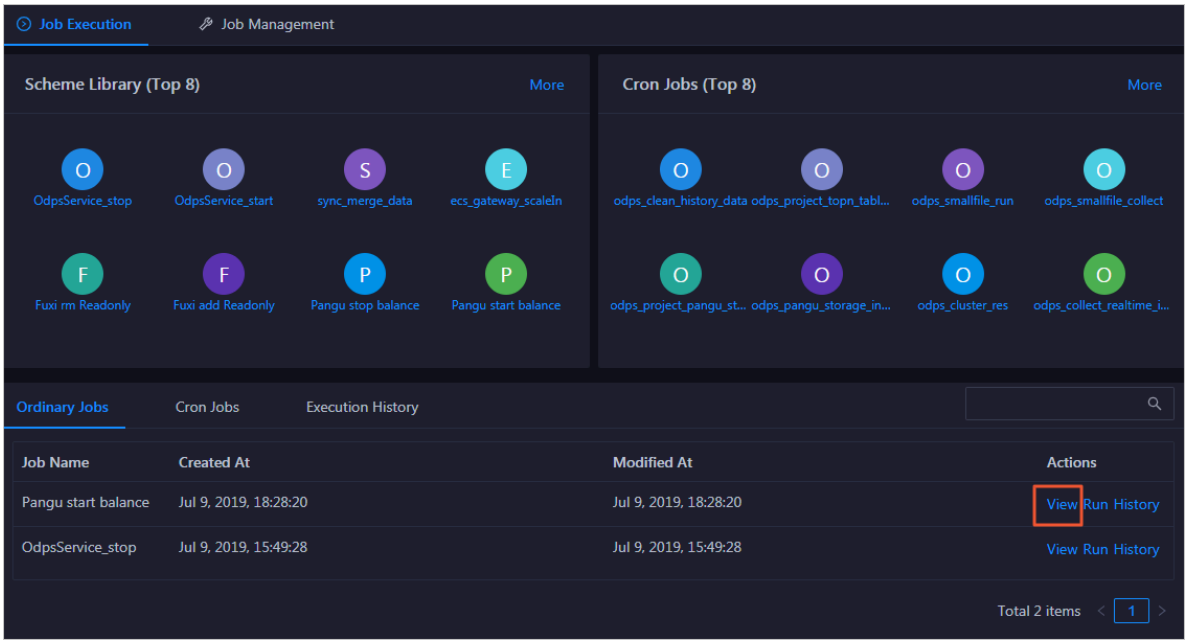

1.10.2.2.6 View the execution history of a job

Apsara Bigdata 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*.

1.10.2.3 Schemes

1.10.2.3.1 Create a scheme from a job

If an ordinary job or a cron job adapts to an O&M scenario of your service, you can save the job as a scheme to create service O&M tasks in similar scenarios.

Prerequisites

An Apsara Bigdata Manager (ABM) administrator account is obtained.

Context

Both cron jobs and ordinary jobs can be used to generate schemes. The procedures for these two types of jobs are the same. This topic uses the procedure for an ordinary job as an example.

> **Notice:**
> **When a cron job is saved as a scheme, no parameters are included.**

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.

4. On the Ordinary Jobs page, click View in the Actions column of an ordinary job.

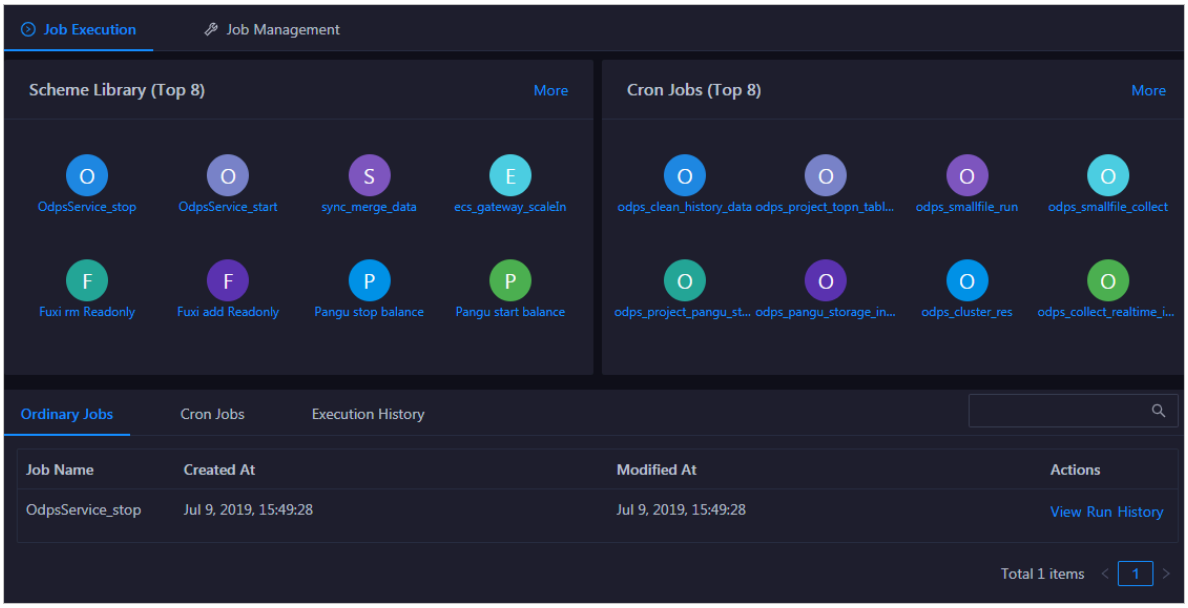

5. On the Job Details page, click Save as Scheme in the upper-right corner. The system prompts that you have saved the scheme.

Result

The new scheme has the same name as the job from which it was created and is listed on the Schemes page.

1.10.2.3.2 View schemes

A scheme is displayed in the scheme list after it is created. Apsara Bigdata Manager (ABM) allows you to view existing schemes in different ways, filter schemes, and search for specific schemes.

Prerequisites

An ABM administrator account is obtained.

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 Jobs page, click Job Management.

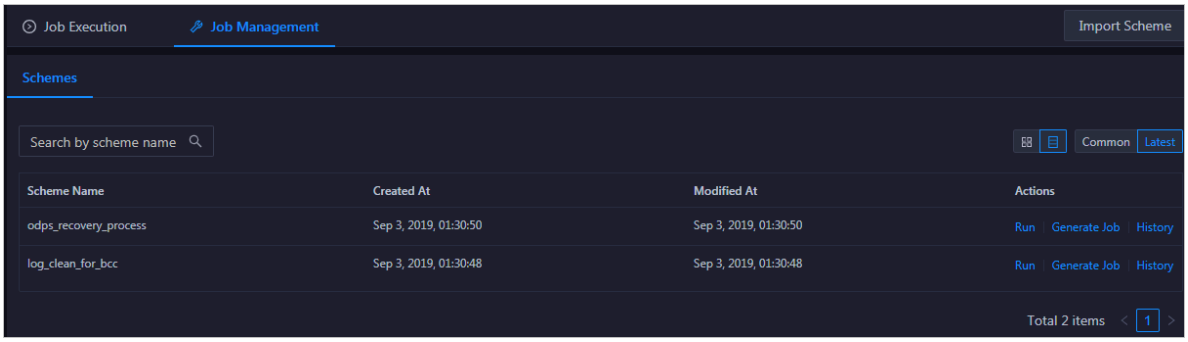

4. If there are too many schemes, you can enter the scheme name in the search bar to search for the required scheme.
5. Change the method for viewing schemes
   - View schemes in list (default): Click in the upper-right corner.
   - View schemes in cards: Click in the upper-right corner.

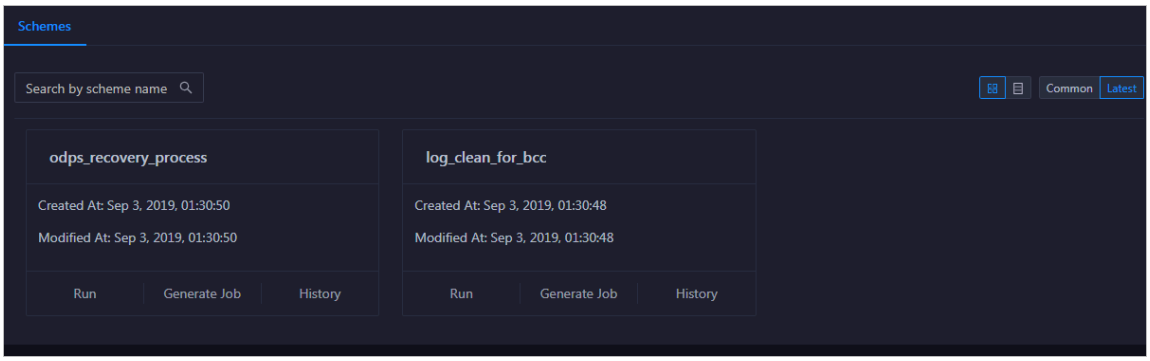

1.10.2.3.3 View the execution history of a scheme

Apsara Bigdata Manager (ABM) allows you to view the execution history of a specified scheme to learn the execution status of it.

Prerequisites

An ABM administrator account is obtained.

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 Jobs page, click Job Management.

4. On the Schemes page, click History in the Actions column of a scheme that has directly run jobs.

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

1.10.2.4 View the execution history

Apsara Bigdata 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.

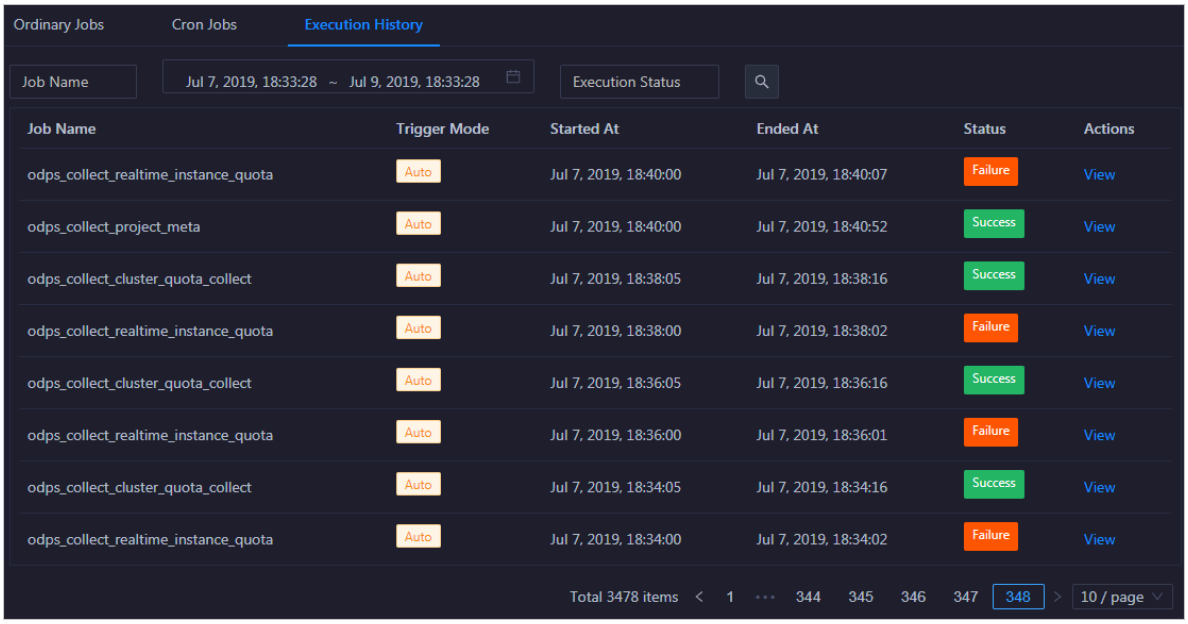

4. If there are too many execution records, filter them by a combination of one or more of the following filter conditions: job name, creator, execution status, and time range. Then, click to search for required records.

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.

Execution status	Feature	Operation
	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.
	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.

Execution status	Feature	Operation
	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.
	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.

Execution status	Feature	Operation
	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.
	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.

1.10.3 Patch management

Apsara Bigdata Manager (ABM) allows you to deploy and roll back upgrade patches for the services that it maintains. It also allows you to view detailed records of patch deployment and rollback by patch package or host.

Prerequisites

- **An ABM account with the required permissions to perform O&M operations on the corresponding service and the corresponding password are obtained.**
- **The patch package in the `tar.gz` format for the service to be upgraded is obtained.**
- **The cluster of the service to be upgraded is running properly.**

Entry

1. *Log on to the ABM console*.
2. **Click Management in the upper-right corner. On the page that appears, click Packages in the left-side navigation pane. The Packages page appears.**

   Description of the Packages page:

   - **Package Management: allows you to manage the patch packages of the service. You can upload, deploy, or delete the packages.**
   - **Package Deployment: displays the deployment history and details.**

   The Package Management page appears by default.

Upload a patch package

This section describes how to upload a patch package for ABM.

1. **Click Upload Package on the Package Management page.**
2. **In the dialog box that appears, select a patch package, and then click Upload. Wait until the uploading is complete.**

   After the patch package is uploaded, the system prompts a success message. The patch package is then displayed in the list.

Deploy a patch package

After a patch package is uploaded, you can deploy it to the corresponding service cluster.

1. **In the patch package list, click Deploy in the Actions column of a patch package.**

2. **In the dialog box that appears, set Cluster and Deployment Mode.**

   The valid values of Deployment Mode include:

   - **All: Deploy the patch package to all hosts where it has not been deployed.**
   - **Phased Release: Deploy the patch package on a random host.**

3. **Click OK.**

   The deployment status of the patch package is Deploying. Patch deployment takes some time. Wait until the patch package is deployed. Refresh the page after the deployment is complete. The deployment status is changed to Deployed.

Handle deployment failures

After you use ABM to deploy a patch for a service, the patch will be automatically bound to the service release (SR) version of the service. If the service is upgraded, the SR version is changed, and the deployment status of the patch package is changed to Deployment Failed (Product Upgraded).

After the service is upgraded, ABM cannot determine whether the new version has fixed the problem to be resolved by the patch. Therefore, the patch automatically becomes invalid. If the service upgrade cannot fix the problem to be resolved by the patch, click Force Deploy to deploy the patch again. If the service upgrade has fixed the problem to be resolved by the patch, click Ignore.

View the deployment history and details

The Deployment Records page displays the deployment information about all patch packages. The Deployment Details page displays the deployment information about all hosts.

1. **Click the Package Deployment tab on the Packages page to view the deployment records.**

   The Deployment Records page displays the deployment records of all patch packages. You can view the name, version, product, cluster, service, service role, application type, deployment mode, and operation type of each patch package. You can also view the users who submitted the deployment requests, the total number of hosts where each patch package needs to be deployed, the number of hosts where each patch package is deployed, the number of hosts where each

patch package fails to be deployed, the number of hosts where the deployment has not finished, and the deployment time.

If too many deployment records exist, you can filter them by service name or package name.

2. Click the Deployment Details tab to view the deployment details.

   The Deployment Details page displays the deployment information about all hosts, including the IP address, patch package name, version, product, cluster, service, service role, deployment progress, deployment status, associated build ID, deployment time, and log details.

   If too many deployment details exist, you can filter them by service name, package name, or deployment status.

Roll back an upgrade patch

After an upgrade patch is deployed, you can roll back the cluster to the version before the deployment if the cluster runs abnormally or encounters other problems.

1. Click Roll Back in the Actions column of the patch package to be rolled back.

   > **Note:**
   > A patch package can be rolled back only when the deployment status is Deployed.

2. In the dialog box that appears, set Cluster to the cluster where the patch package is deployed, and then click OK.

   Refresh the page in the rollback process. The deployment status is changed to Rolling Back. Rollback takes some time. Wait until the patch package is rolled back.

   Refresh the page after the rollback is complete. The deployment status is changed to Rolled Back.

1.10.4 Hot upgrade

Apsara Bigdata Manager (ABM) allows you to upgrade monitoring configuration and items without interrupting the service. On the Hot Upgrades page, you can view

the hot upgrade history and upgrade logs. You can also delete the upgrade packages and upgrade history on this page.

Prerequisites

- **Your ABM account is granted the required permissions to perform O&M operations on ABM.**
- **The monitoring item upgrade package in the `tar.gz` format is obtained.**

Upgrade a monitoring item without interrupting the service

1. *Log on to the ABM console*.
2. **Click Management in the upper-right corner. On the page that appears, click Hot Upgrades in the left-side navigation pane.**
3. **Click Upload File, and then select and upload the obtained tar.gz file.**

   The upload logs are displayed in the Upload Log section of the page in the upload process. After the upload is complete, the page displays the upgrade items for this upgrade package.
4. **Select the monitoring items to be upgraded, and then click OK.**
5. **In the dialog box that appears, click OK to start the upgrade.**

   After the upgrade is complete, the system prompts that the upgrade is successful .

View the hot upgrade history and logs

After the hot upgrade is complete, a hot upgrade record is generated on the File Management page, including the creation time and ID of the record, and the storage address of the upgrade package. When the hot upgrade fails, you can view the hot upgrade logs to locate the fault.

1. **Click the File Management tab on the Hot Upgrades page to view the hot upgrade history.**
2. **Click View Logs in the Actions column of an upgrade record to view the upgrade logs for each monitoring item in this hot upgrade process.**

Delete a hot upgrade record

ABM allows you to delete hot upgrade records, together with the corresponding hot upgrade packages and hot upgrade logs.

1. **Click the File Management tab on the Hot Upgrades page to view the hot upgrade history.**
2. **Click Delete in the Actions column of an upgrade record. In the dialog box that appears, click OK.**

1.10.5 Health management

Apsara Bigdata 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.
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.

1.10.6 Operation auditing

This feature allows you to view the O&M operations of the current service of Apsara Bigdata 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.

1.11 Go to other platforms

Apsara Bigdata Manager (ABM) provides the links to Apsara Stack Operation, Apsara Infrastructure Management Framework, and Apsara Stack Security to facilitate the O&M of big data services.

Prerequisites

An ABM account that works properly and the corresponding password are obtained.

Procedure

1. *Log on to the ABM console*.

2. On the homepage of ABM, click in the upper-left corner, and then click ASO, TIANJI, or YUNDUN in the Site Navigation section. The corresponding platform appears.

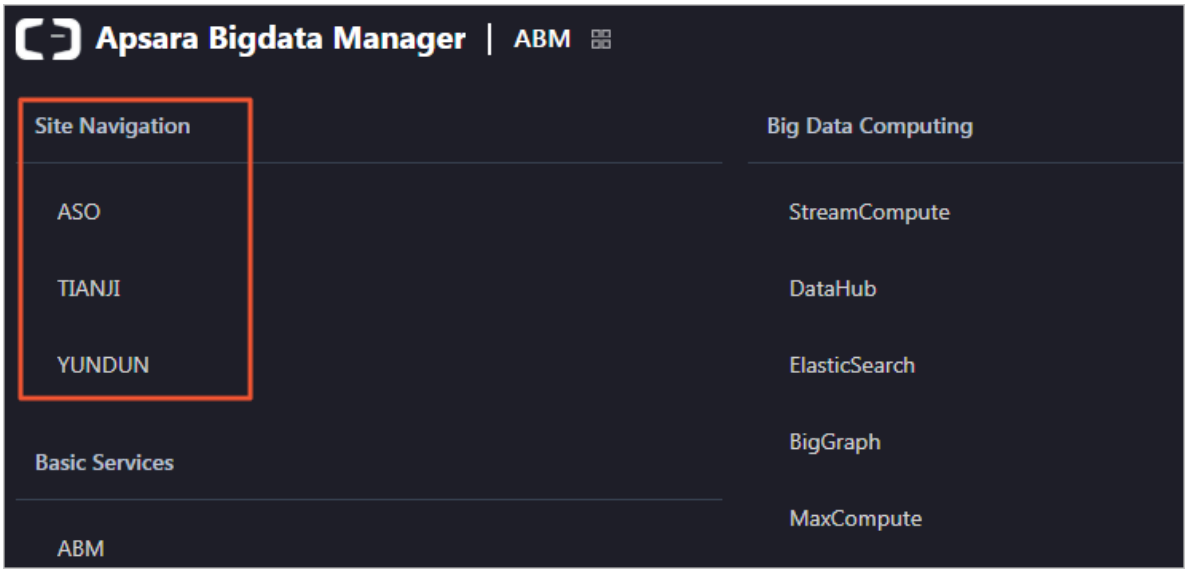

Result

After clicking ASO or TIANJI, you can log on to the corresponding platform without entering the username or password.

After clicking YUNDUN, however, you need to enter your username and password to log on to the platform.

2 MaxCompute

2.1 Concepts and architecture

Figure 2-1: MaxCompute architecture **shows the MaxCompute architecture.**

Figure 2-1: MaxCompute architecture

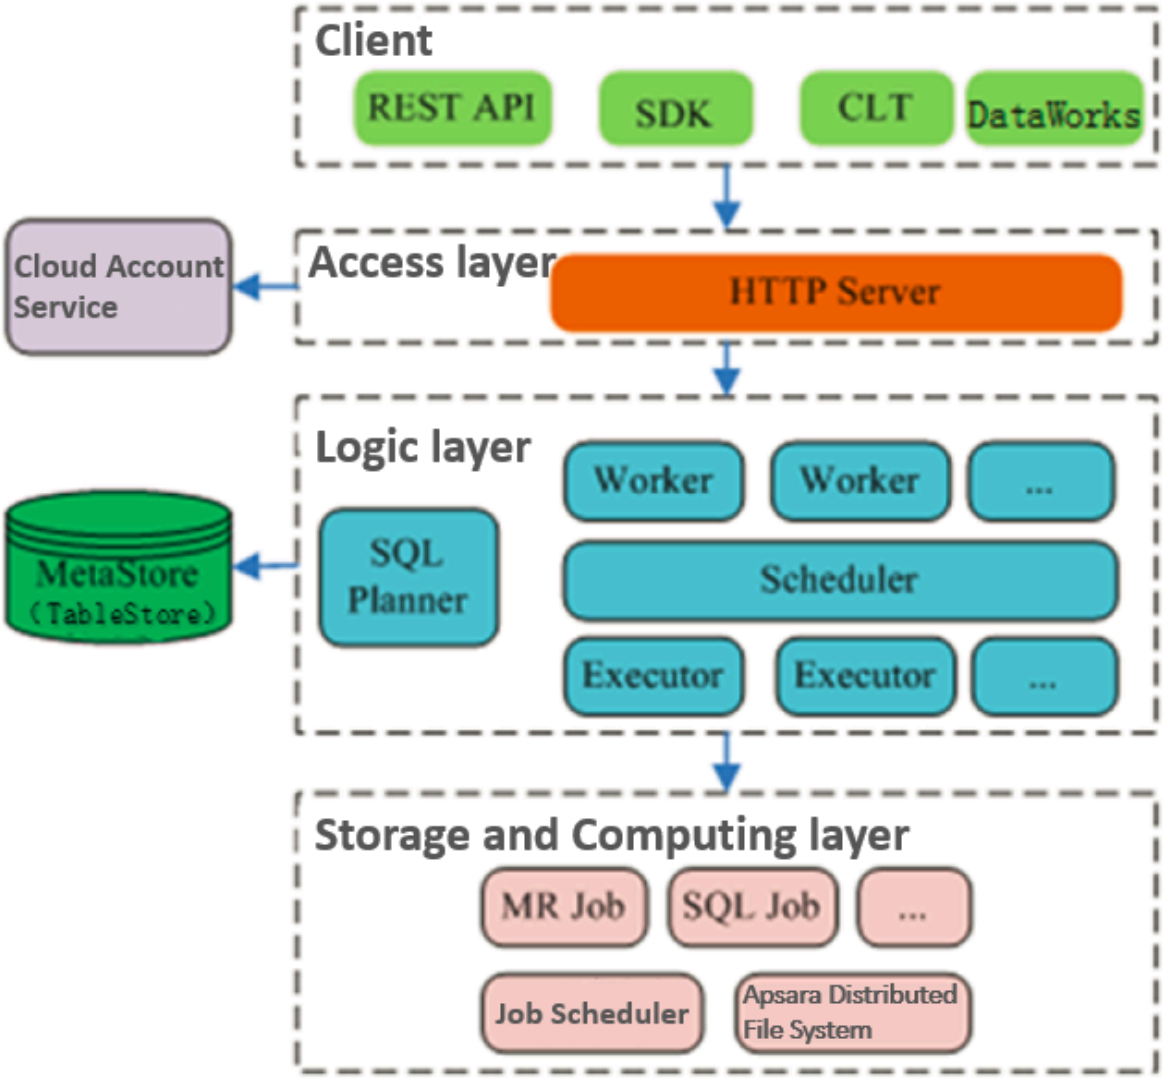

The MaxCompute service is divided into four parts: client, access layer, logic layer, and storage and computing layer. Each layer can be horizontally scaled.

The following methods can be used to implement the functions of a MaxCompute client:

- **API: RESTful APIs are used to provide offline data processing services.**
- **SDK: RESTful APIs are encapsulated within SDKs. SDKs are currently available in programming languages such as Java.**
- **Command line tool (CLT): This client-side tool runs on Windows and Linux. CLT allows you to submit commands to manage projects and use DDL and DML.**
- **DataWorks: DataWorks provides upper-layer visual ETL and BI tools that allow you to synchronize data, schedule tasks, and create reports.**

The access layer of MaxCompute supports HTTP, HTTPS, load balancing, user authentication, and service-level access control.

The logic layer is at the core of MaxCompute and supports project and object management, command parsing and execution logic, and data object access control and authorization. The logic layer contains two clusters: control and compute clusters. The control cluster is designed to manage projects and objects, parse and start queries and commands, and control and authorize access to data objects. The compute cluster executes tasks. Both control and compute clusters can be horizontally scaled as needed. The control cluster has three roles: Worker, Scheduler, and Executor. These roles are described as follows:

- **The Worker role processes all RESTful requests and manages projects, resources, and jobs. Workers forward jobs that need to launch Fuxi tasks (such as SQL, MapReduce, and Graph jobs) to the Scheduler for further processing.**
- **The Scheduler role schedules instances, splits instances into multiple tasks, sorts tasks that are pending for submission, and queries resource usage from FuxiMaster in the compute cluster for throttling. If there are no idle slots in Job Scheduler, the Scheduler stops processing task requests from Executors.**
- **The Executor role is responsible for launching SQL and MapReduce tasks. Executors submit Fuxi tasks to FuxiMaster in the compute cluster and monitor the operating status of these tasks.**

When you submit a job request, the web server at the access layer queries the IP addresses of registered Workers and sends API requests to randomly selected Workers. The Workers then send these requests to the Scheduler for scheduling and throttling. Executors actively poll the Scheduler queue. If the necessary resources are available, the Executors start executing tasks and return the task execution

status to the Scheduler. The following figure shows the MaxCompute job execution process.

Figure 2-2: MaxCompute job execution process

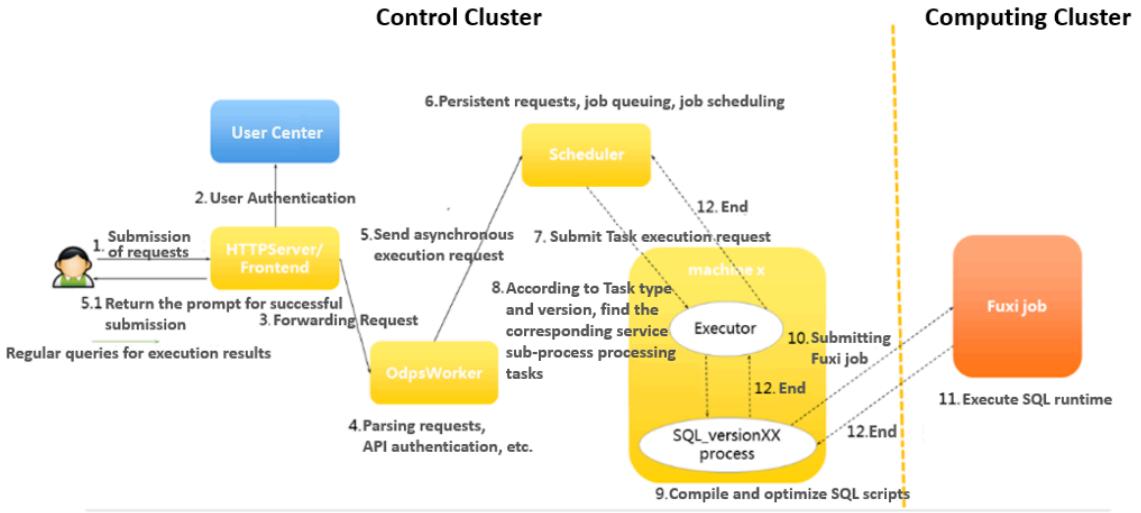

The following concepts are involved in the MaxCompute job execution process:

1. **MaxCompute instance: the instance of a MaxCompute job. A job is anonymous if it is not defined. A MaxCompute job can contain multiple MaxCompute tasks. In a MaxCompute instance, you can submit multiple SQL or MapReduce tasks, and specify whether to run the tasks in parallel or serial mode. This scenario is rarely seen because MaxCompute jobs are not commonly used. In most cases, an instance contains only one task.**
2. **MaxCompute task: a specific task in MaxCompute. Currently, there are almost 20 task types, such as SQL, MapReduce, Admin, Lot, and Xlib. The execution logic varies greatly depending on the task type. Different tasks in an instance are differentiated by their task name. MaxCompute tasks can run in the control cluster. Simple tasks such as metadata modification can run in the control cluster for their entire lifecycles. To run computing tasks, submit Fuxi jobs to the compute cluster.**

3. Fuxi job: a computing model provided by the Job Scheduler module. A Fuxi job corresponds to a Fuxi service. A Fuxi job represents a task that can be completed, while a Fuxi service represents a resident process.
   - The DAG scheduling approach can be used to schedule Fuxi jobs. Each job has a job master to schedule its job resources.
   - For SQL, Fuxi jobs are divided into offline and online jobs. Online jobs evolve from the service mode jobs. An online job is also called a quasi-real-time task . An online job is a resident process that can be executed whenever there are tasks, reducing the time required to start and stop a job.
   - You can submit a MaxCompute task to multiple compute clusters. The primary key name of a Fuxi job is the cluster name followed by the job name.
   - The JSON plan for Job Scheduler to submit a job and the status of a finished job are stored in Apsara Distributed File System.
4. Fuxi task: a sub-concept of Fuxi job. Similar to MaxCompute tasks, different Fuxi tasks represent different execution logics. Fuxi tasks can be linked together as pipes to implement complex logic.
5. Fuxi instance: the instance of a Fuxi task. A Fuxi instance is the smallest unit that can be scheduled by Job Scheduler. During the actual execution process, a task is divided into many logical units to improve the processing speed. Different instances will run on the same execution logic but work with different input and output data.
6. Fuxi worker: an underlying concept of Job Scheduler. A worker represents an operating system process. A worker can be reused by multiple Fuxi instances, but a worker can only handle one instance at a time.

 **Note:**

- InstanceID: the unique identifier of a MaxCompute job. It is commonly used for troubleshooting. You can construct the LogView of the current instance based on the project name and instance ID.
- Service master or job master: a primary node of the service or job type. The primary node is responsible for requesting and scheduling resources, creating work plans for workers, and monitoring workers across their entire lifecycles.

The storage and computing layer of MaxCompute is a core component of the proprietary cloud computing platform of Alibaba Cloud. As the kernel of the Apsara

system, this component runs in the compute cluster independent of the control cluster. The architecture diagram illustrates only the major modules.

2.2 O&M commands and tools

2.2.1 Before you start

Before using MaxCompute O&M commands and tools, you must be aware of the following information:

During the MaxCompute O&M process, the default account is admin. You must run all commands as an admin user. You must use your admin account and sudo to run commands that require sudo privileges.

2.2.2 odpscmd commands

You can use the command line to perform operations and maintenance. You must log on to the command line tool before you can run commands. The specific procedure is as follows:

1. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose Operations > Cluster Operations. In the Cluster search box, enter odps to search for the expected cluster.
2. Click the cluster in the search result. On the Cluster Details page, click the Services tab. In the Services search box, search for odps-service-computer. Click odps-service-computer in the search result.
3. After you access the odps-service-computer service, select ComputerInit# on the Service Details page. In the `Actions` column corresponding to the machine, click Terminal. In the TerminalService window that appears, you can perform subsequent command line operations.

Console command directories and configurations

The MaxCompute client is located in the clt folder under the */apsara/odps_tools* directory of odpsag. The client configuration file is located in the conf directory under the clt folder. The access_id, access_key, end_point, log_view, and tunnel_point parameters are configured by default. You can use the `./clt/bin/odpscmd` command to view information such as the version number in interactive mode. For example, run the `HTTP GET /projects/admin_task_project/system;` command to check the version information of MaxCompute.

Description of client command options

The following figure shows the client command options.

Figure 2-3: Client command options

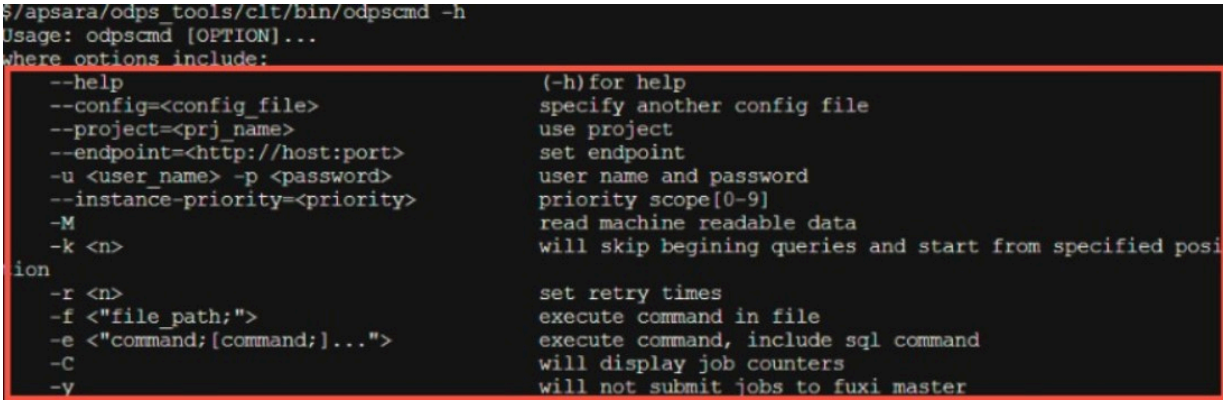

- **-e: The MaxCompute client does not execute SQL statements in interactive mode.**
- **--project, -u, and -p: The client directly uses the specified values for the project, user, and pass parameters. If you do not specify a parameter, the client uses the corresponding value configured in the conf file.**
- **-k and -f: The client directly executes local SQL files.**
- **--instance-priority: This option is used to assign a priority to the current task. Valid values: 0 to 9. A lower value indicates a higher priority.**
- **-r: This option indicates the number of times a failed command will be retried. It is commonly used in scripting jobs.**

Commonly used SQL commands for O&M

The following table lists the commonly used commands.

Table 2-1: Commonly used commands

Command	Description
whoami;	**Allows you to view your Apsara Stack tenant account and endpoint information.**
show p;	**Allows you to view information about all instances that have been run.**

Command	Description
wait <instanceid>;	Allows you to re-generate the LogView and Fuxi job information of a task. To run this command, you must have owner permissions, and the LogView and Fuxi job information must be stored in the same project.
kill <instanceid>;	Allows you to terminate specified instances.
tunnel upload/download;	Allows you to test whether Tunnel is functioning.
desc project <projectname> -extended;	Allows you to view the project usage. · desc extended table: allows you to view table information. · desc table_name partition(pt_spec): allows you to view partition information. · desc resource $resource_name: allows you to view project resource information. · desc project $project_name -extended: allows you to view cluster information.
export <project name> local_file_path;	Allows you to export DDL statements of all tables in a project.
create table tablename (...) ;	Allows you to create a table.
select count(*) from tablename;	Allows you to search for a table.
Explain	Allows you to create plans without submitting Fuxi jobs to view resources required for tasks.
list	Allows you to list tables, resources, and roles.
show	Allows you to view table and partition information.
purge	Allows you to remove all data from the MaxCompute recycle bin directly to the Apsara Distributed File System recycle bin. · purge table <tablename>: allows you to purge a single table. · purge all: allows you to purge all tables from the current project.

2.2.3 Tunnel commands

The client provides Tunnel commands that implement the original functions of the Dship tool. Tunnel commands are mainly used to upload or download data.

Table 2-2: Tunnel commands

Command	Description
tunnel upload	Allows you to upload data to MaxCompute tables. You can upload files or level-1 directories. Data can only be uploaded to a single table or table partition each time. The destination partition must be specified for partitioned tables.
tunnel download	Allows you to download data from MaxCompute tables. You can only download data to a single file. Only data in one table or partition can be downloaded to one file each time. For partitioned tables, the source partition must be specified.
tunnel resume	If an error occurs because of network or Tunnel service faults, you can resume file or directory transmission after interruption. This command only allows you to resume the previous data upload. Every data upload or download operation is called a session. Run the resume command and specify the ID of the session to be resumed.
tunnel show	Allows you to view historical task information.
tunnel purge	Purges the session directory. Sessions from the last three days are purged by default.

Tunnel commands allow you to view help information by using the Help sub-command on the client. The sub-commands of each Tunnel command are described as follows:

Upload

Imports data of a local file into a MaxCompute table. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help upload;
usage: tunnel upload [options] <path> <[project.]table[/partition]>
                upload data from local file
 -acp,-auto-create-partition <ARG>   auto create target partition if
not
                                     exists, default false
 -bs,-block-size <ARG>               block size in MiB, default 100
```

```
 -c,-charset <ARG>                       specify file charset, default
ignore.
                                         set ignore to download raw data
 -cp,-compress <ARG>                     compress, default true
 -dbr,-discard-bad-records <ARG>         specify discard bad records
                                         action(true|false), default false
 -dfp,-date-format-pattern <ARG>         specify date format pattern,
default
                                         yyyy-MM-dd HH:mm:ss
 -fd,-field-delimiter <ARG>              specify field delimiter, support
                                         unicode, eg \u0001. default ","
 -h,-header <ARG>                        if local file should have table
                                         header, default false
 -mbr,-max-bad-records <ARG>             max bad records, default 1000
 -ni,-null-indicator <ARG>               specify null indicator string,
                                         default ""(empty string)
 -rd,-record-delimiter <ARG>             specify record delimiter, support
                                         unicode, eg \u0001. default "\r\n
"
 -s,-scan <ARG>                          specify scan file
                                         action(true|false|only), default
true
 -sd,-session-dir <ARG>                  set session dir, default
                                         D:\software\odpscmd_public\
plugins\ds
                                         hip
 -ss,-strict-schema <ARG>                specify strict schema mode. If
false,
                                         extra data will be abandoned and
                                         insufficient field will be filled
                                         with null. Default true
 -te,-tunnel_endpoint <ARG>              tunnel endpoint
    -threads <ARG>                       number of threads, default 1
 -tz,-time-zone <ARG>                    time zone, default local timezone
:
                                         Asia/Shanghai
Example:
    tunnel upload log.txt test_project.test_table/p1="b1",p2="b2"
```

Parameters:

- **-acp: indicates whether to automatically create the destination partition if it does not exist. No destination partition is created by default.**
- **-bs: specifies the size of each data block uploaded with Tunnel. Default value: 100 MiB (MiB = 1024 * 1024B).**
- **-c: specifies the local data file encoding format. Default value: UTF-8. If this parameter is not set, the encoding format of the downloaded source data is used by default.**
- **-cp: indicates whether to compress the local data file before it is uploaded to reduce network traffic. By default, the local data file is compressed before it is uploaded.**

- **-dbr: indicates whether to ignore dirty data (such as additional columns, missing columns, and columns with mismatched data types).**
  - **If this parameter is set to true, all data that does not comply with table definitions is ignored.**
  - **If this parameter is set to false, an error is returned when dirty data is found, so that raw data in the destination table is not contaminated.**
- **-dfp: specifies the DateTime format. Default value: yyyy-MM-dd HH:mm:ss.**
- **-fd: specifies the column delimiter used in the local data file. Default value: comma (,).**
- **-h: indicates whether the data file contains the header. If this parameter is set to true, Dship skips the header row and starts uploading data from the second row.**
- **-mbr: terminates any attempts to upload more than 1,000 rows of dirty data. This parameter allows you to adjust the maximum allowable volume of dirty data.**
- **-ni: specifies the NULL data identifier. Default value: an empty string ("").**
- **-rd: specifies the row delimiter used in the local data file. Default value: \r\n.**
- **-s: indicates whether to scan the local data file. Default value: false.**
  - **If this parameter is set to true, the system scans the source data first, and then imports the data if the format is correct.**
  - **If this parameter is set to false, the system imports data directly without scanning.**
  - **If this parameter is set to only, the system only scans the source data, and does not import the data after scanning.**
- **-sd: sets the session directory.**
- **-te: specifies the Tunnel endpoint.**
- **-threads: specifies the number of threads. Default value: 1.**
- **-tz: specifies the time zone. Default value: Asia/Shanghai.**

Show

Displays historical records. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help show;
usage: tunnel show history [options]
               show session information
 -n,-number <ARG>   lines
Example:
    tunnel show history -n 5
```

```
    tunnel show log
```

Parameters:

-n: specifies the number of rows to be displayed.

