ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Enterprise

Operations and Maintenance Guide

Product Version: 2006, Internal: V3.12.0

Document Version: 20201106

Alibaba Cloud

Legal disclaimer

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

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

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

Table of Contents

1.O&M overview

This topic describes the Apsara Stack Operations (ASO) system.

Based on the Information Technology Infrastructure Library (ITIL) and IT Service Management (ITSM) standards, the O&M processes and requirements must be abstract and automation is implemented by using intelligent O&M tools for common high-frequency operations. For customized operations, interfaces and multi-level approval must be used to reduce risks.

Apsara Stack adopts the ISO 20000 series standards and references the methods regulated by the Information Technology Service Standards (ITSS) and ITIL frameworks to build the ASO framework. ASO framework shows the ASO framework.

ASO framework

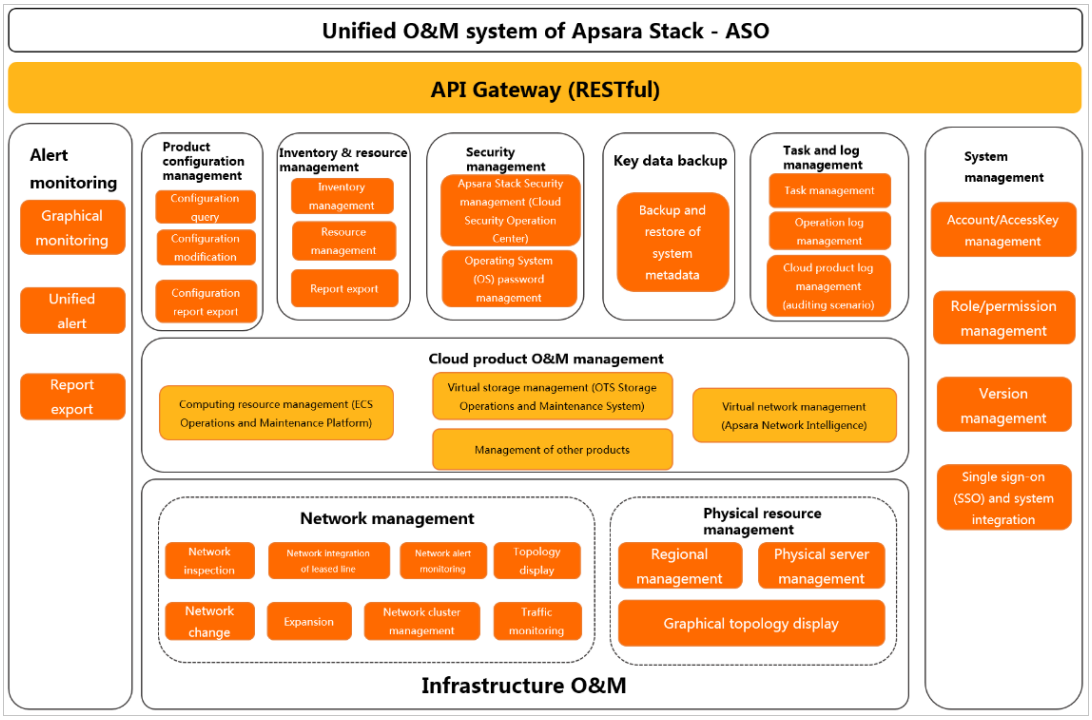

ASO is a unified intelligent O&M platform for Apsara Stack. In the ASO framework, cloud O&M is classified into three layers: infrastructure O&M, cloud product O&M, and business O&M. The ASO framework contains the full lifecycle management methods, management standards, management modes, management supporting tools, management objects, and process-based management methods of IT O&M services.

ASO provides a centralized O&M portal. You can use the O&M portal to have a consistent O&M experience and a unified O&M entrance. ASO supports interconnections with third-party platforms and provides centralized API O&M capabilities to deliver data to third-party systems by using APIs.

ASO performs centralized O&M management, such as automated deployment, upgrade, change, and configuration, on physical devices, operating systems, computing resources, network, storage, databases, middleware, and business applications in the cloud computing environment. ASO also provides the functions of alert monitoring and automatic analysis, diagnosis, and troubleshooting for faults, performance, and configurations. By analyzing, processing, and evaluating the running status and quality of cloud platforms, ASO guarantees the continuous and stable running of cloud computing business applications and provides services and support for O&M service processes to build an improved O&M service management platform.

Based on ITIL/ISO20000, with process-oriented, normalized, and standardized management as the method, adaption to various management modes as the aim, the O&M service management framework uses management supporting tools to implement the systematic management of the overall process of O&M services.

Based on the accumulated O&M experience and data collection of the three-layer system, Alibaba Cloud Apsara Stack aggregates data collected by O&M platforms to the Configuration Management Database (CMDB) of the platform. ASO, the intelligent O&M platform, consolidates, analyzes, and comprehensively processes the data and solidifies rich practical experience and O&M capabilities to the platform O&M tools. By using the design concept of facing to the final status, ASO uses the unified O&M tools for the fault discovery, fault tracking, link display, ITIL process, and self-repaired faults of the platform to realize the ultimate goal of AIOps.

In addition to tools, process assurance and personnel management are essential to the O&M integrity. Apsara Stack provides on-site development supporting services for major problems, on-site services, expert escort services, business consulting services, and business optimization services. Apsara Stack provides the first-line, second-line, and third-line supporting systems to support platform problems of customers and the upgrade channel to support urgent problems of customers. As an autonomous and controllable platform, ASO ensures that technical problems can be effectively solved in a timely manner.

The ASO system defines various entities involved in O&M activities and relationships between these entities. Relevant entities are well organized and coordinated based on the O&M service management system and can provide different levels of O&M services based on the service agreements.

2.Preparations before operations

2.1. Prepare an operations account

Before you perform O&M operations in the ASO console, make sure that you have obtained an operations account from an administrator.

Perform the following steps to create an operations account and grant permissions to the account:

1. Log on to the ASO console as a system administrator.
2. Create a role. For more information, see Role management.
3. Create an operations account and grant the role to the account. For more information, see User management.

> **Note** For a more fine-grained division of the operations role, you can create a basic role as specified in **Appendix > OAM**, grant permissions to the role, and then grant the role to the corresponding operations account as an administrator.

2.2. Log on to the ASO console

This topic describes how to log on to the Apsara Stack Operations (ASO) console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

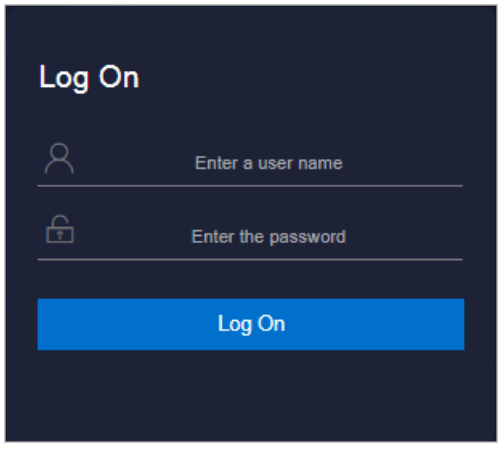

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.

2.3. ASO console overview

After you log on to the ASO console, the homepage appears. This topic describes the basic operations and features of the ASO console.

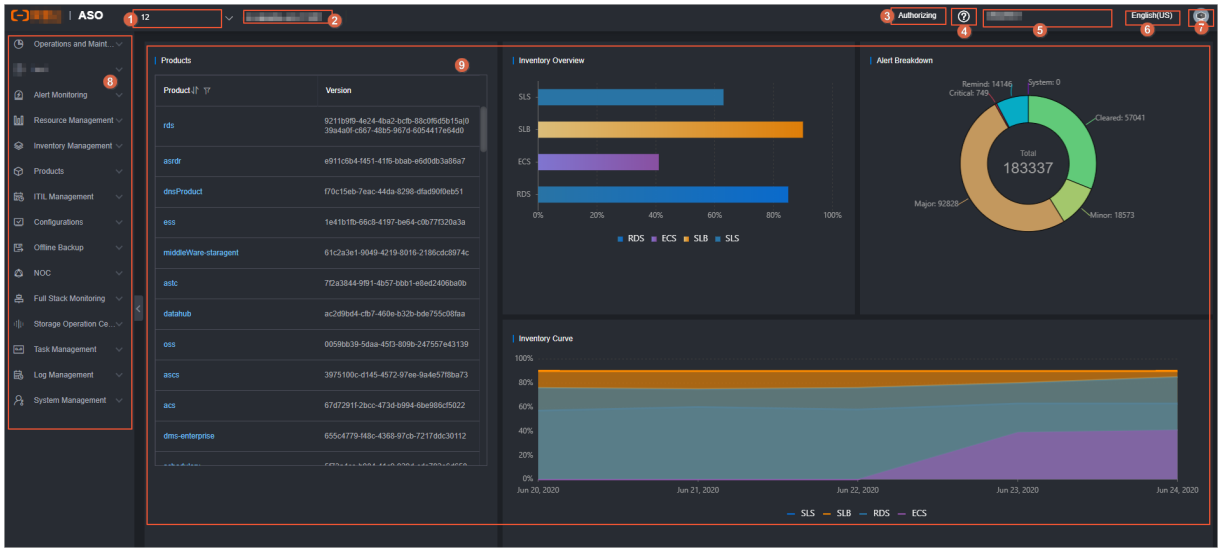

The following table describes the ASO homepage sections.

Section		Description
1	Cloud	Switch the cloud from the drop-down list.
2	Region	Switch the region from the drop-down list and centrally manage each region.
3	Authorization information	Click this section to go to the **Authorization** page and then view the authorization conditions of services.
4	Help center	In the help center, you can view the alert knowledge base and upload other relevant HTML documents.