Resume

Resumes the execution of historical operations (only applicable to data upload).

The following example shows how to use the sub-commands:

```
odps@  project_name>tunnel help resume;
usage: tunnel resume [session_id] [-force]
               resume an upload session
 -f,-force   force resume
Example:
    tunnel resume
```

Download

The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help download;
usage: tunnel download [options] <[project.]table[/partition]> <path>
               download data to local file
 -c,-charset <ARG>                  specify file charset, default
ignore.
                                    set ignore to download raw data
 -ci,-columns-index <ARG>           specify the columns index(starts
from
                                    0) to download, use comma to split
each
                                    index
 -cn,-columns-name <ARG>            specify the columns name to
download,
                                    use comma to split each name
 -cp,-compress <ARG>                compress, default true
 -dfp,-date-format-pattern <ARG>    specify date format pattern,
default
                                    yyyy-MM-dd HH:mm:ss
 -e,-exponential <ARG>              When download double values, use
                                    exponential express if necessary.
                                    Otherwise at most 20 digits will be
                                    reserved. Default false
 -fd,-field-delimiter <ARG>         specify field delimiter, support
                                    unicode, eg \u0001. default ","
 -h,-header <ARG>                   if local file should have table
header,
                                    default false
    -limit <ARG>                    specify the number of records to
                                    download
 -ni,-null-indicator <ARG>          specify null indicator string,
default
                                    ""(empty string)
 -rd,-record-delimiter <ARG>        specify record delimiter, support
                                    unicode, eg \u0001. default "\r\n"
 -sd,-session-dir <ARG>             set session dir, default
                                    D:\software\odpscmd_public\plugins\
dshi
                                    p
```

```
 -te,-tunnel_endpoint <ARG>         tunnel endpoint
    -threads <ARG>                  number of threads, default 1
 -tz,-time-zone <ARG>               time zone, default local timezone:
                                    Asia/Shanghai
usage: tunnel download [options] instance://<[project/]instance_id> <
path>
              download instance result to local file
 -c,-charset <ARG>                  specify file charset, default
ignore.
                                    set ignore to download raw data
 -ci,-columns-index <ARG>           specify the columns index(starts
from
                                    0) to download, use comma to split
each
                                    index
 -cn,-columns-name <ARG>            specify the columns name to
download,
                                    use comma to split each name
 -cp,-compress <ARG>                compress, default true
 -dfp,-date-format-pattern <ARG>    specify date format pattern,
default
                                    yyyy-MM-dd HH:mm:ss
 -e,-exponential <ARG>              When download double values, use
                                    exponential express if necessary.
                                    Otherwise at most 20 digits will be
                                    reserved. Default false
 -fd,-field-delimiter <ARG>         specify field delimiter, support
                                    unicode, eg \u0001. default ","
 -h,-header <ARG>                   if local file should have table
header,
                                    default false
    -limit <ARG>                    specify the number of records to
                                    download
 -ni,-null-indicator <ARG> specify null indicator string, default
                                    ""(empty string)
 -rd,-record-delimiter <ARG>        specify record delimiter, support
                                    unicode, eg \u0001. default "\r\n"
 -sd,-session-dir <ARG>             set session dir, default
                                    D:\software\odpscmd_public\plugins\
dshi
                                    p
 -te,-tunnel_endpoint <ARG>         tunnel endpoint
    -threads <ARG>                  number of threads, default 1
 -tz,-time-zone <ARG>               time zone, default local timezone:
                                    Asia/Shanghai
Example:
    tunnel download test_project.test_table/p1="b1",p2="b2" log.txt
    tunnel download instance://test_project/test_instance log.txt
```

Parameters:

- **-c: specifies the local data file encoding format. Default value: UTF-8.**
- **-ci: specifies the column index (starting from 0) for downloading. Separate multiple entries with commas (,).**
- **-cn: specifies the names of columns to be downloaded. Separate multiple entries with commas (,).**

- **-cp, -compress: indicates whether to compress the data file before it is uploaded to reduce network traffic. By default, a data file is compressed by it is uploaded.**
- **-dfp: specifies the DateTime format. Default value: yyyy-MM-dd HH:mm:ss.**
- **-e: allows you to express the values as exponential functions when you download Double type data. If this parameter is not set, a maximum of 20 digits can be retained.**
- **-fd: specifies the column delimiter used in the local data file. Default value: comma (,).**
- **-h: indicates whether the data file contains a header. If this parameter is set to true, Dship skips the header row and starts downloading data from the second row.**

  

  Note:

  `-h=true` **and** `threads>1` **cannot be used together.**
- **-limit: specifies the number of files to be downloaded.**
- **-ni: specifies the NULL data identifier. Default value: an empty string ("").**
- **-rd: specifies the row delimiter used in the local data file. Default value: \r\n.**
- **-sd: sets the session directory.**
- **-te: specifies the Tunnel endpoint.**
- **-threads: specifies the number of threads. Default value: 1.**
- **-tz: specifies the time zone. Default value: Asia/Shanghai.**

Purge

Purges the session directory. Sessions from the last three days are purged by default. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help purge;
usage: tunnel purge [n]
               force session history to be purged.([n] days before,
default
               3 days)
Example:
```

```
tunnel purge 5
```

2.2.4 LogView tool

2.2.4.1 Before you start

You must confirm the LogView process status before using LogView. If the process status is off, you must start the LogView process.

The procedure for querying the process status and starting the process is as follows:

1. **Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose Operations > Cluster Operations. In the Cluster search box, enter odps to search for the expected cluster.**
2. **Click the cluster in the search result. On the Cluster Details page, click the Services tab. In the Service search box, search for odps-service-console. Click odps-service-console in the search result.**
3. **After you access the odps-service-console service, select LogView# on the Service Details page. In the `Actions` column corresponding to the machine, click Terminal to open the TerminalService window.**
4. **Run the following command to find the Docker container where LogView resides:**

   ```
   docker ps|grep logview
   ```

5. **Run the following commands to view the LogView process status:**

   ```
   ps -aux|grep logview
   ```

   ```
   netstat -ntulp|grep 9000
   ```

6. **If the process status is off, run the following command to start the process:**

   ```
   /opt/aliyun/app/logview/bin/control start
   ```

The following sections describe what is LogView and how to use LogView to perform basic operations.

2.2.4.2 LogView introduction

LogView is a tool for checking and debugging a job submitted to MaxCompute. LogView allows you to check the running details of a job.

LogView functions

LogView allows you to check the running status, details, and results of a job, and the progress of each phase.

LogView endpoint

Take the odpscmd client as an example. After you submit an SQL task on the client, a long string starting with logview is returned.

Figure 2-4: A long string starting with logview

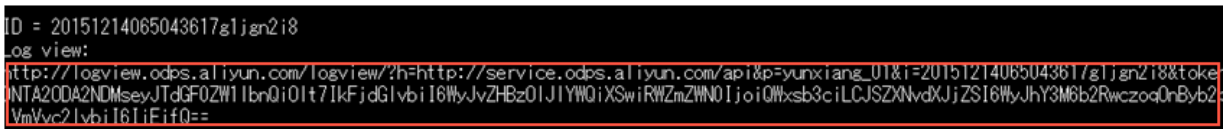

Enter the string with all carriage return and line feed characters removed in the address bar of the browser.

Composition of a LogView string

A LogView string consists of five parts, as shown in the following figure.

Figure 2-5: Composition of a LogView string

2.2.4.3 Preliminary knowledge of LogView

For complex SQL queries, you must have an in-depth knowledge of the relationships between MaxCompute tasks and Fuxi instances before you can understand LogView .

In short, a MaxCompute task consists of one or more Fuxi jobs. Each Fuxi job consists of one or more Fuxi tasks. Each Fuxi task consists of one or more Fuxi instances.

Figure 2-6: Relationships between MaxCompute tasks and Fuxi instances

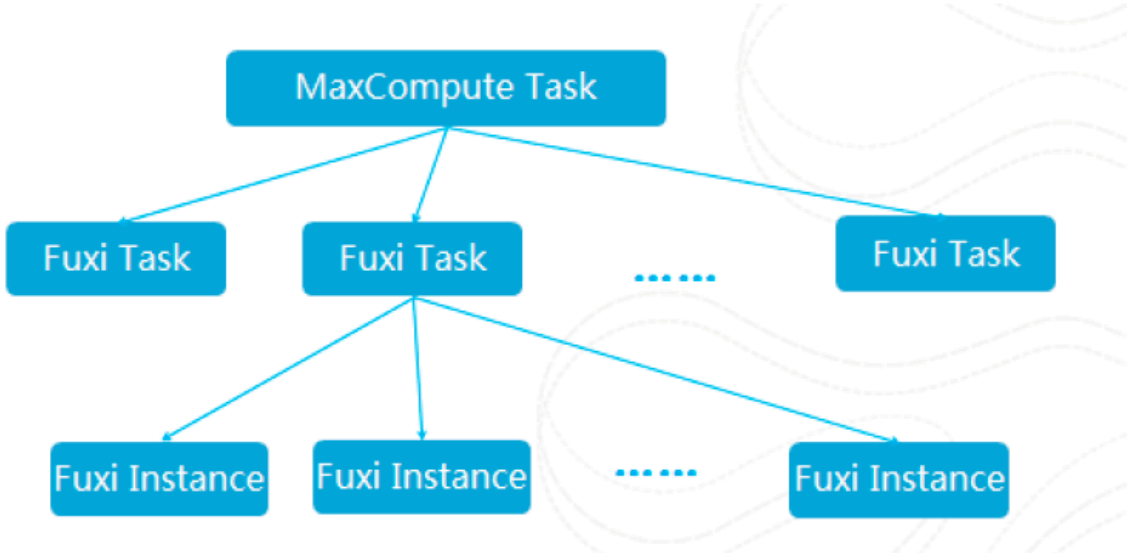

The following figures show the relevant information in LogView.

MaxCompute Instance

Figure 2-7: MaxCompute Instance

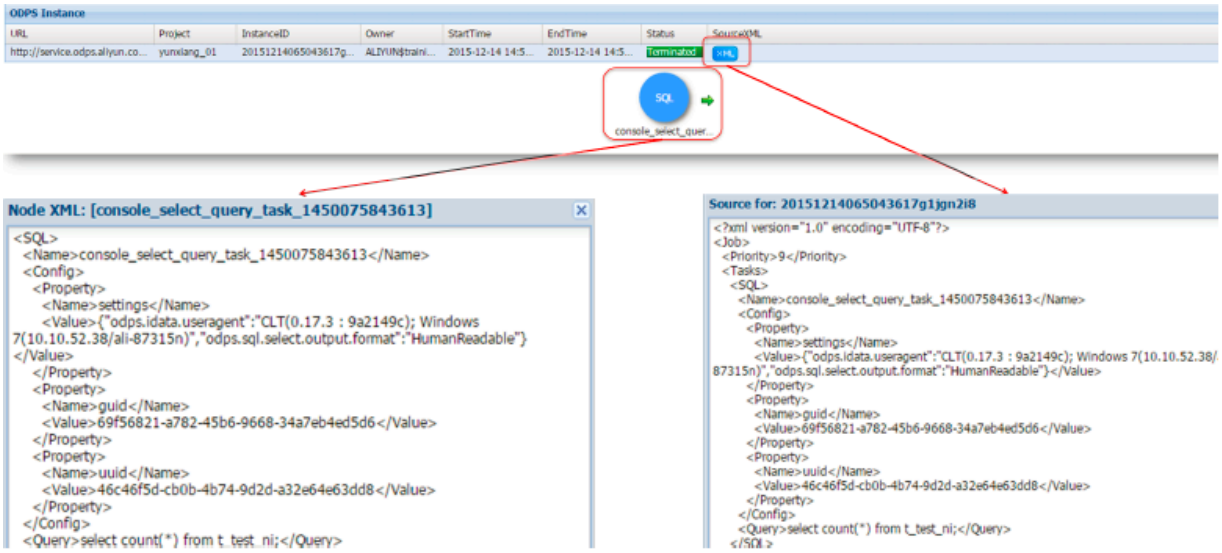

MaxCompute Task

Figure 2-8: MaxCompute Task

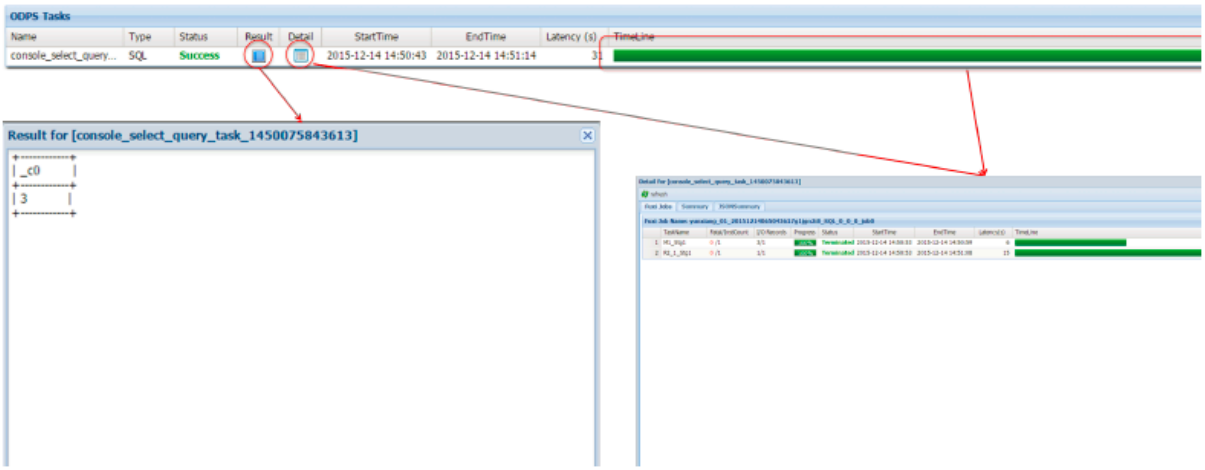

Task Detail - Fuxi Job

Figure 2-9: Task Detail - Fuxi Job(1)

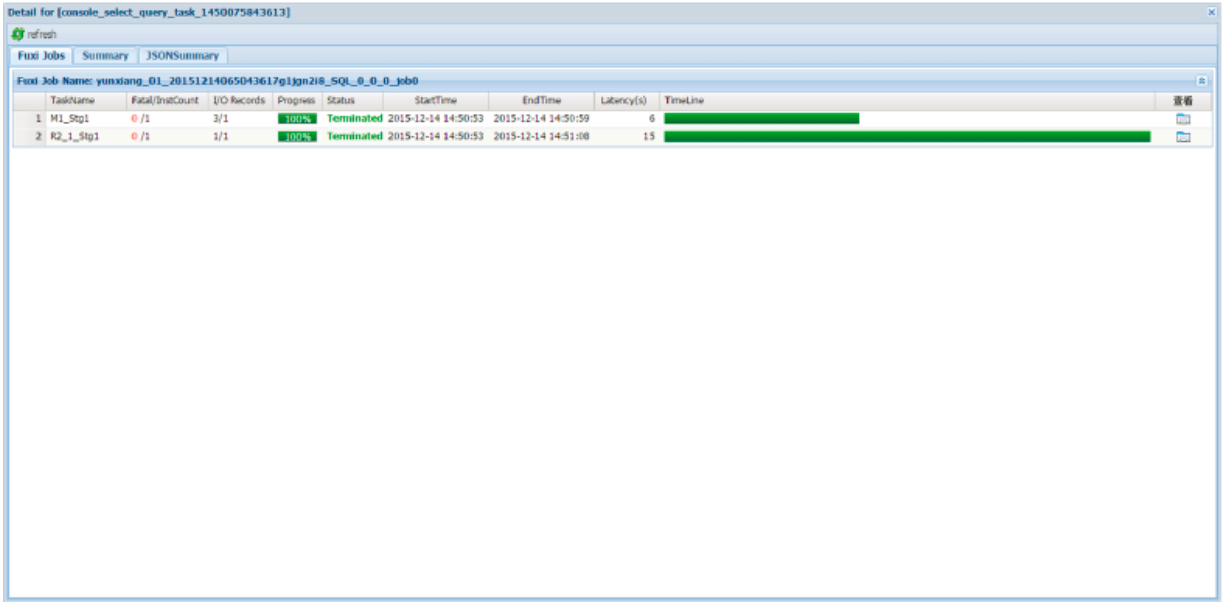

Figure 2-10: Task Detail - Fuxi Job(2)

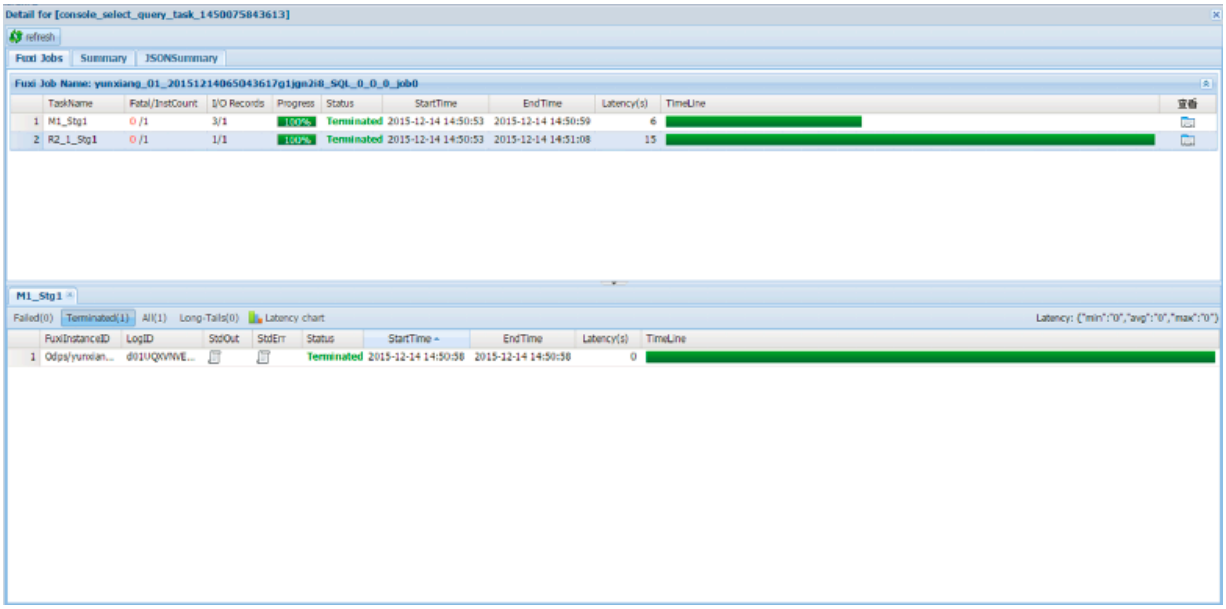

Task Detail - Summary

Figure 2-11: Task Detail - Summary

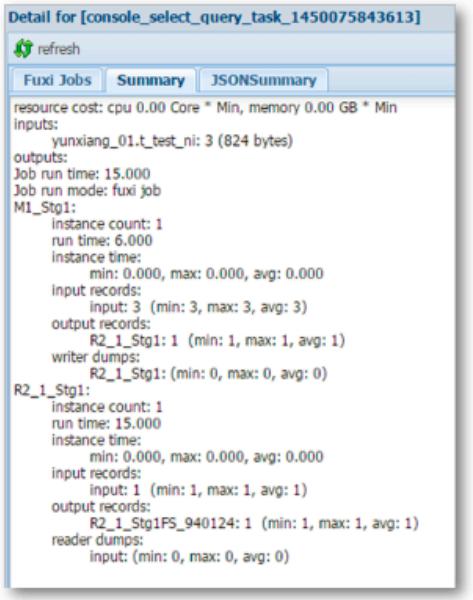

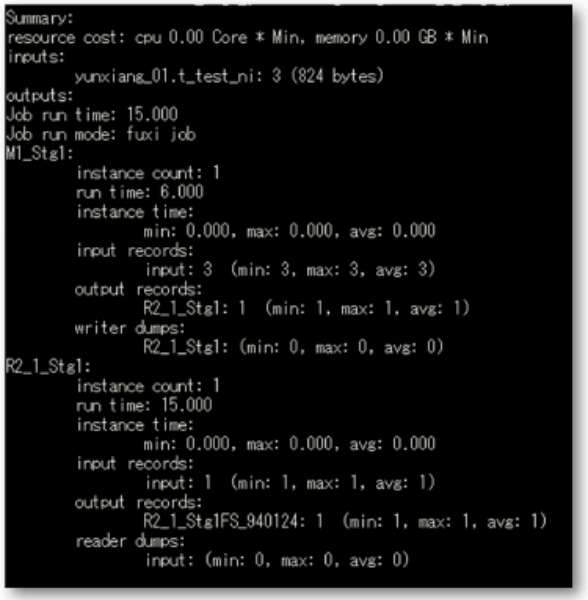

Task Detail - JSONSummary

Figure 2-12: Task Detail - JSONSummary

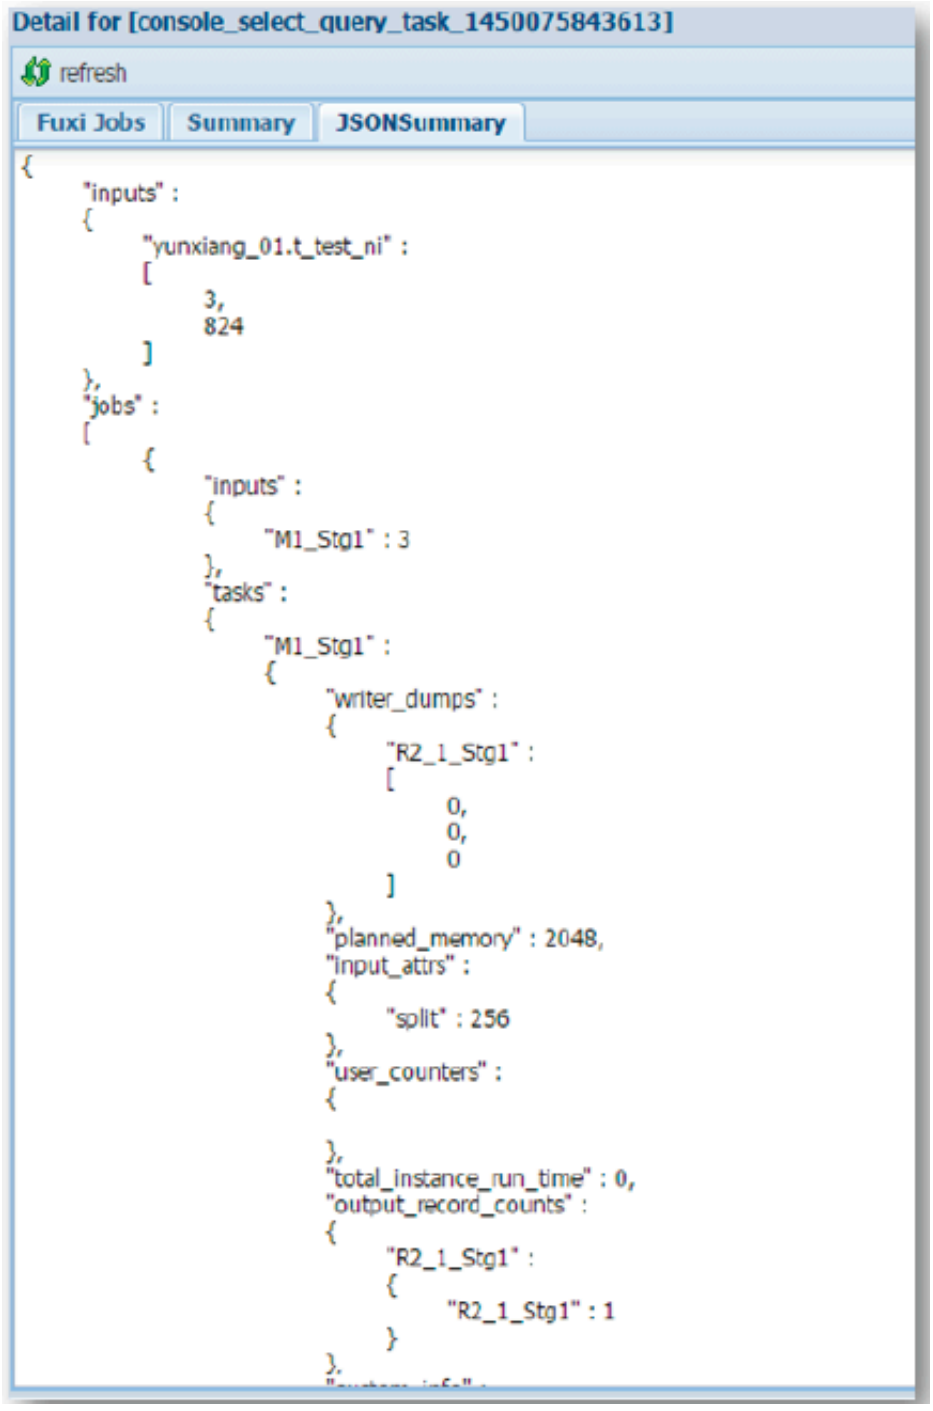

2.2.4.4 Basic operations and examples

View each point in time in the life cycle of a job.

Figure 2-13: View each point in time in the life cycle of a job

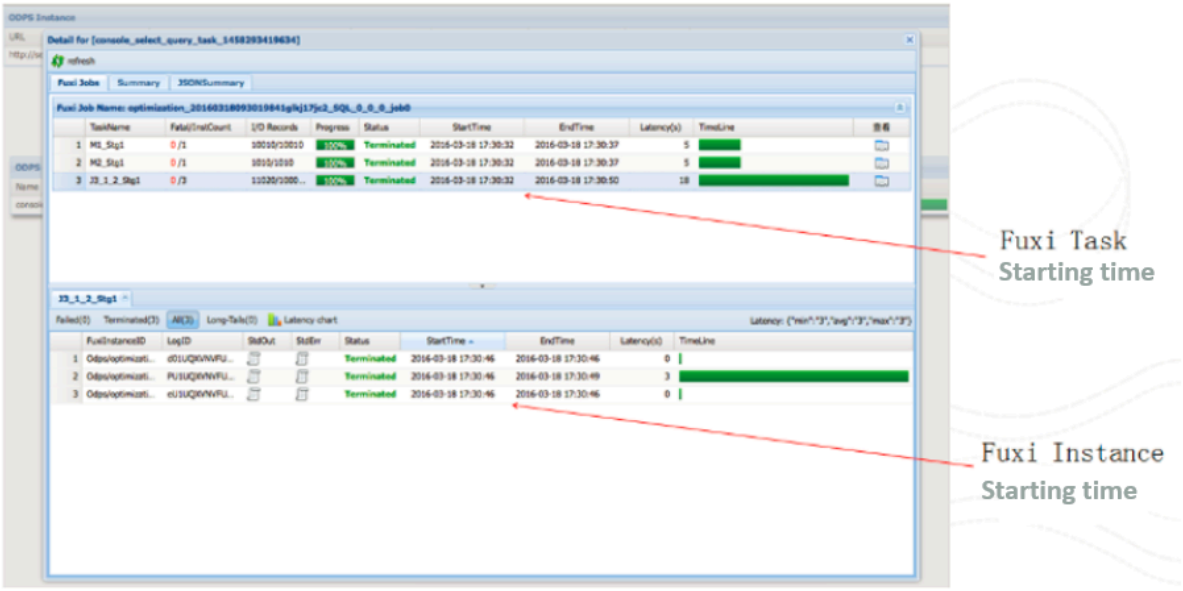

View the time it takes for Job Scheduler to schedule an instance.

Figure 2-14: View the time it takes for Job Scheduler to schedule an instance

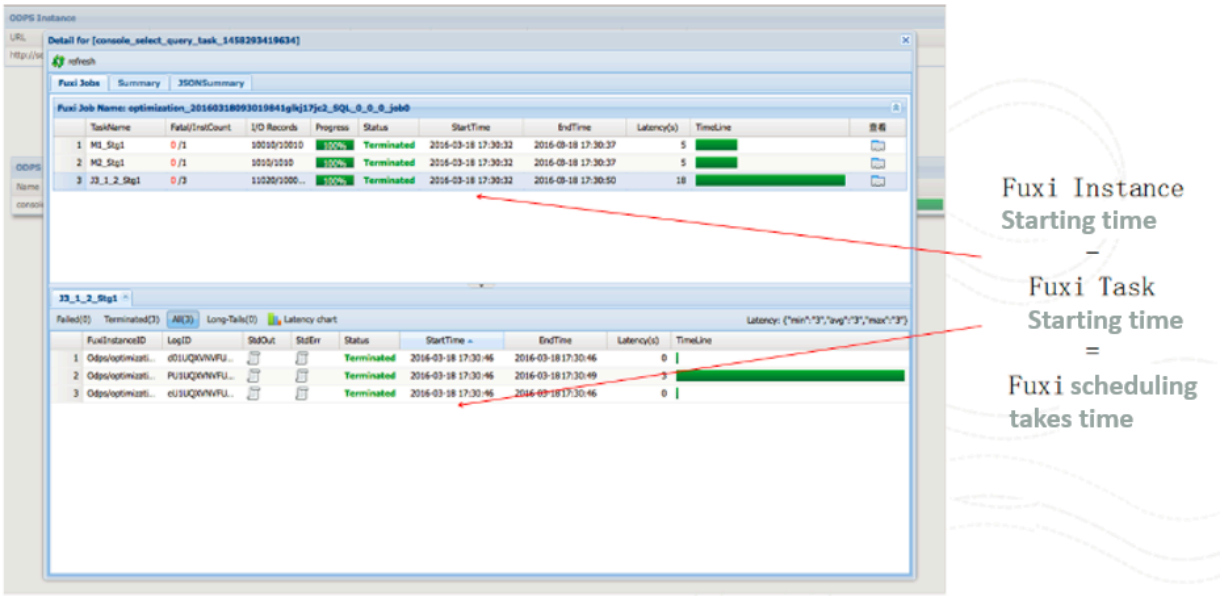

View the polling interval.

Figure 2-15: View the polling interval

```
odps@ optimization>select * from skew a join small b on a.key=b.key;

ID = 20160318092653630gstax6jc2
Log view:
http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=optimization&i=20160318092653630gs
tax6jc2&token=d05vbmZoWUpSRkhCVllzUHdGM3I1SEFoeEFVPSxPRFBTX09CTzoxMDExODIyNTI0ODIzNDU5LDE0NTg4OTgwMTMseyJTdGF0ZW1
lbnQiOlt7IkFjdGlvbiI6WyJvZHBzOlJlYWQiXSwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2RwczoqOn8yb2plY3RzL29wdGlt
aXphdGlvbi9pbnN0YW5jZXMvMjAxNjAzMTgwOTI2NTM2MzBnc3RheDZqYzIiXX1dLCJWZXJzaW9uIjoiMSJ9
2016-03-18 17:27:05 M1_Stg1_job0:0/0/1[0%]      M2_Stg1_job0:0/0/1[0%]  J3_1_2_Stg1_job0:0/0/3[0%]
2016-03-18 17:27:10 M1_Stg1_job0:0/1/1[100%]    M2_Stg1_job0:0/1/1[100%]        J3_1_2_Stg1_job0:0/0/3[0%]
2016-03-18 17:27:16 M1_Stg1_job0:0/1/1[100%]    M2_Stg1_job0:0/1/1[100%]        J3_1_2_Stg1_job0:0/0/3[0%]
Summary:
resource cost: cpu 0.02 Core * Min, memory 0.03 GB * Min
```

After a MaxCompute instance is submitted, odpscmd polls the execution status of the job at a specified interval of approximately 5s.

Check for data skews

Figure 2-16: Check for data skews

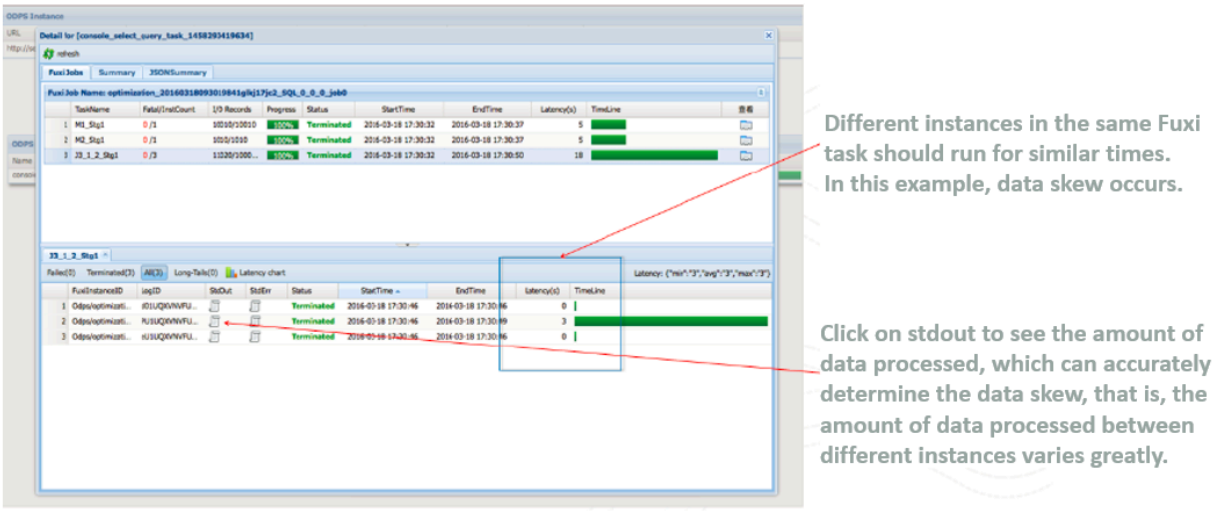

View the UDF and MR debugging information

Figure 2-17: View the UDF and MR debugging information

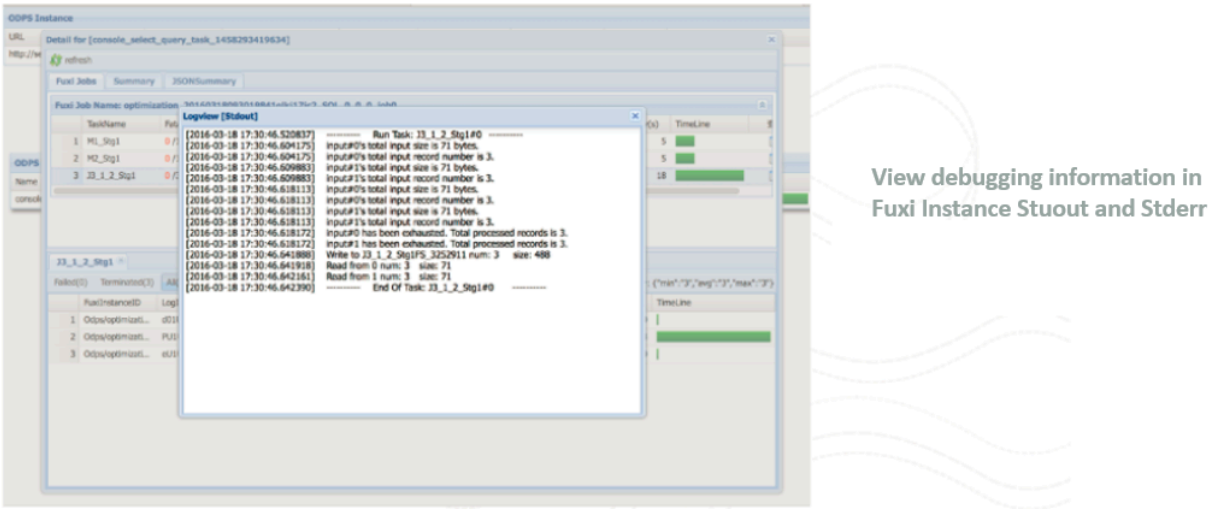

View the task status - Terminated

Figure 2-18: View the task status - Terminated

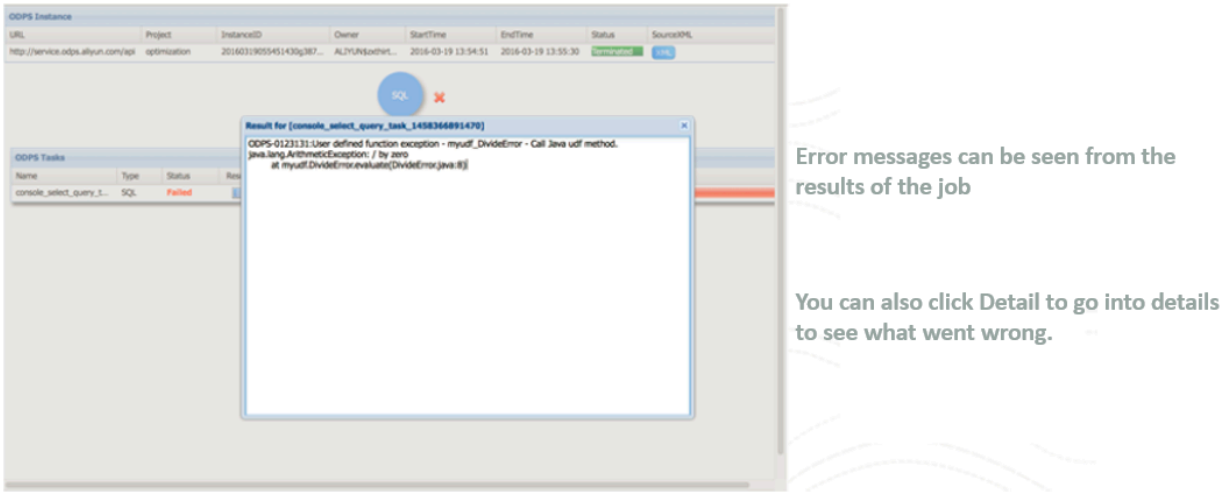

2.2.4.5 Best practices

Locate LogView based on the instance ID

After you submit a job, you can press Ctrl+C to return to odpscmd and perform other operations. You can run the `wait <instanceid>;` command to locate LogView and obtain the job status.

Figure 2-19: Locate LogView based on the instance ID

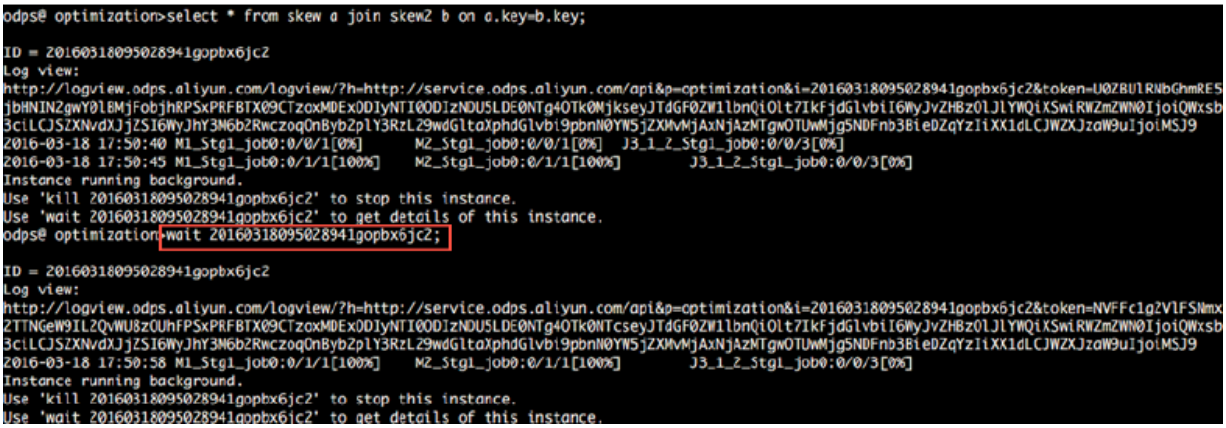

Locate running tasks

After you exit the control window, you can run the `show p;` command to locate currently running tasks and historical tasks.