Section		Description
5	Current user	The name of the current logon user.
6	Language	Move the pointer over this section and select a target language to switch the language.
7	Current user information	Move the pointer to this section and select an item to view the personal information of the current user, modify the password, configure logon parameters, or log off from the ASO console.
8	Left-side navigation pane	Select an O&M operation.
9	Operation area	The information display and operation area.

3.System settings

3.1. Default operations roles

This topic describes the default roles of Apsara Stack Operations (ASO) and their responsibilities.

For quick management, the following roles are preset in ASO: Operation Administrator Manager (OAM) super administrator, system administrator, security officer, security auditor, and multi-cloud configuration administrator. The following table describes these roles and their responsibilities.

Role	Responsibility
OAM super administrator	The administrator of OAM, with the root permissions of the system.
System administrator	Manages platform nodes, physical devices, and virtual resources, backs up, restores, and migrates product data, and searches for and backs up system logs.
Security officer	Manages permissions, security policies, and network security, and reviews and analyzes security logs and activities of auditor officers.
Security auditor	Audits, tracks, and analyzes operations of the system administrator and the security officer.
Multi-cloud configuration administrator	Manages multi-cloud operations, and adds, deletes, and modifies multi-cloud configurations.

3.2. ITIL Management

3.2.1. Overview

Information Technology Infrastructure Library (ITIL) manages the incidents and problems generated during the daily system operations, which allows operations engineers to better maintain the network stability, improve the performance indicators quickly, reduce operation and maintenance costs, and finally enhance the user satisfaction.

ITIL has the following functions:

- Dashboard

  The **Dashboard** section displays the summary of incidents and problems and the corresponding data in specific days.

- Services

  The **Services** section is used to record, diagnose, resolve, and monitor the incidents and problems generated during the operations. Multiple types of process transactions are supported.

  You can submit the incidents and problems generated when using the system to the service request platform and receive the information about the problem processing.

  - Incident management: used to recover from exceptions and guarantee the normal production by a series of recovery operations, including diagnosis, processing, resolution, and confirmation. Incident management provides a unified mode and standardizes the process for incident processing, and supports automatically collecting or manually recording the incident information.
  - Problem management: Temporarily resolved incidents or incidents whose root cause is not clear can be transformed to problems for further analysis and thorough troubleshooting. Incidents aim to resume the production, whereas problems aim to be completely solved to make sure the problems do not recur. Problem management allows you to find the root cause of incidents, thoroughly troubleshoot the incidents, and reduce repeated incidents.

- Version control

  The **Version Control** section displays the version information of Apsara Stack products.

- Process template configuration

  By configuring the operations process template, operations engineers can select the corresponding type from the catalogue based on the actual Operations & Maintenance (O&M) operations and assign tasks according to different types of process templates.

- CAB/ECAB configuration

  The change management process has the **CAB Audit** and **ECAB Audit** phases. Therefore, you must configure the CAB or ECAB.

3.2.2. Dashboard

The Dashboard module allows you to view the summary of incident, problem, and change requests: the total numbers of incident, problem, and change requests, the numbers of new and closed incident, problem, and change requests, and their change trend. You can also view the distribution of request fulfillment and the information of version management.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Dashboard**.

   The data for the last seven days is displayed on the **Dashboard** page by default.

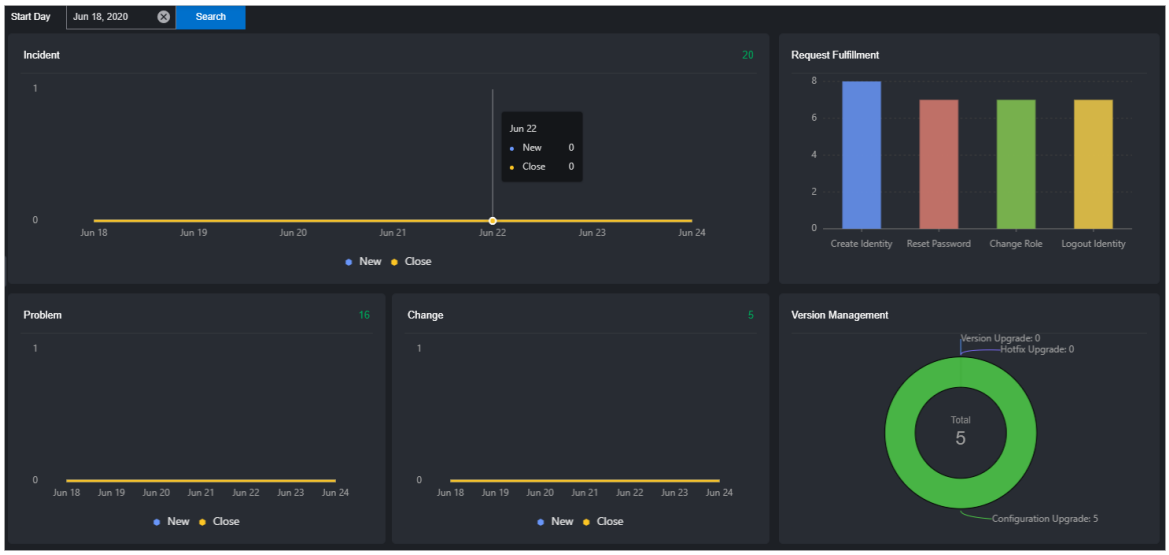

2. In the upper part of the page, set a time range, and then click **Search**.

3.2.3. Services

3.2.3.1. Basic functions

3.2.3.1.1. Overview

This topic describes the basic features of requests and tasks.

Services consists of the requests and tasks modules.

- Requests

  A request is the complete process of an incident or problem request. For example, the process of an incident request may consist of diagnose, resolve, and confirm phases.

- Tasks

  A task is an operation of a phase in the processing of an incident request or problem request. For example, the diagnose phase in the incident request processing can be considered as a task.

3.2.3.1.2. Request management

This topic describes how to create, filter, and view details of requests.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**.The **Request** tab appears.
2. On the Request tab, perform the following operations:
   - Create a request

     Click + New and then select a request type. Configure the parameters and then click Confirm to create a request. This topic describes how to create an incident request and a problem request. For more information, see Create an incident request and Create a problem request.

     Requests are classified into the following types based on the processing status:

     - : in processing, indicating the requests that are waiting to be processed.
     - : closed, indicating the requests that have the whole process completed.
     - : recycle bin, indicating the requests that have been recycled.

   - Filter requests

Select a request type from the drop-down list. The requests of that type are displayed in the list.

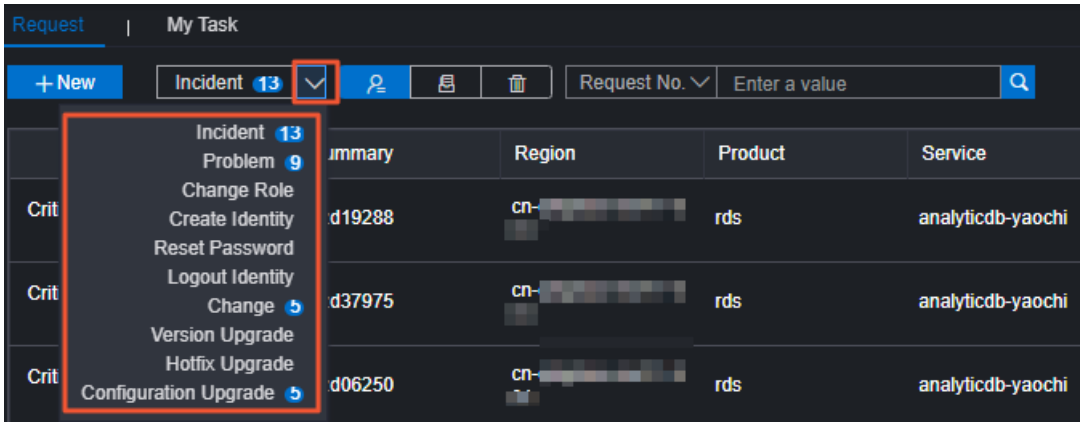

- Search for requests

  Select **Request No.** or **Summary** from the second drop-down list, enter the information in the search box, and then click the Search icon.

- View request details

  Find the target request and then click **Detail**. The request details page contains the following sections:

  - **Top navigation bar**: buttons for the request processing. For more information, see Manage incident requests and Manage problem requests.

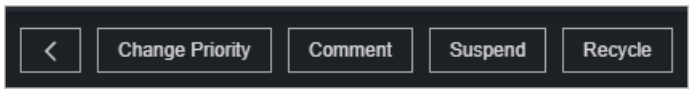

  - **Request Flow**: the current processing flow of the request.

  - **Basic Information**: the basic information of the request, which is generally the information configured when you create the request.

- **Track**: the phases of the request processing and their points in time.

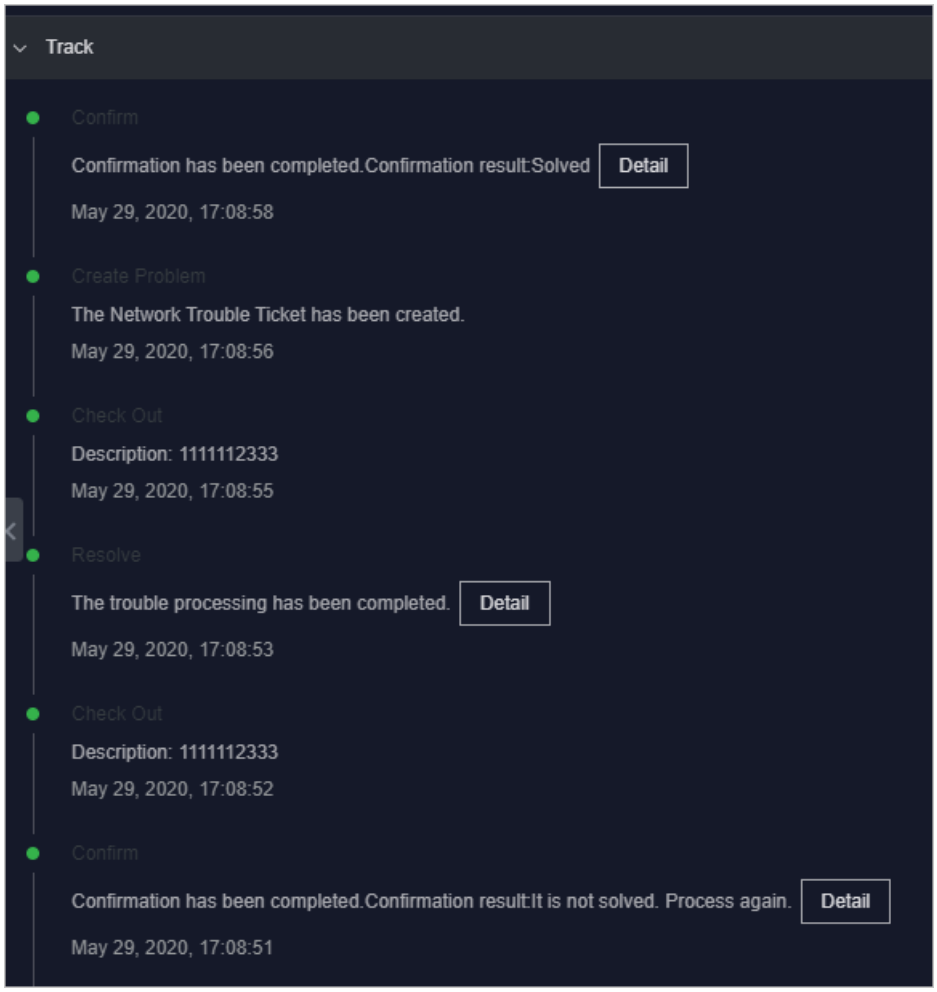

- **Detail Tabs**: the task list and comments related to the request.

3.2.3.1.3. Task management

After a request is created, the request enters the diagnose phase. In the diagnose phase, the system automatically generates a task. Each task corresponds to a processing phase.

Context

Tasks are classified into the following types:

- : my task, indicating tasks that are waiting to be processed by you.
- 13: task pool, indicating a collection of tasks that are not assigned to relevant responsible

persons. You can check out the tasks in the task pool to make the tasks exclusive to you. Others cannot process the tasks that you have checked out. You can view the checked out tasks under .

- : processed by me, indicating the historical tasks that have been processed by you. After you process the tasks under , they are displayed under .

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**.
2. Click the **My Task** tab.

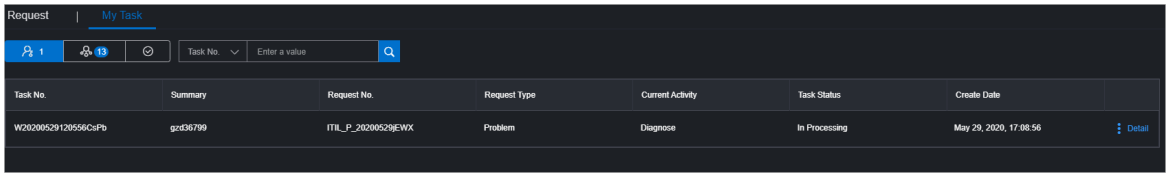

3. On the My Task tab, perform the following operations:
   - Search for tasks

     Select **Task No.**, **Request No.**, or **Summary** from the drop-down list, enter the information in the search box, and then click the Search icon.

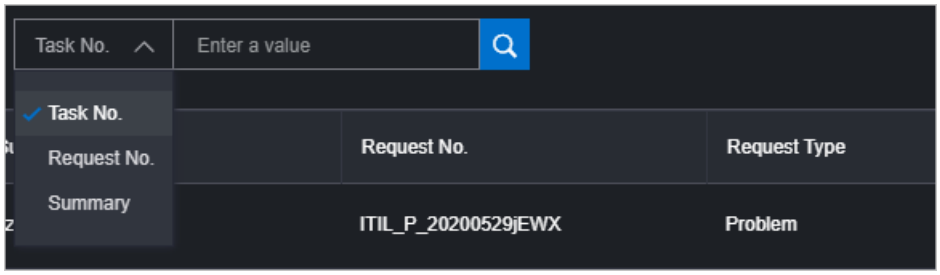

   - View task details

     Find the target task and then click **Detail**. On the task details page, you can view the request details related to the task. For more information, see the "View request details" section of Manage requests.

3.2.3.2. Manage incidents

3.2.3.2.1. Create an incident request

An incident is a system exception that affects the normal business. Incident management is used to recover from exceptions and guarantee normal business by performing a series of recovery operations including diagnosis, resolution, and confirmation. If the system has an exception, you can create an incident request to track the progress of the incident.

Context

ITIL management allows you to create incident requests by using one of the following methods:

- Automatically create incident requests

The incident information comes from the alert information in ASO. The Alert module transfers the alert information to the ITIL Management module to generate incident requests based on the actual conditions, such as the alert level and the alert filtering.

- Manually create incident requests

  You can manually create incident requests to supplement the automated system. For example, if the incident is not automatically recognized, you can manually create an incident request. This topic describes how to manually create an incident request.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**. The **Request** tab appears.
2. Click the + New icon and then select **Incident** from the drop-down list. On the page that appears, configure the parameters.

Parameter	Description
Report Object	The person who is required to process the request.
Callback Email	The email of the person who submits the request.
Callback Telephone	The telephone number of the person who submits the request.
Region	The region to which the request belongs.
Product	The product to which the request belongs. Select a specific product from the drop-down list.
Service Name	The service related to the product to which the request belongs. Select a service from the drop-down list.
Happen Date	The time when the request was sent.
Priority	The priority of processing this request. The priority indicates the urgency of the request. The higher the urgency, the higher the priority must be. Based on the urgency, the priority of requests are divided into the following levels in descending order: ○ Critical ○ Major ○ Minor ○ Remind ○ Cleared ○ System
Alarm Code	The alert ID.
Summary	The summary of the request.
Description	The detailed description about the request.

Parameter	Description
Suggestion	Optional. Suggestions to process the request.

3. Click **Confirm**.

3.2.3.2.2. Manage incident requests

After you create an incident request, you can change the priority of, comment on, suspend, resume, recycle, restore, and delete the created incident request.

Prerequisites

An incident request is created. For more information about how to create an incident request, see Create an incident request.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**.The **Request** tab appears.
2. Select **Incident** from the drop-down list next to the + New icon. The incident requests are displayed in the list.

3. Find the incident request that you want to manage, and then click **Detail**.
4. On the request details page, perform the following operations:
   - Change the priority

     Click **Change Priority** in the upper part of the page. In the dialog box that appears, select the new priority. Perform this operation to temporarily adjust the priority or correct an error in priority.

     > **Note** You can only change the priority of an incident request that is in the **Diagnose** phase.

   - Comment on the incident request

     Click **Comment** in the upper part of the page. In the dialog box that appears, enter the comment for the incident request. Perform this operation for scenarios that require collaborations. For example, users can comment on an incident request to share the information with each other and guide each other when they process the same incident.

   - Suspend the incident request

     Click **Suspend** in the upper part of the page. In the dialog box that appears, enter the **Remarks**. Perform this operation for incident requests that currently do not require to be processed.

   - Resume the incident request

     Click **Resume** in the upper part of the page. In the dialog box that appears, enter the **Remarks**. Perform this operation for suspended incident requests that require to be processed.

   - Recycle the incident request

Perform this operation for incident requests in the processing ( ) list. Click Recycle to cancel or logically delete the incident request. When the incident request is recycled, it is displayed in the recycle bin ( ) list.

- Restore the incident request

  Perform this operation for incident requests in the recycle bin ( ) list. Click Restore to restore the recycled incident request. When the incident request is restored, it is displayed in the processing ( ) list and restored to the status before the request is recycled.

- Delete the incident request

  Perform this operation for incident requests in the recycle bin ( ) list. Click Delete to delete the incident request. When the incident request is deleted, it is physically deleted and cannot be restored.

3.2.3.2.3. Manage incident tasks

When you create an incident request, it is divided into different tasks based on the incident processing flow. Different tasks are to be processed by different people in charge.

Context

The processing of an incident task consists of the following steps:

- Diagnose: When an incident request is created, the system automatically goes to the Diagnose phase and analyzes the reason of the incident.
- Resolve: The system goes to the Resolve phase after the Diagnose phase. The incident is repaired in this phase.
- Confirm: The system goes to the Confirm phase after the Resolve phase and reviews whether the incident was processed in a reasonable manner. If Temporary Solution is selected in the Diagnose phase, or an incident requires further analysis, you can create a problem request in this phase to track the incident processing.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**.
2. Click the **My Task** tab.
3. Click the icon.

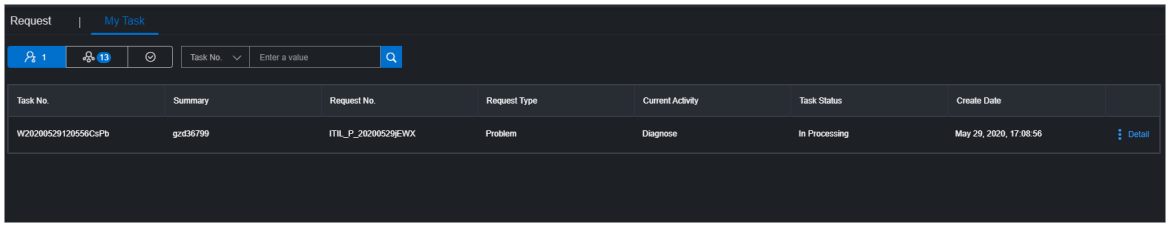

> **Note** To check out a task in the task pool to the current username for processing, perform the following operations: Click the 13 icon and then click **Detail** corresponding to the task. Click **Check Out**. In the dialog box that appears, enter the **Description** and then click **OK**.