Figure 2-20: Locate running tasks

2.2.5 Apsara Bigdata Manager

Apsara Bigdata Manager (ABM) supports O&M on big data products from the perspective of business, services, clusters, and hosts. You can also upgrade big data products, customize alert configurations, and view the O&M history in ABM.

On-site Apsara Stack engineers can use ABM to easily manage big data products through actions such as viewing resource usage, checking and handling alerts, and modifying configurations.

For more information about how to log on to Apsara Bigdata Manager, see related documentation.

2.3 Routine O&M

2.3.1 Configurations

MaxCompute configurations are stored in the `/apsara/odps_service/deploy/env.cfg` **directory in odpsag. The configuration file contains the following content:**

```
odps_worker_num=3
executor_worker_num=3
hiveserver_worker_num=3
replication_server_num=3
messager_partition_num=3
```

You can modify these parameter values based on your requirements and start the corresponding MaxCompute services based on the configured values. For more information, see *Restart a MaxCompute service***.**

If you add `xstream_max_worker_num=3` **at the end of the configuration file, XStream will be started with three running workers.**

2.3.2 Routine inspections

1. **On the Cluster Operations page in Apsara Infrastructure Management Framework, check whether all machines have reached the desired state.**
   a. **Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose Operations > Cluster Operations. In the Cluster search box, enter odps to search for the expected cluster.**
   b. **Based on the information in the** `Status`, `Machine Status`, **and** `Server Role Status` **columns, check whether all machines have reached the desired state. The following figure shows that some machines have not reached the desired state.**
   c. **Click the exceptions in the** `Machine Status` **and** `Server Role status` **columns to view the exception details.**

2. **Go to the** */home/admin/odps/odps_tools/clt/bin/odpscmd -e* **directory and run the following command:**

```
select count(*) from datahub_smoke_test;
```

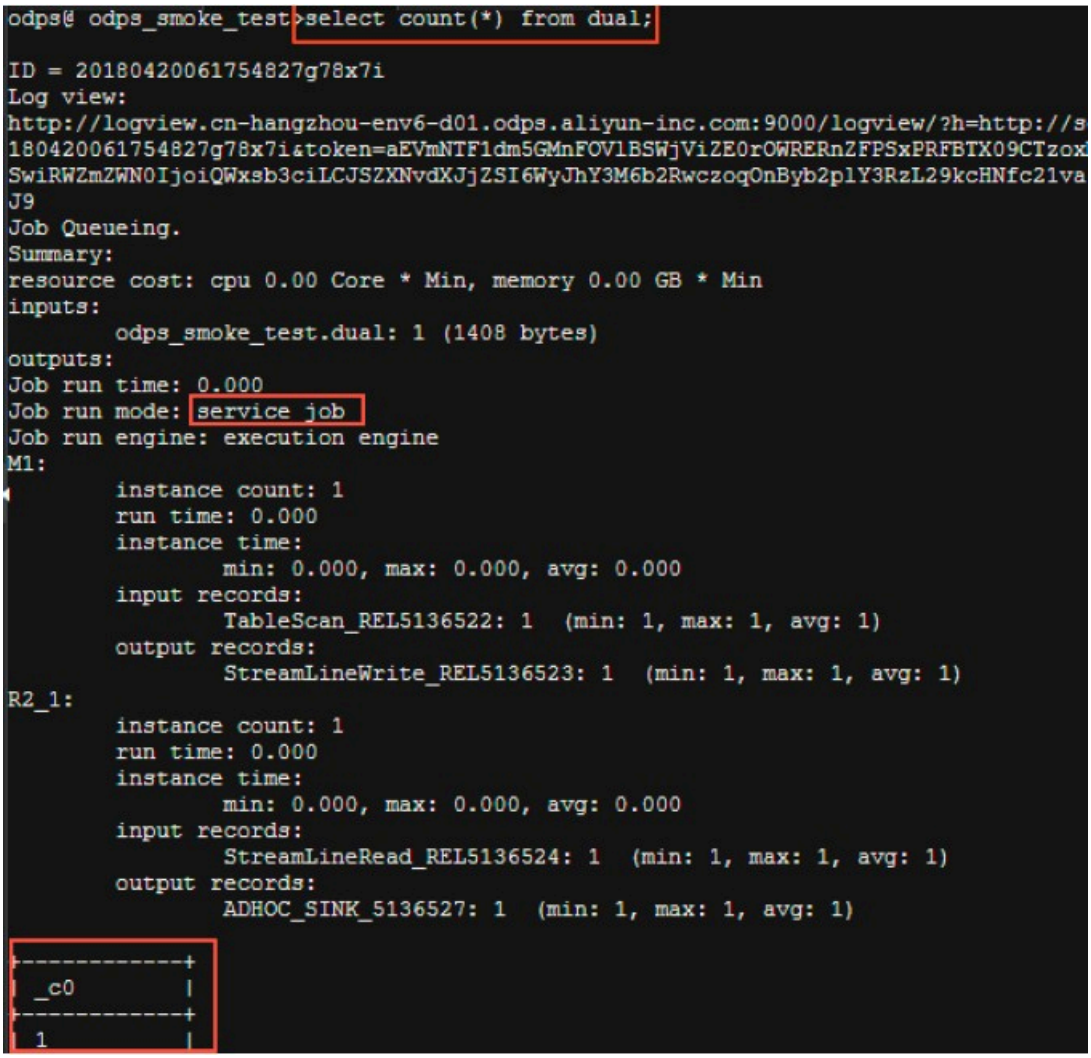

As shown in the following figure, `fuxi job` **is running. The command output indicates that the cluster functions properly.**

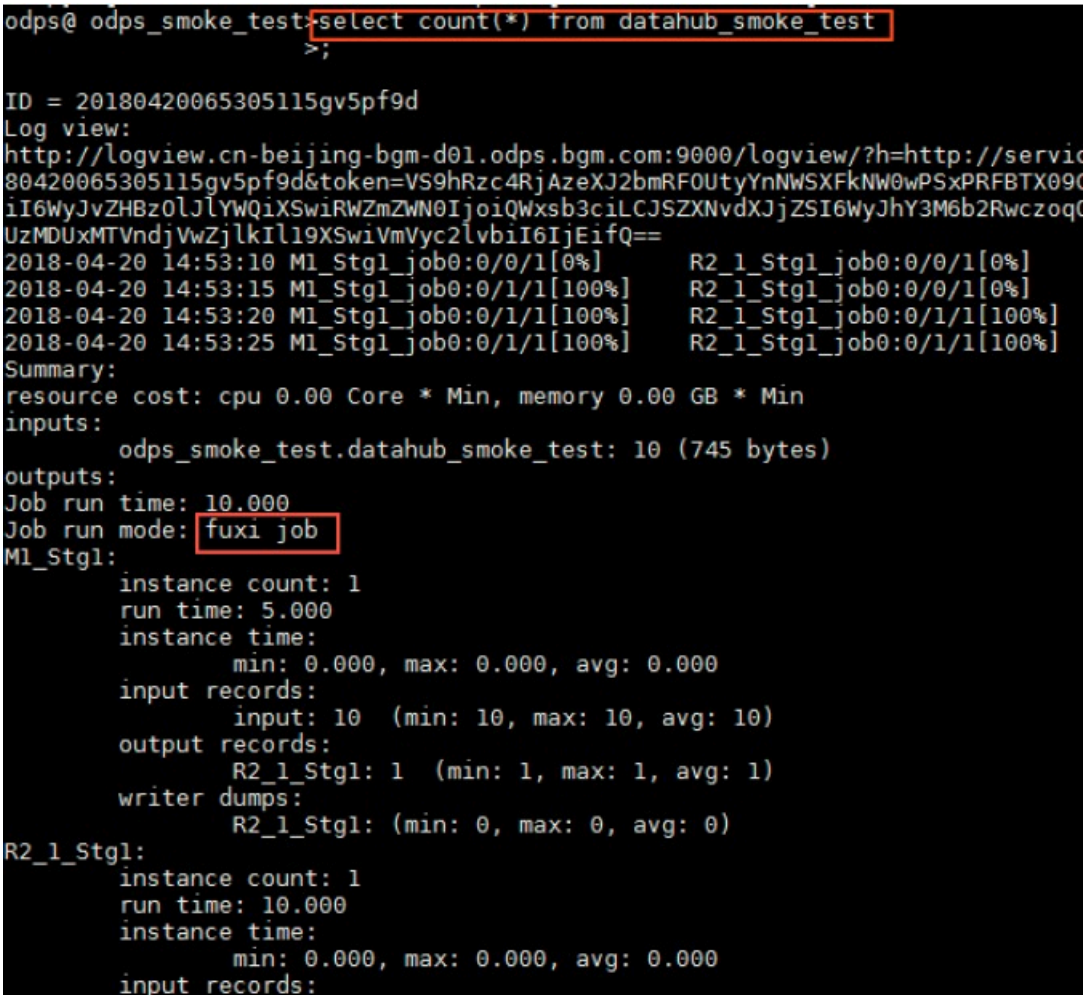
odps@ odps_smoke_test>select count(*) from datahub_smoke_test
>;
ID = 20180420065305115gv5pf9d
Log view:
http://logview.cn-beijing-bgm-d01.odps.bgm.com:9000/logview/?h=http://servic
80420065305115gv5pf9d&token=VS9hRzc4RjAzeXJ2bmRFOUtyYnNWSXFkNW0wPSxPRFBTX09
iI6WyJvZHBzOlJlYWQiXSwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2Rwczoq
UzMDUxMTVndjVwZjlkIl19XSwiVmVyc2lvbiI6IjEifQ==
2018-04-20 14:53:10 M1_Stg1_job0:0/0/1[0%] R2_1_Stg1_job0:0/0/1[0%]
2018-04-20 14:53:15 M1_Stg1_job0:0/1/1[100%] R2_1_Stg1_job0:0/0/1[0%]
2018-04-20 14:53:20 M1_Stg1_job0:0/1/1[100%] R2_1_Stg1_job0:0/1/1[100%]
2018-04-20 14:53:25 M1_Stg1_job0:0/1/1[100%] R2_1_Stg1_job0:0/1/1[100%]
Summary:
resource cost: cpu 0.00 Core * Min, memory 0.00 GB * Min
inputs:
odps_smoke_test.datahub_smoke_test: 10 (745 bytes)
outputs:
Job run time: 10.000
Job run mode: fuxi job
M1_Stg1:
instance count: 1
run time: 5.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:
input: 10 (min: 10, max: 10, avg: 10)
output records:
R2_1_Stg1: 1 (min: 1, max: 1, avg: 1)
writer dumps:
R2_1_Stg1: (min: 0, max: 0, avg: 0)
R2_1_Stg1:
instance count: 1
run time: 10.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:

3. **Run the following commands to check whether the following workers exist and whether they have been restarted recently:**

   a. `r swl Odps/MessagerServicex`

```
$r swl Odps/MessagerServicex
WorkerName                                  | LastUpdateTime           | pid   | planned | loaded | unloaded
MessageServerRole@101h05215.cloud.h07.amtest1284 | Mon Apr  9 16:49:03 2018 | 24697 | 1       | 1      | 0
MessageServerRole@101h11210.cloud.h13.amtest1284 | Mon Apr  9 16:48:37 2018 | 15149 | 1       | 1      | 0
MessageServerRole@101h08109.cloud.h09.amtest1284 | Mon Apr  9 16:49:03 2018 | 23586 | 1       | 1      | 0
```

   b. `r swl Odps/OdpsServicex`

```
$r swl Odps/OdpsServicex
WorkerName                                  | LastUpdateTime           | pid   | planned | loaded | unloaded
RecycleWorker@101h08114.cloud.h09.amtest1284    | Mon Apr  9 17:05:42 2018 | 52905 | 0       | 0      | 0
OdpsWorker@101h08114.cloud.h09.amtest1284       | Mon Apr  9 17:05:42 2018 | 52904 | 0       | 0      | 0
OdpsWorker@101h11010.cloud.h11.amtest1284       | Mon Apr  9 17:04:06 2018 | 4454  | 0       | 0      | 0
ExecutorWorker@101h08114.cloud.h09.amtest1284   | Mon Apr  9 17:05:42 2018 | 52903 | 0       | 0      | 0
ExecutorWorker@101h11010.cloud.h11.amtest1284   | Mon Apr  9 17:04:22 2018 | 6524  | 0       | 0      | 0
SchedulerWorker@101h08114.cloud.h09.amtest1284  | Mon Apr  9 17:05:47 2018 | 53609 | 0       | 0      | 0
WorkflowWorker@101h08114.cloud.h09.amtest1284   | Mon Apr  9 17:05:48 2018 | 53610 | 0       | 0      | 0
```

   c. `r swl Odps/HiveServerx`

```
$r swl Odps/HiveServerx
WorkerName                                  | LastUpdateTime           | pid   | planned | loaded | unloaded
AuthServer@101h08114.cloud.h09.amtest1284    | Tue Apr 10 18:05:54 2018 | 23585 | 0       | 0      | 0
HiveServer@101h11010.cloud.h11.amtest1284    | Mon Apr  9 17:03:07 2018 | 1696  | 1       | 1      | 0
HiveServer@101h08114.cloud.h09.amtest1284    | Tue Apr 10 18:06:02 2018 | 23587 | 2       | 2      | 0
CatalogServer@101h08114.cloud.h09.amtest1284 | Tue Apr 10 18:05:55 2018 | 23586 | 1       | 1      | 0
```

   d. `r swl Odps/QuotaServicex`

```
$r swl Odps/QuotaServicex
WorkerName                                  | LastUpdateTime           | pid   | planned | loaded | unloaded
QuotaWorkerRole@101h08114.cloud.h09.amtest1284 | Mon Apr  9 16:55:32 2018 | 32814 | 0       | 0      | 0
```

   e. `r swl Odps/ReplicationServicex`

```
$r swl Odps/ReplicationServicex
WorkerName                                  | LastUpdateTime           | pid   | planned | loaded | unloaded
ReplicationServer@101h05215.cloud.h07.amtest1284 | Mon Apr  9 16:49:12 2018 | 26594 | 0       | 0      | 0
ReplicationServer@101h11210.cloud.h13.amtest1284 | Mon Apr  9 16:48:51 2018 | 26859 | 0       | 0      | 0
ReplicationServer@101h11215.cloud.h13.amtest1284 | Mon Apr  9 16:49:18 2018 | 3453  | 0       | 0      | 0
ReplicationMaster@101h11010.cloud.h11.amtest1284 | Mon Apr  9 16:50:21 2018 | 34315 | 0       | 0      | 0
```

4. **Run the following command to check for errors:**

```
puadmin lscs |grep -vi NORMAL|grep -vi DISK_OK
```

```
$puadmin lscs |grep -vi NORMAL|grep -vi DISK_OK

The pangu disk status:
Total Disk Size:681225 GB
Total Free Disk Size:635009 GB
Total File Size:1093 GB
Total UnReserved Disk Space4Piops:0 GB
Total Disk Space4Piops:0 GB
Total UnReserved Disk Iops4Piops:0
Total Disk Iops4Piops:0
TotalChunkNumber:26074944       NonTempChunkNumber:26074030     NonTempChunkDataSize:1093 GB    TempChunkNumber:914     TempChunkDataSize:0 GB

No.     Rack    UsableChunkserver/TotalChunkserver      UsableDisk/TotalDisk    TotalDiskSize   TotalFreeDiskSize
1       101g15  2/2                                     23/23                   128427 GB       119672 GB
2       101h05  1/1                                     11/11                   61421 GB        57318 GB
3       101h08  2/2                                     23/23                   150763 GB       140758 GB
4       101h11  5/5                                     57/57                   340612 GB       317859 GB

Number of Racks:                        4
Number of Usable Racks(Having at least one disk with Free Disk Size > 15GB):    4
Notice!: Total Disk Size of 101h11 >= 1/3 of Total Disk Size of the Cluster, three replicas may not locate in different racks
```

5. **Run the following commands to check the data integrity:**

   a. `puadmin fs -abnchunk -t none`

   ```
   $puadmin  fs -abnchunk  -t  none
   Master Address: nuwa://localcluster/sys/pangu/master
   ChunkId Type    FoundTime
   ```

   b. `puadmin fs -abnchunk -t onecopy`

   ```
   $puadmin  fs -abnchunk  -t  onecopy
   Master Address: nuwa://localcluster/sys/pangu/master
   ChunkId Type    FoundTime
   ```

   c. `puadmin fs -abnchunk -t lessmin`

   ```
   $puadmin  fs -abnchunk  -t  lessmin
   Master Address: nuwa://localcluster/sys/pangu/master
   ChunkId Type    FoundTime
   ```

6. **Log on to the machine where Apsara Name Service and Distributed Lock Synchronization System resides.**

   ```
   echo srvr | nc localhost 10240 | grep Mode
   ```

   Example:

   ```
   tj_show -r nuwa.NuwaZK#>/tmp/nuwa;pssh -h /tmp/nuwa -i "echo srvr |
   nc localhost 10240 | grep Mode"
   ```

   ```
   $tj_show -r nuwa.NuwaZK#>/tmp/nuwa;pssh -h /tmp/nuwa -i "echo srvr | nc localhost 10240 | grep Mode"
   [1] 15:59:01 [SUCCESS] vm010036016093
   Mode: follower
   [2] 15:59:02 [SUCCESS] vm010036032042
   Mode: leader
   [3] 15:59:02 [SUCCESS] vm010036024022
   Mode: follower
   ```

7. **Run the following commands to check whether Apsara Distributed File System functions properly:**

```
puadmin gems
```

```
puadmin gss
```

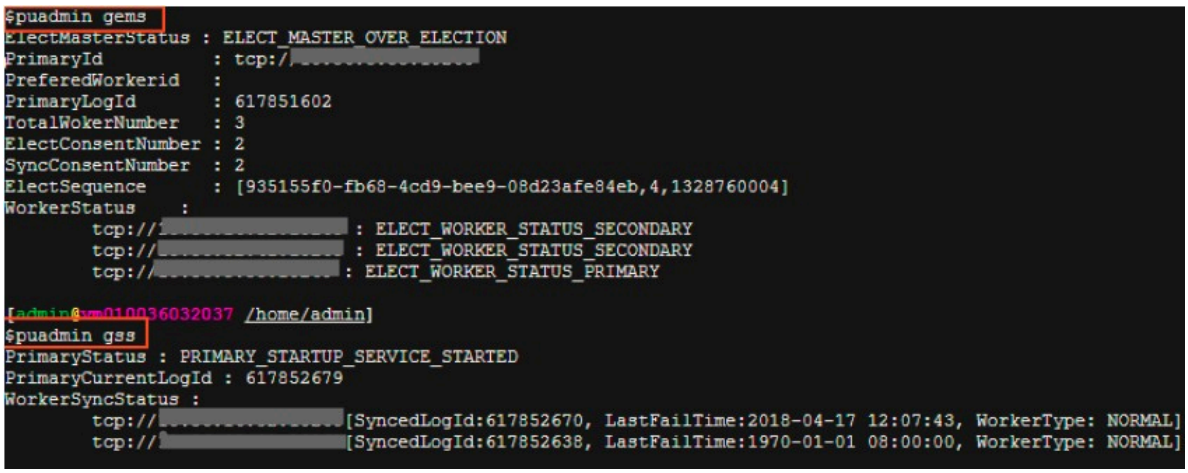

8. **Perform daily inspections in Apsara BigData Manager to check disk usage.**

2.3.3 Shut down a chunkserver, perform maintenance, and then clone the chunkserver

Prerequisites

- **A customer has asked to fix a faulty instance of odps_cs and clone a new one.**
- **You must inform the customer that this operation will temporarily render a chunkserver in the cluster unavailable, but will not affect the overall operation of the service.**
- **All MaxCompute services have reached the desired state and are functioning properly.**
- **All services on the OPS1 server have reached the desired state and are functionin g properly.**
- **You must ensure that the disk space available is sufficient for data migration triggered when a node goes offline.**
- **If the primary node exists on the machine to be brought offline, you must ensure that services are switched from the primary node to the secondary node.**

Procedure

1. **In Apsara Infrastructure Management Framework, find ComputerInit# in the** `odps-service-computer` **service of the odps cluster, and open the corresponding**

TerminalService window. Run the following commands to check the data integrity of Apsara Distributed File System:

```
puadmin abnchunk fs -t none
-- Check for any missing files. If no output is displayed, no files
are missing.
puadmin abnchunk fs -t onecopy
-- Check whether each file has only one copy. If no output is
displayed, each file has only one copy.
puadmin abnchunk fs -t lessmin
-- Check whether the number of files is smaller than the minimum
number of backups. If no output is displayed, the number of files is
 smaller than the minimum number of backups.
```

2. Add the machine to be shut down to a Job Scheduler blacklist.
   a. Run the following command to enable the blacklisting function of Job Scheduler (ignore this step if the function has been enabled):

      ```
      /apsara/deploy/rpc_caller --Server=nuwa://localcluster/sys/fuxi
      /master/ForClient --Method=/fuxi/SetGlobalFlag --Parameter={\"
      fuxi_Enable_BadNodeManager\":false}
      ```

   b. Run the following command to check the hostnames in the existing blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
      ```

   c. Run the following command to add the machine to be shut down to the blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster add $hostname
      ```

   d. Run the following command to check whether the machine to be shut down is already included in the blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
      ```

3. Shut down the machine, perform maintenance, and then restart the machine.

   Note:
   Do not compromise the system during maintenance.

4. Run the following commands to remove the Job Scheduler blacklist:

   ```
   /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster remove $hostname
   ```

```
/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
```

5. **Set the status of rma to pending for the faulty machine.**

   a. **Log on to the OPS1 server. Set the status of the rma action to pending for the faulty machine. The hostname of the faulty machine is m1.**

      Run the following command:

      ```
      curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1" -d
       '{"action_name":"rma", "action_status":"pending"}'
      ```

      The command output is as follows:

      ```
      {
          "err_code": 0,
          "err_msg": "",
          "data": [
              {
                  "hostname": "m1"
              }
          ]
      }
      ```

   b. **Run the following command to configure the audit log:**

      ```
      curl "http://127.0.0.1:7070/api/v5/AddAuditLog?object=/m/m1&category
      =action" -d '{"category":"action", "from":"tianji.HealingService#",
      "object":"/m/m1", "content": "{\n \"action\" : \"/action/rma\",\n \"
      description\" : \"/monitor/rma=error, mtime: 1513488046851649\",\n
       \"status\" : \"pending\"\n}\n" }'
      ```

      The mtime parameter, which represents action_description@mtime, is set to 1513488046851649 in the example. Set the parameter to the current system time when you configure the audit log. Run the following command to query the mtime value:

      ```
      curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1&attr=
      action_name,action_status,action_description@mtime"
      ```

      The command output is as follows:

      ```
      {
          "err_code": 0,
          "err_msg": "",
          "data": {
              "action_description": "",
              "action_description@mtime": 1516168642565661,
              "action_name": "rma",
              "action_name@mtime": 1516777552688111,
              "action_status": "pending",
      ```

```
        "action_status@mtime": 1516777552688111,
        "hostname": "m1",
        "hostname@mtime": 1516120875605211
    }
}
```

6. **Wait for approval.**
   a. **Wait until the status of the rma action becomes approved or doing on the machine. Check the action status.**

      Run the following command to obtain the machine information:

      ```
      curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1"
      ```

      Command output:

      A large amount of information is returned. You can locate the following keyword: "action_status": "pending".

   b. **Check the SR approval status on the machine. pending indicates that the SR is being approved. approved, doing, or done indicates that the SR has been approved. If no action was taken, the SR was not approved.**

      Run the following query command:

      ```
      curl http://127.0.0.1:7070/api/v5/GetMachineInfoPackage?hostname=m1&
      attr=sr.id,sr.action_name,sr.action_status
      ```

      Command output: A large amount of information is returned. You can also view items in the doing state on the webpage.

7. **Shut down the machine when the status of rma becomes approved or doing. After the maintenance is completed, start the machine.**

   Note:
   If you need to clone the machine after the maintenance is completed, proceed with the next step. Otherwise, skip the next step.

8. **Clone the machine.**
   a. **After the maintenance is completed, run the following command to clone the machine on the OPS1 server:**

      ```
      curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1&
      action_name=rma&action_status=doing" -d '{"action_name":"clone", "
      action_status":"approved", "action_description":"", "force":true}'
      ```

      The command output is as follows:

      ```
      {
      "err_code": 0,
      "err_msg": "",
      "data": [
      {
      "hostname": "m1"
      }
      ]
      }
      ```

   b. **Access the clone container. Run the following commands to check the clone status and confirm whether the clone operation takes effect.**
      A. **Run the following command to query the clone container:**

         ```
         docker ps|grep clone
         ```

         The command output is as follows:

         ```
         18c1339340ab reg.docker.god7.cn/tianji/ops_service:1f147fec48
         83e082646715cb79c3710f7b2ae9c6e6851fa9a9452b92b4b3366a ops.
         OpsClone__.clone.1514969139
         ```

      B. **Run the following command to log on to the container:**

         ```
         docker ps|grep clone
         ```

      C. **Run the following command to query the clone task:**

         ```
         /home/tops/bin/python /root/opsbuild/bin/opsbuild.py acli list --
         status=ALL -n 10000 | vim -
         ```

9. **Run the following command to restore the machine status:**

   ```
   curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1&
   action_name=rma" -d '{"action_name":"rma","action_status":"done", "
   force":true}'
   ```

10. **Check the machine status through the command or Apsara Infrastructure Management Framework. If the status is GOOD, the machine is normal.**

    Run the following command to check the machine status:

    ```
    curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1&attr=
    state,hostname"
    ```

    

11. **Check whether the cluster has reached the desired state. Ensure that all services on the machine being brought online have reached the desired state.**
12. **Run the following commands to remove the Job Scheduler blacklist:**

    ```
    /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster remove $hostname
    /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
    ```

2.3.4 Adjust the virtual resources of the Apsara system in MaxCompute

Prerequisites

All MaxCompute services have reached the desired state and are functioning properly.

Procedure

1. **Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose Operations > Cluster Operations. In the Cluster search box, enter odps to search for the expected cluster.**

2. **Click the cluster in the search result. On the Cluster Details page, click the Cluster Configuration tab. In the left-side file list, find the role.conf file in the fuxi directory.**

Figure 2-21: role.conf file

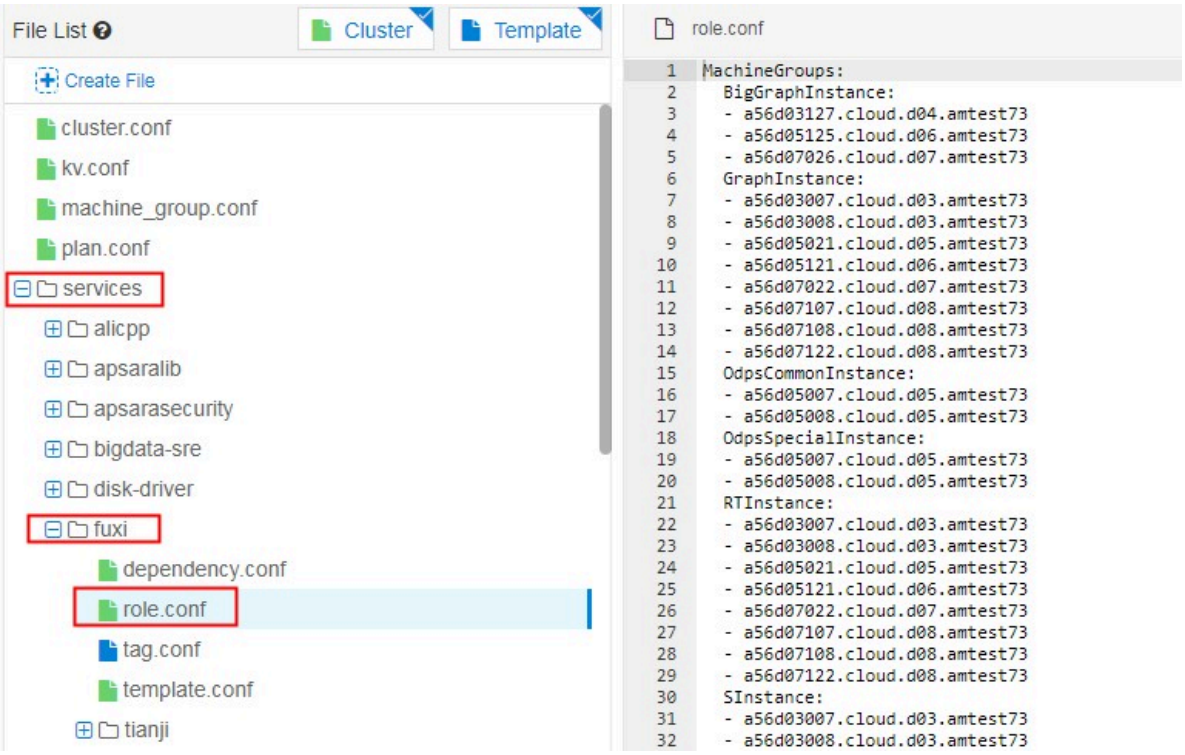

3. **Adjust the machine tags on the right and click Preview and Submit.**

Figure 2-22: Adjust machine tags

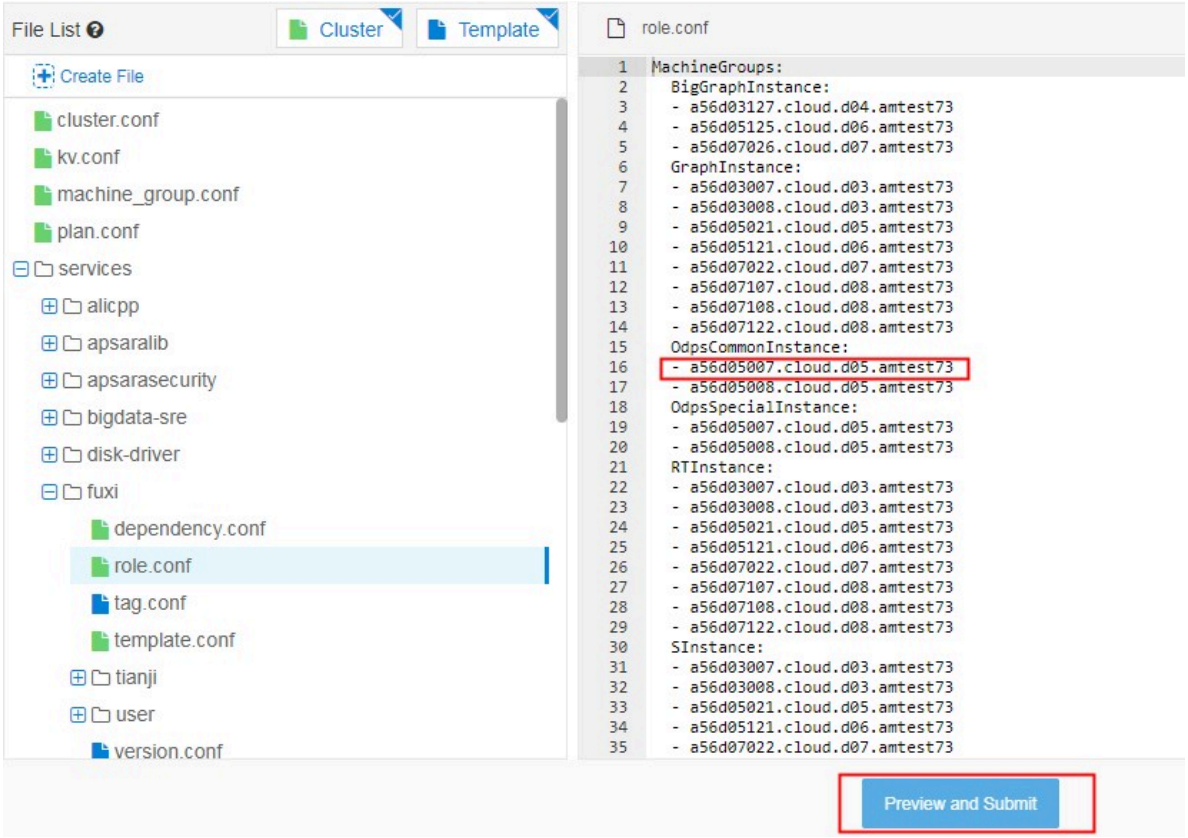

4. **In the Confirm and Submit dialog box that appears, enter the change description and click Submit.**

   Figure 2-23: Submit

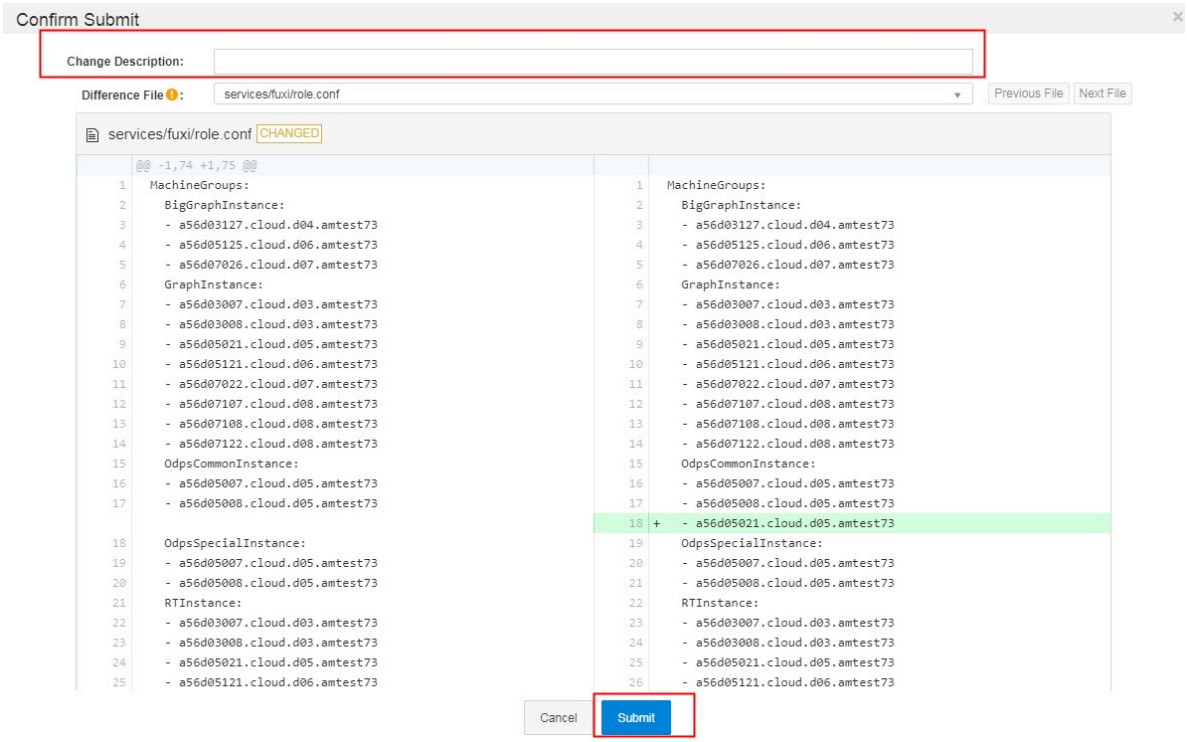

5. **The cluster starts rolling and the changes start to take effect.**

> **Note:**
> **You can check the task status in the operation log. If the changes take effect, the status becomes Successful.**

6. **After the changes are made, run the `r ttrl` command in the TerminalService window to confirm the changes.**

2.3.5 Shut down a chunkserver for maintenance without compromising the system

Prerequisites

Check that all MaxCompute services have reached the final status and are functioning properly.

Procedure

1. **In Apsara Infrastructure Management Framework, locate ComputerInit# in the `odps-service-computer` service of the odps cluster, and open the corresponding**

TerminalService window. Run the following commands to check the data integrity of Apsara Distributed File System:

```
puadmin abnchunk fs -t none
-- Check for any missing files. If no output is displayed, no files
are missing.
puadmin abnchunk fs -t onecopy
-- Check whether each file has only one copy. If no output is
displayed, each file has only one copy.
puadmin abnchunk fs -t lessmin
-- Check whether the number of files is smaller than the minimum
number of backups. If no output is displayed, the number of files is
 smaller than the minimum number of backups.
```

2. Add the machine to be shut down to a Job Scheduler blacklist.
   a. Run the following command to enable the blacklisting function of Job Scheduler (ignore this step if the function has been enabled):

      ```
      /apsara/deploy/rpc_caller --Server=nuwa://localcluster/sys/fuxi
      /master/ForClient --Method=/fuxi/SetGlobalFlag --Parameter={\"
      fuxi_Enable_BadNodeManager\":false}
      ```

   b. Run the following command to check the hostnames in the existing blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
      ```

   c. Run the following command to add the machine to be shut down to the blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster add $hostname
      ```

   d. Run the following command to check whether the machine to be shut down is already included in the blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
      ```

3. Shut down the machine for maintenance and then restart the machine.

   Note:
   Do not compromise the system during maintenance.

4. Run the following commands to remove the Job Scheduler blacklist:

   ```
   /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster remove $hostname
   /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
   ```

Expected results

During the shutdown of Pangu_chunkserver, Apsara Distributed File System will keep trying to read data, and SQL tasks will remain in the running state. The tasks

are completed after seven to eight minutes, or after the machine resumes operation .

2.3.6 Restart MaxCompute services

Procedure

1. **Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose Operations > Cluster Operations. In the Cluster search box, enter odps to search for the expected cluster.**
2. **Click the cluster in the search result. On the Cluster Details page, click the Services tab. In the Service search box, search for odps-service-computer. Click odps-service-computer in the search result.**
3. **After you access the odps-service-computer service, select ComputerInit# on the Service Details page. In the `Actions` column corresponding to the machine, click Terminal. In the TerminalService window that appears, you can perform subsequent command line operations.**
4. **Run the following command to obtain the number of machines:**

```
tj_show -r fuxi.Tubo#
```

5. **Divide the number of machines by 3 to obtain the workernum value.**

> **Note:**
> **The workernum value ranges from 1 to 3.**

6. **Modify workernum in *`vim /apsara/odps_service/deploy/env.cfg`*.**

```
odps_worker_num = 2
executor_worker_num = 2
hiveserver_worker_num = 2
replication_server_num = 2
messager_partition_num = 2
-- The values here are used as an example. Set these values as
needed.
```

7. **Restart Hive and MaxCompute.**

```
/apsara/odps_service/deploy/install_odps.sh restart_hiveservice
-- Restart Hive.
/apsara/odps_service/deploy/install_odps.sh restart_odpsservice
-- Restart MaxCompute.
```

```
r swl Odps/OdpsServicex
r swl Odps/HiveServerx
```

```
-- Check the service update status and time after restart.
```

8. **Restart the messager service.**

```
cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployme
ssagerservice
-- Restart the messager service.
```

```
r swl Odps/MessagerServicex
-- Check the service update status and time after restart.
```

9. **Restart the quota service.**

```
cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployqu
otaservice
-- Restart the quota service.
```

```
r swl Odps/QuotaServicex
-- Check the service update status and time after restart.
```

10. **Restart the replication service.**

```
cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployre
plicationservice
-- Restart the replication service.
```

```
r swl Odps/ReplicationServicex
-- Check the service update status and time after restart.
```

11. **Restart the service mode.**

```
r plan Odps/CGServiceControllerx >/home/admin/servicemode.json
r sstop Odps/CGServiceControllerx
r start /home/admin/servicemode.json
-- Restart the service mode.
```

```
r swl Odps/CGServiceControllerx
-- Check the CGServiceControllerx service update status and time
after restart.
```

2.4 Common issues and solutions

2.4.1 View and allocate MaxCompute cluster resources

This topic describes how to view the storage and computing resources in a MaxCompute cluster. This topic also describes the quota group-related concepts, relationships between a quota group and a MaxCompute project, and quota group division policies.

Resources that can be allocated to projects in a MaxCompute cluster

- **Storage resources: The total sum of storage resources available in a MaxCompute cluster is limited and can be calculated based on the number of compute nodes in the entire cluster. The storage capacity in a MaxCompute cluster is managed through Apsara Distributed File System. You can run Apsara Distributed File System commands to view the total storage capacity, such as the current storage usage statistics. The following metrics are available for measuring storage resources:**
  - **Storage capacity metric: indicates the total size of files that can be stored in a cluster. You can calculate the total file size in a cluster based on the following formula: Total file size in a cluster = Number of machines * (Size of a single disk * (Number of disks on a single machine – 1)) * System security level * System compression ratio/Number of distributed replicas.**

    

    Note:

    - **Based on the standard TPC-H test data set, the ratio of the original data size to the compressed data size is 3:1. The ratio varies depending on the characteristics of business data.**
    - **Typically, three replicas are stored in a distributed manner.**
    - **Security level: The default value is 0.85 in the MaxCompute system. You can set a custom security level as required. For example, when the business data increases rapidly and reaches 85% of the total storage quota, the**

security level is low. You must scale out the system as required or delete unnecessary data.

How to view the storage capacity of a MaxCompute cluster

- **Run the `puadmin lscs` command on the cluster AG. The total disk size, total free disk size, and total file size are displayed at the end of the command output.**

  Figure 2-24: Capacity information

  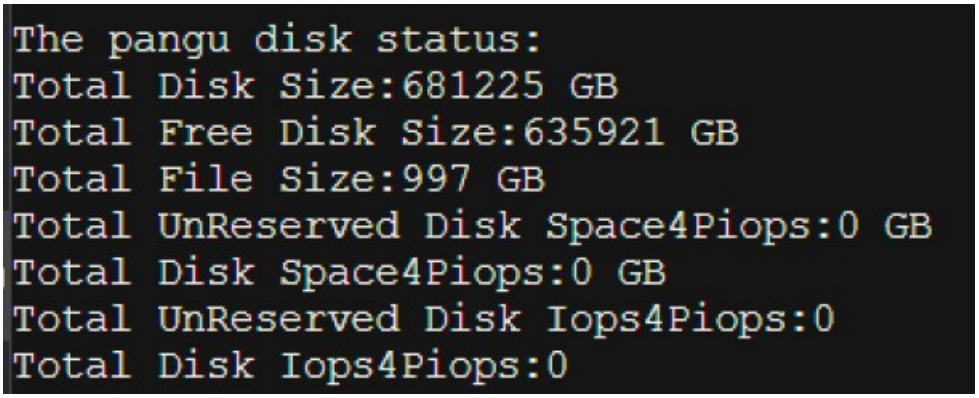

  Note:

  Parameters:

  - **Total Disk Size: the total amount of physical space. Each file is stored in three copies. The logical space is one third the size of the physical space.**
  - **Total Free Disk Size: the total size of available disks, excluding recycle bins on chunkservers.**
  - **Total File Size: the total amount of physical space used by Apsara Distributed File System files, including the /deleted/ directory.**

- **Run the following command on the cluster AG to view the storage capacity used by all projects:**

  ```
  pu ls -l pangu://localcluster/product/aliyun/odps/
  ```

  Example:

  ```
  pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -
  A 4
  ```

```
-- View the capacity used by a single project, such as adsmr.
```

Figure 2-25: Project capacity information

```
$pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -A 4
pangu://localcluster/product/aliyun/odps/adsmr/
Length         : 551267930
FileNumber     : 570
DirNumber      : 143
Pinned         : 0
```

Note:

Parameters:

- **Length: the logical length used by a project. The physical length required is three times the logical length.**
- **FileNumber: the number of files used.**
- **DirNumber: the number of directories used.**

- **File size metric: The total size of files that can be stored in a cluster is limited based on the memory capacity of PanguMaster. The existence of a large number of small files or an improper number of files in a cluster can also affect the stability of the cluster and its services.**

  The Apsara Distributed File System index files, including the information of Apsara Distributed File System files and directories, are stored in the PanguMaster memory. Each file in PanguMaster corresponds to a file node. Each file node uses XXX bytes of memory, each level of directory uses XXX bytes of memory, and each chunk uses XXX bytes of memory. A large file is split into multiple chunks in Apsara Distributed File System. Therefore, the

factors that affect PanguMaster memory usage include the number of files, directory hierarchy, and number of chunks.

If the size of the original files in Apsara Distributed File System is large, the memory usage of PanguMaster is relatively low. When a large number of small files exist, the memory usage of PanguMaster is relatively high.

We recommend that you perform the following operations to reduce the memory usage of PanguMaster:

- **Reduce or even delete empty directories which occupy memory, and reduce the number of directory levels.**
- **Do not create directories. A directory is created automatically when you create a file.**
- **Store multiple files in a directory. However, a maximum of 100,000 files can be stored.**
- **Decrease the length of file names and directory names to reduce the memory usage and network traffic in PanguMaster.**
- **Reduce the number of small tables and files. We recommend that you use Tunnel to upload and commit MaxCompute tables only when the table data size reaches 64 MB.**

The following figure shows the numbers of files that can be stored in Apsara Distributed File System for different PanguMaster memory capacities.

Figure 2-26: Numbers of files that can be stored for different PanguMaster memory capacities

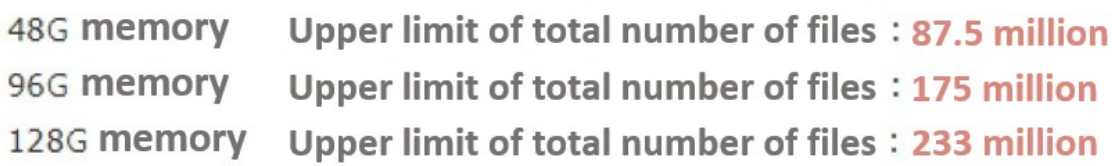

How to view the number of files stored in a MaxCompute cluster

- **Run the `pu quota` command on the cluster AG to view the total number of files stored in a MaxCompute cluster.**

  Figure 2-27: Total number of files

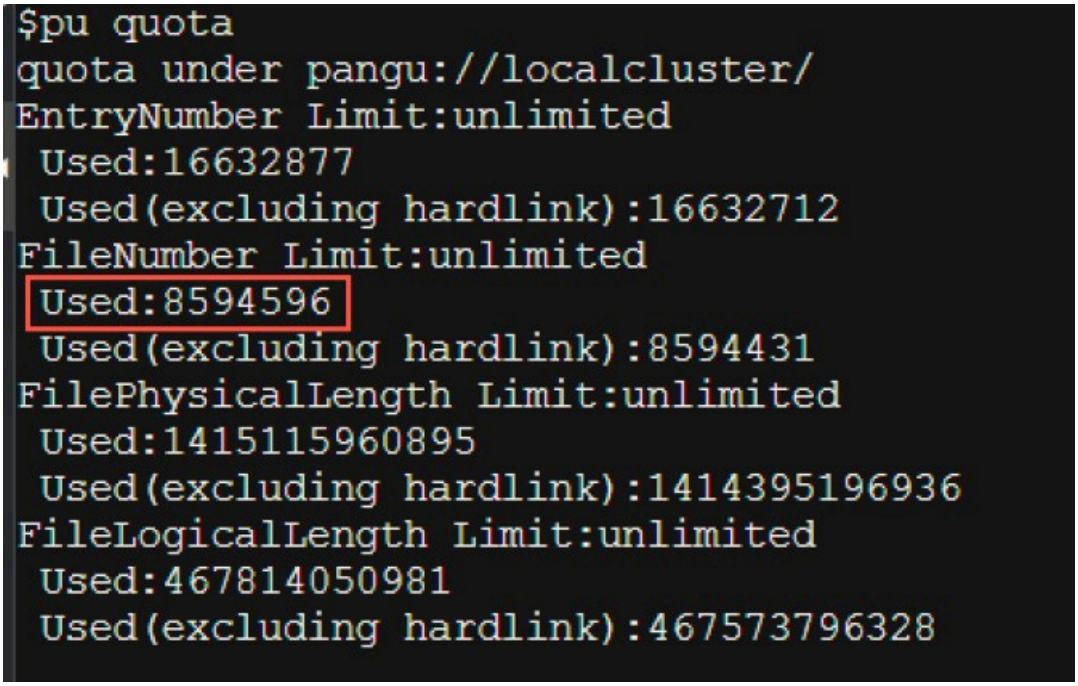

- **This example uses the adsmr project to demonstrate how to view the number of files. Run the following command on the cluster AG to view the number of files for a single project in a MaxCompute cluster:**

```
pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -
A 4
```

  Figure 2-28: Number of files for a single project

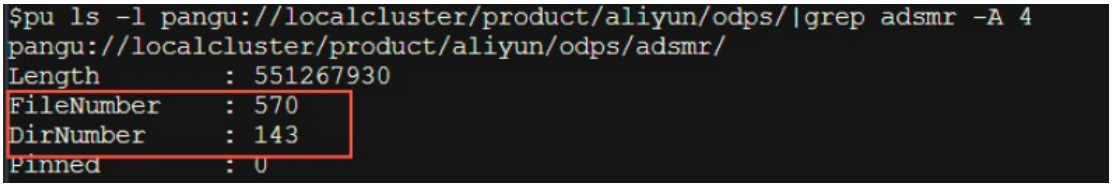

  Note:

  Parameters:

  - **FileNumber: the number of files used.**
  - **DirNumber: the number of directories used.**
  - **FileNumber + DirNumber = Number of files for the current project.**

- **Computing resources: CPU and memory are typically referred to as computing resources in a MaxCompute cluster. The total amount of computing resources is calculated based on the following formula: Total amount of computing resources = (Number of CPU cores + Memory size of each machine) * Number of machines. For example, each machine has 56 CPU cores. One core on each machine is used by the system. The remaining 55 cores are managed by the distributed scheduling system and are scheduled for use by the MaxCompute service. The memory (aside from the chunk of memory for system overhead) is allocated by Job Scheduler. Typically, 4 GB of memory is allocated per CPU core in each MaxCompute task. The ratio varies depending on MaxCompute tasks.**

  How to view computing resources

  - **Run the `r ttrl` command on the cluster AG to view all computing resources.**

    Figure 2-29: All computing resources

```
$r ttrl
total tubo in cluster=13

detail table for every machine:
Machine Name                    | CPU    | Memory    | Other
          .cloud.    .amtest1284 | 6,300  | 170,453   | GraphInstance:8 RTInstance:4 SInstance:99
          .cloud.    .amtest1284 | 6,300  | 234,014   | BigGraphInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453   | GraphInstance:8 RTInstance:4 SInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453   | ElasticSearchInstance:5
          .cloud.    .amtest1284 | 6,300  | 234,014   | BigGraphInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453   |
          .cloud.    .amtest1284 | 6,300  | 170,453   | GraphInstance:8 RTInstance:4 SInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453   | OdpsSpecialInstance:20 OdpsCommonInstance:20
          .cloud.    .amtest1284 | 6,300  | 170,453   | ElasticSearchInstance:5
          .cloud.    .amtest1284 | 6,300  | 170,453   | ElasticSearchInstance:5
          .cloud.    .amtest1284 | 6,300  | 234,014   | BigGraphInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453   | OdpsSpecialInstance:20 OdpsCommonInstance:20
          .cloud.    .amtest1284 | 6,300  | 170,453   | GraphInstance:8 RTInstance:4 SInstance:99
Total                             | 81,900 | 2,406,572 | NA
```

Note:

In the command output, the domain name, total CPU capacity (Unit: U. 100 U = 1 core), and total memory (Unit: MB) of each Tubo machine, as well as

the role of each Tubo machine in Job Scheduling System are listed in four columns.

- **Run the `r tfrl` command on the cluster AG to view the remaining computing resources.**

Figure 2-30: Remaining computing resources

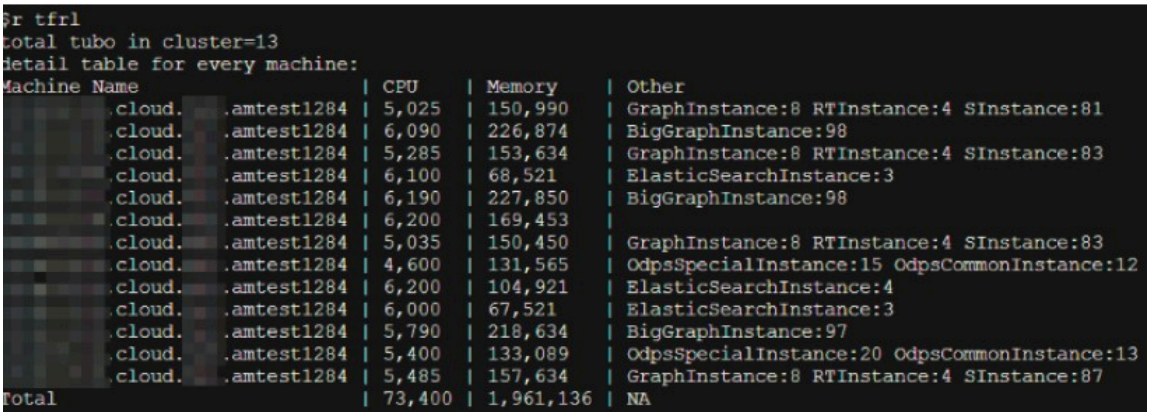

Note:

In the command output, the domain name, total CPU capacity (Unit: U. 100 U = 1 core), and total memory (Unit: MB) of each Tubo machine, as well as the role of each Tubo machine in Job Scheduling System are listed in four columns.

- **Run the `r cru` command on the cluster AG to view the resources used by all running jobs in MaxCompute.**

Figure 2-31: Resources used by all running jobs

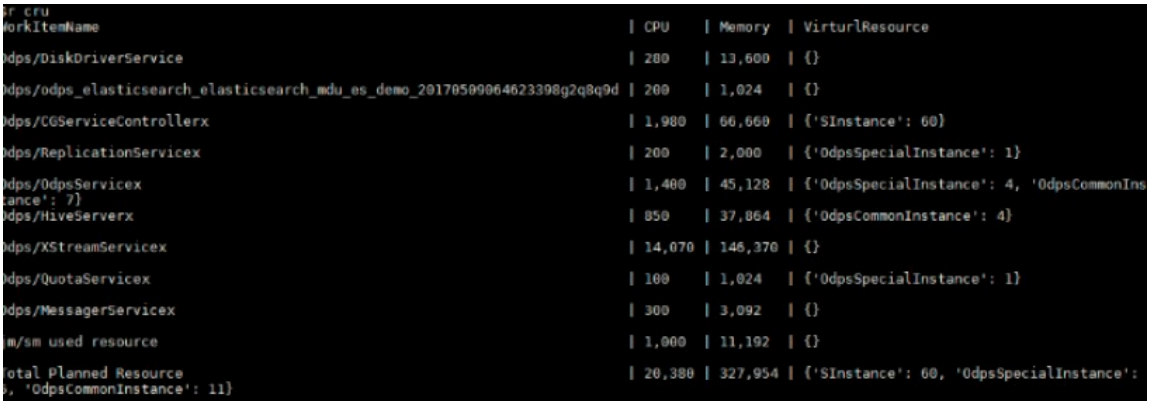

Note:

The name, total CPU capacity, total memory of each job, as well as the number of Fuxi instances started in the role of each job in Job Scheduling System are listed in four columns.

How to allocate project resources in a MaxCompute cluster

- **Storage resource allocation: Based on the characteristics of a project, the space size and file size limit are configured when you create the project.**

  If the following error messages are displayed, the file size limit of the project has been exceeded. In this case, you must organize the data in the project by deleting unnecessary table data or increasing the storage resource quota.

  Figure 2-32: Error messages

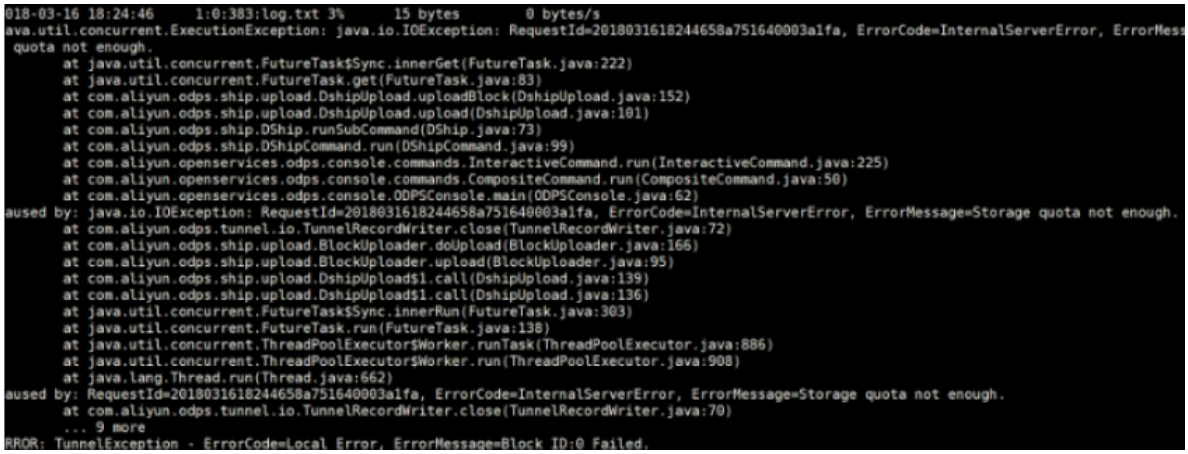

  Notice:
  The sum of the storage capacity of all projects cannot exceed the total allowable storage capacity of a service. Similarly, the total file size of all projects cannot exceed the total allowable file size. Therefore, you must properly allocate the storage space and file size limit by project and make timely adjustment based on your business requirements.

- **Computing resource allocation: division of quota groups.**
  - **What is a quota group?**

    A MaxCompute cluster allows you to divide computing resources into different quota groups, and schedule them as required. A quota group represents a certain amount of CPU and memory resources. MinQuota and MaxQuota are used for CPU and memory configurations. MinQuota is the minimum quota allowed for the quota group, and MaxQuota is the maximum quota allowed for

the quota group. For example, MinCPU=500 indicates that the quota group has been assigned at least 500/100=5 cores. MaxCPU=2000 indicates that the quota group has been assigned at least 2000/100=20 cores.

MaxCompute uses a FAIR scheduling policy and a first-in-first-out (FIFO) scheduling policy by default. The difference between the FAIR and FIFO scheduling polices lies in the keys by which tasks in waiting queues are sorted . If each schedule unit has its own priority, both FAIR and FIFO scheduling policies allocate high-priority schedule units first. If all schedule units share the same priority, the FIFO scheduling policy sorts the schedule units by the time when they are submitted. The earlier they are submitted, the higher priority they have. The FAIR scheduling policy sorts the scheduling units by the slotNum allocated to them. The smaller the slotNum is, the higher priority they have. For the FAIR policy group, this can basically ensure that the same amount of resources are assigned to schedule units with the same priority.

You can run the `r quota` command on the cluster AG to view quota group settings.

Figure 2-33: View quota group settings

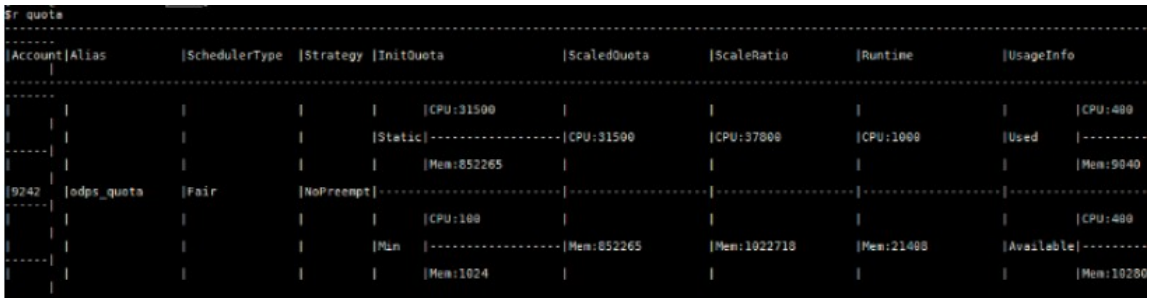

You can run the following command on the cluster AG to create and modify a quota as needed:

```
sh /apsara/deploy/rpc_wrapper/rpc.sh setquota -i $QUOTAID -a $
QUOTANAME -t fair -s $max_cpu_quota $max_mem_quota -m $min_cpu_qu
ota $min_mem_quota
```

Note:

The command with `$QUOTAID` is used to modify a quota. The command without $QUOTAID is used to create a quota.

Figure 2-34: Create a quota

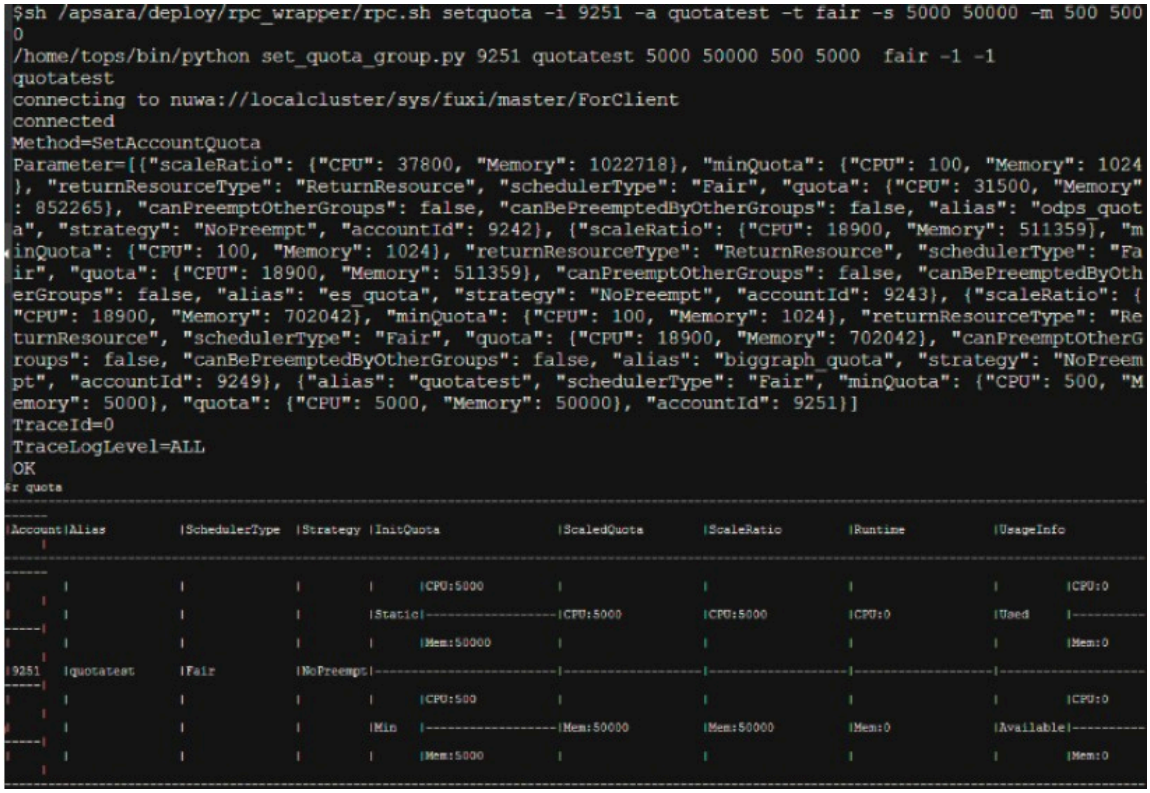

Figure 2-35: Modify a quota

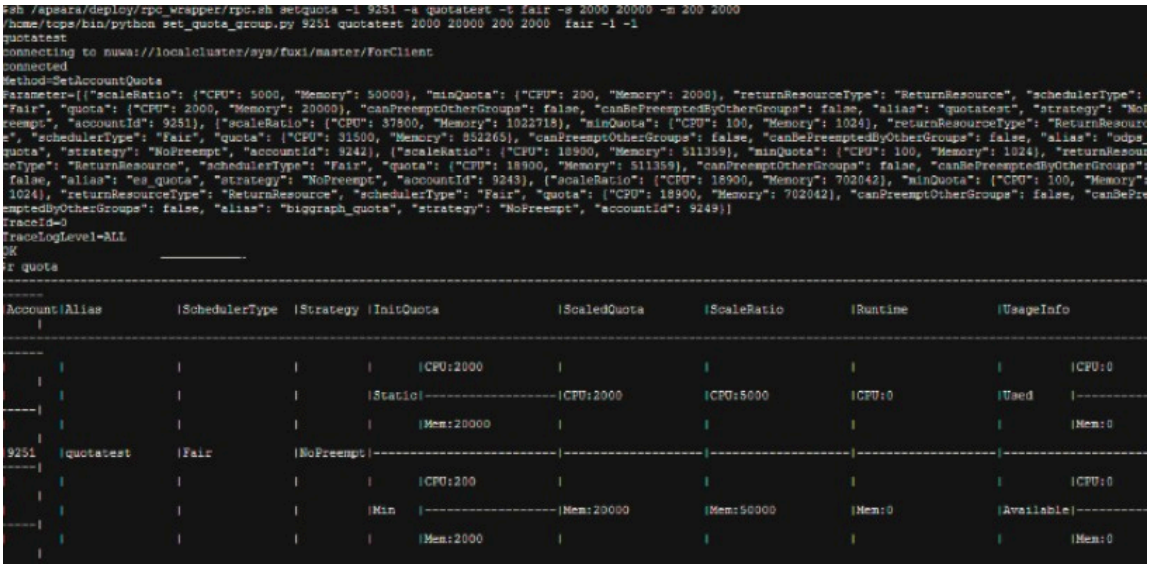

- **How to divide quota groups**

  To divide quota groups correctly, you must understand the relationship between a MaxCompute project and a quota group.

You can select the quota group to which a project belongs upon project creation or modify the quota group after project creation.

Resources in a quota group can be used by all running tasks of all projects in this quota group. Therefore, the project tasks in the same quota group may be affected during peak hours. That is, one or several large tasks may take up all resources in the quota group, while other computing tasks can only wait for resources.

For example, in the following two figures, the first figure shows that a lot of jobs are waiting for resources (in red box). However, a lot of cluster resources are left unused. You can check the quota usage. In the second figure, quota 9243 is only allocated with 5000U, all of which are in use. The CPU quota for 9243 is used up, but there are still pending tasks in 9243. In this case, even if

there are unused cluster resources, the tasks under this quota cannot have resources allocated to them.

Figure 2-36: Jobs waiting for resources

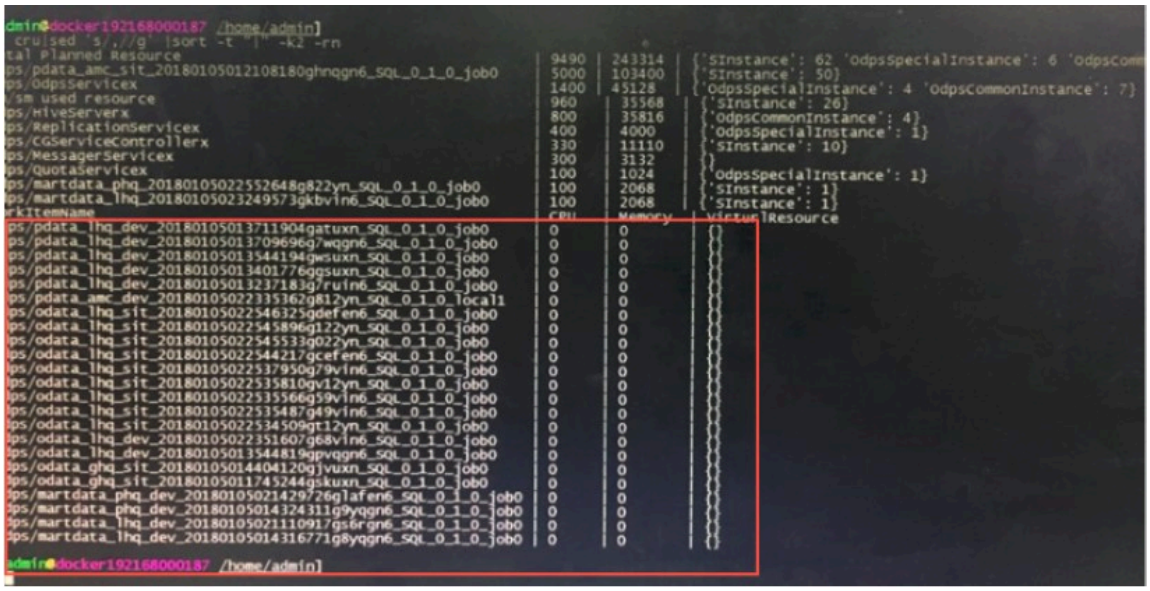

Figure 2-37: Quota used up

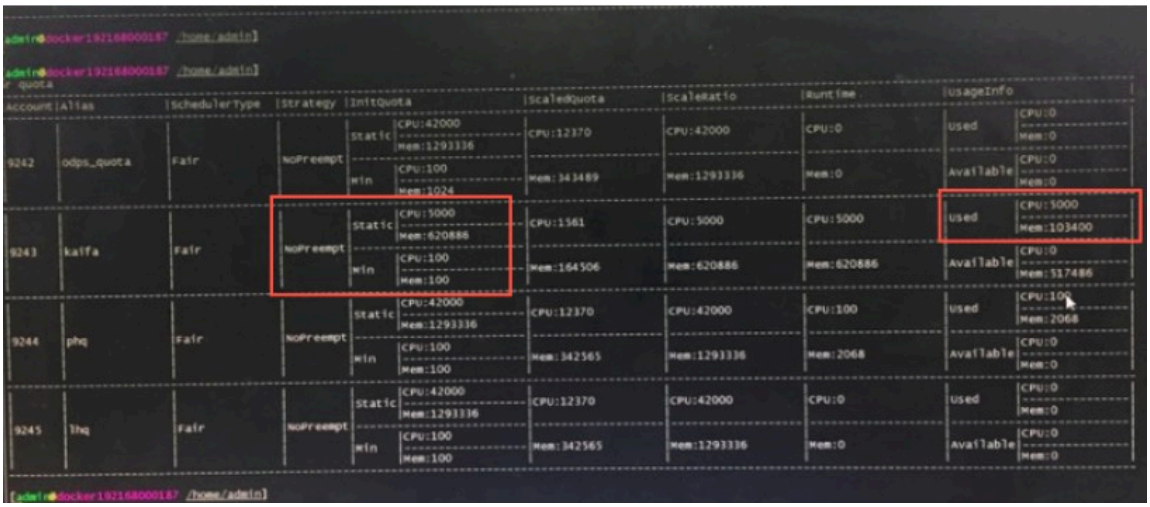

You must divide quota groups based on the following general principles:

- **You must plan quota groups in a way that they do not mutually interfere with each other in a large resource pool, and avoid overly fine-grained division of resource groups. For example, some large tasks cannot be**

scheduled due to quota group limits, or occupy a quota group for an extended period of time, which affects other tasks in the group.

- You must consider the configured MinQuota and MaxQuota when dividing quota groups.
- You can oversell the resources in your cluster, that is, the sum of MaxQuotas of all quota groups can be greater than the total amount of cluster resources. However, the oversell ratio cannot be too high. If the oversell ratio is too high, a quota group with a running project may perpetually occupy a large amount of resources.
- When dividing quota groups, you must consider the priorities of tasks, task execution duration, amount of task data, and characteristics of computing types.
- Properly configure quota groups for peak hours. We recommend that you configure a separate quota group for tasks that are important and time-consuming.
- The division of quota groups and the selection and configuration of projects are conducted based on a resource pre-allocation policy, which needs to be adjusted in a timely manner, based on actual requirements.

2.4.2 Common issues and data skew troubleshooting

Scenario 1: how to determine whether a job has stopped running due to insufficient resources

Symptom: The job does not progress as expected.

Figure 2-38: Symptom

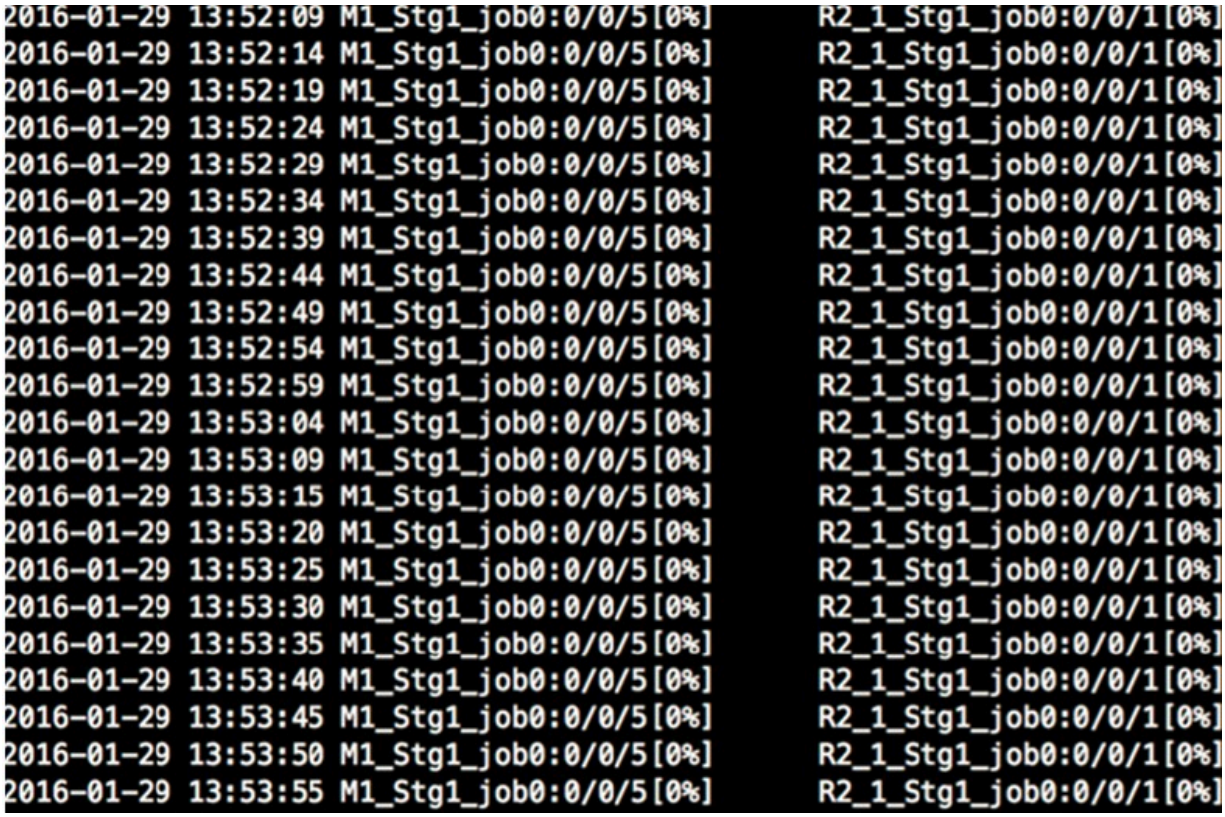

Cause: The issue is typically caused by insufficient resources. You can use LogView to determine the status of job resources (task instance status).

- **Ready: indicates that instances are waiting for Job Scheduler to allocate resources. Instances can resume operation after they obtain the necessary resources.**
- **Wait: indicates that instances are waiting for dependent tasks to complete.**

The task instances in the Ready state shown in the following figure indicate that there are insufficient resources to run these tasks. After an instance obtains the necessary resources, its status changes to `Running`**.**

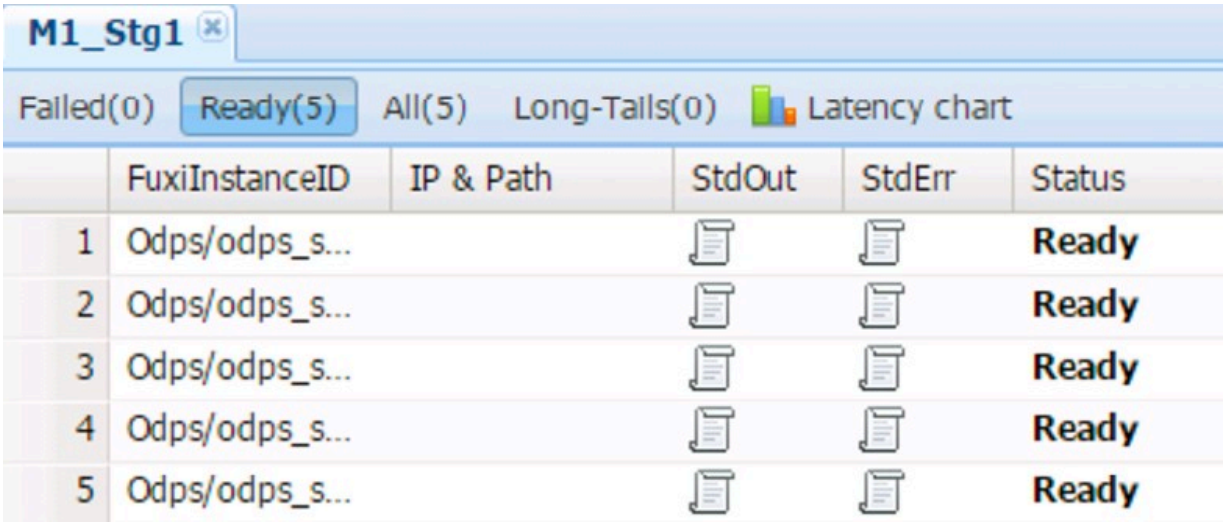

Solution:

- **If there are insufficient resources during peak hours, you can reschedule the tasks to run during off-peak hours.**
- **If the computing quotas are insufficient, check whether the quota group of the project has sufficient computing resources.**
- **If computing resources in the cluster are occupied for long periods of time, you can develop a computing quota allocation policy to scale the quota as necessary.**
- **We recommend that you do not run abnormally large jobs to prevent the jobs from occupying resources for extended periods of time.**
- **You can enable SQL acceleration, so that you can run small jobs without requesting resources from Job Scheduler.**
- **You can use the First-In First-Out (FIFO) scheduling policy.**

Scenario 2: how to find the root cause of a job that has been running for an extended period of time

Symptom: The MaxCompute job execution progress has remained at 99% for a long period of time.

Cause: The running time of some Fuxi instances in the MaxCompute job is significantly longer than that of other Fuxi instances.

Figure 2-39: Cause analysis

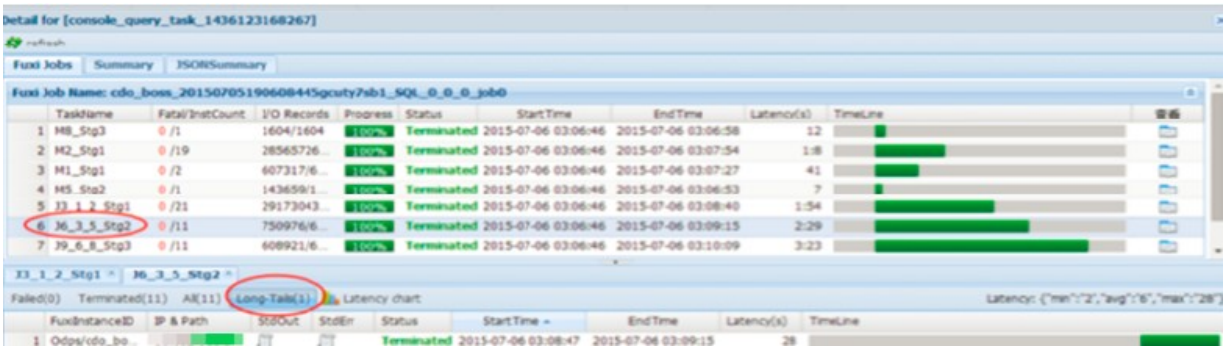

Further analysis: Analyze the job summary in LogView, and calculate the difference between the max and avg values of input and output records of a slow task. If the max and avg values differ by several orders of magnitude, it can be initially determined that the job data is skewed.

Figure 2-40: Further analysis

```
R2_1_Stg1:
        instance count: 1
        run time: 12.000
        instance time:
                min: 0.000, max: 0.000, avg: 0.000
        input records:
                input: 15  (min: 15, max: 15, avg: 15)
        output records:
                R2_1_Stg1FS_11934: 15  (min: 15, max: 15, avg: 15)
```

Solution: If there are slow Fuxi instances on a particular machine, check whether a hardware failure has occurred on the machine.

Scenario 3: How to improve the concurrency of MaxCompute jobs

Fault locating: The concurrency of Map tasks depends on the following factors:

- **Split size and merge limit.**

  Map takes a series of data files as inputs. Larger files are split into partitions based on the odps.sql.mapper.split.size value, which is 256 MB by default. An instance is started for each partition. However, starting an instance requires resources and time. Small files can be merged into a single partition based on the odps.sql.mapper.merge.limit.size value and be processed by a single instance to

improve instance utilization. The default value of odps.sql.mapper.merge.limit.size is 64 MB. The total size of small files merged cannot exceed this value.

- Instances cannot process data across multiple partitions.