4. In the task list, find the task that you want to manage and then click **Detail**.
5. On the task details page that appears, click **Diagnose** in the upper part of the page. In the **Diagnose** dialog box, configure the parameters and then click **OK**.

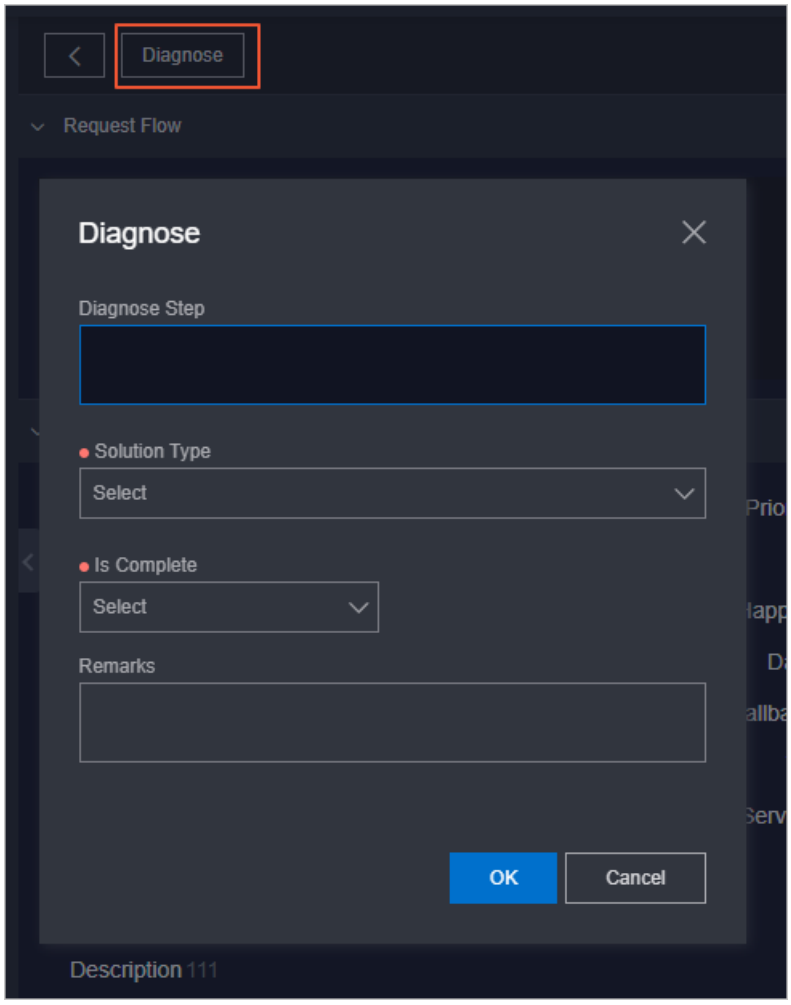

Parameter	Description
Diagnose Step	Analyzes the task steps.
Solution Type	Select **Temporary Solution** or **Permanent Solution**. If you select **Temporary Solution**, you may have to create a problem request in the Confirm phase to further troubleshoot and find the root cause.

Parameter	Description
Is Complete	Select Yes or No to indicate whether the task processing is complete. If **No** is selected, the system goes to the Resolve phase. In some cases, the incident may have been processed after it was reported because of a time difference. In this case, you can directly select **Yes** and configure the resolved date. The system will then skip the Resolve phase and go directly to the Confirm phase.
Remarks	The information about the task.

6. The system goes to the Resolve phase after the Diagnose phase. After the incident is processed offline, click **Resolve** in the upper part of the page. In the **Resolve** dialog box, configure **Resolve Date** and **Handle Step**. Click **OK**.

   The Resolve phase consists of troubleshooting and solving the incident. ITIL only tracks this step in a standardized way and processes the log records.

7. The system goes to the Confirm phase after the Resolve phase. This phase reviews the processing result of the incident. Click **Confirm** in the upper part of the page.
8. In the **Confirm** dialog box, select the review result from the **Is Pass** drop-down list, and then click **OK**.

   The review results have the following status:

   - **Solved**: The incident has been completely solved.
   - **It is not solved. Analyze again**: The incident cannot be solved because of an error in the cause analysis. The task is sent back to the Diagnose phase to restart the processing until the incident can be solved.
   - **It is not solved. Process again**: The reason of the incident is clear. The incident cannot be solved because the incident cannot be effectively processed. The task is sent back to the Resolve phase to restart processing until the incident can be solved.

3.2.3.3. Manage problems

3.2.3.3.1. Create a problem request

If the system has a problem that requires further troubleshooting, you can create a problem request to track the progress of the problem.

Context

Temporarily resolved incidents or incidents whose root cause is not clear can be transformed to problems for further analysis and thorough troubleshooting. Problem management allows you to find the root causes of incidents, thoroughly troubleshoot the incidents, and reduce repeated incidents.

Compared with the incident processing, problem processing takes more time. The occurrence rate of repeated incidents is used to determine whether the problem management is good. A lower occurrence rate indicates more effective problem processing.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**. The **Request** tab appears.
2. Click the + New icon and then select **Problem** from the drop-down. On the page that appears, configure the parameters.

Parameter	Description
Report Object	The person who is required to process the request.
Callback Email	The email of the person who submits the request.
Callback Telephone	The telephone number of the person who submits the request.
Region	The region to which the request belongs.
Product	The product to which the request belongs. Select a specific product from the drop-down list.
Service Name	The service related to the product to which the request belongs. Select a service from the drop-down list.
Happen Date	The time when the request was sent.
Priority	The priority of processing this request. The priority indicates the urgency of the request. The higher the urgency, the higher the priority must be. Based on the urgency, the priority of requests are divided into the following levels in descending order: ○ Critical ○ Major ○ Minor ○ Remind ○ Cleared ○ System
Alarm Code	The alert ID.
Summary	The summary of the request.
Description	The detailed description about the request.
Suggestion	Optional. Suggestions to process the request.

3. Click **Confirm**.

3.2.3.3.2. Manage problem requests

After you create a problem request, you can change the priority of, comment on, suspend, and resume the created problem request.

Prerequisites

A problem request is created. For more information about how to create a problem request, see Create a problem request.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**. The **Request** tab appears.
2. Select **Problem** from the drop-down list next to the  icon. The problem requests are displayed in the list.
3. Find the problem request that you want to manage, and then click **Detail**.
4. On the request details page, perform the following operations:
   - Change the priority

     Click **Change Priority** in the upper part of the page. In the dialog box that appears, select the new priority. Perform this operation to temporarily adjust the priority or correct an error in priority.

     > **Note** You can only change the priority of an problem request that is in the **Diagnose** phase.

   - Comment on the problem request

     Click **Comment** in the upper part of the page. In the dialog box that appears, enter the comment for the problem request. Perform this operation for scenarios that require collaborations. For example, users can comment on a problem request to share the information with each other and guide each other when they process the same problem.

   - Suspend the problem request

     Click **Suspend** in the upper part of the page. In the dialog box that appears, enter the **Remarks**. Perform this operation for problem requests that currently do not require to be processed.

   - Resume the problem request

     Click **Resume** in the upper part of the page. In the dialog box that appears, enter the **Remarks**. Perform this operation for suspended problem requests that require to be processed.

   - Recycle the problem request

     Perform this operation for problem requests in the processing () list. Click Recycle to cancel or logically delete the problem request. When the problem request is recycled, it is displayed in the recycle bin () list.

   - Restore the problem request

     Perform this operation for problem requests in the recycle bin () list. Click Restore to restore the recycled problem request. When the problem request is restored, it is displayed in the processing () list and restored to the status before the request is recycled.

   - Delete the problem request

Perform this operation for problem requests in the recycle bin ( ) list. Click Delete to delete the problem request. When the problem request is deleted, it is physically deleted and cannot be restored.

3.2.3.3.3. Manage problem tasks

When you create a problem request, it is divided into different tasks based on the problem processing flow.

Context

The processing of a problem task consists of the following steps:

- Diagnose: During this phase, the cause of the problem is analyzed.
- Resolve: After the Diagnose phase, the system enters the Resolve phase. During this phase, the problem is repaired.
- Confirm: After the Resolve phase, the system enters the Confirm phase. During this phase, the system reviews whether the problem was processed in a reasonable manner.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Services**.
2. Click the **My Task** tab.
3. Click the icon.

> **Note** To check out a task in the task pool to the current username for processing, perform the following operations: Click the 13 icon and then click **Detail** corresponding to the task. Click **Check Out**. In the dialog box that appears, enter the **Description** and then click **OK**.

4. In the task list, find the task that you want to manage and then click **Detail**.
5. On the task details page, click **Diagnose** in the upper part of the page. In the **Diagnose** dialog box, configure the parameters and then click **OK**.

Parameter	Description
Diagnose Step	Analyzes the task steps.
Solution Type	Select **Temporary Solution** or **Permanent Solution**. If you select **Temporary Solution**, you may have to create a problem request in the Confirm phase to further troubleshoot and find the root cause.
Is Complete	Select Yes or No to indicate whether the task processing is complete. If **No** is selected, the system goes to the Resolve phase. In some cases, the problem may have been processed after it was reported because of a time difference. In this case, you can directly select **Yes** and configure the resolved date. The system will then skip the Resolve phase and go directly to the Confirm phase.

Parameter	Description
Remarks	The information about the task.

6. The system goes to the Resolve phase after the Diagnose phase. After the problem is processed offline, click **Resolve** in the upper part of the page. In the **Resolve** dialog box, configure **Resolve Date** and **Handle Step**. Click **OK**.

   The Resolve phase consists of troubleshooting and solving the problem. ITIL only tracks this step in a standardized way and processes the log records.

7. The system goes to the Confirm phase after the Resolve phase. This phase reviews the processing result of the problem. Click **Confirm** in the upper part of the page.
8. In the **Confirm** dialog box, select the review result from the **Is Pass** drop-down list, and then click **OK**.

   The review results have the following status:

   - **Solved**: indicates that the fault has been completely solved.
   - **It is not solved. Analyze again**: The problem cannot be solved because of an error in the cause analysis. The task is sent back to the Diagnose phase to restart the process until the problem can be solved.
   - **It is not solved. Process again**: The reason of the problem is clear. The problem cannot be solved because the problem cannot be effectively processed. The task is sent back to the Resolve phase to restart the process until the problem can be solved.

3.2.4. Version control

The Version Control module allows you to view the version information and history versions of Apsara Stack products.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > Version Control**. The **Version Control** page appears.

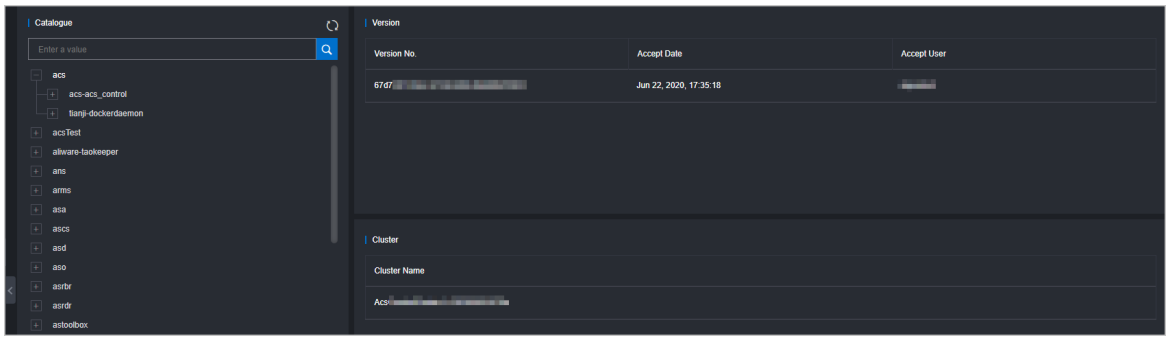

2. Select a node in the tree structure, or enter a name in the search box and click the search icon. The version and cluster information is displayed on the right.

   Note Before you perform a search, click the icon to synchronize the information to the ASO console.

3.2.5. Configure process templates

By configuring the operations process templates, operations engineers can select the corresponding type from the catalog based on the actual O&M operations and assign tasks based on different types of process templates.

In the left-side navigation pane, choose **ITIL Management > Process Template Configuration**. On the **Process Template Configuration** page, you can view the following sections: **Process**, **Process Template**, and **Regulation**.

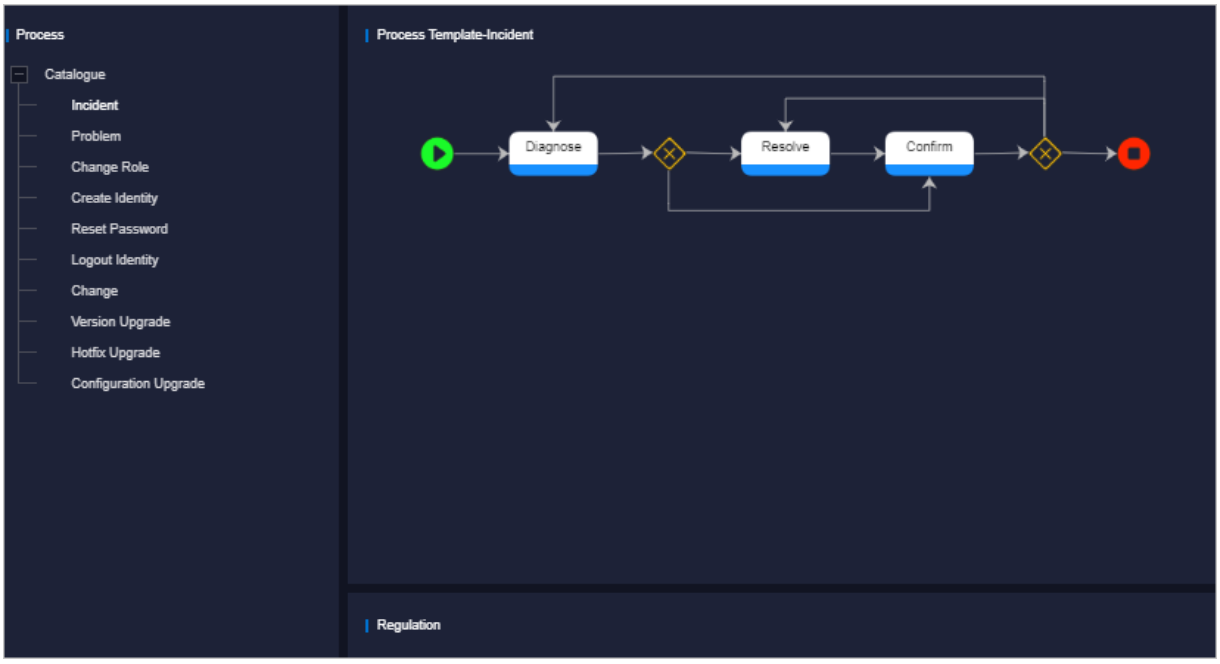

Processes

The following processes are supported:

- Incident
- Problem
- Change Role
- Create Identity
- Reset Password
- Logout Identity
- Change
- Version Upgrade
- Hotfix Upgrade
- Configuration Upgrade

Process templates

After you select a process, the corresponding process template is displayed in the **Process Template** section. The following section describes the nodes in the process:

-  is the start node of the process. A process usually starts with the request creation.

-  indicates the gateway. The gateway defines the process trend in different branches. In the BPMN specification, gateways are classified into different types, such as inclusive gateway, exclusive gateway, parallel gateway, and hybrid gateway. This is an exclusive gateway, indicating that multiple routes have only one valid path.
-  is the end node of the process. A process usually ends with archiving.
- 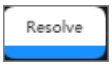

indicates the phase. A phase is usually composed of roles with specific functions.
-  is the route, indicating the process trend. A phase contains one or more egress and ingress routes.

The templates can be classified into the following types:

- Incidents and problems

  Incident and Problem. The whole process has the following phases: Record, Diagnose, Resolve, Confirm, and Close.
- Request fulfillment

  Change Role, Create Identity, Reset Password, and Logout Identity. The whole process has the following phases: Record, Approve, Handle, and Close.
- Change management

  Change, Version Upgrade, Hotfix Upgrade, and Configuration Upgrade. The whole process has the following phases: Record, Preliminary Approval, Modify Information, CAB Audit, ECAB Audit, Schedule Arrangement, Task Execution, Task Confirmation, Review, and Close.

Regulations

Each phase in the process template involves one or more tasks and each task corresponds to a handler. A regulation defines how to assign tasks to correct handlers.

The following regulations are supported:

- Assign by role
- Assign by user
- Assign by owner
- CAB/ECAB configuration

In practice, you can click a phase in the process template to configure the regulation.

> **Note** By default, if no regulations are configured in this phase, all users can view the current task in the task pool.

- Assign by role

  Select **Assign by Role** and then select roles from the drop-down list.
  - By default, if no role is selected, all users can view the current task in the task pool.
  - If the selected role has only one user, only that user can view the current task on the My Task tab.

  - If the selected role has more than one user, all users that have the selected role can view the current task in the task pool.

- Assign by user

  Select **Assign by User** and then select users from the drop-down list.

  - By default, if no user is selected, all users can view the current task in the task pool.
  - If only one user is selected, only that user can view the current task on the My Task tab.
  - If more than one user is selected, all of the selected users can view the current task in the task pool.

- Assign by owner

  If **Assign by Owner** is selected, only the user who creates the process request can view the current task on the My Task tab. The person who creates the request is the owner of the request.

- CAB/ECAB configuration

  CAB/ECAB Configuration only applies if you click the CAB Audit and ECAB Audit phases in a change management process. CAB stands for Change Advisory Board and ECAB stands for Emergency Change Advisory Board.

  Click **CAB/ECAB Configuration** to go to the **CAB/ECAB Configuration** page. For more information about how to configure CAB or ECAB, see Configure CAB or ECAB.

3.2.6. Configure CAB or ECAB

The change management process has the CAB Audit and ECAB Audit phases. Therefore, you must configure the CAB or ECAB.

Context

CAB and ECAB are terms of ITIL specifications. CAB stands for Change Advisory Board and ECAB for Emergency Change Advisory Board.

In all process templates, the CAB configuration of the CAB Audit phase is similar to the ECAB configuration of the ECAB Audit phase. This topic describes how to configure CAB.

If no rule is configured, all users can generate the current task in My Task by default. When one or more users are configured, each user can generate the current task in My Task, and the task can go to the next phase only after all users configured in this phase complete the current task.

Procedure

1. In the left-side navigation pane, choose **ITIL Management > CAB/ECAB Configuration**.
2. Click the **CAB Configuration** tab.
3. Select one or more users in the left section and then click » to add them to the right section.

   Users in the right section are the current CAB configuration.