  A partition is mapped to a folder in Apsara Distributed File System. You must run at least one instance to process data in a partition. Instances cannot process data across multiple partitions. In a partition, you must run instances based on the preceding rule.

Typically, the number of instances for Reduce tasks is 1/4 of that for Map tasks. The number of instances for Join tasks is the same as that for Map tasks, but cannot exceed 1,111.

You can use the following methods to increase the number of concurrent instances for Reduce and Join tasks:

```
set odps.sql.reducer.instances = xxx
```

```
set odps.sql.joiner.instances = xxx
```

Scenarios that require higher concurrency:

- A single record only contains a small amount of data.

  Because a single record contains a small amount of data, there are many records in a file of the same size. If you split data into 256 MB chunks, a single Map instance needs to process a large number of records, reducing concurrency.
- Dump operations occur in the Map, Reduce, and Join stages.

  Based on the preceding job summary analysis, the displayed dump information indicates that the instance does not have sufficient memory to sort data in the Shuffle stage. Improving concurrency can reduce the amount of data processed by a single instance to the amount of data that can be handled by the memory, eliminate disk I/O time consumption, and improve the processing speed.
- Time-consuming UDFs are used.

  The execution of UDFs is time-consuming. If you execute UDFs concurrently, you can reduce the UDF execution time of an instance.

Solution:

- **You can decrease the following parameter values to improve the concurrency of Map tasks:**

```
odps.sql.mapper.split.size = xxx
odps.sql.mapper.merge.limit.size = xxx
```

- **You can increase the following parameter values to improve the concurrency of Reduce and Join tasks:**

```
odps.sql.reducer.instances = xxx
odps.sql.joiner.instances = xxx
```

Note: Improving concurrency will result in a greater amount of resources being consumed. We recommend that you take cost into account when improving concurrency. An instance takes an average of 10 minutes to complete after optimization, improving overall resource utilization. We recommend that you optimize jobs in critical paths so that they consume less time.

Scenario 4: how to resolve data skew issues

Different types of data skew issues in SQL are resolved in different ways.

- **GROUP BY data skew**

  The uneven distribution of GROUP BY keys results in data skew on reducers. You can set the anti-skew parameter before executing SQL tasks.

```
set odps.sql.groupby.skewindata=true
```

  After this parameter is set to true, the system automatically adds a random number to each key when running the Shuffle hash algorithm and prevents data skew by introducing a new task.

- **DISTRIBUTE BY data skew**

  Using constants to execute the DISTRIBUTE BY clause for full sorting of the entire table will result in data skew on reducers. We recommend that you do not perform this operation.

- **Data skew in the Join stage**

  Data is skewed in the Join stage when the Join keys are unevenly distributed. For example, a key exists in multiple joined tables, resulting in a Cartesian explosion

of data in the Join instance. You can use one of the following solutions to resolve data skew in the Join stage:

- **When a large table and a small table are joined, use MapJoin instead of Join to optimize query performance.**
- **Use a separate logic to handle a skewed key. For example, when a large number of null values exist in the key, you can filter out the null values or execute a CASE WHEN statement to replace them with random values before the Join operation.**
- **If you do not want to modify SQL statements, configure the following parameters to allow MaxCompute to perform automatic optimization:**

```
set odps.sql.skewinfo=tab1:(col1,col2)[(v1,v2),(v3,v4),...]
set odps.sql.skewjoin=true;
```

- **Data skew caused by multi-distinct**

  Multi-distinct syntax aggravates GROUP BY data skew. You can use the GROUP BY clause with the COUNT function instead of multi-distinct to alleviate the data skew issue.

- **UDF OOM**

  Some jobs report an OOM error during runtime. The error message is as follows: `FAILED: ODPS-0123144: Fuxi job failed - WorkerRestart errCode:9,errMsg:SigKill(OOM), usually caused by OOM(out of memory)`**. You can fix the error by configuring the UDF runtime parameters. Example:**

```
odps.sql.mapper.memory=3072;
set odps.sql.udf.jvm.memory=2048;
```

```
set odps.sql.udf.python.memory=1536;
```

The related data skew settings are as follows:

```
set odps.sql.groupby.skewindata=true/false
```

Description: allows you to enable GROUP BY optimization.

```
set odps.sql.skewjoin=true/false
```

Description: allows you to enable Join optimization. It is effective only when odps.sql.skewinfo is set.

```
set odps.sql.skewinfo
```

Description: allows you to set detailed information for Join optimization. The command syntax is as follows:

```
set odps.sql.skewinfo=skewed_src:(skewed_key)[("skewed_value")]
src a join src_skewjoin1 b on a.key = b.key;
```

Example:

```
set odps.sql.skewinfo=src_skewjoin1:(key)[("0")]
-- The output result for a single skewed value of a single field is
as follows: explain select a.key c1, a.value c2, b.key c3, b.value c4
from src a join src_skewjoin1 b on a.key = b.key;
```

```
set odps.sql.skewinfo=src_skewjoin1:(key)[("0")("1")]
```

```
-- The output result for multiple skewed values of a single field is
as follows: explain select a.key c1, a.value c2, b.key c3, b.value c4
from src a join src_skewjoin1 b on a.key = b.key;
```

Scenario 5: how to configure common SQL parameters

Map settings

```
set odps.sql.mapper.cpu=100
```

Description: allows you to set the number of CPUs used by each instance in a Map task. Default value: 100. Valid values: 50 to 800.

```
set odps.sql.mapper.memory=1024
```

Description: allows you to set the memory size of each instance in a Map task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

```
set odps.sql.mapper.merge.limit.size=64
```

Description: allows you to set the maximum size of control files to be merged. Unit: MB. Default value: 64. You can set this variable to control the inputs of mappers. Valid values: 0 to Integer.MAX_VALUE.

```
set odps.sql.mapper.split.size=256
```

Description: allows you to set the maximum data input volume for a Map task. Unit: MB. Default value: 256. You can set this variable to control the inputs of mappers. Valid values: 1 to Integer.MAX_VALUE.

Join settings

```
set odps.sql.joiner.instances=-1
```

Description: allows you to set the number of instances in a Join task. Default value: -1. Valid values: 0 to 2000.

```
set odps.sql.joiner.cpu=100
```

Description: allows you to set the number of CPUs used by each instance in a Join task. Default value: 100. Valid values: 50 to 800.

```
set odps.sql.joiner.memory=1024
```

Description: allows you to set the memory size of each instance in a Join task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

Reduce settings

```
set odps.sql.reducer.instances=-1
```

Description: allows you to set the number of instances in a Reduce task. Default value: -1. Valid values: 0 to 2000.

```
set odps.sql.reducer.cpu=100
```

Description: allows you to set the number of CPUs used by each instance in a Reduce task. Default value: 100. Valid values: 50 to 800.

```
set odps.sql.reducer.memory=1024
```

Description: allows you to set the memory size of each instance in a Reduce task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

UDF settings

```
set odps.sql.udf.jvm.memory=1024
```

Description: allows you to set the maximum memory size used by the UDF JVM heap. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

```
set odps.sql.udf.timeout=600
```

Description: allows you to set the timeout period of a UDF. Unit: seconds. Default value: 600. Valid values: 0 to 3600.

```
set odps.sql.udf.python.memory=256
```

Description: allows you to set the maximum memory size used by the UDF Python API. Unit: MB. Default value: 256. Valid values: 64 to 3072.

```
set odps.sql.udf.optimize.reuse=true/false
```

Description: When this parameter is set to true, each UDF function expression can only be calculated once, improving performance. Default value: true.

```
set odps.sql.udf.strict.mode=false/true
```

Description: allows you to control whether functions return NULL or an error if dirty data is found. If the parameter is set to true, an error is returned. Otherwise, NULL is returned.

MapJoin settings

```
set odps.sql.mapjoin.memory.max=512
```

Description: allows you to set the maximum memory size for a small table when running MapJoin. Unit: MB. Default value: 512. Valid values: 128 to 2048.

```
set odps.sql.reshuffle.dynamicpt=true/false
```

Description:

- **Dynamic partitioning scenarios are time-consuming. Disabling dynamic partitioning can accelerate SQL.**
- **If there are few dynamic partitions, disabling dynamic partitioning can prevent data skew.**

Scenario 6: how to check the storage usage of a single project

Launch the MaxCompute console as a project owner and run the `desc project <project_name>-extended;` command to view the following information.

Figure 2-41: Storage information

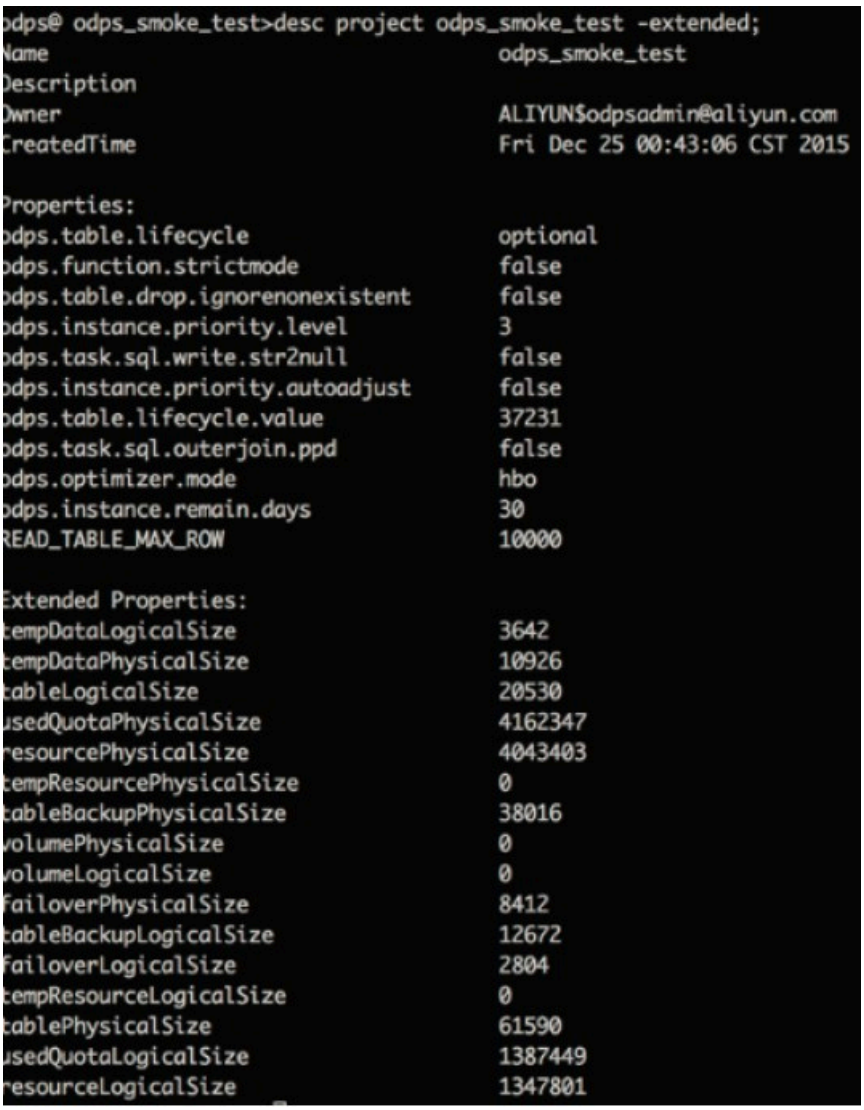

The preceding figure shows the capacity-related storage information of the project. The relationship between the physical and logical values of the related metrics is: Physical value of a metric = Logical value of the metric * Number of replicas.

3 DataWorks

3.1 Basic concepts and structure

3.1.1 What is DataWorks (base)?

DataWorks, also known as base, is a visual workflow development platform that applies MaxCompute as its compute and storage engine. This platform is integrated with a hosted scheduling system, an administration system, and a synchronization system that can handle massive data. You can schedule your tasks by specifying a particular time and task relationships. You can also use the monitoring and management tools to ensure the punctual and accurate execution of millions of tasks. In addition, you are provided with a global overview of each workflow in the form of a directed acyclic graph (DAG).

3.1.2 Functions of base

Data collection

The data synchronization feature enables you to synchronize tables in a source database to a destination database using the data synchronization feature provided by base. Tables can be synchronized between heterogeneous data sources.

Data analysis

Write Shell, MapReduce (MR), or SQL code, and then submit the code to MaxCompute for computing.

Workflow

In base, you can combine task nodes of different types into a workflow. A workflow can contain data sync nodes, SQL nodes, Shell nodes, and MR nodes.

Task scheduling

You can run the tasks periodically with different cycles.

3.1.3 Introduction to data analytics

Scenario 1: data synchronization and analysis

Scenario 1 shows an example of data analytics in typical scenarios.

1. **You can collect data from various databases, and send the data to MaxCompute by using DataWorks.**
2. **You can log on to DataWorks, create SQL, MapReduce, and shell nodes, and commit the nodes to MaxCompute for data analysis.**
3. **You can use DataWorks to synchronize the analysis results from MaxCompute to the databases from which you collect data.**

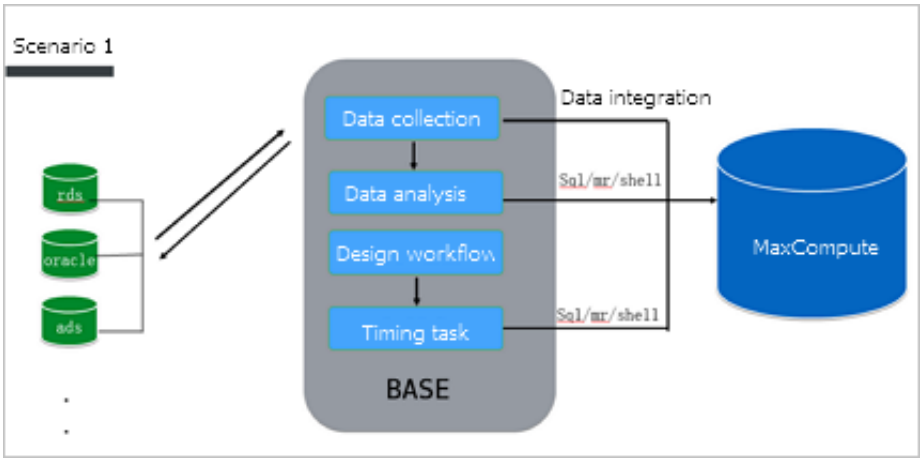

Scenario 2: data synchronization

DataWorks supports data synchronization between various databases. You can synchronize data by using DataWorks.

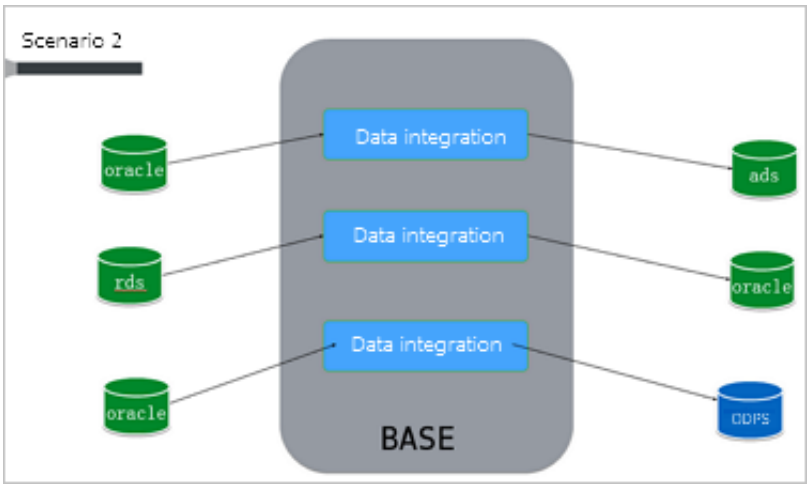

3.1.4 Architecture of DataWorks in Apsara Stack V3

Figure 3-1: The core architecture of base

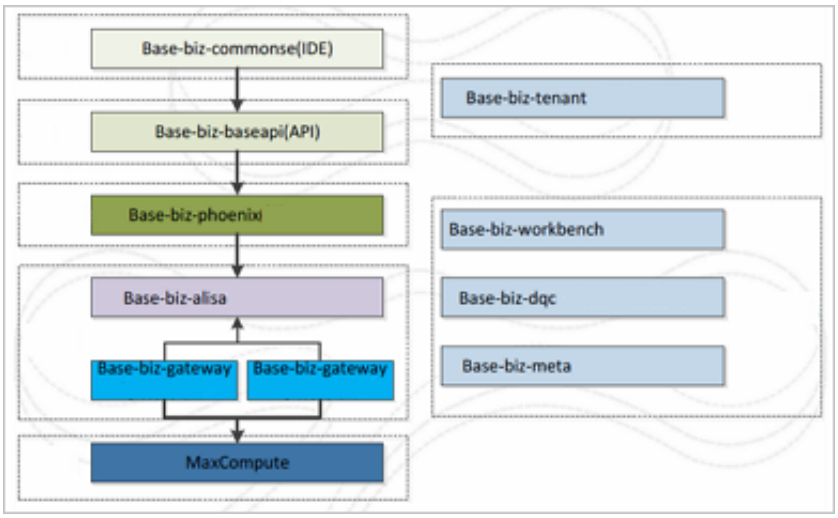

Services in DataWorks play an important role for node scheduling and running. You can perform all O&M operations for DataWorks of Apsara Stack V3 in the Apsara

Infrastructure Management Framework. DataWorks consists of the following services.

Figure 3-2: base components

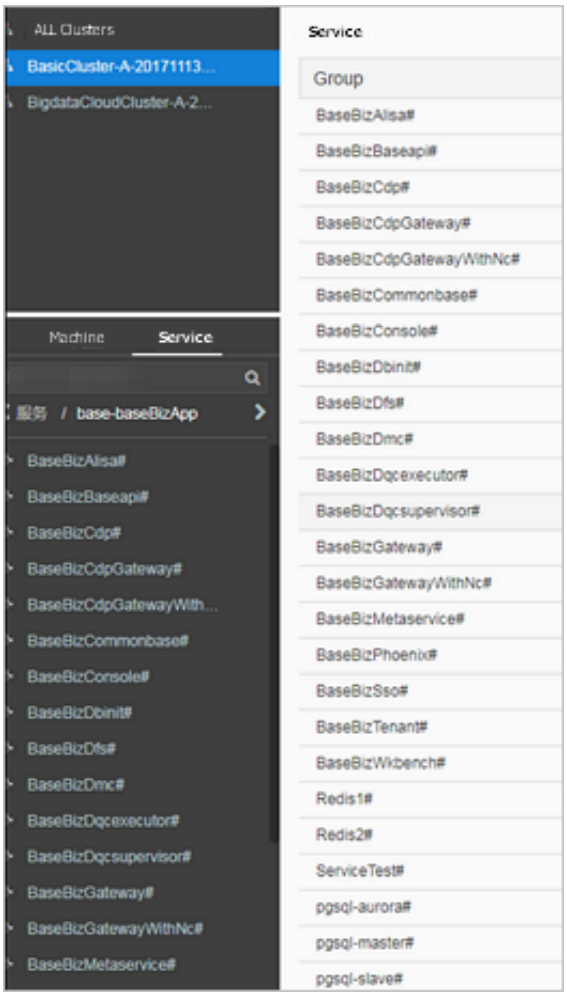

All services in DataWorks are deployed on Docker containers. You can log on to a host, and run the docker ps command to view the containers on which the services are deployed.

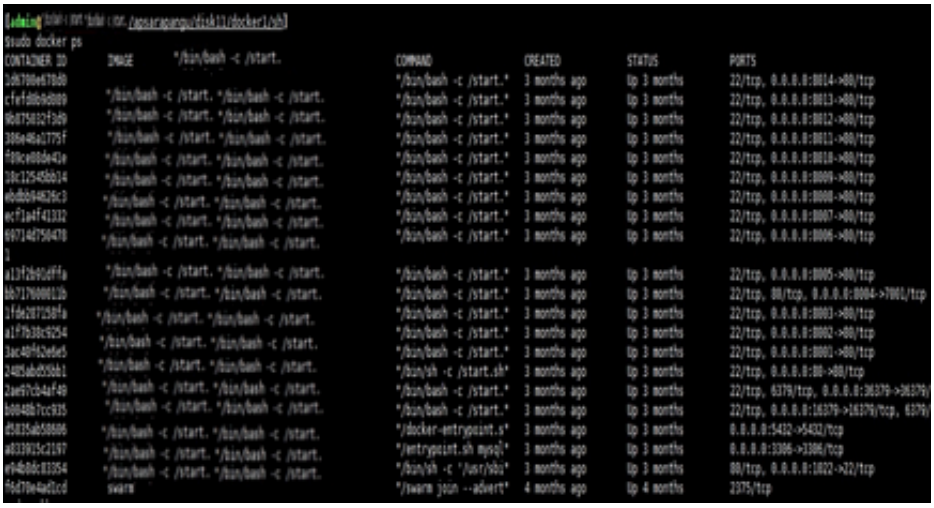

The internal structure of each service (except gateway) is shown in *Figure 3-3: Internal structure*.

Figure 3-3: Internal structure

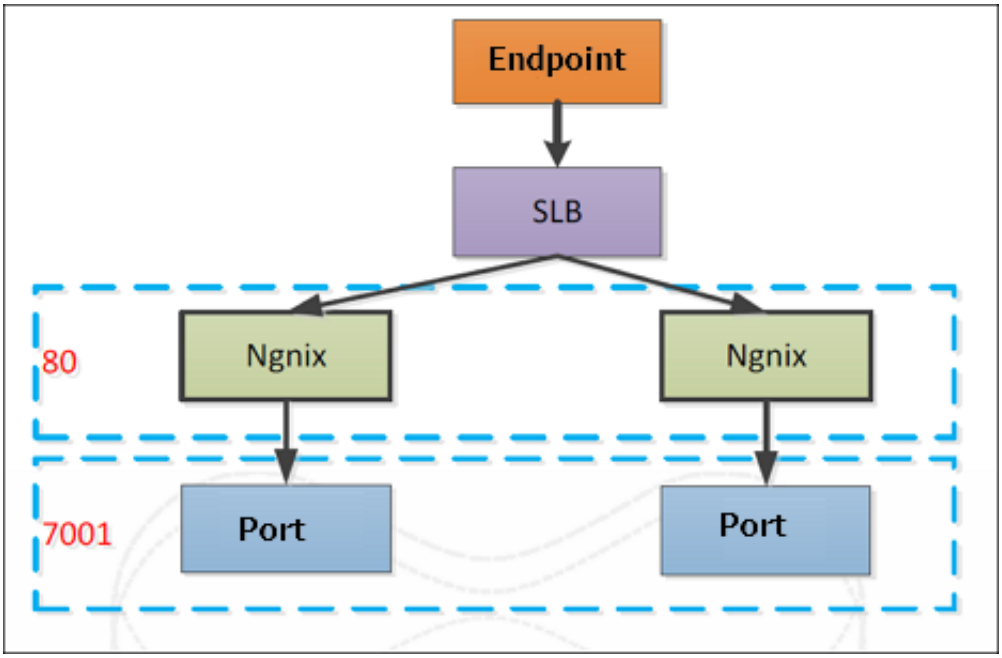

3.1.5 Directory of each service

base-biz-gateway

This service receives tasks from the development platform and the scheduling system, and proceeds to run the tasks.

- **logs: The directory stores operational logs of the gateway service.**
- **taskinfo: The directory stores the code and logs of tasks.**

- **target: The directory is the home directory of the gateway service, which includes service code, scripts for starting and stopping the service, and configuration files.**

cdp

This service handles data synchronization tasks.

- **logs: The directory stores operational logs of the cdp service.**
- **conf: The directory stores configuration files of the cdp service.**
- **bin: The directory stores the script for starting the service.**

The directory structure of other services

The following example shows the directory structure of the alisa service.

- **logs: The directory stores operational logs.**
- **conf: The directory stores configuration files.**
- **bin: The directory stores the script for starting the service.**

3.2 Common administration tools and commands

3.2.1 Find the container that runs the service

In Apsara Infrastructure Management Framework V3, select base from the project drop-down list, and then select BasicCluster.

Double-click baseBizApp in the lower part of the left-side navigation pane to view all services.

You can find the VM host that runs the service by double-clicking the service name. All services are deployed in containers. Therefore, you can run the `docker exec -it [container ID] bash` command to enter the container.

3.2.2 Cluster resource list

In Apsara Infrastructure Management Framework, select base from the project drop-down list. Select BasicCluster from the project list, move the pointer over the More icon next to BasicCluster, and select Dashboard from the menu to go to the Cluster Dashboard page.

On the Cluster Dashboard page, you can find the cluster resource list.

The Result column of the cluster resource list contains the details of each application. You can obtain the database logon information of a service from the Result column.

3.2.3 Commands to restart services

Enter the container that runs the service as an admin user, and then run the following commands to restart services.

Note:

Only admin users can run the following commands to restart the service.

- **To restart the base-biz-cdp service, run the** `/home/admin/cdp_server/bin/appctl.sh restart` **command.**
- **To restart the base-biz-gateway service, run the** `/home/admin/alisatasknode/target/alisatasknode/bin/serverctl restart` **command.**
- **To restart other services, run the** `/home/admin/base-biz-[application name]/bin/jbossctl restart` **command.**

  For example, to restart the base-biz-alisa service, run `/home/admin/base-biz-alisa/bin/jbossctl restart`**.**

3.2.4 View logs of a failed node

Log on to the DataWorks console. Click the DataWorks icon in the upper-left corner and select Operation Center from the menu.

On the Dashboard page of Operation Center, you can view the statistics of the running status of nodes and node instances. Click Failed in the upper-left corner to view the list of nodes that failed to run.

In the failed node list that appears, find the target node and choose More > View Runtime Log in the Actions column to view the runtime log of the node.

3.2.5 Rerun a task

If you want to rerun a failed task, select the task in the Administration console and click Rerun.

3.2.6 Terminate a task

If you want to terminate a running task, select the task in Administration, and then click Terminate.

Note:

Only running tasks can be terminated.

3.2.7 Filter tasks in the administration center

You can choose Administration > Task List and filter the tasks to maintain.

3.2.8 Commonly used Linux commands

top: You can run this command to view the system load.

The load average section shows the average system load over the last 5, 10, and 15 minutes. The system is overloaded if any of the average load divided by the number of logical CPUs is greater than five.

du: You can run this command to view the file size.

Run the `du -sh [file name]` command to view the size of the file. Run the `du -sh *` command to list the sizes of all files in the current directory.

ps: You can run this command to view system processes.

Run the `ps -ef` command to view all processes that are running in the system.

grep: You can run this command to print lines which match a specified string.

Run the following command to print log file lines that match a specified string.

```
grep ["string"] [file_name]
```

Run the following command to print first few lines in a log file.

```
grep -C [NUM] ["string"] [file_name]
```

Note:

C is uppercase, and NUM is the number of lines you want to print.

Run the following command to print last few log file lines that match a specified string.

```
grep -A [NUM] ["string"] [file_name]
```

kill: You can run this command to terminate a process.

Run `kill -9 [process ID]` to terminate the process.

Docker commands

`docker ps -a`: You can run this command to list all containers.

`docker logs [container ID]`: You can run this command to veiw the container logs.

`docker exec -it [container ID] bash`: You can run this command to enter the container.

3.2.9 View the slots usage of each resource group

Scenario: When a large amount of tasks are waiting for resources, you need to view the slots usage of each resource group.

Log on to the alisa database. In the Cluster Resource list, find and right-click the base-biz-alisa service that has a type of db, and then click Show More. The database logon address, username, and password are displayed. Connect to the database using MySQL statements.

Run the following command to view the top 10 longest running tasks.

```
select task_id,gateway,slot,create_time from alisa_task where status=2
 order by create_time limit 10;
```

Run the following command to view the top 10 tasks that occupy most slots.

```
select task_id,gateway,slot,create_time from alisa_task where status=2
 order by slot desc limit 10;
```

Run the following command to view the number of tasks that run in each slot. You can learn which tasks occupy a large number of slots.

```
select slot,count(*) from alisa_task where status=2 group by slot;
```

Run the following command to view the slots usage of each resource group.

```
select exec_target,sum(slot) from alisa_task where status=2 group by
exec_target;
```

Run the following command to view the status of each gateway node. If either the live value or the active_type value of a node is 1, the server does not work properly.

```
select * from alisa_node;
```

3.3 Process daily administration operations

3.3.1 Daily check

3.3.1.1 Check the service status and the basic information of the servers

Log on to the Apsara Infrastructure Management Framework console and select base from the project drop-down list. Hover over the vertical dots next to BasicCluster, and then click Dashboard. On the Dashboard page that appears, check whether servers are in GOOD status and services are in the desired status. If you find any issues, troubleshoot specific servers and services or contact an O&M engineer.

The blue column indicates the number of servers in GOOD status. If an orange column appears, errors occur on some servers.

3.3.1.2 Check the postgres database

1. **In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster. On the Service tab in the lower part of the left-side navigation pane, find baseBizApp.**
2. **Double-click baseBizApp, and then double-click psql-master.**
3. **Open the terminal window of the VM server.**
4. **Run the `docker ps|grep master` command to view the container ID.**
5. **Run the `docker exec -it [container ID] bash` command to enter the container.**
6. **Run the `psql -h127.0.0.1 -Uphoenix_prod -ddpphoenix -p3320` command, and enter the password pgsql to connect to the postgres database. Run the following statement in the database.**

   ```
   select to_char(to_timestamp(next_fire_time/1000), 'YYYY-MM-DD HH24:
   MI:SS') from qrtz_triggers;
   ```

   View the result.

   If the result contains 00:00:00 of the current day, the service is running properly . If not, ask for Alibaba Cloud technical support.
7. **Run the following command in the database.**

   ```
   select pid ,(now() - xact_start) as time , state,query from
   pg_stat_activity where state ! = 'idle' order by time desc;
   ```

   In the result, if the stat value is active, the service is running properly. If not, contact Alibaba Cloud Customer Support.

3.3.1.3 Check the status of each gateway server

1. **In the Apsara Infrastructure Management Framework console, open the dashboard page of BasicCluster.**
2. **In the cluster resource list, find and right-click the base-biz-alisa service that has a type of db, and then click Show More. The database logon address, username, and password are displayed.**

3. **Connect to the database using MySQL statements and run the following statement.**

```
Select * from alisa_node;
```

In the result that is returned, if either the active_type value or the live value is -1 or 0, the service does not run properly. In this case, contact Alibaba Cloud Customer Support.

3.3.1.4 Check the case test report

1. **Log on to the Apsara Infrastructure Management Framework console, and enter base in the search box on the Service (S) tab.**
2. **In the search result, select base-baseBizApp to open the Dashboard page of the service instance.**
3. **In the Service Monitoring List, click Details.**

   If the Failed Cases tab contains any record, contact Alibaba Cloud Customer Support.

3.3.2 View logs of the services

Logs of the gateway service are stored in `/home/admin/alisatasknode/logs/alisatasknode.log`**.**

Logs of the cdp services are stored in `/home/admin/cdp_server/logs/cdp_server.log`**.**

Logs of other services are stored in `/home/admin/base-biz-[service name]/base-biz-[service name].log`**.**

For example, the logs of the base-biz-phoenix service are stored in `/home/admin/base-biz-phoenix/base-biz-phoenix.log`**.**

3.3.3 Scale out the node cluster that runs the base-biz-gateway service

Prerequisites

Check whether the current environment meets the requirements for scale-out, such as disk space, file ownership and permissions, file execution path, software version, and any other necessary scale-out conditions.

- Before you scale out the BasicCluster cluster, make sure that it reaches the desired state and functions as expected.
- Save a screenshot of the key initial configurations for the cluster.
- Check for IP address conflicts. If you want to use a new buffer cluster for the scale-out, make sure that the IP addresses that Deployment Planner assigns to the servers in the cluster are not used in the current environment. This can avoid exceptions arising from IP address conflicts after the scale-out.
- Check the clone_mode parameter.

   **Note:**
  Apsara Infrastructure Management Framework of V3.3 and later versions supports cloning protection. Before scaling out the cluster, you need to set the clone_mode parameter to normal. After the scale-out process is complete, you need to set this parameter to block.

  Choose Apsara Infrastructure Management Framework > Operations > Cluster Operations > Global Clone Switch.

  In the Global Clone Switch dialog box that appears, select normal, and then click OK.

Procedure

Add a buffer cluster

 **Note:**
You can use idle servers in an existing buffer cluster for the scale-out operation, without adding a new buffer cluster. This method is applicable if the host, memory, CPU, and disk size of the idle servers match those of current servers that run the base-biz-gateway service. In this case, start from moving the idle servers to the default cluster.

In the scale-out procedure, use the actual parameter values and IP addresses instead of the specific parameter values in this guide.

 **Note:**
When you plan to scale out the cluster with Deployment Planner, make sure that the name of the new buffer cluster is different from that of any existing buffer cluster.

1. Copy and paste _tianji_imports to the /apsarapangu/disk3/u_disk/ directory of the ops1 server, and run the following command in the tianji_zhuque_sdk directory.

```
./tianji_zhuque_exchanger.py import --skip_packages -o ${desired
 state in the Apsara Infrastructure Management Framework} -c
tianji_dest.conf
```

2. Log on to the Apsara Infrastructure Management Framework. In the left-side navigation pane, locate the buffer cluster in the cluster list. Then, move the pointer over the More icon next to the buffer cluster, and select Cluster Operations and Maintenance Center from the shortcut menu to view the status of servers in the buffer cluster.
3. Run the following commands on the ops1 server to check scale-out information by calling API operations.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_clt machinestatus -c buffer --config clt2.conf
```

Scale in the buffer cluster

Note:

You can use the default cluster to scale out the cluster that runs base-biz-cdp and base-biz-gateway services.

1. Make sure that the value of the scalable tag value is true for the new buffer cluster.
2. Log on to the ops1 server, and then run the following commands to scale in the buffer cluster.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
```

```
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf (You can ignore the
message indicating that the symbolic link already exists.)
```

Scale-in command

```
./tianji_ops_tool.py contract_nc -c [buffer cluster name] -l [
hostname of the server to be removed], [hostname of the server to be
 removed],.... --config clt2.conf -s [SRG name]
```

Parameters

- **-c: the name of the buffer cluster that you scale in, which starts with buffer-cluster. This parameter is required.**
- **-l: a list of server hostnames that are included in the scale-in operation. Separate multiple hostnames with commas (,). This parameter is required.**
- **-s: the name of the SRG where the servers reside. You can find the SRG name in the machine_group.conf file of the buffer cluster. This parameter is required. If you want to remove the server, use this method to find the SRG name of the server.-config: the tianji_clt configuration file. This parameter is required.**

Note:

Chinese characters are not supported in the command line.

3. **Check whether the operation takes effect in the Apsara Infrastructure Management Framework.**

   On the Cluster Operations page, make sure that the servers are removed.
4. **Run the following commands to view the scaling information by calling API operations.**

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf (You can ignore the
message indicating that the symbolic link already exists.)
./tianji_clt machinestatus -c default --config clt2.conf
```

5. **On the Cluster Configuration page of the buffer cluster, check whether the server is deleted from the machine_group.conf file. If the server still exists in the machine_group.conf file, delete the server, and then submit a rolling task.**

Add servers to the BasicCluster cluster, and specify the SRG name where these servers reside.

1. **Check whether the clone mode for the BasicCluster cluster is set to Real Clone.**

2. **Run the following commands to perform scaling. A rolling task is triggered after running the command.**

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
```

```
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf (You can ignore the
message indicating that the symbolic link already exists.)
```

To add servers to the cluster that runs the base-biz-gateway service, run the following command:

```
./tianji_ops_tool.py expand_nc -c [the name of a BasicCluster
cluster] -s BaseGwGroup -l [machine1,machine2] --config clt2.conf
```

To add servers to the cluster that runs the base-biz-cdpgateway service, run the following command:

```
./tianji_ops_tool.py expand_nc -c [the name of a BasicCluster
cluster] -s BaseCdpGwGroup -l [machine1,machine2] --config clt2.conf
```

Parameters

- **-c: the name of a BasicCluster cluster. The name starts with BasicCluster.**
- **-l: a list of server hostnames that are included in the scale-out operation. Separate multiple hostnames with commas (,).**

> **Note:**
> **Chinese characters are not supported in the command line.**

3. **You can run the following command to call an API operation to check the cluster to which the servers belong and the uplink information of the servers. This process may take a few minutes.**

```
curl http://127.0.0.1:7070/api/v3/column/m. *?m.id=[machine hostname
]
```

4. **Log on to the OpsClone container, and run the following command to view the clone status:**

```
/home/tops/bin/python /root/opsbuild/bin/opsbuild.py acli list --
status=ALL -n 10000 | vim -
```

5. **Check the rolling task status in the Apsara Infrastructure Management Framework.**

Export the file that contains the information of desired state

After you complete the scale-out, export the file that contains the information of recent desired state to Deployment Planner. This ensures the success of subsequent scale-in and scale-out operations.

Verify the scale-out operation

1. **View the heartbeat log.**

   Open the terminal of the added server, log on to the gateway container, and then run the `tail -f /home/admin/alisatasknode/logs/heartbeat.log` command.

   If the heartbeat log is updated every five seconds, the heartbeat function is running as expected.
2. **Query the database.**

   In the Apsara Infrastructure Management Framework, open the dashboard page of the BasicCluster cluster. In the cluster resource list, find the base-biz-alisa service of the db type, right-click the result field, and then click Show More. You can find the database logon credentials. Connect to the database by using a MySQL command, and run the `select * from alisa_node;` command. The information of all gateway servers is displayed.

   Check the values of the live field and the active_type field for the added server. If both the two values are 1, the server is added.
3. **Verify that the server reaches the desired state on the Cluster Operation and Maintenance Center page.**

3.3.4 Scale in the base-biz-gateway cluster

Prerequisite

If a server in the base-biz-gateway cluster fails, you can repair and restart the server to redeploy the server.

If you want to remove a healthy server from the base-biz-gateway cluster, follow the instructions in this topic.

Note:
Before removing a healthy server, perform an on-site check to guarantee that the following conditions are met:

- **No business applications are running on the server.**

- **The hostname of the server is correct.**

Procedure

Perform checks before the scale-in

1. **Perform an on-site check.**

   Collect the detailed information of the server to be removed and the cluster that contains the server.
2. **Make sure that the value of the scalable tag is true for the service resource group (SRG) of the server to be removed. If the value is false, change it to true and submit a rolling task.**

   Log on to Apsara Infrastructure Management Framework. In the left-side navigation pane, choose BasicCluster > Cluster Configuration File > machine_group.conf. In this file, verify that the value of the scalable tag is true for the SRG of the server to be removed.

Stop the base-biz-gateway service

1. **Log on to the server to be removed and run the** `ps -ef|grep gateway` **command to obtain the container ID of the base-biz-gateway service.**
2. **Run the** `docker exec -it [container ID] bash` **command to enter the container.**
3. **Switch to the admin account and run the** `/home/admin/alisatasknode/target/alisatasknode/bin/servervtl stop` **command.**
4. **Run the** `ps -ef|grep java` **command to check whether any process is running on the server. If any process is running, run the** `kill -9 [process ID]` **command to terminate the process.**
5. **Delete the program directories from the server.**

   Clean up the disks of the server. Skip this step if you want to clone the server.