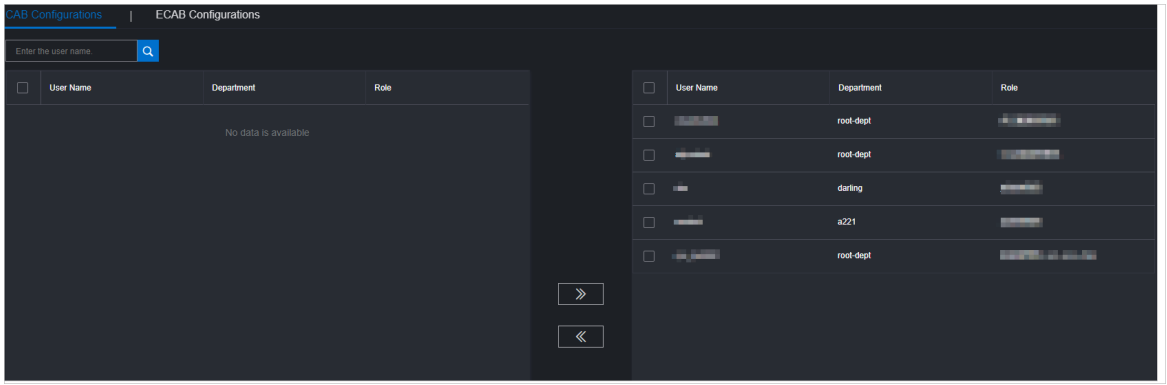

Note

- You can use the search box in the upper-left corner to search for users. Fuzzy search is supported.
- You can select one or more users in the right section and then click to remove them from the right section.

3.3. Configurations

3.3.1. Overview

The Configurations module allows you to modify the related configuration items of each product as required. To modify a configuration item of a product, you can modify the configuration value in Apsara Stack Operations (ASO) and then apply the modifications. To restore the configuration value of a modified configuration item, you can roll back the configuration value with one click.

You can also manage the kernel configurations and scan the configuration values of kernel configurations for a host.

3.3.2. Modify a configuration item of a product

You can modify a configuration item of a product.

Procedure

1. In the left-side navigation pane, choose **Configurations > Configuration Items** to go to the **Configuration Items** page.
2. In **Product** or **Configuration Name** search box, enter the name of the product or configuration item. Click **Search** to check whether the configuration item exists in the list.
   - If the configuration item exists in the list, you can perform the following operations:

     Click **Get** in the **Actions** column to load the actual data from the product to the ASO console.

     Click **Modify** in the **Actions** column. In the **Modify Configurations** dialog box that appears, modify the values and then click OK to modify the configuration item in the ASO console.
   - If the configuration item does not exist in the list, you can perform the following operations:

     You must add a configuration item in the following way:

a. Click **Add** in the upper-right corner of the page.

b. In the **Add Configuration** dialog box that appears, specify the information of the configuration item, such as **Product**, **Configuration Name**, **Default Value**, and **Data Source Type**.

c. Click **OK**.

Then, the created configuration item is displayed in the list. You can search for and modify this configuration item.

3. After the configuration item is modified, click **Apply** in the **Actions** column to apply the modifications.
4. (Optional)To import or export configuration items as a file, click **Import** or **Export** in the upper-right corner.

> **Note** To import configuration items as a file, we recommend that you export a file before the import and then complete the configurations based on the format in the exported file.

3.3.3. Restore the value of a modified configuration item

To restore the value of a modified configuration item, you can roll back the configuration item with one click.

Procedure

1. In the left-side navigation pane, choose **Configurations > Restore** to go to the **Restore** page.

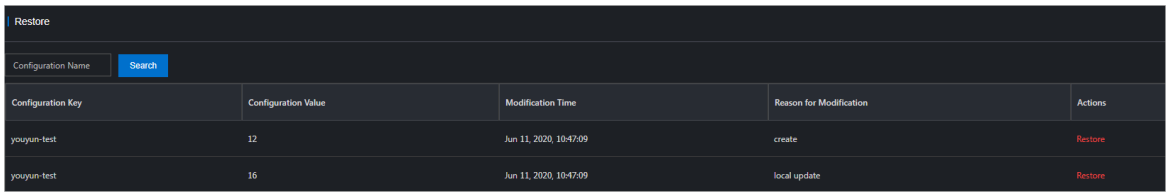

2. In the **Configuration Name** search box, enter the name of the configuration item that you want to roll back and then click **Search**. All modification records of the configuration item appear in the list.
3. Find the target record, and then click **Restore** in the **Actions** column.
4. In the message that appears, click **OK**.

3.3.4. Manage kernel configurations

You can add, modify, or delete a kernel configuration item.

Procedure

1. In the left-side navigation pane, choose **Configurations > Kernel Configuration** to go to the **Kernel Configurations** page.

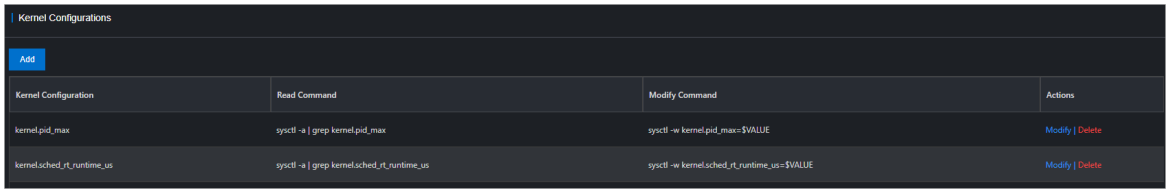

2. Perform the following operations:
   - Add a kernel configuration item

     In the upper part of the page, click **Add**. In the Add Configuration dialog box that appears, set **Configuration Name**, **Read Command**, and **Modify Command**. Then, click **Submit**.
   - Modify a kernel configuration item

     Find the kernel configuration item to be modified. Click **Modify** in the **Actions** column. Modify the values of **Kernel Configuration**, **Read Command**, and **Modify Command**. Then, click **Save**.
   - Delete a kernel configuration item

     Find the kernel configuration item to be deleted. Click **Delete** in the **Actions** column. In the message that appears, click **OK**.

3.3.5. Scan kernel configurations

You can scan kernel configuration items for a host to obtain their values.

Prerequisites

Before the scan, make sure that the following conditions are met:

- The kernel configuration items to be scanned exist in the list. For more information about how to add a kernel configuration item, see Manage kernel configurations.
- The hostname or IP address of the host is obtained from Apsara Infrastructure Management Framework.

Procedure

1. In the left-side navigation pane, choose **Configurations > Kernel Configuration Actions** to go to the **Kernel Configuration Actions** page.

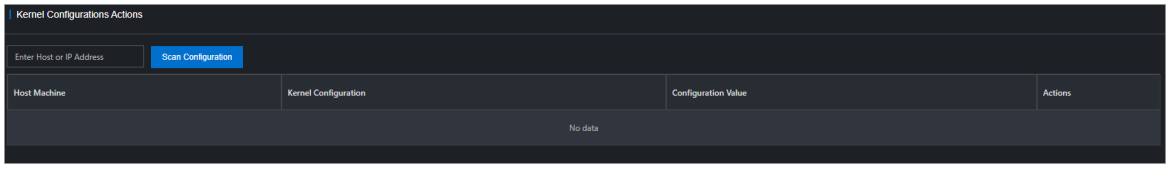

2. Enter the hostname or IP address in the search box and then click **Scan Configuration**.The scan results are displayed in the list.
3. (Optional)To modify the value of a scanned configuration item, click **Modify** in the Actions column and modify **Configuration Value** in the dialog box that appears. Click **Save**.After the modification, click **Apply** to apply the new value of the kernel configuration item to the host. To read the value of the kernel configuration item from the host again, click **Get**.

3.4. System Management

3.4.1. Overview

The System Management module allows you to centrally manage the departments, roles, and users involved in the ASO console. This makes it easy to grant different resource access permissions to different users. As the core for centralized permission management, the System Management module integrates the features such as department management, role management, logon policy management, user management, and password management.

3.4.2. Role management

You can add custom roles in the ASO console to more efficiently grant permissions to users.

Context

A role is a set of access permissions. You can assign different roles to different users to meet requirements for system access control. Roles are classified into basic roles and user-created roles. The basic roles, also known as atomic roles, are preset by the OAM system and cannot be modified or deleted by users. The user-created roles can be modified and deleted.

Procedure

1. In the left-side navigation pane, choose **System Management > Roles**.

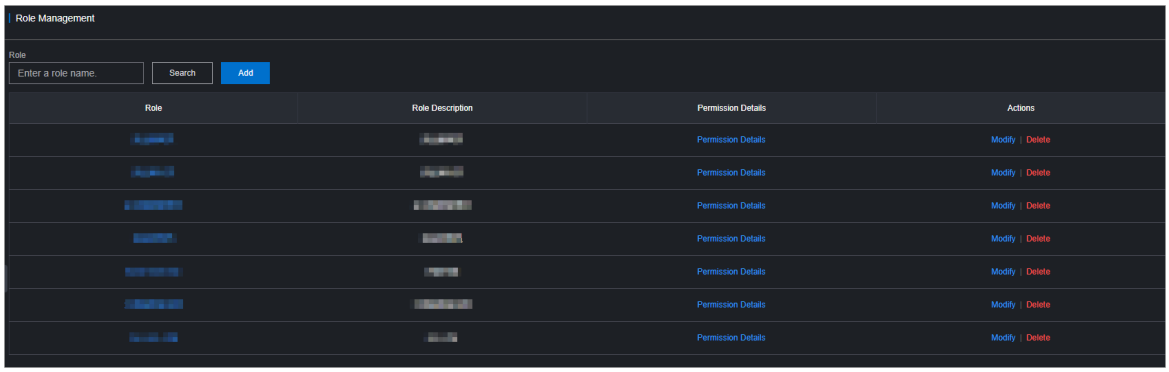

2. On the **Role Management** page that appears, perform the following operations:
   - Query roles

     > **Note** To query roles, you must have the ASO security officer role or system administrator role.

     In the upper-left corner of the page, enter a role name in the **Role** field, and then click **Search** to view the role information in the list.
   - Add a role

     > **Note** To add a role in the ASO console, you must have the ASO security officer role.

     Click **Add** in the upper part of the page. In the **Add** dialog box that appears, specify **Role Name**, **Role Description**, and **Basic Role**, and then click **OK**.
   - Modify a role

> **Note** To modify a user in the ASO console, you must have the ASO security officer role.

Find the role that you want to modify, and then click **Modify** in the **Actions** column. In the **Modify Role** dialog box that appears, modify the information, and then click **OK**.

- Delete a role

> **Notice** Before you delete a role, make sure that the role is not bound to any user. Otherwise, the role cannot be deleted.

Find the role that you want to delete, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

3.4.3. Department management

Department management allows you to create, modify, delete, and search for departments.

Context

By default, after ASO is deployed, a root department is created. You can create departments under the root department.

Departments are displayed in a hierarchy and you can create sub-departments under each level of departments. Up to five levels of departments can be created.

Procedure

1. In the left-side navigation pane, choose **System Management > Departments**.On the **Department Management** page, you can view the tree structure of all created departments, and the user information under each department.
2. (Optional)In the upper-left corner of the page, enter a role name in the search box and click the search icon to find the target department.
3. You can perform the following operations:
   - Add a department

     In the left catalog tree, select the department to which you want to add a department. Click the  icon in the upper part of the page and select **Add Department**. In the **Add Department** dialog box that appears, set **Department Name**, select a department administrator that have a proper role, and then click **OK**. Then, you can view the created department in the navigation tree.

     > **Note** When you add a department, you can select one or more department administrators.

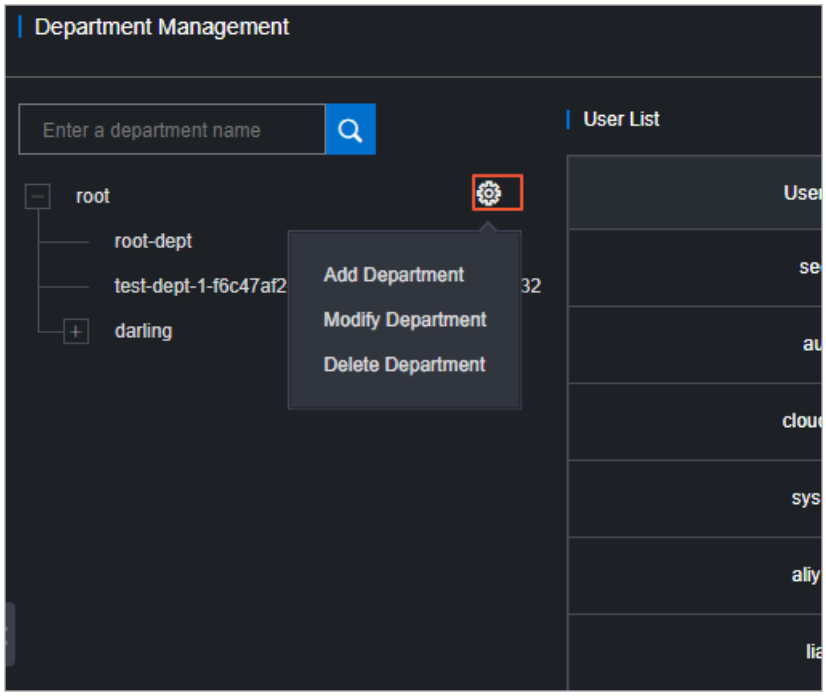

- Modify a department

  In the left catalog tree, select the department to be modified. Click the icon in the upper part of the page and select **Modify Department**. In the **Modify Department** dialog box that appears, set **Department Name**, select another department administrator with a proper role associated, and then click **OK**.

- Delete a department

  > **Notice** Before you delete a department, make sure that no users exist in the department. Otherwise, the department cannot be deleted.

  In the left catalog tree, select the department to be deleted. Click the icon in the upper part of the page and select **Delete Department**. In the message that appears, click **OK**.

- Add a user to a department

  In the left catalog tree, select the department to which you want to add a user. In the **User List** section on the right, click **Create User**. In the **Create User** dialog box that appears, enter user information, and then click **OK**.

  Then, you can view the added user information on the **Users** tab after you choose **System Management > Users**.

- Add a user group to a department

  In the left catalog tree, select the department to which you want to add a user group. In the **User Groups** section on the right, click **Create User Group**. In the **Create User Group** dialog box that appears, enter a user group name, and then click **OK**.

  Then, you can view the added user group information after you choose **System Management > Manage User Group**.

3.4.4. Region management

In multi-region scenarios, you can bind a department to a region as a system administrator. After that, users in the department can manage and view resources in the region.

Context

In multi-region scenarios, a region is managed by its own administrator. After an administrator logs on to the ASO console, the administrator can only manage resources in the authorized region.

Relationship between departments and regions:

- A department can be bound to multiple regions.
- A region can be bound to multiple departments.

Procedure

1. In the left-side navigation pane, choose **System Management > Region Management**.
2. (Optional)In the upper-left corner of the page, enter a department name and click the search icon.
3. Click the target department in the tree on the left and select one or more regions in the **Regions** list on the right.
4. Click **Update Association**.

3.4.5. Logon policy management

You can configure logon policies to control the logon time and IP addresses of users as an administrator.

Context

A default policy has been defined. You can configure logon policies to better control the read and write permissions of users and improve the system security.

Procedure

1. In the left-side navigation pane, choose **System Management > Logo Policy Management**.

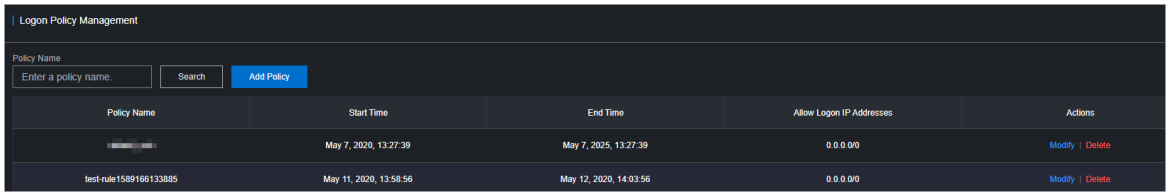

2. On the **Logon Policy Management** page that appears, perform the following operations:
   - Query policies

     In the upper-left corner of the page, enter a policy name in the **Policy Name** field, and then click **Search** to view the policy information in the list.
   - Add a policy

     Click **Add Policy** in the upper part of the page. In the Add Policy dialog box that appears, specify **Policy Name**, **Start Time**, **End Time**, and **IP addresses prohibited for logon**. Click **OK**.
   - Modify a policy

Find the policy that you want to modify, and then click **Modify** in the **Actions** column. In the **Update Policy** dialog box that appears, modify the information, and then click **OK**.

- Delete a policy

  Find the policy that you want to delete, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

3.4.6. User management

You can create users and assign different user roles to meet different requirements for system access control as an administrator.

Prerequisites

Before you create a user, make sure that the following requirements are met:

- A department is created. For more information, see Department management.
- A custom role is created if needed. For more information, see Role management.

Procedure

1. In the left-side navigation pane, choose **System Management > Users**.

   The **Users** tab appears.

2. On the **Users** tab, perform the following operations:

   - Query users

     > **Note** To search for users, you must have the security officer role or system administrator role.

     In the upper-left corner of the tab, configure the **User Name**, **Role**, and **Department** parameters, and then click **Search** to view the user information in the list.

   - Add a user

     > **Note** To add a user, you must have the security officer role.

     Click **Add** in the upper part of the tab. In the **Add User** dialog box that appears, configure the information, such as **User Name** and **Password**, and then click **OK**.

     The added user is displayed in the user list. The value of the **Primary Key Value** parameter is used for authentication when other applications call application API operations in ASO.

   - Modify a user

     > **Note** To modify a user, you must have the ASO security officer role.

     Find the user to be modified, and then click **Modify** in the **Actions** column. In the **Modify User** dialog box that appears, modify the parameters, and then click **OK**.

   - Delete a user

     Find the user to be deleted, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

> **Note** Deleted users are displayed on the **Recycled** tab. To restore a deleted user, click the **Recycled** tab. Find the user to be restored, click **Cleared** in the **Actions** column, and then click **OK**.

- Bind a logon policy

  Select a user in the user list. Click **Bind Logon Policy** to bind a logon policy to the user.

- Query personal information of the current user

  Move the pointer over the profile picture in the upper-right corner of the page, and select **Personal Information**. In the **Personal Information** dialog box that appears, view the personal information of the current user.

  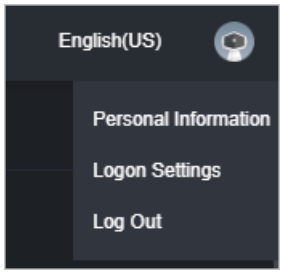

- Configure logon settings

  Move the pointer over the profile picture in the upper-right corner of the page, and select **Logon Settings**. In the **Logon Settings** dialog box that appears, configure Logon Timeout, Multiple-Terminal Logon Settings, Maximum Allowed Password Retries, Account Validity, and Logon Policy. Click **Save**.