   ```
   #rm -rf /home/admin/*
   ```

   ```
   #rm -rf /opt/taobao/tbdpapp/
   ```

Move servers from the base-biz-gateway cluster to the default cluster in Apsara Infrastructure Management Framework

1. **Log on to the ops1 server and run the following commands to remove a server from the base-biz-gateway cluster:**

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf (After you run
this command, if a message appears indicating that a symbolic link
already exists, proceed with the next command.)
./tianji_ops_tool.py contract_nc -c [clusterName] -l [machineList]
 --config tianji_clt.conf -s [SRGname]
```

   The parameters are described as follows:

   - **-c: Required. Set this parameter to the name of the base cluster to be scaled in . To obtain the cluster name, choose Operations > Cluster Operations in the top navigation bar and select base from the Project drop-down list.**
   - **-l: Required. Set this parameter to the hostname of the server to be removed. Separate multiple hostnames with commas (,).**
   - **-s: Required. Set this parameter to the SRG name of the server to be removed . Find the machine_group.conf file among the configuration files of the base cluster. In this file, find the SRG of the server to be removed.**
   - **-config: Required. Set this parameter to tianji_clt.conf.**

2. **After you run the preceding command, check whether the scale-in operation succeeds in Apsara Infrastructure Management Framework.**

   Go to the Cluster Operation and Maintenance Center of the base cluster.

3. **On the Cluster Operation and Maintenance Center page, check the number of servers that are being removed.**
4. **Click the number next to Machine: in: to identify the status of the servers that are being removed.**

   If the scale-in operation succeeds, the number of servers that are being removed decreases to zero. Otherwise, check the server status on this page.

   You can follow the preceding steps to scale in a node cluster by moving servers to the default cluster in Apsara Infrastructure Management Framework. The following section describes how to remove servers from Apsara Infrastructure Management Framework.

Remove servers from Apsara Infrastructure Management Framework

1. **In the top navigation bar, choose Operations > Machine Operations.**

2. **On the Machine Operations page that appears, click Machine Online/Offline in the upper-right corner.**
3. **In the Machine Online/Offline dialog box that appears, click Remove Machine.**
4. **On the Remove Machine tab, search for the server to be removed by hostname in the left-side Enter Machine List section. You can only remove servers in the default cluster.**
5. **Confirm the information of the server and click Clear Machines to remove it.**

Verify the server removal result

1. **Check whether the server is moved to the default cluster in Apsara Infrastructure Management Framework.**

   In the top navigation bar, choose Operations > Machine Operations. On the Machine Operations page that appears, search for the target server by hostname and check whether it is in the default cluster.
2. **Check whether the server is removed from the default cluster.**

   In the top navigation bar, choose Operations > Machine Operations. On the Machine Operations page that appears, search for the server by hostname. If you cannot find the server in the search results, the server is removed.
3. **To check whether the server is removed from the default cluster, run the following command on the ops1 server to call the GetMachineInfo operation:**

   ```
   curl http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=$hostname
   ```

3.3.5 Restart the base-biz-alisa service

Procedure

1. **Click C on the top of the left-side navigation pane in the Apsara Infrastructure Management Framework. Select base from the project drop-down list, and select BasicCluster from the cluster list.**
2. **Then, double-click base-baseBizApp to view the base-baseBizApp service list. Find and double-click BaseBizAlisa to view servers that run the base-biz-alisa service.**
3. **Select one of the servers, and choose More > Terminal to open the Terminal Service page. In the upper part of the left-side navigation pane, click Add and then run the `docker ps|grep alisa` command to obtain the container ID.**

4. **Run the `docker exec -it [container ID] bash` command to enter the container.**
5. **Switch to the admin account and run the `/home/admin/base-biz-alisa/bin/jbossctl restart` command to restart the service.**

   If you see NGINX start Done in the command output, the base-biz-alisa service is restarted.

3.3.6 Restart the base-biz-phoenix service

Procedure

1. **Click C on the top of the left-side navigation pane in the Apsara Infrastructure Management Framework. Select base from the project drop-down list, and select BasicCluster from the cluster list. Then, double-click base-baseBizApp to view the base-baseBizApp service list. Find and double-click BaseBizPhoenix to view servers that run the base-biz-phoenix service.**
2. **Select one of the servers, and choose More > Terminal to open the Terminal Service page. In the upper part of the left-side navigation pane, click Add and then run the `docker ps|grep phoenix` command to obtain the container ID.**
3. **Run the `docker exec -it [container ID] bash` command to log on to the container.**
4. **Switch to the admin account and run the `/home/admin/base-biz-phoenix/bin/jbossctl restart` command to restart the service.**

   If you see NGINX start Done in the command output, the phoenix service is restarted.

3.3.7 Restart base-biz-tenant

Procedure

1. **In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster. In the lower part of the left-side navigation pane, double-click BaseBizTenant in the service list, and then the host that runs the service appears.**
2. **Open the terminal window of the vm host, and run `docker ps|grep phoenix` to find the container ID.**
3. **Run `docker exec -it [container ID] bash` to enter the container.**

4. **Switch to the admin account and run `/home/admin/base-biz-tenant/bin/jbossctl restart` to restart the service.**

   After you run the command, if the status is OK and the command output ends with NGINX start Done, the tenant service is restarted successfully.

3.3.8 Restart base-biz-gateway

Procedure

1. **In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and then select BasicCluster from the search result.**
2. **On the Service tab in the lower part of the left-side navigation pane, double-click base-baseBizApp, double-click BaseBizCdpGateway, and then the host that runs the service appears.**
3. **Open the terminal window of the host, and run the `docker ps|grep gateway` command to find the container ID.**
4. **Run the `docker exec -it [container ID] bash` command to enter the container.**
5. **Switch to the admin account, and run the `/home/admin/alisatasknode/target/alisatasknode/bin/serverctl restart` command to restart the service.**
6. **After the service is restarted, run the `ps -ef|grep java` command to check whether the process is started.**

   Note:
   This method can only be used where the gateway service is deployed in a Docker container.

For the service deployed on a physical server

If the service is deployed on a physical server, use the following method to restart the service.

1. **In the Apsara Infrastructure Management Framework console, open the Dashboard page of BasicCluster. In the cluster resource list, find and right-click the base-biz-alisa service that has a type of db, and then click Show More. The database logon address, username, and password are displayed.**

2. Run the `select * from alisa_node;` command in the database to view the information of all gateway servers, and use the node IP address to find and maintain the gateway server.
3. In the terminal window of the server, switch to the admin account, and then run the `/home/admin/alisatasknode/target/alisatasknode/bin/serverctl restart` command.

3.3.9 Restart the base-biz-api service

Procedure

1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster. On the Service tab in the lower part of the left-side navigation pane, double-click baseBizApp, double-click BaseBizCdpGateway, and then the host that runs the service appears.
2. Open the terminal window of the host, and run the `docker ps|grep gateway` command to find the container ID.
3. Run the `docker exec -it [container ID] bash` command to enter the container.
4. Switch to the admin account, and run the `/home/admin/alisatasknode/target/alisatasknode/bin/serverctl restart` command to restart the service.
5. After the service is restarted, run the `ps -ef|grep java` command to check whether the process is started.

> **Note:**
> The above method can only be used where the gateway service is deployed in a Docker container.

3.3.10 Restart the base-redis service

Procedure

1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster.
2. On the Service tab in the lower part of the left-side navigation pane, double-click base-baseBizApp, and you can find redis1 and redis2.
3. Open the terminal window of the VM host, and run the `docker ps|grep redis` command to find the container ID.

4. **Run the `docker exec -it [container ID] bash` command to enter the redis container.**
5. **Run the following commands to restart the redis service.**

```
/etc/init.d/ssh restart
```

```
/etc/init.d/redis-sentinel restart
```

3.3.11 Restart DataWorks Data Service

Procedure

1. **In the Apsara Infrastructure Management Framework console, search for dataworks-dataservice on the S tab.**
2. **Hover over the vertical dots next to BasicCluster, and then click Operations to open the Operations page to view the details of dataworks-dataservice.**
3. **Click the service instance name to open Service Instance Dashboard, and then find Service Role List.**
4. **If you want to restart the server, select BaseBizDataServiceServer#. If you want to restart the Web application, select BaseBizDataServiceWeb#. Hover over the vertical dots next to the service name, and then click Details to open the Service Role Dashboard page, and then find the virtual machine in the Server Information area.**
5. **Open the terminal window of the VM host, and run the `docker ps|grep dataservice` command to find the container ID.**
6. **Run the `docker exec -it [container ID] bash` command to enter the container.**
7. **Switch to the admin account, and run the `/home/admin/data-service-web/bin/jbossctl restart` command to restart the service.**

   If you are restarting the server, run the `/home/admin/data-service-server/bin/jbossctl restart` command.
8. **After you run the command, if the status is OK and the command output displays [OK] -- SUCCESS at the end, the dataservice service is restarted successfully.**

3.3.12 Restart DataWorks Data Management

Procedure

1. **In the Apsara Infrastructure Management Framework console, select base from the project drop-down list, and select BasicCluster.**
2. **In the lower part of the left-side navigation pane, double-click BaseBizAlisa in the service list, and then find BaseBizDmc.**
3. **Double-click BaseBizTenant, and then the host that runs the service appears.**
4. **Open the terminal window of the vm host, and run `docker ps|grep dmc` to find the container ID.**
5. **Run `docker exec -it [container ID] bash` to enter the container.**
6. **Switch to the admin account.**
7. **Run `/home/admin/base-biz-dmc/bin/jbossctl restart` to restart the service.**
8. **After you run the command, if the status is OK, and the command output ends with NGINX start Done, the Data Management (DMC) service is restarted successfully.**

3.4 Common issues and solutions

3.4.1 Nodes remain in the Pending (Resources) state

Symptom

After you log on to the DataWorks console and click Operation Center in the upper-right corner of the console, the following issue occurs on the Dashboard page that appears: The instances of many recurring nodes remain in the Pending (Resources) state for a long period of time.

Causes

The issue may occur due to any one of the following four reasons:

- **A gateway server is overloaded or offline and its status value is -1 in the database.**
- **The slots that handle concurrent jobs are fully occupied.**
- **The disk on a gateway server is full.**
- **The system time of servers in the base cluster is out of sync with the time of the Network Time Protocol (NTP) server.**

Solutions

To resolve this issue, follow these steps:

- **Check the status of a gateway server in the database.**
  1. **Log on to the database that hosts the base-biz-alisa service. In Apsara Stack V3 , you can find the database endpoint from the resource list of the base cluster in Apsara Infrastructure Management Framework.**
  2. **Run the `select * from alisa_node;` command to check the values of the active_type and live fields.**

     If the value of the live field is -1, the server is offline. If the value of the active_type field is -1, the server is overloaded.

     

     Note:

     In either case, use SSH to connect to the gateway server and then check the server load and heartbeat.

     - **Run the `tail -f/home/admin/alisatasknode/logs/heartbeat.log` command to check the heartbeat of the gateway server.**

       If the heartbeat log is updated every five seconds, the heartbeat is normal. Otherwise, check the configuration files for an error.
     - **Run the top command to display the load of the gateway server.**

       The status of the server becomes -1 in the database as a result of the high load. In this case, check whether the CPU and memory are overloaded. You can find out the high-load processes in the output of the top command.

       You can run the `ps -ef|greppid` command to view processes of the specified node and identify which process causes the high load. To terminate a process, run the `kill -9 [process ID]` command. After the load drops, check whether the status of the server resumes to 1.
- **Check whether the slots that handle concurrent jobs are fully occupied.**

  Log on to the database that hosts the base-biz-alisa service and run the following statements:

```
select group_name,max_slot from alisa_group where group_name like '%
default%';
```

```
select exec_target,sum(slot) from alisa_task where status=2 group by
 exec_target;
```

Compare the query results of the two statements.

- **The first statement returns the maximum number of slots that can be assigned in each resource group.**
- **The second statement returns the number of slots that are occupied in each resource group.**

If the query results of the two statements are the same or almost the same, all resource groups run out of slots. In this case, if a large number of nodes are running, the subsequent nodes do not run until the preceding nodes are completed.

Run the following statement to list the top 10 nodes that require the longest runtime:

```
select task_id,gateway,slot,create_time from alisa_task where status
=2 and create_time>current_time order by create_time desc limit 10;
```

Log on to the gateway server and run the `ps -ef|grep task_id` command.

> **Note:**
> **Replace task_id in this command with one of the node IDs that are returned by the preceding SELECT statement. You can obtain the node name from the command output.**

Then, you can troubleshoot the node. If required, run the `kill -9` command to terminate the node and release resources immediately. Otherwise, new nodes can start only after the existing nodes are completed.

- **Check whether the disk on a gateway server is full.**

  Log on to the gateway server and run the `df -h` command to check whether the disk attached to /home/admin is full. If the disk is full, run the `du-sh` command to identify the files in the /home/admin directory that consume a large amount of space. You can manually remove some large log files from the /home/admin/alisatasknode/taskinfo/ directory.

- **Check the system time of servers in the base cluster against the time of the NTP server.**

  1. **Log on to the database that hosts the base-biz-alisa service and run the `select now();` command to view the current time of the database.**
  2. **Check the system time of servers in the base cluster against the time of the database.**
  3. **Run the date command on the servers to check whether the system time of each server is synchronized with the time of the database. If the time difference is greater than 30 seconds, the base-biz-alisa service may fail. In this case, synchronize the system time of servers in the base cluster with the time of the NTP server.**

     

     Note:
     In Apsara Stack V3, you can find the servers of the base cluster in the service list in the Apsara Infrastructure Management Framework console and follow the proceeding steps to resolve the issue.

- **Rename the phoenix folder to change it to a .bak file and restart the base-biz-alisa service.**

  If the issue persists after you perform the preceding steps, run the following command on the gateway server:

  ```
  cd /home/admin/alisatasknode/taskinfo/prevDay/phoenix/
  ```

  

  Note:
  Replace prevDay in this command with the date of the previous day in the format YYYYMMDD, for example, 20180306.

  In this directory, run the `mkdir test` command. If the error message "Cannot create directory too many links" appears, the issue occurs because the number

of subdirectories in the directory has reached the maximum and you cannot create more subdirectories. To resolve this issue, follow these steps:

1. **Rename the /home/admin/alisatasknode/taskinfo/20180306/phoenix directory as /home/admin/alisatasknode/taskinfo/20180306/phoenix.bak.**
2. **Run the following command to restart the base-biz-alisa service:**

```
sudo su admin -c "/home/admin/alisatasknode/target/alisatasknode/
bin/serverctlrestart"
```

Note:

This is a rare problem which tends to occur when a gateway server uses the third extended (ext3) file system.

3.4.2 An out-of-memory (OOM) error occurs when synchronizing data from an Oracle database

Description

During the data synchronization from an Oracle database to MaxCompute or other platforms, an `java.lang.OutOfMemoryError: Java heap space` error is displayed in the task log.

Cause

This issue is often caused by a large volume of data in the data synchronization task , which causes a JVM OOM error.

Solution

Set a low fetchsize value.

Use MySQL statements to connect to the cdp database, and modify the template configuration of the Oracle reader plug-in by changing the fetchsize value from 1024 to 128. Run the following statement:

```
update t_plugin_template set template=replace(template,'1024','128')
where name='oracle' and type='reader';
```

Rerun the task after the fetchsize value is changed. To reset the fetchsize value, run the following statement:

```
update t_plugin_template set template=replace(template,'128','1024')
where name='oracle' and type='reader';
```

3.4.3 A task does not run at the specified time

Description

A periodic task does not run, and no data is displayed in the overview.

Solution

1. **Check whether periodic scheduling is enabled in this workspace.**

   On the Workspace Configuration page in Workspace Management, ensure that the periodic scheduling is enabled.
2. **If it is enabled, check whether the phoenix service runs properly.**

   Connect to the phoenix database and run the following statement.

   ```
   select to_char(to_timestamp(next_fire_time/1000),'YYYY-MM-DDHH24:MI:
   SS') from qrtz_triggers;
   ```

   If the output contains 00:00:00 of the next day, the service is running properly. If not, you need to check whether the time of the two base-biz-phoenix containers are different.

   If the two containers have the same system time, you need to switch to the admin account and run the `/home/admin/base-biz-phoenix/bin/jbossctl restart` **command to restart the phoenix service, and then check the time again.**
3. **After the time is corrected, you can run tasks that failed to run on the previous day.**

   Run the following command in either of the phoenix containers. Note that you can run this command only once.

   ```
   curl -v -H "Accept:application/json"-H "Content-type: application/
   json"-X POST -d'{"opCode":11,"opSEQ":12345,"opUser":"067605","name
   ```

```
":"SYSTEM","bizdate":"2017-04-2300:00:00","gmtdate":"2017-04-2400:00
:00"}' http://localhost:7001/engine/2.0/flow/create_unified_daily
```

Note:
bizdate refers to the previous day, and gmtdate refers to the current day. Modify the command if needed before running it.

3.4.4 The test service of base is not in the desired status

1. **On the S tab, select base-baseBizApp.**
2. **Select the cluster in the lower part of the left-side navigation pane, and then open the dashboard.**
3. **View the report of service monitoring.**

   Analyze the causes of the failed test based on the log.

3.4.5 The Data Management page does not display the number of tables and the usage of tables

Description

The Data Management page is blank.

Solution

1. **Log on to the Apsara Infrastructure Management Framework console, select odps from the project drop-down list, and then open the HybridOdpsCluster dashboard page.**
2. **Find the accesskey type base_admin service in the Cluster Resource area.**
3. **Right-click the result field, and click Show More to view the username and the password.**
4. **Log on to DataWorks.**

   Note:
   To log on to DataWorks, enter the domain name of base in the browser. By default, the domain name is ide.[your Apsara Stack second-level domain].
5. **Select the base_meta workspace, and go to Administration.**

   Rerun all failed tasks, and then check whether the Data Management page is displayed properly. If the task fails again, contact Alibaba Cloud Customer Support.

3.4.6 Logs are not automatically cleaned up

Description

Logs are not cleaned up automatically because of an error.

Solution

Follow the following steps to clean up the logs manually.

1. **Establish a terminal session to the VM.**
2. **Run the following command to clean up real-time analysis logs.**

```
find /home/admin/cai/logs/cronolog/ -mtime +7 -type f -exec rm -rf
 {} \;
find /home/admin/dw-realtime-analysis/logs/ -mtime +7 -type f -exec
rm -rf {} \;
```

3. **Run the following command to clean up base-biz-diide application logs.**

```
find /home/admin/cai/logs/cronolog/ -mtime +7 -type f -exec rm -rf
 {} \;
find /home/admin/base-biz-diide/logs/ -mtime +7 -type f -exec rm -rf
 {} \;
```

4. **Run the following command to clean up base-biz-cdp application logs.**

```
find /home/admin/cai/logs/cronolog/ -mtime +7 -type f -exec rm -rf
 {} \;
find /home/admin/base-biz-cdp/logs/ -mtime +7 -type f -exec rm -rf
 {} \;
```

3.4.7 The real-time analysis service is not in the desired status

Description

The real-time analysis service is not in the desired status.

Solution

1. **On the S tab, select dataworks-realtime.**
2. **Open the dashboard page of the cluster in the lower part of the left-side navigation pane.**
3. **View the report of service monitoring.**

   View the log to find out what caused the failed test.

4 Realtime Compute

4.1 Job status

4.1.1 Overview

StreamCompute allows you to view the real-time running information and instantaneous values of a job. You can also determine whether a job is running properly and whether the job performance meets expectations based on the job status.

4.1.2 Task status

A task can be in one of the following seven statuses: created, running, failed, completed, scheduling, canceling, and canceled. You can determine whether a job is running properly based on the task status.

4.1.3 Health score

To help you quickly locate job performance issues, Realtime Compute offers a health check feature.

If the health score of a job is lower than 60, lots of data has been piled up on the current task node and data processing performance needs to be optimized. To optimize the performance, you can enable *automatic resource configuration* **or** *manually reconfigure the resources*. **You can optimize the performance based on your business requirements.**

4.1.4 Job instantaneous values

Table 4-1: Job parameters

Name	Description
Consumed compute time	Indicates the computing performance of a job.

Name	Description
Input TPS	Indicates the number of data blocks that are read from the source per second. For Log Service, multiple data records can be included in a log group and the log group functions as the basic unit of measurement for data. In this scenario, the number of blocks indicates the number of log groups that are read from the source per second.
Input RPS	Indicates the number of data records that are read from the source table per second.
Output RPS	Indicates the number of data records that are written into result tables per second.
Input BPS	Indicates the data transmission rate per second, which is measured in bytes per second.
CPU usage	Indicates the CPU usage of the job.
Start time	Indicates the start time of the job.
Running duration	Indicates the duration during which the job has been running.

4.1.5 Running topology

A running topology shows the execution of the underlying computational logic of Realtime Compute. Each component corresponds to a task. Each dataflow starts with one or more sources and ends in one or more result tables. The dataflows resemble arbitrary directed acyclic graphs (DAGs). For more efficient distributed execution, Realtime Compute chains operator subtasks together into tasks if possible. Each task is executed by one thread.

Chaining operators together into tasks provides the following benefits:

- Reduces the thread-to-thread handover.
- Reduces the message serialization and deserialization.
- Reduces the data handover in the buffer zone.
- Increases overall throughput while decreasing latency.

An operator indicates the computational logic, and a task is a collection of multiple operators.

View mode

The underlying computational logic is visualized in a view, as shown in *Figure 4-1: View mode*, **to offer you a more intuitive display.**

Figure 4-1: View mode

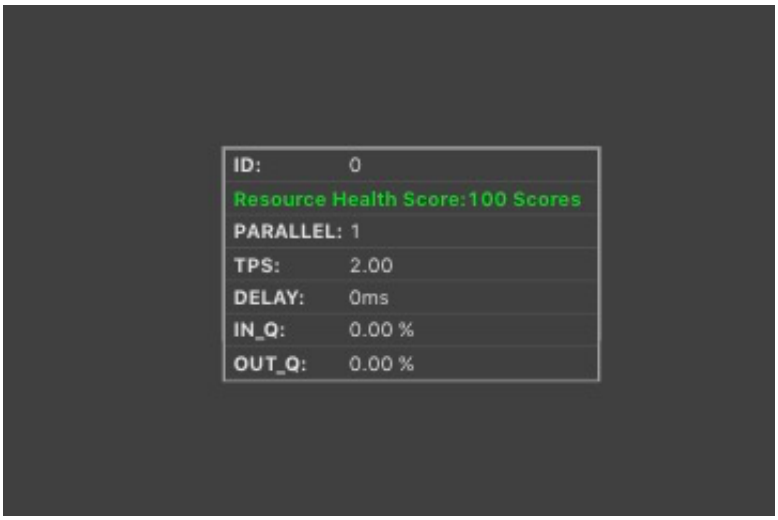

You can view the detailed information about a task by moving the pointer over the task. *Table 4-2: Parameter description* **describes the task parameters.**

Table 4-2: Parameter description

Parameter	Description
ID	The task ID in the running topology.
PARALLEL	The parallelism, which is the number of operator subtasks.
CPU	The CPU usage of a parallelism.
MEM	The memory usage of a parallelism.
TPS	The amount of data read from the inputs, which is measured in blocks per second.
LATENCY	The compute time consumed on the task node.
DELAY	The processing delay on the task node.
IN_Q	The percentage of input queues for the task node.

Parameter	Description
OUT_Q	The percentage of output queues for the task node.

You can also click a task node to access its details page. On this page, you can view its subtasks, as shown in *Figure 4-2: Task details page*.

Figure 4-2: Task details page

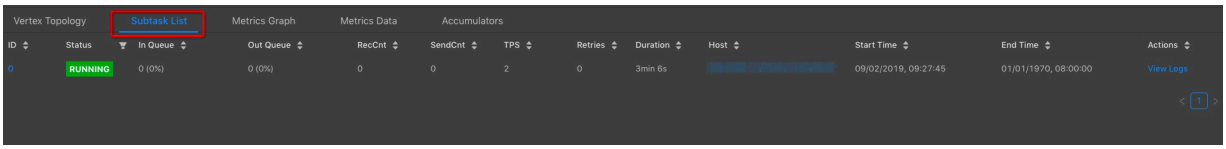

The Curve Charts tab provides curve charts to show the metrics of each task, as shown in *Figure 4-3: Curve charts for task metrics*.

Figure 4-3: Curve charts for task metrics

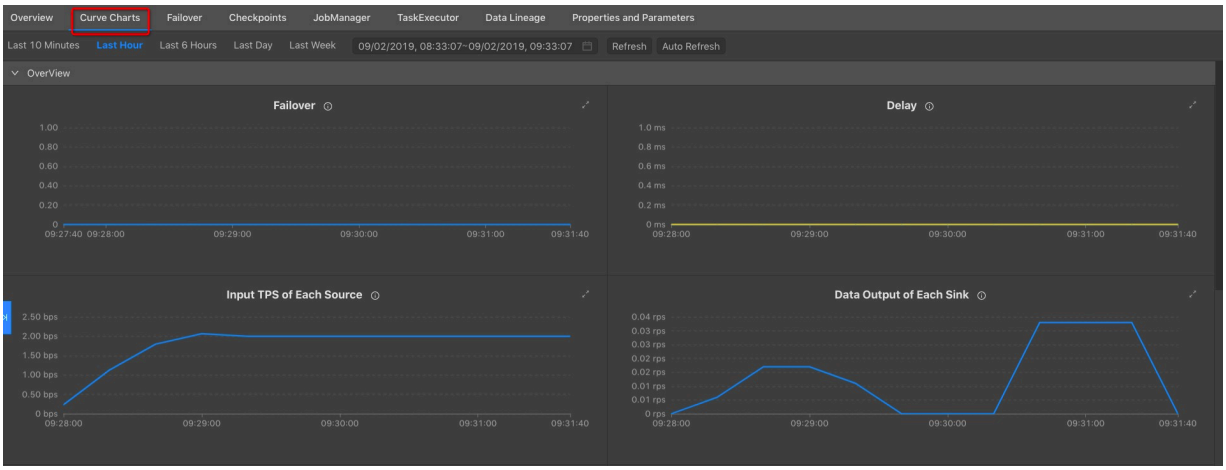

List mode

In addition to the view mode, Realtime Compute also allows you to view each task in the list mode, as shown in *Figure 4-4: List mode*.

Figure 4-4: List mode

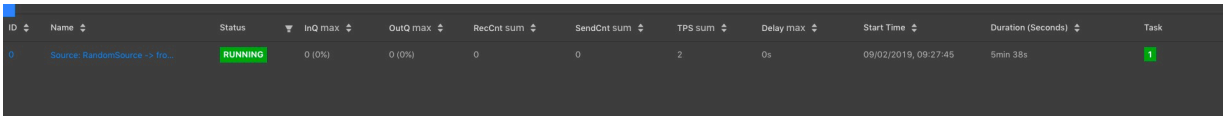

Table 4-3: Parameter description **describes the task parameters.**

Table 4-3: Parameter description

Parameter	Description
ID	The task ID in the running topology.
Name	The name of the task.
Status	The status of the task.
INQ max	The maximum percentage of input queues for the task node.
OUTQ max	The maximum percentage of output queues for the task node.
RecvCnt sum	The total amount of data that is received by the task node.
SendCnt sum	The total amount of data that is sent from the task node.
TPS sum	The total amount of data that is read from the inputs per second.
Delay max	The longest processing delay on the task node.
Task	The status of each parallelism on the task node.
StartTime	The start time of the task node.
Durations(s)	The running duration of the task node.

4.2 Curve charts

4.2.1 Overview

On the Curve Charts tab of the Realtime Compute development platform, you can view the key metrics of a job. This allows you to easily analyze the performance of a

job. Currently, we are working on intelligent and automatic diagnosis by developing in-depth intelligent analysis algorithms based on the job running information.

Figure 4-5: Curve Charts tab

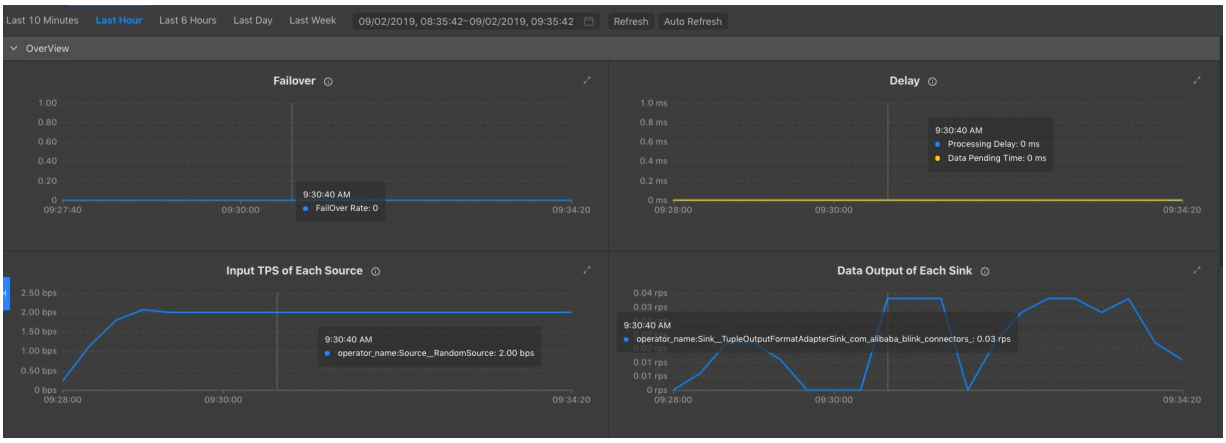

Note:

- **The metrics shown in this figure are displayed only when the job is in the running status.**
- **The metrics are asynchronously collected in the background, which results in delays. The metrics can be collected and displayed only after a job has been running for more than 1 minute.**

4.2.2 Overview

Failover rate

The failover rate indicates the percentage of the number of times that errors or exceptions occur on the current job. The failover rate curve allows you to easily analyze the issues of the current job.

Processing delay

The processing delay refers to the time interval between the current processing time and the time of reading data in the Realtime Compute service. If the time of reading data is not specified, the upstream DataHub or LogHub assigns the system timestamp to the data. The processing delay shows the timeliness of Realtime Compute end-to-end processing. For example, if the current processing time is 05:00 and the timestamp of the stored data is 01:00, the data to be processed was stored at 01:00, which is 4 hours earlier than the current processing time. In this scenario, the processing delay is 4 hours. The processing delay is used to monitor

the data processing progress. If the source data fails to flow into DataHub because of certain faults, the processing delay increases accordingly.If the source data fails to enter DataHub because of certain faults, the processing delay increases accordingly. The following table shows the processing delay.

Figure 4-6: Processing delay

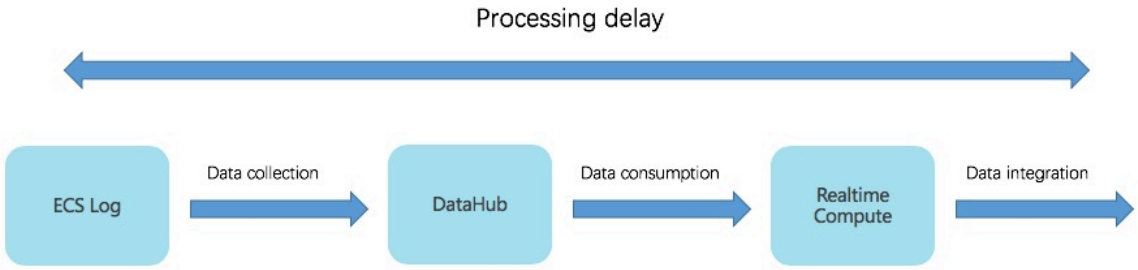

The processing delay can be categorized into the following three types:

- **Shortest delay: indicates the shortest processing delay of shards among data sources.**
- **Longest delay: indicates the longest processing delay of shards among data sources.**
- **Average delay: indicates the average processing delay of shards among data sources.**

Input TPS of each source

Realtime Compute collects statistics about the streaming data inputs of each Realtime Compute job to help you easily view the input transactions per second (TPS). The input TPS describes the amount of data that is read from the source table , which is measured in blocks per second. Unlike the TPS, records per second (RPS) indicates the number of data records parsed based on the data blocks that are read from the source table.

For example, in Log Service, N log groups are read per second and M log records are parsed based on the N log groups. In this example, the input TPS is N, and the output RPS is M.

Data outputs of each sink

Realtime Compute collects statistics about data outputs of each Realtime Compute job to help you easily view the output RPS.

Note:

The outputs show all data outputs rather than streaming data outputs.

As an administrator, if you find that no data output is detected, you must check whether data inputs from the upstream exist. You also need to check whether data outputs in the downstream exist.

Input RPS of each source

Realtime Compute collects statistics about the streaming data inputs of each Realtime Compute job to help you easily view the input data records per second . As an administrator, if you find that no data output is detected, you must check whether data inputs from the source exist.

Input BPS of each source

Realtime Compute collects statistics about the streaming data inputs of each Realtime Compute job to help you easily view the input data bytes per second (BPS). The input BPS indicates the amount of data that is read from the source table per second.

CPU usage

The CPU usage describes the CPU resources consumed by a Realtime Compute job. Realtime Compute provides the following two metrics to reflect the CPU usage:

- **The number of CPUs that you have applied for.**
- **The CPU usage of the current job at the specified time, which is shown in the curve chart.**

Memory usage

The memory usage describes the memory resources consumed by a Realtime Compute job. Realtime Compute provides the following two metrics to reflect the memory usage:

- **The size of memory space that you have applied for.**
- **The memory usage of the current job at the specified time, which is shown in the curve chart.**

Dirty data from each source

Realtime Compute allows you to view the dirty data from each source through the corresponding curve chart.

4.2.3 Advanced view

Realtime Compute offers a fault tolerance mechanism to consistently recover the state of data streaming applications. The central part of the fault tolerance mechanism is drawing consistent snapshots of the distributed data stream and the state. These snapshots act as consistent checkpoints to which the system can fall back when a failure occurs.

One of the core concepts of distributed snapshots is the barrier. Barriers are inserted into data streams and flow together with the data streams to the downstream. Barriers never overtake records, and the dataflow is strictly in line. A barrier separates the records in the data stream into two sets of records.

- **One set of records is sorted into the current snapshot.**
- **The other set of records is sorted into the next snapshot.**

Each barrier carries the ID of the snapshot that covers the records before the barrier. Barriers are a lightweight mechanism. They do not interrupt the flow of the stream. Multiple barriers from different snapshots can be in the stream at the same time. This means that multiple snapshots may be created concurrently.

Figure 4-7: Barriers

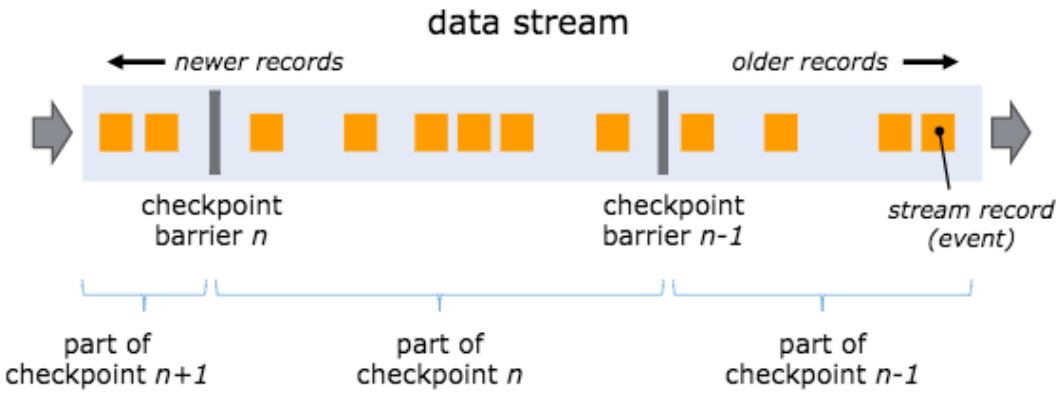

Stream barriers are injected into the dataflow at the stream sources. The point where the barrier for snapshot n is injected is the position in the source stream, up to which the snapshot covers the data. This point is indicated by Sn. The barriers

then flow downstream. When an intermediate operator has received a barrier for snapshot n from all of its input streams, it emits a barrier for snapshot n into all of its outgoing streams. When a sink operator has received the barrier n from all of its input streams, it acknowledges that snapshot n to the checkpoint coordinator. A sink operator is the end of a streaming directed acyclic graph (DAG). After all sinks have acknowledged a snapshot, the snapshot is considered completed.

Figure 4-8: Barrier mechanism

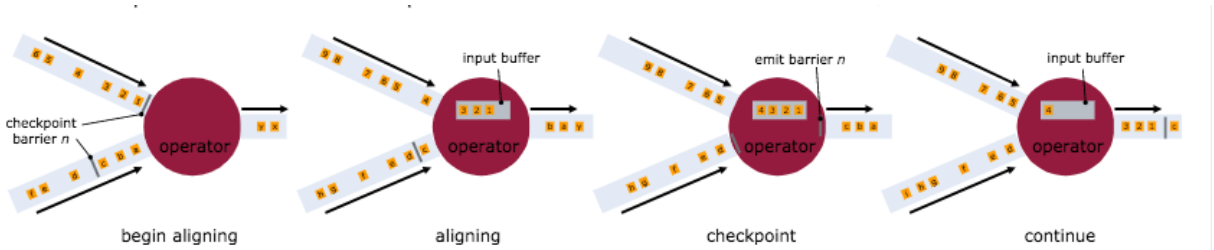

Checkpoint parameters

- **Checkpoint Duration**

  This parameter indicates the time spent on saving the state for each checkpoint. The duration is measured in milliseconds.
- **CheckpointSize**

  This parameter indicates the state size of a checkpoint, which is measured in MiB .
- **checkpointAlignmentTime**

  This parameter indicates the time spent on receiving and acknowledging the barrier n from all incoming streams. When a sink operator (the end of a streaming DAG) has received the barrier n from all of its input streams, it acknowledges that snapshot n to the checkpoint coordinator. After all sinks have acknowledged a snapshot, the snapshot is considered completed. The time consumed by the acknowledgement is included in the checkpoint alignment time .
- **CheckpointCount**
- **Get**

  This parameter indicates the longest time that a subtask spends on performing a GET operation on the RocksDB within a specified period.

- **Put**

  This parameter indicates the longest time that a subtask spends on performing a PUT operation on the RocksDB within a specified period.

- **Seek**

  This parameter indicates the longest time that a subtask spends on performing a SEEK operation on the RocksDB within a specified period.

- **State Size**

  This parameter indicates the state size of a job. If the size increases excessively fast, you need to check and resolve potential issues.

- **CMS GC Time**

  This parameter indicates the garbage collection (GC) time that is consumed by the underlying container that runs the job.

- **CMS GC Rate**

  This parameter indicates how often the garbage collection is performed in the underlying container that runs the job.

4.2.4 Processing delay

Top 15 source subtasks with the longest processing delay

This metric describes the processing delays of each parallelism of a source.

4.2.5 Throughput

Task Input TPS

This indicates the data inputs of all tasks for the job.

Task Output TPS

This indicates the data outputs of all tasks for the job.

4.2.6 Queue

Input Queue Usage

This indicates the input data queues of all tasks for the job.

Output Queue Usage

This indicates the output data queues of all tasks for the job.

4.2.7 Tracing

The available parameters for advanced users are as follows:

- **Time Used In Processing Per Second**

  This parameter indicates the time that a task spends on processing the data of each second.

- **Time Used In Waiting Output Per Second**

  This parameter indicates the time that a task spends on waiting for outputs of each second.

- **TaskLatency**

  This parameter indicates the computing delay of each task for a job. This delay is indicated by the interval between the time when data enters a task node and the time when data processing is completed on the task node. You can view the delay from the corresponding curve chart.

- **WaitOutput**

  This parameter indicates the time that a task spends on waiting for outputs. You can view the waiting time from the corresponding curve chart.

- **WaitInput**

  This parameter indicates the time that a task spends on waiting for inputs. You can view the waiting time from the corresponding curve chart.

- **Source Latency**

  This parameter indicates the delay of each parallelism for a data source. You can view the delay from the corresponding curve chart.

4.2.8 Process

Process MEM Rss

You can view the memory usage of each process from the curve chart.

Memory NonHeap Used

You can view the non-heap memory usage of each process from the curve chart.

CPU Usage

You can view the CPU usage of each process from the curve chart.

4.2.9 JVM

Memory Heap Used

This indicates the Java Virtual Machine (JVM) heap memory usage of the job.

Memory NonHeap Used

This indicates the JVM non-heap memory usage of the job.

Threads Count

This indicates the number of threads for the job.

GC (CMS)

This indicates how often garbage collection (GC) is performed for the job.

4.3 FailOver

On the FailOver tab of the Realtime Compute development platform, you can check whether the job is running properly.

Latest FailOver

On the Latest FailOver tab, you can view the running errors of the job.

FailOver History

On the FailOver History tab, you can view the previous running errors of the job.

4.4 CheckPoints

Realtime Compute offers a fault tolerance mechanism to consistently recover the state of data streaming applications. The central part of the fault tolerance mechanism is drawing consistent snapshots of the distributed data stream and the state. These snapshots act as consistent checkpoints to which the system can fall back when a failure occurs.

Completed Checkpoints

On this tab, you can view the checkpoints that have been created. *Table 4-4: Parameter description* **describes the parameters for the created checkpoints.**

Table 4-4: Parameter description

Parameter	Description
ID	The ID of the checkpoint.
StartTime	The start time when the checkpoint is created.
Durations(ms)	The time that is spent on creating the checkpoint.

Task Latest Completed Checkpoint

On this tab, you can view the detailed information about the latest checkpoint. *Table 4-5: Parameter description* describes the parameters for the latest checkpoint.

Table 4-5: Parameter description

Parameter	Description
SubTask ID	The ID of the subtask.
State Size	The state size of the checkpoint.
Durations(ms)	The time that is spent on creating the checkpoint.

4.5 JobManager

After a Realtime Compute cluster is started, one JobManager and one or more TaskManagers are started. A client submits jobs to the JobManager, and the JobManager assigns the tasks of jobs to TaskManagers. During task execution , TaskManagers report the heartbeats and statistics to the JobManager. The TaskManagers exchange the data streams.

Similar to Storm Nimbus, a JobManager schedules jobs and functions as a coordinator to create checkpoints for tasks. A JobManager receives resources, such as jobs and JAR files, from a client. Then, the JobManager generates an optimized execution plan based on these resources and assigns tasks to TaskManagers.

4.6 TaskExecutor

After a Realtime Compute cluster is started, one JobManager and one or more TaskManagers are started. A client submits jobs to the JobManager, and the

JobManager assigns the tasks of jobs to TaskManagers. During task execution , TaskManagers report the heartbeats and statistics to the JobManager. The TaskManagers exchange the data streams.

The number of slots is specified before a TaskManager is started. A TaskManage r executes each task in each slot, and each task can be considered as a thread. A TaskManager receives tasks from the JobManager, and then establishes a Netty connection with its upstream to receive and process data.

TaskExecutor shows the detailed information about each TaskManager.

4.7 Data lineage

On the Data Lineage tab of the Realtime Compute development platform, you can view the dependencies of a job, including its relationship with its source table and result table. The topology on this tab allows you to easily and clearly view the complex dependencies of a job.

Data sampling

Realtime Compute provides the data sampling feature for source tables and result tables of jobs. The data to be sampled is the same as the data on the Developmen t page. The data sampling feature allows you to check data at any time on the Administration page to facilitate fault locating. In the topology, click the button on the right side of the table name to enable the data sampling feature.

4.8 Properties and Parameters

The Properties and Parameters page provides detailed information about the current job, including the current running information and running history.

Job Code

On this tab page, you can preview the SQL job. You can also click Edit Job to go to the Development page.

Resource Configuration

On this tab page, you can view the resources that have been configured for the current job, including the CPU, memory, and parallelism.

Properties

On this tab page, you can view the basic running information of the current job. *Table 4-6: Job properties* **describes the basic job properties that are displayed on this tab page.**

Table 4-6: Job properties

No.	Field and Description
1	Job Name: indicates the name of the job.
2	Job ID: indicates the ID of the job.
3	Referenced Resources: indicates the resources that are referenced by the job.
4	Execution Engine: indicates the engine of the job.
5	Last Operated By: indicates the user who last operates the job.
6	Action: indicates the action that is last performed.
7	Created By: indicates the user who creates the job.
8	Created At: indicates the time when the job is created.
9	Last Modified By: indicates the user who last modifies the job.
10	Last Modified At: indicates the time when the job is last modified.

Running Parameters

On this tab page, you can view the underlying checkpoints, start time, and running parameters of the job.

History

On this tab page, you can view the detailed information about all versions of the job , including the start time, end time, and the user who operates the job.

Parameters

On this tab page, you can view additional job parameters, such as the separator used in the debugging file.

4.9 Improve performance by automatic configuration

Background

To improve user experience, the Realtime Compute team offers the automatic configuration feature.

This feature optimizes the configuration of resources and parallelism for each operator of a job when the operators, data sources, and data sinks of Realtime Compute jobs are running properly. The automatic configuration feature also helps to globally improve job performance and handle issues, such as low throughput and data piling up on the upstream nodes.

This feature can optimize job performance in the following scenarios, but cannot address the performance bottlenecks of Realtime Compute jobs. To address the performance bottlenecks, contact the technical support team of Apsara Stack or your administrator.

- **The performance of data sources or sinks needs to be improved.**
  - **Data sources. For example, the partitions of a DataHub source table are insufficient or the message queue (MQ) throughput is low. In this scenario, you need to increase the partitions of the data source table.**
  - **Data sinks. For example, an ApsaraDB for RDS deadlock occurs.**
- **The performance of user-defined extensions (UDXs) needs to be improved, such as user-defined functions (UDFs), user-defined aggregation functions (UDAFs), and user-defined table functions (UDTFs).**

Improve the performance of a new job

1. **After you write the SQL statements and the statements pass the syntax check, click Publish. The Publish New Version dialog box appears.**
   - **If you select Automatic CU Configuration, the automatic configuration algorithm determines the number of compute units (CUs) based on the system default configuration to optimize resource configuration. If automatic configuration is performed for the first time, the algorithm determines**

the number based on empirical values. We recommend that you perform automatic configuration after a job has been running for more than 10 minutes. In most cases, the resources are optimally allocated after you perform automatic configuration three to five times.

- If you select Use Latest Manually Configured Resources, the latest saved resource configuration is used, no matter whether the resources are configured automatically or manually.

> **Note:**
> We recommend that you select Automatic CU Configuration. If you are performing automatic configuration for the first time, use the default number of CUs.

2. After you configure the resources for the job, click Next to check the data, and then click Publish to publish the job. Note that the default number of CUs is used.
3. Start the job.

   The following section uses an example to describe how to improve job performance by using the automatic configuration feature. In this example, the default number of CUs for the job is 71. Note that the job must run for more than 10 minutes before automatic configuration is performed.
4. Improve the performance by using the automatic configuration feature.

   > **Note:**
   > Optimize the resource configuration. In this example, you can specify 40 CUs and select automatic configuration. You can increase or decrease the number of CUs based on the job running information. We recommend that you set the number of CUs to a value that is greater than or equal to 1 and 50% of the default number of CUs. For example, if the default number of CUs is 71, we recommend that you set the number of CUs to a value that is greater than or equal to 35.5 (71 CUs × 50% = 35.5 CUs). If the specified CUs cannot meet the throughput requirements of the job, you can increase the number of CUs. We recommend that you increase the number of CUs by more than 30% each time. For example, if 10 CUs were last specified, you can specify 13 CUs. If the result does not meet your needs, you can perform automatic configuration for several times and increase or decrease the number of CUs based on the job running information.

5. View the result of performance improvement.

> **Note:**
> If you are performing automatic configuration on a new job, do not select Use Latest Manually Configured Resources. Otherwise, an error message is displayed.

Improve the performance of an existing job

The following figure shows the procedure for improving the performance of an existing job.

Figure 4-9: Procedure

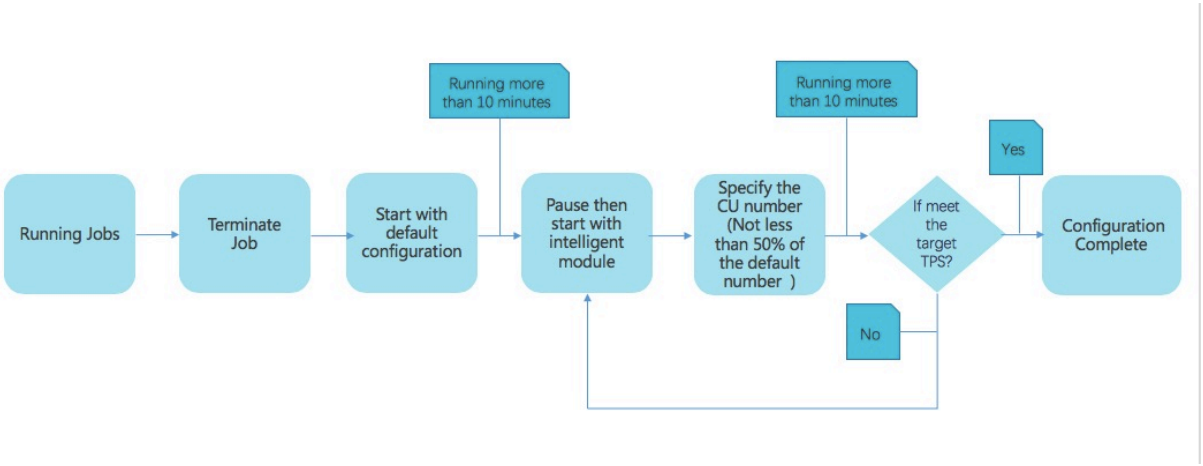

Before performing automatic configuration on an existing job, check whether stateful operations are involved. This is because the saved state information of a job may be cleared during the automatic configuration process.

If a job is changed, for example, an SQL statement is modified or the Realtime Compute version is changed, the automatic configuration may fail. The reason is that these changes may lead to topology changes, which further results in certain issues. These issues include: 1. Curve charts do not display the latest data. 2. The state cannot be used for fault tolerance. In this scenario, resource configuration cannot be optimized based on the job running history, and an error occurs while performing automatic configuration. To perform automatic configuration on a job that has been changed, perform steps 1 to 5 from the previous section on the changed job.

To perform automatic configuration on an existing job, perform steps 1 to 5 for a new job, and resume the job with the latest configuration.

Restrictions

The result of automatic configuration may be compromised in the following scenarios:

- **The target job runs only for a short period. In this scenario, the useful information collected during data sampling is limited. This reduces the accuracy of the results calculated based on the automatic configuration algorithm. We recommend that you perform automatic configuration after the curves, such as Input RPS of Each Source, have been stable for 2 to 3 minutes.**
- **The target job has encountered a failover. This reduces the accuracy of the results calculated based on the automatic configuration algorithm. We recommend that you check and handle failovers before performing automatic configuration.**
- **Only a small amount of data is available for the target job. This reduces the accuracy of the results calculated based on the automatic configuration algorithm. We recommend that you trace more historical data.**
- **The configuration obtained by using the automatic configuration feature is not always better than that from the last time. If the automatic configuration feature cannot meet your needs for improving the job performance,** *manually configure the resources*.

Recommendations

- **Before performing automatic configuration on a job, ensure that the job has been running stably and properly for more than 10 minutes. This helps to collect accurate job running information for the automatic configuration algorithm.**
- **You may need to perform automatic configuration for three to five times before the job performance is significantly improved.**
- **Before performing automatic configuration on a job, you can specify the start offset to read data from the past or even pile up large amounts of data for a job. This allows you to easily and quickly view performance improvement results.**

Method for determining the effectiveness of automatic configuration

The automatic configuration feature for Realtime Compute is enabled based on a JSON configuration file. After performing automatic configuration, you can view the JSON configuration file to check whether this feature is running properly. You can view the JSON configuration file on either of the following tabs:

- **Configuration Comparison tab under Properties on the Development page.**
- **Resource Configuration tab under Properties and Parameters on the Administration page.**

The configurations in the JSON file are described as follows:

```
"autoConfig" : {
    "goal": {  // The goal of automatic configuration.
        "maxResourceUnits": 10000.0,  // The maximum number of CUs for
 a Blink job. The value cannot be modified, and you can ignore this
item when checking whether the feature is running properly.
        "targetResoureUnits": 20.0  // The number of CUs, which you
have specified.
    },
    "result" : {  // The results of automatic configuration. We
recommend that you pay special attention to this item.
      "scalingAction" : "ScaleToTargetResource",  // The action of
automatic configuration. *
      "allocatedResourceUnits" : 18.5, // The total resources.
      "allocatedCpuCores" : 18.5,      // The total CPU cores.
      "allocatedMemoryInMB" : 40960    // The total memory size.
      "messages" : "xxxx"  // We recommend that you pay special
attention to the displayed messages. *
    }
}
```

- **The InitialScale value of the scalingAction parameter indicates that automatic configuration is performed for the first time. The ScaleToTargetResource value of the scalingAction parameter indicates that automatic configuration is not performed for the first time.**
- **If no message is displayed, the automatic configuration feature is running properly. If certain messages are displayed, you need to analyze the messages and handle the issues. Messages are categorized into the following two types:**
  - **Warning: Messages of this type indicate that the feature is running properly , but you need to pay attention to potential issues, such as insufficient partitions of source tables.**
  - **Error or exception: Messages of this type indicate that the automatic configuration has failed. The following error message is usually displayed:**

Previous job statistics and configuration will be used. The automatic configuration for a job fails in either of the following two scenarios:

- **The job or Realtime Compute version has been modified. In this scenario, the previous running information cannot be used for automatic configurat ion.**
- **The "xxxException" message is displayed. This message indicates that an error occurred while performing automatic configuration. You can analyze the error based on the job running information and logs. If the available information cannot help you to analyze the error, contact our technical support and development teams.**

Error messages

IllegalStateException:

If the following error messages are displayed, the state cannot be used for fault tolerance. To resolve this issue, terminate the target job, clear its state, and then specify the start offset to re-read the data.

If you cannot migrate the target job to a backup node and you are concerned that online business may be interrupted, click Properties on the right side of the Development page, roll back the target job to the earlier version, and then specify the start offset to re-read the data during off-peak hours.

```
java.lang.IllegalStateException: Could not initialize keyed state
backend.
    at org.apache.flink.streaming.api.operators.AbstractStreamOperator
.initKeyedState(AbstractStreamOperator.java:687)
    at org.apache.flink.streaming.api.operators.AbstractStreamOperator
.initializeState(AbstractStreamOperator.java:275)
    at org.apache.flink.streaming.runtime.tasks.StreamTask.initialize
Operators(StreamTask.java:870)
    at org.apache.flink.streaming.runtime.tasks.StreamTask.initialize
State(StreamTask.java:856)
    at org.apache.flink.streaming.runtime.tasks.StreamTask.invoke(
StreamTask.java:292)
    at org.apache.flink.runtime.taskmanager.Task.run(Task.java:762)
    at java.lang.Thread.run(Thread.java:834)
Caused by: org.apache.flink.api.common.typeutils.SerializationExcepti
on: Cannot serialize/deserialize the object.
    at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSe
condaryState.deserializeStateEntry(AbstractRocksDBRawSecondaryState.
java:167)
    at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRe
storeOperation.restoreRawStateData(RocksDBIncrementalRestoreOperation.
java:425)
    at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRe
storeOperation.restore(RocksDBIncrementalRestoreOperation.java:119)
    at com.alibaba.blink.contrib.streaming.state.RocksDBKeyedStateBac
kend.restore(RocksDBKeyedStateBackend.java:216)
```

```
    at org.apache.flink.streaming.api.operators.AbstractStreamOperator
.createKeyedStateBackend(AbstractStreamOperator.java:986)
    at org.apache.flink.streaming.api.operators.AbstractStreamOperator
.initKeyedState(AbstractStreamOperator.java:675)
    ... 6 more
Caused by: java.io.EOFException
    at java.io.DataInputStream.readUnsignedByte(DataInputStream.java:
290)
    at org.apache.flink.types.StringValue.readString(StringValue.java:
770)
    at org.apache.flink.api.common.typeutils.base.StringSerializer.
deserialize(StringSerializer.java:69)
    at org.apache.flink.api.common.typeutils.base.StringSerializer.
deserialize(StringSerializer.java:28)
    at org.apache.flink.api.java.typeutils.runtime.RowSerializer.
deserialize(RowSerializer.java:169)
    at org.apache.flink.api.java.typeutils.runtime.RowSerializer.
deserialize(RowSerializer.java:38)
    at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSe
condaryState.deserializeStateEntry(AbstractRocksDBRawSecondaryState.
java:162)
    ... 11 more
```

4.10 Improve performance by manual configuration

4.10.1 Overview

You can manually configure resources to improve job performance using one of the following methods:

- **Optimize resource configuration. You can modify the resources to improve the performance by reconfiguring parameters, such as parallelism, core, and heap_memory.**
- **Improve performance based on job parameter settings. You can specify the job parameters such as miniBatch to improve the performance.**
- **Improve upstream and downstream data storage based on parameter settings. You can specify related parameters to optimize the upstream and downstream storage for a job.**

More details about these three methods are described in the following sections . After parameters are reconfigured to improve the performance of a job, the corresponding job must be re-published and started or resumed to apply the new configuration. The detailed process is provided in the following section.

4.10.2 Optimize resource configuration

Problem analysis

1. **The percentage of input queues at task node 2 has reached 100%. Large amounts of data have piled up at task node 2, which results in the piling up of output queues at task node 1 in the upstream.**
2. **You can click task node 2 and find the subtask where the percentage of input queues has reached 100%. Then, click View TaskExecutor Logs to view the detailed information.**
3. **On the TaskExecutor page, you can view the CPU and memory usage. You can increase the number of CPU cores and expand the memory based on the current usage to handle the large amounts of data that have piled up.**

Performance improvement

1. **On the Development page of the StreamCompute development platform, click Properties.**
2. **Click Configure Resources to enter the page for editing resources.**
3. **Find the group (if any) or operator that corresponds to task node 2. You can modify the parameters of one operator or multiple operators in one group at a time.**
   - **Modify the parameters of multiple operators in a group.**
   - **Modify the parameters of an operator.**
4. **After modifying the parameters, click Apply and Close the Page in the upper-right corner of the page.**

> **Note:**
>
> **If the resources of a group have increased but the performance is not improved, you need to separately analyze each operator in the group and find the abnormal operators. Then, you can modify the resources for the abnormal operators for performance tuning. To separately analyze each operator in a group, click the target operator and change the value of its chainingStrategy parameter to HEAD. If the value is already set to HEAD, click the next operator and change the value of its chainingStrategy parameter to HEAD. The values of the chainingStrategy parameter are as follows:**
>
> - **ALWAYS: indicates that operators are chained into a group.**

- **NEVER: indicates that operators are not chained.**
- **HEAD: indicates that operators are separated from a group.**

Principles and recommendations

You can modify the following parameters:

- **parallelism**
  - **Source**

    Set the parallelism parameter based on the number of source table partitions . For example, if the number of sources is 16, set the parallelism parameter to 16, 8, or 4. Note that the maximum value is 16.
  - **Operators**

    Set the parallelism parameter based on the estimated queries per second (QPS). For tasks with low QPS, set the parallelism parameter for the operators to the same value as that for the sources. For tasks with high QPS, set the parallelism parameter to a larger value, such as 64, 128, or 256.
  - **Sinks**

    Set the parallelism parameter for the sinks to a value that is two or three times the number of downstream sink partitions. However, if the specified parallelism limit is exceeded, a write timeout or failure occurs. For example, if the number of downstream sinks is 16, the maximum value of the parallelism parameter for sinks is 48.
- **core**

  This parameter indicates the number of CPU cores. The default value is 0.1. Set this parameter based on CPU usage. We recommend that you set this parameter to a value whose reciprocal is an integer. The recommended value is 0.25.
- **heap_memory**

  This parameter indicates the heap memory size, whose default value is 256 MB. The value is determined based on the actual memory usage. You can click GROUP on the resource editing page to modify the preceding parameters.
- **For the task nodes that use the GROUP BY operator, you can configure the state_size parameter.**

  This parameter specifies the state size. The default value is 0. If the operator state is used, set the state_size parameter to 1. In this case, the corresponding

job requests extra memory for this operator. The extra memory is used to store the state. If the state_size parameter is not set to 1, the corresponding job may be killed by YARN.

Note:

- **The state_size parameter must be set to 1 for the following operators: GROUP BY, JOIN, OVER, and WINDOW.**
- **General users only need to focus on the core, parallelism, and heap_memory parameters.**
- **For each job, we recommend that you assign 4 GB memory for each core.**

4.10.3 Improve performance based on job parameter settings

The miniBatch parameter can be used to optimize only GROUP BY operators. During the streaming data processing of Flink SQL, the state is read each time a data record arrives for processing, which consumes large amounts of high I/O resources. After the miniBatch parameter is set, the state is read only once for data records with the same key, and the output contains only the latest data record. This reduces the frequency of reading state and minimizes the data output updates. The settings of the miniBatch parameter are described as follows:

1. **The allowed delay for a job.**

```
blink.miniBatch.allowLatencyMs=5000
```

2. **The size of a batch.**

```
blink.miniBatch.size=1000
```

4.10.4 Optimize upstream and downstream data storage based on parameter settings

In Realtime Compute, each data record can trigger read and write operations on source and result tables. This brings considerable challenges for upstream and downstream data storage performance. To address these challenges, you can set batch size parameters to specify the number of data records that are read from a source table or written into a result table at a time. The following table describes the available batch size parameters.

Table 4-7: Parameter description

Object	Parameter	Description	Value
DataHub source table	batchReadSize	The number of data records that are read at a time.	Optional. Default value: 10.
DataHub result table	batchSize	The number of data records that are written at a time.	Optional. Default value: 300.
Log Service source table	batchGetSize	The number of log groups that are read at a time.	Optional. Default value: 10.
ApsaraDB for RDS result table	batchSize	The number of data records that are written at a time.	Optional. Default value: 50.

Note:

To complete batch data read and write settings, add the above parameters to the parameter list WITH in DDL statements for the corresponding data storage. For example, add `batchReadSize='500'` to the parameter list WITH in DDL statements for the DataHub source table.

4.10.5 Apply new configuration

After resources are reconfigured for a job, you must restart or resume the job to apply the new configuration. Perform the following operations:

1. **Publish the job of the new version. In the Publish New Version dialog box, select Use Latest Configuration.**
2. **Suspend the job.**
3. **Resume the job. In the Resume Job dialog box, select Resume with Latest Configuration. Otherwise, the resource configuration cannot take effect.**
4. **After resuming the job, choose Administration > Overview > Vertex Topology to check whether the new configuration has taken effect.**

Note:

We do not recommend that you terminate and restart a job to apply the new configuration. After a job is terminated, its status is cleared. In this case, the computing result may be inconsistent with the result that is obtained if you suspend and resume the job.

4.10.6 Concepts

- **Global**

  isChainingEnabled: indicates whether the chaining is enabled. Use the default value (true).

- **Nodes**
  - **id: specifies the unique ID of a node. The ID is automatically generated and does not need to be changed.**
  - **uid: specifies the UID of a node, which is used to calculate the operator ID. If this parameter is not specified, the value of id is used.**
  - **pact: specifies the type of a node, such as the data source, operator, and data sink. Use the default value.**
  - **name: specifies the name of a node, which can be customized.**
  - **slotSharingGroup: Use the default value.**
  - **chainingStrategy: specifies the chaining strategy. The options include HEAD, ALWAYS, and NEVER. Use the default value.**
  - **parallelism: specifies the number of parallel subtasks. The default value is 1. You can increase the value based on the data volume.**
  - **core: specifies the number of CPU cores. The default value is 0.1. The value is configured based on the CPU usage. We recommend that you set this parameter to a value whose reciprocal is an integer. The recommended value is 0.25.**
  - **heap_memory: specifies the heap memory size. The default value is 256 MB. Set this parameter based on the memory usage.**
  - **direct_memory: specifies the JVM non-heap memory size. We recommend that you use the default value (0).**
  - **native_memory: specifies the JVM non-heap memory size for the Java Native Interface (JNI). The default value is 0. The recommended value is 10 MB.**

- **Chain**

  A Flink SQL task is a directed acyclic graph (DAG) that contains many nodes, which are also known as operators. Some upstream and downstream operators can be combined to form a chain when they are running. The CPU capacity of a chain is set to the maximum CPU capacity among operators in the chain. The memory size of a chain is set to the total memory size of operators in the chain. For example, after node 1 (256 MB, 0.2 cores), node 2 (128 MB, 0.5 cores), and node 3 (128 MB, 0.25 cores) are combined to form a chain, the CPU capacity of the chain is 0.5 cores and the memory is 512 MB. The prerequisite for chaining operators is that the operators to be chained must have the same parallelism settings. However, some operators cannot be chained, such as GROUP BY operators. We recommend that you chain operators to improve the efficiency of network transmission.

5 Apsara Big Data Manager (ABM)

5.1 Routine maintenance

5.1.1 Perform routine maintenance

You can perform routine maintenance on Apsara Big Data Manager (ABM) through the Apsara Infrastructure Management Framework console.

Apsara Infrastructure Management Framework

1. **Log on to the ABM console.**
2. **Click in the upper-left corner, and then click TIANJI to log on to the Apsara Infrastructure Management Framework console.**
3. **Go to the Clusters page in the ABM console and verify that all containers are in their final state.**
4. **Go to the Dashboard page in the ABM console and verify that alerts have not been generated.**

Metrics and alert handling

- **Hardware monitoring**

  The system retains logs for 30 days and automatically deletes old logs. If a disk alert is triggered when a large volume of logs exhaust disk space, contact technical support.
- **System exception**

  If a system exception is thrown during the inspection, handle the exception in the ABM console. If the exception message is unclear, contact technical support.

5.1.2 View the ABM operating status

ABM monitors its own health and operating metrics. You need to regularly handle ABM alerts and view ABM operating metrics to evaluate system downtime risks in the future.

View ABM operating metrics

In ABM, click O&M on the top and click Clusters. The Overview tab appears.

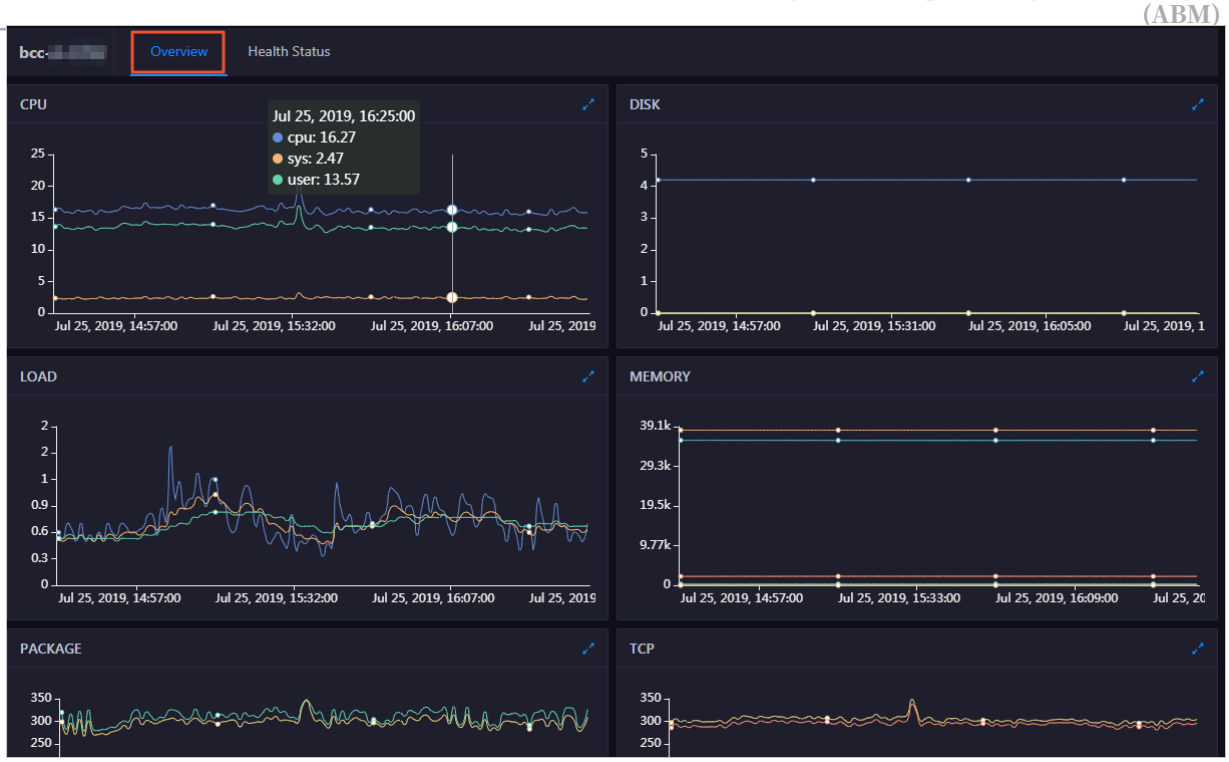

The Overview tab displays tendency charts for cluster metrics, including the CPU, memory, disk, load, package, TCP, and disk root directory usage. You need to regularly view and record these metrics to evaluate system downtime risks in the future.

Handle ABM alerts

ABM cluster alerts are classified into Critical, Warning, and Exception alerts. You need to handle these alerts in time, especially Critical and Warning alerts.

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

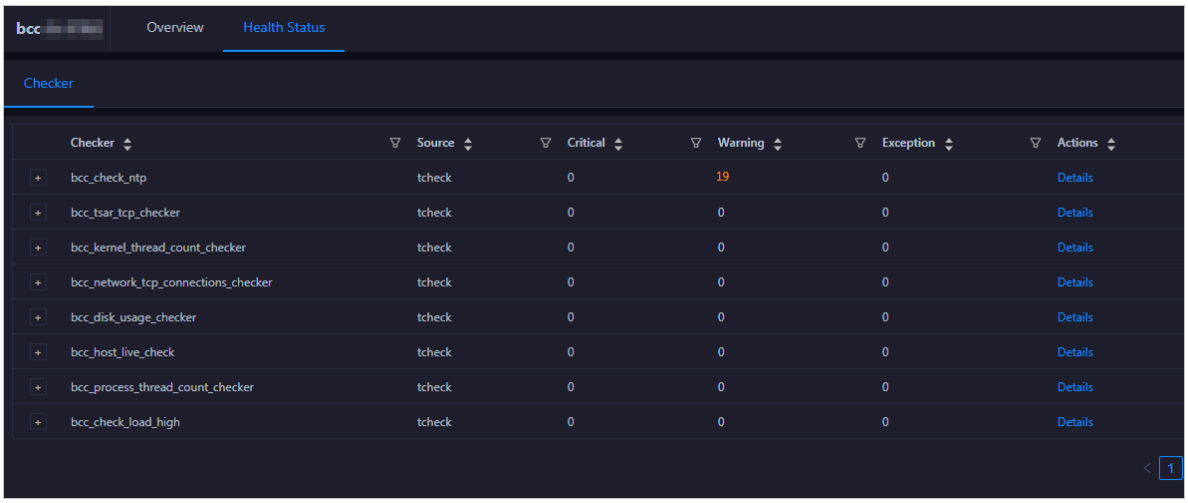

The Health Status tab displays all check items and the alerts that were generated during the check.

2. Click the Fold icon for a check item with alerts. All hosts on which the check item was performed appear.

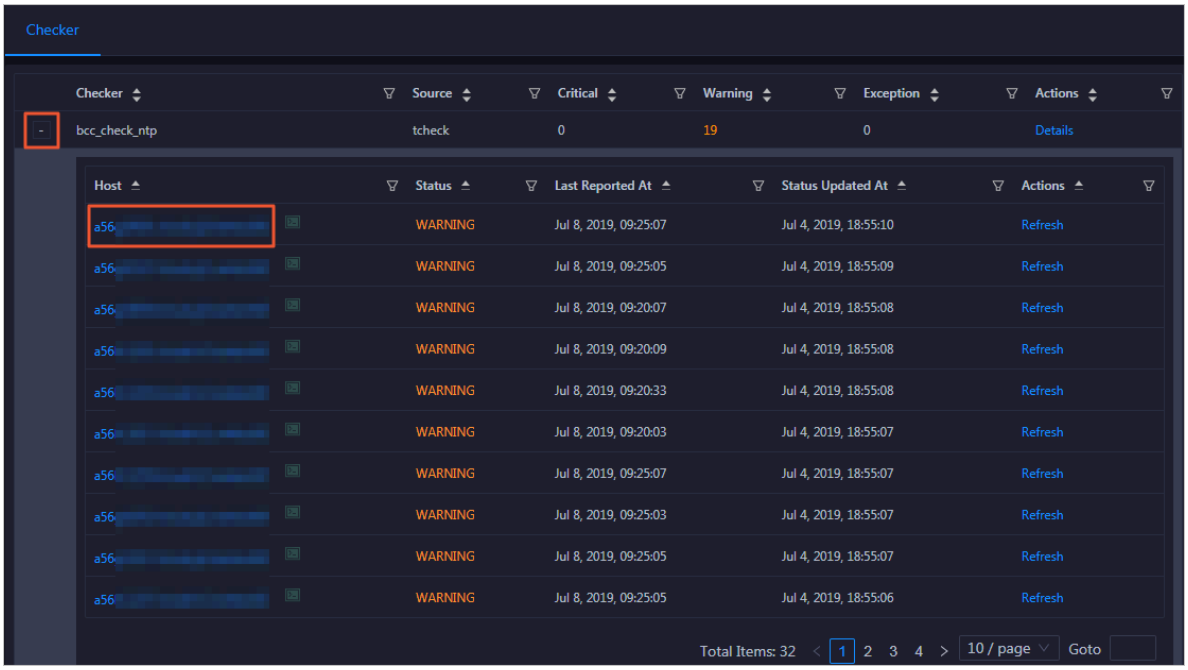

3. Click a host. In the dialog box that appears, click Details for an alert. The alert cause appears on the right.

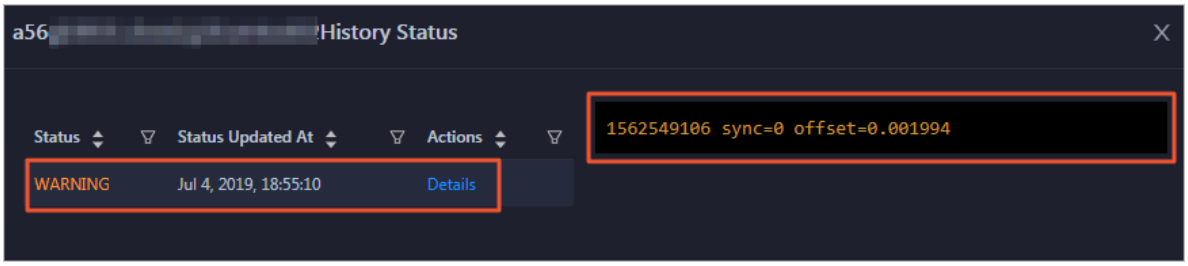

4. Click Details for a check item with an alert and view the fix method for the alert in the dialog box that appears.

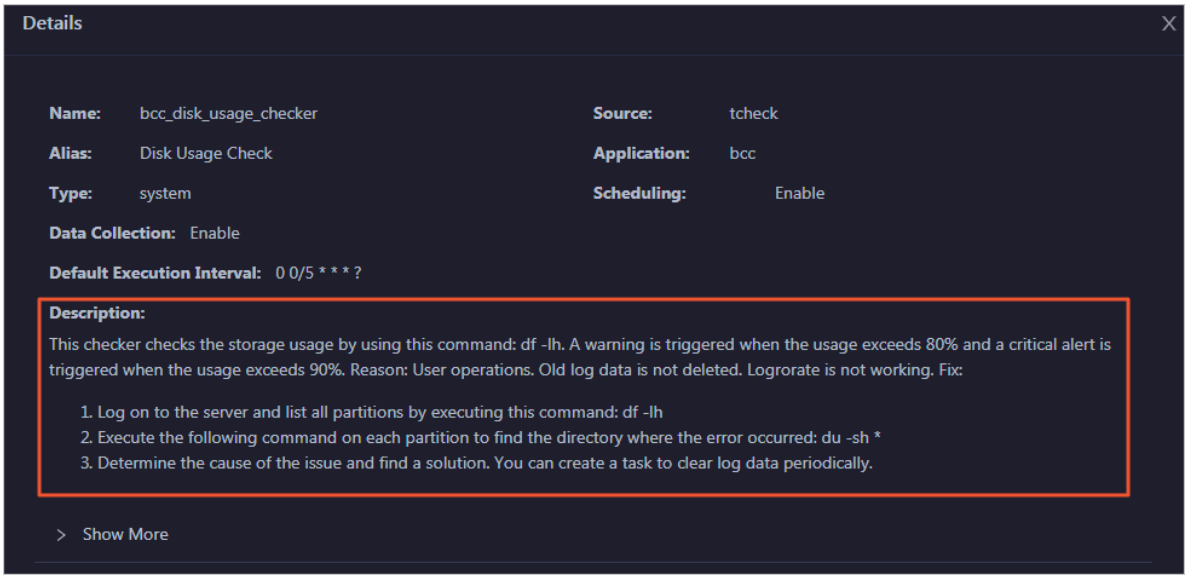

5. Handle the alert based on the fix method.

   You may need to log on to the host when handling the alert. For more information, see *Log on to a host*.

6. After the alert is handled, click Refresh for the host to perform the check again in real time. In this way, you can check whether the alert is cleared.

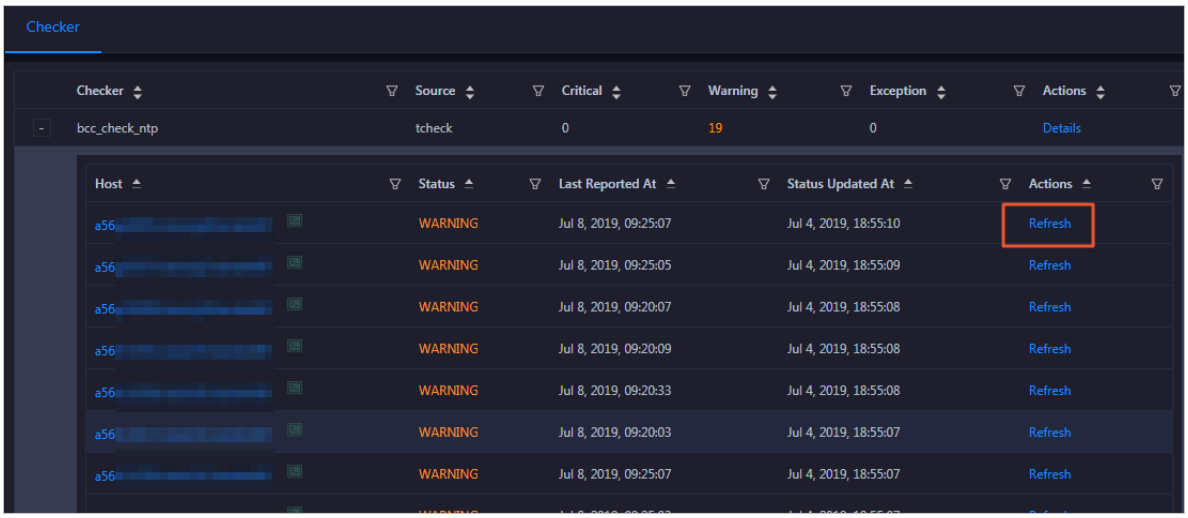

Log on to a host

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 the Fold icon for a check item.

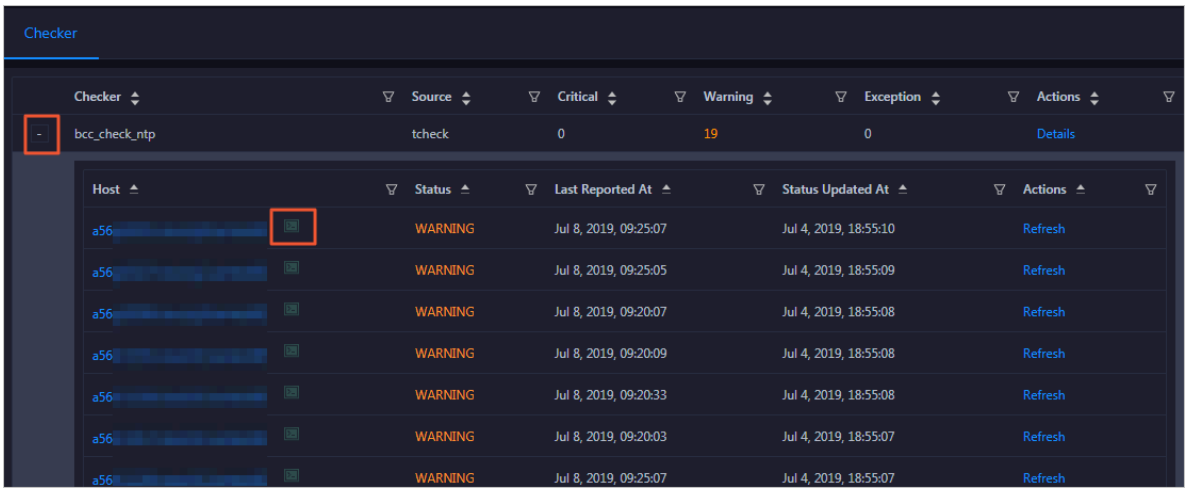

2. Click the Logon icon for a host. The TerminalService page appears.

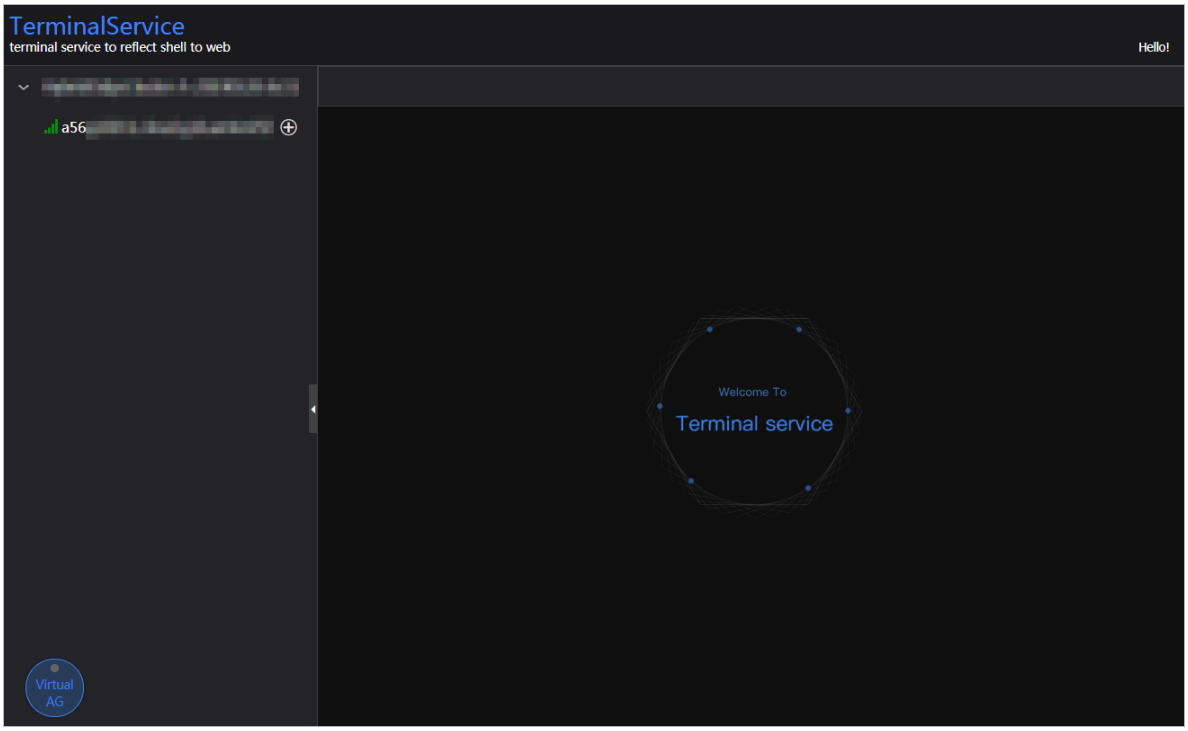

3. On the TerminalService page, select the host on the left to log on to it.

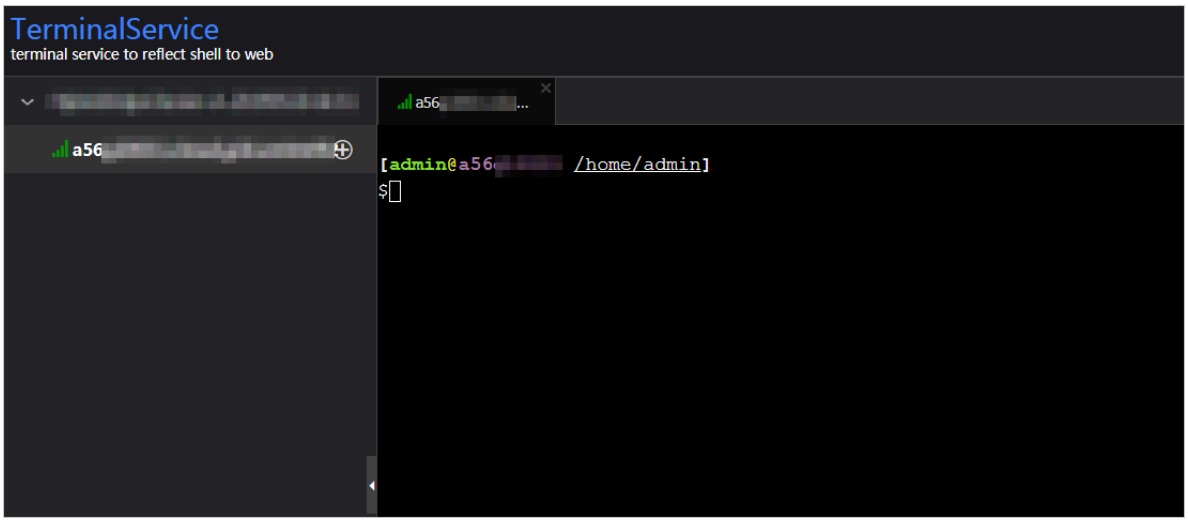

5.1.3 Troubleshooting

Common failures

- **Logon failure**

  If you failed to log on to ABM, clear the cache and cookies in your web browser, and then try again.

  Based on the logon failure message that appears, check whether the following issues exist:

  - **The password that you entered is incorrect.**
  - **Your account has been locked.**
  - **Your account has been disabled.**
- **Other failures**

  Contact technical support.

5.2 Backup and restore

Back up data

ABM uses a high-availability database. You do not need to manually back up data. To obtain full backup data, contact technical support.

Restore data

You do not need to restore data for ABM.

6 Machine Learning Platform for AI

6.1 Query server and application information

6.1.1 Apsara Stack Machine Learning Platform for AI

6.1.1.1 Query server information

Machine Learning Platform for AI is deployed based on Apsara Infrastructure Management Framework. Its application information and database information can be found by accessing the corresponding Apsara Infrastructure Management Framework address. This topic describes how to query server information.

Procedure

1. **Open Chrome and ensure that you can access internal services through the network.**
2. **Enter the username and password to log on to the homepage of Apsara Infrastructure Management Framework.**

   Notice:
   To avoid logon failures, make sure that your network is connected and the hosts have been bound.
3. **Click the C and search for pai. Hover over the dots next to PaiCluster-20170630-c34b, and choose Dashboard from the shortcut menu.**
4. **Query the server information for an application, such as the server where PaiDmscloud runs.**
   a) **Find the service instance and click Details. The instance detail page appears.**
   b) **Find the role list and click Details. The role detail page appears.**
   c) **The IP address of the server is displayed in the server information list. You can click Terminal to manage the server on the terminal management page.**

6.1.1.2 Log on to a server

Machine Learning Platform for AI is deployed based on Apsara Infrastructure Management Framework. Its application information and database information

can be found by accessing the corresponding Apsara Infrastructure Management Framework address. This topic describes how to log on to a server.

Context

Each module is deployed on two servers with the same application package and configuration. You can log on to the back-end server through the server IP address and perform operations.

Procedure

1. **Ensure that the network is connected and the IP address of the jump server has been obtained.**
2. **Log on to the jump server.**
3. **Switch to the root account.**
4. **All applications are deployed by using a Docker container. You can run the following command to view the current container:**

   ```
   sudo docker ps
   ```

5. **Run the following command to go to the container:**

   ```
   sudo docker exec -ti container_id /bin/bash
   ```

   The application log is stored in the *`/home/admin/logs/${app}`* **path.**

6.1.1.3 Query configurations

Prerequisites

Log on to the server of an application and go to the application container to view the configuration of the application.

Procedure

1. **View the application configuration in the** *`/home/admin/{app}/target/exploded/BOOTINF/classes/application.yml`* **file.**

   

   Note:
   In the preceding file path, {app} indicates the component name, such as pai-dms.

2. **View the application log in the** */home/admin/pai-dms/* **path.**

   The pai-dms.log, err_pai-dms.log, java.log, and access.log files store the application log, error log, framework log, and access log, respectively.

3. **Log on to a database.**

   a) **Query the database information of modules from the Dashboard cluster information of Apsara Infrastructure Management Framework. Find the corresponding result column and click More from the shortcut menu to obtain db_host, db_port, db_name, db_password, and db_user of the application.**

   b) **Run the following command to connect to the database through a MySQL client:**

   ```
   mysql -h$db_host -P$db_port -u$db_user -p$ db_password -D$ db_name
   ```

6.1.1.4 Restart an application service

The application structures and directories of the PaiCap, PaiDmscloud, and PaiJcs modules are almost the same. You can restart an application service in either of the following ways:

- **Log on to the container and run the following command to restart the service:**

  ```
  sudo -u admin /home/admin/pai-dms/bin/appclt.sh restart
  ```

- **Run the following command on the server to restart the container:**

  ```
  sudo docker restart $container_id
  ```

Run the following command to check whether the service is restarted:

```
curl localhost/status.taobao
```

6.1.2 Online model service

6.1.2.1 Query online model service information

Check the online model service status

Online model services are deployed in the Kubernetes cluster. Log on to the master node in the Kubernetes cluster and run the following command to query the service deployment status:

```
kubectl get pod -n eas-system
```

If no errors occur, all pods in the `STATUS` **column display** `Running`**.**

If not, run the following command to perform troubleshooting:

```
kubectl describe pod ${pod_name} -n eas-system
```

View the online model service configurations

1. **Log on to the homepage of Apsara Infrastructure Management Framework.**
2. **Click the C tab and search for pai. Hover over the dots next to the PAI cluster, and choose Dashboard from the shortcut menu.**
3. **Search for the `eas-sentinel` role and log on to the VM from the terminal.**
4. **Run the `docker ps |grep eas-sentinel` command to view the ID of the container for the sentinel.**
5. **Run the `docker logs ${sentinelcontainerid}` command to view the output log, which contains the configuration information of the online model service.**

6.1.2.2 Log on to the online model service container

Prerequisites

Ensure that the network is connected and the IP address of the jump server has been obtained.

Procedure

1. **Log on to the jump server.**
2. **Switch to the root account.**
3. **All applications are deployed with a container. Run the following command to log on to the current pod:**

   ```
   kubectl exec -ti ${pod_name} -n ${pod_namespace} - bash
   ```

6.1.2.3 Restart a pod

Procedure

1. **Log on to the master node in the Kubernetes cluster.**
2. **Run the `kubectl get` command to find the corresponding *`pod name`*.**
3. **Run the following command to restart the pod:**

   ```
   kubectl delete ${pod_name}
   ```

6.1.3 GPU cluster and task information

6.1.3.1 Query GPU cluster information

Prerequisites

You must deploy the deep learning service before querying the GPU cluster information. Deep learning tasks are performed in the GPU cluster. You can log on to ApsaraAG of the GPU cluster to query the GPU cluster status.

Procedure

1. **Log on to the homepage of Apsara Infrastructure Management Framework.**
2. **Click the C tab and search for `PAIGPU`. Move the pointer over the dots next to the deployed GPU cluster. Log on to the cluster O&M center.**
3. **Select `pai-deep_learning` from the Service drop-down list and `ApsaraAG#` from the Service Role drop-down list. Log on to the VM from the terminal.**
4. **Run the `r ttrl` command to view all GPU workers in the current GPU cluster.**

   If the `Other` column displays `FUXI_GPU:200`, the worker has two GPUs. If the column displays `FUXI_GPU:800`, the worker has eight GPUs.

6.1.3.2 Query GPU task information

Procedure

1. **Perform steps 1 through 3 in *Query GPU cluster information* and log on to ApsaraAG of the GPU cluster.**
2. **Run the `r al` command to view the running tasks.**
3. **Run the `r wwl WorkItemName` command to view the status of a task and the allocated resources.**

   `WorkItemName`: specifies the values in the first column displayed by the `r al` command.
4. **Run the `r cru` command to view the resources allocated to the current cluster, including CPU, memory, and FUXI_GPU resources.**
5. **Notice:**

Use caution when performing this step.

Run the `r jstop WorkItemName` command to stop a Fuxi task.

`WorkItemName`: specifies the values in the first column displayed by the `r al` command.

6.2 Maintenance and troubleshooting

6.2.1 Machine Learning Platform for AI maintenance

6.2.1.1 Run ServiceTest

After ServiceTest is run, the automated test case is executed.

1. **Log on to the homepage of Apsara Infrastructure Management Framework and choose Tasks > Deployment Summary from the top navigation bar. The Deployment Summary page appears.**
2. **On the Deployment Summary page, click Deployment Details. The Deployment Details page appears.**
3. **Move the pointer over the row in which the project name is PAI. Click Details, and click ServiceTest# to go to the server list page.**
4. **On the machine learning list page, click Terminal to access TerminalService.**
5. **Run the `sudo docker ps -a` command to find the ServiceTest instance of PAI, as shown in the following figure.**

Figure 6-1: ServiceTest instance

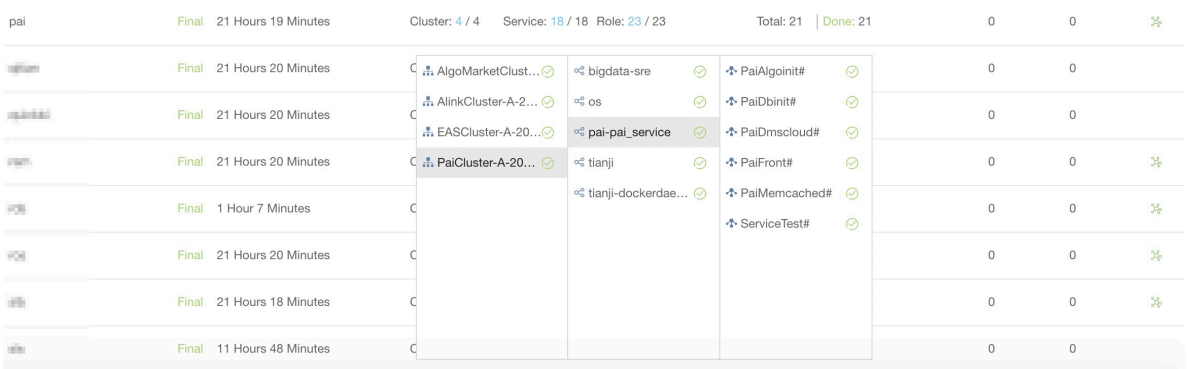

6. **Run the `sudo docker restart e90f70353031` command to restart the ServiceTest service, as shown in the following figure.**

   Figure 6-2: Restart the ServiceTest service

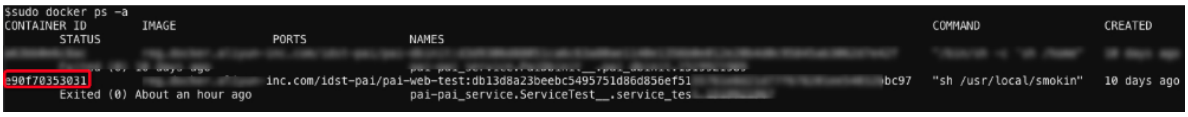

   The test case is executed when the service_test service is restarted. After the execution, you can view the log information.

7. **Run the `sudo docker logs e90f70353031 --tail 1000` command to view the log. Only the last 1,000 rows are displayed.**
8. **After the test case is executed, the testing results for all algorithms are displayed, as shown in the following figure.**

   Figure 6-3: Testing results

   - **PASS: The algorithm is running properly.**
   - **SKIP or FAIL: The algorithm fails.**

6.2.1.2 Common faults and solutions

6.2.1.2.1 Maintenance commands

nc, telent, curl, ping, mysql

`docker images` : shows all images on a server.

`docker ps`: shows the running images on a server.

`docker exec -ti containerID /bin/bash`

`docker log containerID`: shows the container log.

`curl http://localhost/status.taobao`: determines whether the SpringBoot service is started.

6.2.1.2.2 pai.xx.xx access failures

Procedure

1. **Run the `ping pai.xx.xx` command to check whether the domain name has been translated to the corresponding VIP.**

   If the domain name cannot be resolved properly, contact the on-site engineer to check the network configurations.

2. **Run the `curl http://ip/status.taobao` command to check whether all service modules are running normally.**

   If the status.taobao module fails the check, perform the following operations:

   a. **Log on to the server to check whether the container is active.**
   b. **Go to the container and run the following command to check whether the service process is active:**

      ```
      ps -lef | grep java
      ```

   c. **View the *`/home/admin/{app}/logs/err_pai-dms.log`* file to locate causes, such as dependent tenant service request timeout, dependent OCS timeout, and database connection exceptions.**

      We recommend that you view the log after checking all items in the checklist to verify whether the malfunction was not caused by a component exception.

3. **Verify whether ApsaraDB for RDS is accessible.**

   a) **Run the following command to check whether the port is active:**

      ```
      nc -v -z $rds_host $port
      ```

   b) **Run the following command to check whether the database is accessible:**

      ```
      mysql -h$Host -P$Port -u$user -p$password
      ```

4. **Verify whether the caching service is functioning properly.**

   Run the following command to check whether port 11211 is active:

   ```
   nc -v -z $ocs_host 11211
   ```

   Search for ocs_host as follows:

   a. **Search for the dmscloud instance, as shown in the following figure.**

   

   b. **Run the `sudo docker inspect b6ead0fa1d58 | grep ocs` command to view the ocs_host information, as shown in the following figure.**

   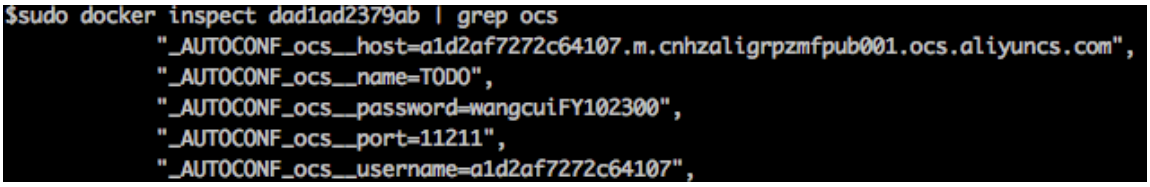

   `host` is a list of servers on which OCS (caching service) is deployed. `port` indicates the port number.

   Machine Learning Platform for AI in Apsara Stack typically uses the built-in memcached service as the dependent caching service. If port 11211 is inaccessible, log on to the server and run the following command to restart the memcached service:

   ```
   docker restart containerid
   ```

6.2.1.2.3 Experiment failures

We recommend that you run a Machine Learning Platform for AI experiment in Google Chrome version 66 or later. Google Chrome is the only supported browser.

- **Components cannot be dragged and dropped.**

  Clear cookies and caches, and then retry. Check the version of Chrome. If the problem persists, it is due to a service failure. Log on to the container to view the log.

- **An error message is displayed while an algorithm is running.**

  If an error message is displayed, the task has been submitted to MaxCompute. Check the parameters and source data against the user guide and algorithm descriptions to locate the error.

6.2.1.2.4 Other failures

If a problem persists after you have checked all items by referring to *pai.xx.xx access failures*, troubleshoot the underlying dependency services, including MaxCompute and DataWorks (tenants and metadata).

- **MaxCompute: Make sure that MaxCompute can pass the pai_console test.**
- **DataWorks: Make sure the configured domain name is accessible, and verify the application log.**

If no errors are found, restart the service.

6.2.2 Online model service maintenance (must be activated separately)

Node maintenance

Online model service nodes are Kubernetes nodes. You can run the `kubectl get node` command to view all nodes in a cluster. A healthy node is in the Ready state. When a node is not in the Ready state, the one of the following errors may have occurred:

- **Node failures**

  There are many reasons that may cause a node to fail. Typically, a node fails when the kernel crashes or the disk does not have sufficient space. If the node can be restarted properly, it rejoins the cluster after it is restarted. If the node cannot be restarted properly, contact the corresponding ECS support personnel.
- **Docker daemon exceptions**

  A Docker daemon exception rarely occurs. Docker daemon exceptions are typically caused by storage issues. Run the `systemctl restart docker` command to restart the Docker daemon.

Online model service maintenance

- **A service cannot be created or deleted.**
  - **If Error 500 is returned while an operation is called, the configurations of the eas-ui component are incorrect. Contact Apsara Stack delivery engineers.**
  - **If a creation or deletion operation is called but no response is returned in a timely manner, the jobworker of the service does not work properly. Check**

whether the KVStore for Redis service in the cluster is normal. If not, restart the pod for KVStore for Redis.

- **The system fails to read the monitoring data.**

  Check whether the influxdb-0 pod under `eas-system` is created properly. If the pod is not in the running state, an influxdb out of memory error has occurred. You can expand the influxdb-0 memory.

Service maintenance

- **Service creation failures.**

  The request is sent but the service creation result displays Failed. A model error has caused a crash. The system then fails to create the model. Check whether the model code contains any null pointers or has any other problems.
- **The system fails to obtain the monitoring data.**

  Check whether the influxdb-0 of each service is normal. The service cannot be created because a persistent volume cannot be created. Check whether the Apsara Stack environment has sufficient disk space. If influxdb-0 runs properly but you cannot obtain the monitoring data, restart the influxdb-0 pod.

6.2.3 GPU cluster maintenance (deep learning must be activated separately)

Node maintenance

A deep learning node is a server where a GPU cluster runs.

1. **Perform steps 1 through 3 in *Query GPU cluster information* and log on to ApsaraAG of the GPU cluster.**
2. **Run the `r ttrl` command to view all nodes that support deep learning tasks.**

- **Node failures**

  There are many reasons that may cause a node to fail. Typically, a node fails when the kernel crashes or the disk does not have sufficient space. If the node can be restarted properly, it rejoins the cluster after it is restarted. If the node cannot be restarted properly, contact the corresponding service support team.

- **Docker daemon exceptions**

  A Docker daemon exception rarely occurs. Docker daemon exceptions are typically caused by storage issues. Run the `systemctl restart docker` command to restart the Docker daemon.

Service maintenance

Failure to allocate resources to a task

Perform the following steps for troubleshooting:

1. **Perform steps 1 through 3 in *Query GPU cluster information* and log on to ApsaraAG of the GPU cluster.**
2. **Run the `r quota` command to view the quota information of the GPU cluster.**
3. **Run the `r cru` command to view the resources allocated to each task in the current cluster.**
4. **Run the `r al` command to view all tasks submitted to the cluster.**
5. **Run the `r wwl WorkItemName` command to view the status of a specific task.**
   - **If only ChildMaster is displayed, no resources are allocated to the worker.**
   - **If worker name is displayed but no hostname is displayed, service resuming is pending or has failed. Log on to the server of the ChildMaster and locate the error. You can also contact the service support team.**
6. **Run the `r ttrl` command to check the value of FUXI_GPU in the Other column. If the value is 200, the worker has two GPUs. If the value is 800, the worker has eight GPUs.**
7. **Log on to a GPU worker in the worker list obtained in Step 3 over SSH. Run the `nvidia-smi` command to view the GPU status. If an exception occurs, contact the relevant service support personnel.**

7 Quick BI

7.1 Introduction to O&M and tools

7.1.1 Introduction to operations and maintenance

Quick BI Operations and Maintenance (O&M) Guide provides step-by-step instructions to explain the O&M process for Quick BI. With the guide, You can perform daily operations, such as monitoring and maintaining Quick BI, and detecting, troubleshooting, and resolving issues. These operations can help ensure that Quick BI is available, stable, and secure.

You can use the Apsara Infrastructure Management Framework to troubleshoot the unavailability issues of Quick BI.

7.1.2 Troubleshoot Quick BI issues by using the Apsara Infrastructure Management Framework

The Apsara Infrastructure Management Framework is a tool that allows you to perform O&M tasks on Quick BI. You can use the Apsara Infrastructure Management Framework to troubleshoot the service unavailability issue of Quick BI.

Prerequisites

Log on to the Apsara Infrastructure Management Framework.

Procedure

1. **Find the Quick BI project in the Apsara Infrastructure Management Framework.**

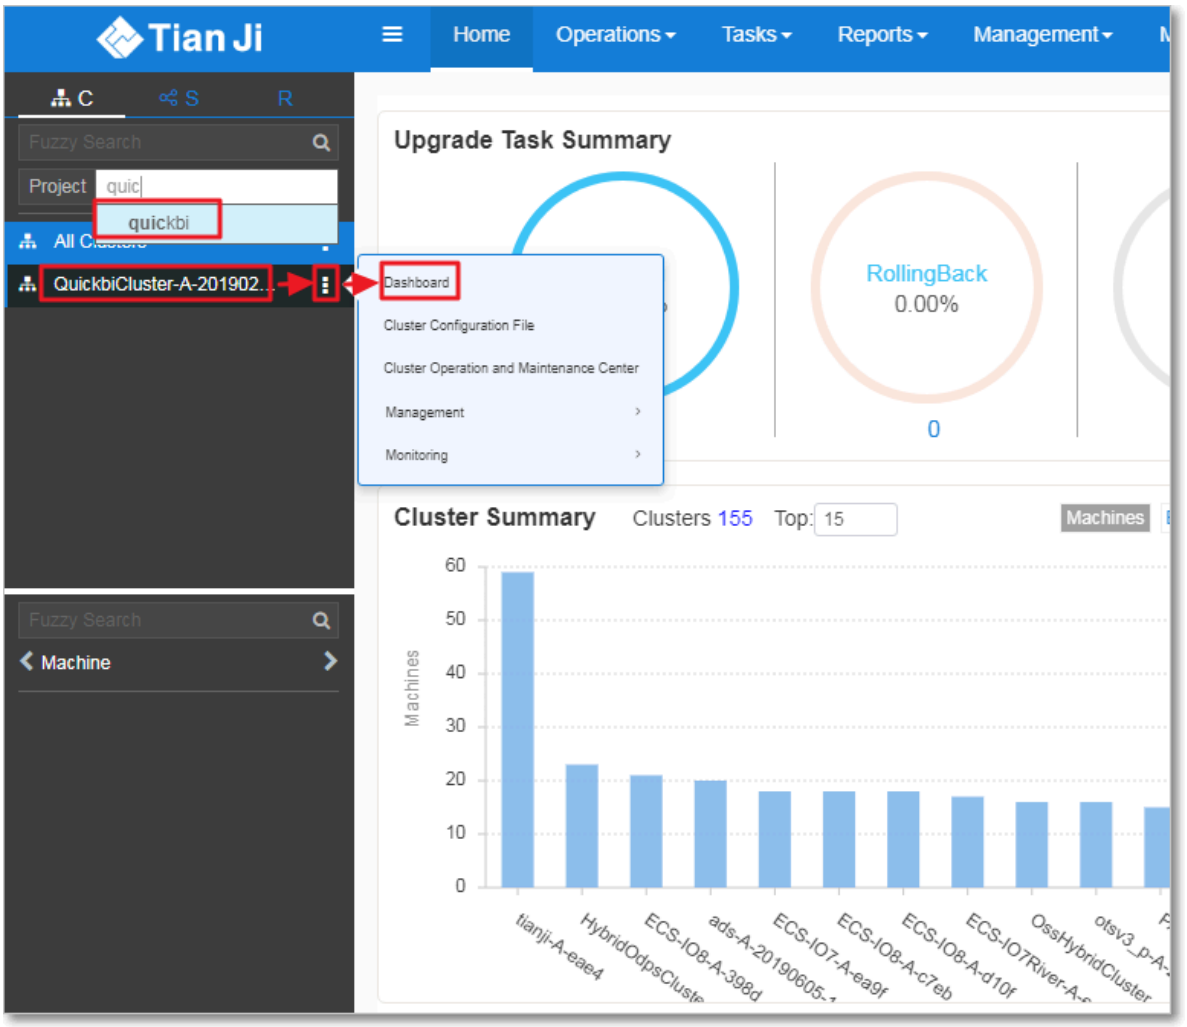

2. **On the Dashboard page, view the cluster status of Quick BI. Check whether the Quick BI cluster is at the desired state. If the cluster is at the desired state, the**

system works as expected. If the cluster is not at the desired state, go to the next step.

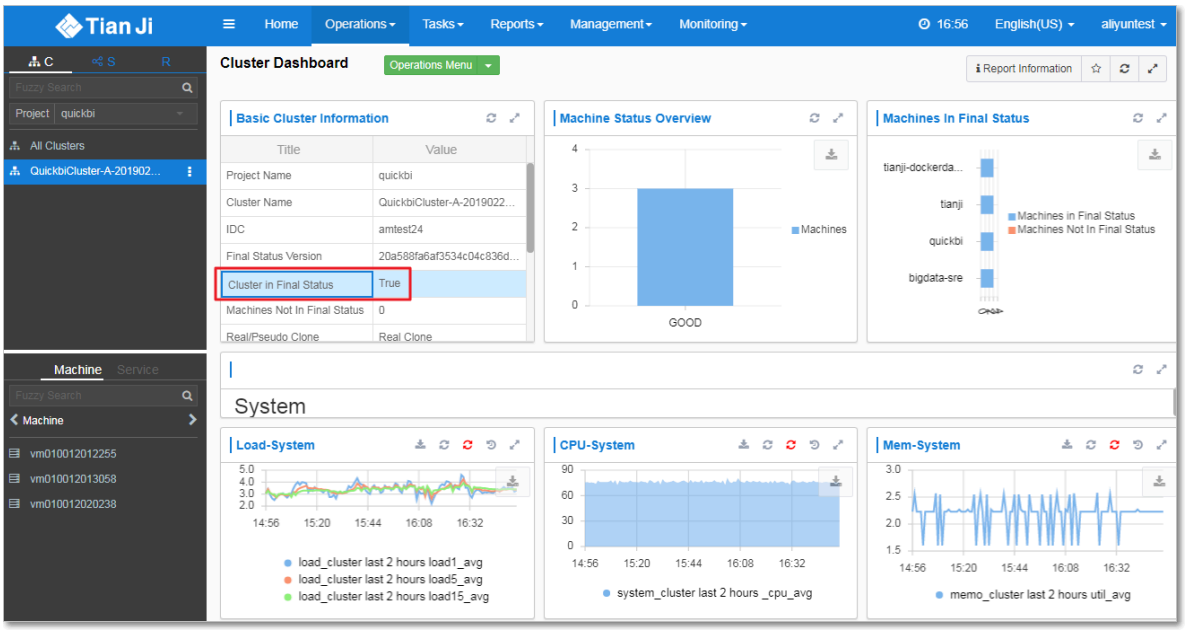

3. On the Dashboard page, find the Service Instances section, and view the service instance details of Quick BI.

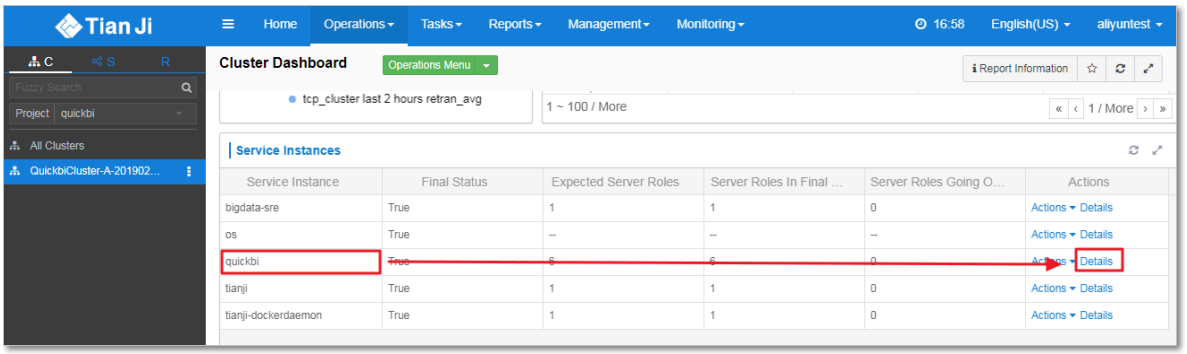

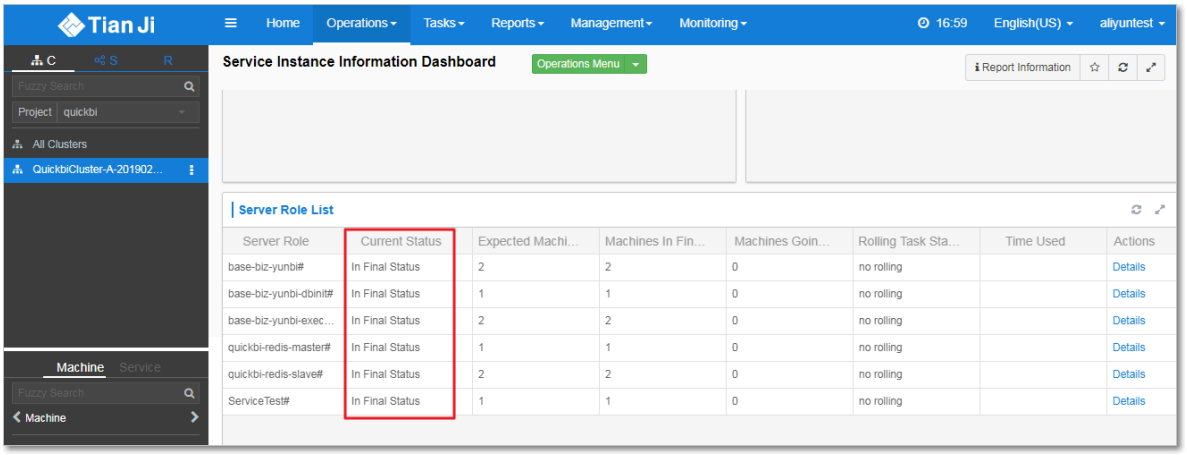

4. **If a service instance is not at the desired state, you need to follow these steps to troubleshoot the issue.**

   Dependencies exist between service roles. If an upstream service role has not reached the desired state, the downstream service role cannot reach the desired state. We recommend that you first troubleshoot the upstream service role. The following figure *Figure 7-1: Relationship between Quick BI service roles* **shows the relationship between Quick BI service roles.**

   Figure 7-1: Relationship between Quick BI service roles

   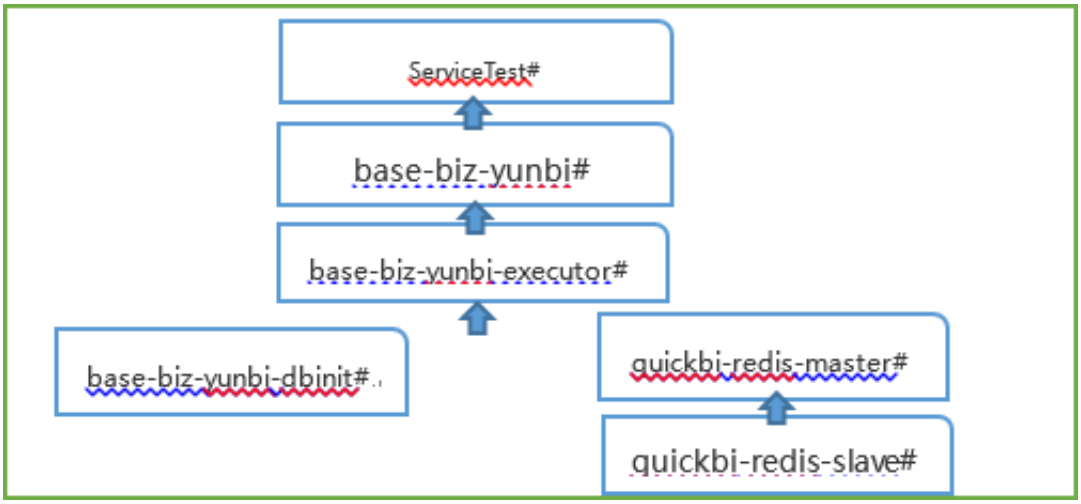

   For example, if the base-biz-yunbi-executor# service role do not reach the desired state, the base-biz-yunbi# and ServiceTest# service roles cannot reach the desired state. You must first ensure that the base-biz-yunbi-executor# service role reaches the desired state. After the base-biz-yunbi-executor# service role reaches the desired state, the base-biz-yunbi# and ServiceTest# service roles will enter the desired state one by one excluding unexpected issues.

7.2 Routine maintenance

7.2.1 Introduction to Quick BI components

You can use container monitoring and periodical detection to check whether service roles related to Quick BI components are at the desired state. You can use these methods to manage and maintain Quick BI. This topic describes Quick BI

operations and maintenance (O&M) components, related service roles, and the description about each component.

Quick BI O&M components, related service roles, and the description of each component

Component	Service role	Description
Database initialization components	base-biz-yunbi-dbinit#	Allows you to initialize Quick BI metadata. The service role must be at the desired state before Quick BI can run as expected.
Cache components	quickbi-redis-master#	Allows you to cache Quick BI data to improve query performance.
	quickbi-redis-slave#	
Runtime components	base-biz-yunbi-executor#	Allows you to perform operations, such as retrieving table metadata and data from data sources.
Web service components	base-biz-yunbi #	Provides Web services. The service role provides Web services that allow frontend clients to visit Quick BI Web pages.
Automated testing components	ServiceTest#	Allows you to check the availability of Quick BI by running batch test cases.

Note:
When you deploy or update Quick BI, the ServiceTest# service role is automatically started.

7.2.2 Database initialization components

This topic describes how to troubleshoot issues when you perform container monitoring on database initialization components.

In the Apsara Infrastructure Management Framework, you need to check whether the base-biz-yunbi-dbinit# service role is at the desired state.

Note:

The service role that is related to database initialization components must be at the desired state before Quick BI is running as expected. If the check result indicates that the service role is not at the desired state, we recommend that you contact Quick BI Technical Support.

7.2.3 Cache components

This topic describes how to detect and troubleshoot issues when you perform container monitoring on cache components.

Container monitoring

In the Apsara Infrastructure Management Framework, you need to check whether the quickbi-redis-master# and quickbi-redis-slave# service roles are at the desired state.

Note:

You can also check the redis process. If the redis process exists, it means that the preceding service roles are at the desired state.

Quick BI is unavailable if the check result indicates that the linked service roles are not at the desired state. Cause: The redis process is interrupted or not started.

Solution: You need to restart the linked service roles. You need to restart the quickbi-redis-master# service role and then restart the quickbi-redis-slave# service role.

Periodical detection

You can check the service availability based on the exit status that is returned after you run the `/checkRedis.sh` script. Quick BI is available if the value of the exit status is 0. Otherwise, Quick BI is unavailable. You can use the preceding script to check whether the redis process exists. The redis process exists if the value of the returned exit status is 0. Otherwise, the redis process does not exist. The detection interval is one second.

7.2.4 Runtime components

This topic describes how to detect and troubleshoot issues when you perform container monitoring on runtime components.

Container monitoring

In the Apsara Infrastructure Management Framework, you need to check whether the base-biz-yunbi-executor# service role is at the desired state.

Quick BI is unavailable if the check result indicates that the linked service role is not at the desired state. Cause: The runtime component process is interrupted or not started.

Solution: You need to restart the base-biz-yunbi-executor# service role.

Periodical detection

You can visit *http://container:7001/checkpreload.htm* **at regular intervals to call the HTTP service. Quick BI is available if a status code of 200 is returned. Otherwise, Quick BI is unavailable. The detection interval is one second.**

> **Note:**
> **The container in the preceding HTTP link is a variable. You must replace the variable with an IP address that is used by the base-biz-yunbi# service role.**

7.2.5 Web service components

This topic describes how to detect and troubleshoot issues when you perform container monitoring for Web service components.

Container monitoring

Check whether the base-biz-yunbi# service role is at the desired state.

Quick BI is unavailable if the check result indicates that the linked service role is not at the desired state. Cause:

- **The Java process is interrupted or not started. Symptom: You cannot visit** *http://container:7001/checkpreload.htm***.**
- **No HTTPS certificate is issued and port 443 is inaccessible. Symptom: You cannot visit** *https://container/checkpreload.htm***.**

> **Note:**

The container in the preceding link is a variable. You must replace the variable with an IP address that is used by the base-biz-yunbi# service role.

Solutions:

- **If the Java process is interrupted or not started, you need to restart the base-biz-yunbi# service role.**
- **If no HTTPS certificate is issued, you need to restart the base-biz-yunbi# service after the HTTPS certificate is issued.**

Periodical detection

You can visit *https://container/checkpreload.htm* **at regular intervals to call an HTTPS service. Quick BI is available if a value of 200 is returned. Otherwise, Quick BI is unavailable. The detection interval is five minutes.**

Note:
The container in the preceding HTTPS link is a variable. You must replace the variable with an IP address that is requested by the base-biz-yunbi# service role.

7.2.6 Automated testing components

This topic describes how to detect and troubleshoot executor issues when you perform container monitoring on automated testing components.

Container monitoring

In the Apsara Infrastructure Management Framework console, check whether the ServiceTest# server role is at desired state.

If the server role is not at desired state, a service error occurs. Causes:

- **The service is unavailable. Symptom: You cannot visit** *https://container/checkpreload.htm* **or log on to the Quick BI console.**

  Note:
  "container" in the link is a variable. You must replace it with the IP address that is used by the base-biz-yunbi# server role.
- **The service is available but an error is detected. Symptom: You can log on to the Quick BI console and search data. However, a logon error is reported in the Apsara Infrastructure Management Framework console. You can view the error message provided in the Description column.**

Solutions:

- **If the service is unavailable, check whether other server roles are at desired state . If they are not, handle the problem.**
- **If the service is available but an error is detected, contact Quick BI technical support and provide error information.**

Periodical detection

You can execute test cases at regular intervals to check the availability of Quick BI . A service is available if the linked server role is at desired state. Otherwise, the service is unavailable. The detection interval is 30 minutes.