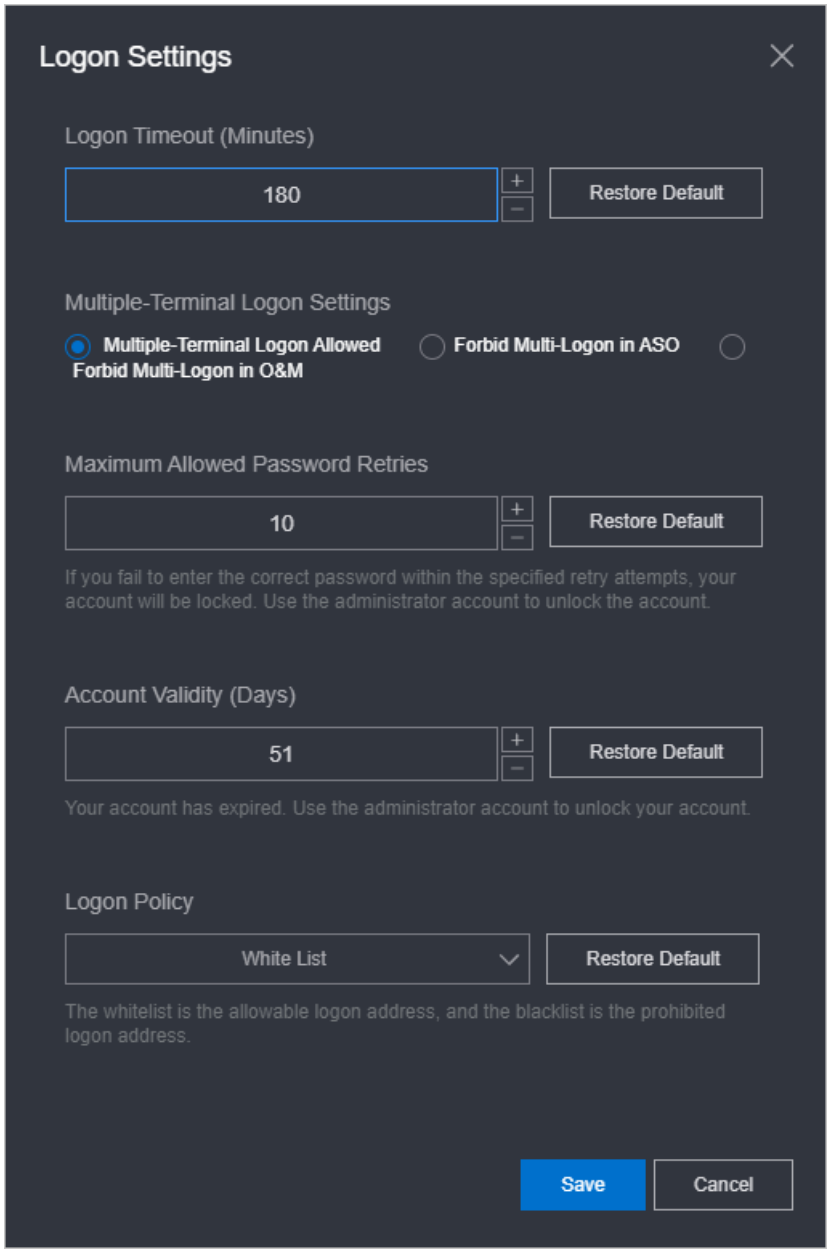

3.4.7. User group management

You can add multiple users to a user group and add the same roles to them as an administrator for centralized management.

Create a user group

1. In the left-side navigation pane, choose **System Management > Manage User Group**.
2. In the upper-right corner of the Manage User Group page, click **Create User Group**.
3. In the **Create User Group** dialog box that appears, enter a user group name, select a department, and then click **OK**.

Modify the name of a user group

1. In the left-side navigation pane, choose **System Management > Manage User Group**.

2. (Optional)In the upper-left corner of the page, configure the Department Name, User Group Name, User Name, and Role Name parameters, and then click **Search**. If you have defined filter conditions, you can click **Clear** to remove the conditions with one click.
3. In the user group list, find the target user group and click **Modify User Group** in the **Operation** column.
4. In the dialog box that appears, modify the user group name.
5. Click **OK**.

Add a user to a user group

1. In the left-side navigation pane, choose **System Management > Manage User Group**.
2. (Optional)In the upper-left corner of the page, configure the Department Name, User Group Name, User Name, and Role Name parameters, and then click **Search**. If you have defined filter conditions, you can click **Clear** to remove the conditions with one click.
3. In the user group list, find the target user group and click **Add User** in the **Operation** column.
4. In the dialog box that appears, select one or more users from the **Users in the Department** list and add the users to the Users to Add list.
5. Click **OK**. After that, the newly added user is displayed in the **User** column corresponding to the user group.

Add a role to a user group

You can add only one role to a user group.

1. In the left-side navigation pane, choose **System Management > Manage User Group**.
2. (Optional)In the upper-left corner of the page, configure the Department Name, User Group Name, User Name, and Role Name parameters, and then click **Search**. If you have defined filter conditions, you can click **Clear** to remove the conditions with one click.
3. In the user group list, find the target user group and click **Add Role** in the **Operation** column.
4. From the **Role** drop-down list, select a role.
5. Click **OK**.

   After that, the newly added role is displayed in the **Role** column corresponding to the user group. All users in the user group are granted the permissions of this role.

Delete a role

1. In the left-side navigation pane, choose **System Management > Manage User Group**.
2. (Optional)In the upper-left corner of the page, configure the Department Name, User Group Name, User Name, and Role Name parameters, and then click **Search**. If you have defined filter conditions, you can click **Clear** to remove the conditions with one click.
3. In the user group list, find the target user group and click **Delete Role** in the **Operation** column.
4. In the message that appears, click **OK**. After that, the deleted role is not displayed in the **Role** column corresponding to the user group. The users in the user group do not have the permissions of the role.

Delete a user

1. In the left-side navigation pane, choose **System Management > Manage User Group**.

2. (Optional)In the upper-left corner of the page, configure the Department Name, User Group Name, User Name, and Role Name parameters, and then click **Search**. If you have defined filter conditions, you can click **Clear** to remove the conditions with one click.
3. In the user group list, find the target user group and click **Delete Users** in the **Operation** column.
4. In the dialog box that appears, select one or more users from the **Users in the Group** list and add the users to the Users to Delete list.
5. Click **OK**. After that, the deleted user is not displayed in the **User** column corresponding to the user group.

Delete a user group

> **Notice** Before you delete a user group, make sure that users or roles are added to the user group.

1. In the left-side navigation pane, choose **System Management > Manage User Group**.
2. (Optional)In the upper-left corner of the page, configure the Department Name, User Group Name, User Name, and Role Name parameters, and then click **Search**. If you have defined filter conditions, you can click **Clear** to remove the conditions with one click.
3. In the user group list, find the target user group and click **Delete User Group** in the **Operation** column.
4. In the message that appears, click **OK**.

3.4.8. Two-factor authentication

To improve the security of user logon, you can configure two-factor authentication for users.

Context

ASO supports the following authentication methods. You can use one of the following authentication methods:

- Google two-factor authentication

  This authentication method uses a password and mobile app to provide two layers of protection for accounts. You can obtain the logon key after you configure users in ASO, and then enter the key in the Google Authenticator app of your mobile phone. The app dynamically generates a verification code for logon based on the time and key.

- USB key authentication

  If you use this authentication method, you must install the drive and browser controls (only Windows + IE 11 environment is supported) based on the third-party manufacturer instructions. The third-party manufacturer provides the USB key hardware and the service for authentication and verification of certificates. The USB key contains the serial number and certificate information. You must bind the user account and the serial number on the management page of the two-factor authentication, and configure the authentication server provided by the third-party manufacturer. Then, you can enable the USB key authentication for the user.

If the USB key authentication is enabled for the account, upon logon, the ASO frontend will call the browser controls, read the certificate in the USB key, obtain the random code from the backend, encrypt the information, and then send the information to the backend. The backend calls the authentication server to parse the encrypted strings, verifies the certificate and serial number, and then completes the other logon processes if the verification is successful.

- PKI authentication

  If you use this authentication method, you must enable ASO HTTPS mutual authentication and change the certificate provided by the user. The third-party manufacturer makes the certificate and verifies the certificate at the backend. After HTTPS mutual authentication is enabled, the request carries the Client certificate upon logon and is passed to the backend. The backend calls the DNS and verification services of the third-party manufacturer for verification. The certificate includes the name and ID card number of a user. Therefore, bind the name and ID card number with a user account when you configure the authentication method in ASO.

Both USB key authentication and PKI authentication depend on the authentication server provided by the third-party manufacturer to verify the encrypted information or certificate provided upon logon. Therefore, you must add the authentication server configurations before you use these two authentication methods.

Google two-factor authentication is implemented based on public algorithms. Therefore, no third-party authentication service is required, and you are not required to configure the authentication server.

Procedure

1. In the left-side navigation pane, choose **System Management > Two Factor Authentication**.
2. On the Two Factor Authentication page, you can perform the following operations:
   - Google two-factor authentication
     a. Set **Current Authentication Method** to **Google Two-Factor Authentication**.
     b. Click **Add User** in the upper-right corner of the page. In the Add User dialog box, enter a username and click OK. The added user is displayed in the user list.
     c. Find the user for whom you want to enable Google two-factor authentication, and then click **Create Key** in the **Actions** column. After the **Added** message appears, **Show Key** is displayed in the **Actions** column. Click **Show Key**, and the key is displayed in plain text.
     d. Enter the key in the Google Authenticator app on your mobile phone. The app dynamically generates a verification code for logon based on the time and key. With two-factor authentication enabled, you are required to enter the verification code on your app when you log on to the system.

        > **Note** The Google Authenticator app and server generate the verification code by using public algorithms and based on the time and key, and can work offline without connecting to the Internet or Google server. Therefore, you must keep the key safe.

     e. To disable two-factor authentication, click **Delete Key** in the **Actions** column.
   - USB key authentication
     a. Set **Current Authentication Method** to **USB Key Authentication**.

b. In the upper-right corner of the **Authentication Server Configuration** section, click **Add Server**. In the Add User dialog box, specify the **IP Address** and **Port** parameters for the server, and then click OK. The added server is displayed in the server list. Click **Test** to test the connectivity of the authentication server.

c. In the upper-righter corner of the **User List** section, click **Add User**. The added user is displayed in the user list.

d. Find the user for whom you are about to enable the USB key authentication, and then click **Bind Serial Number** in the **Actions** column. In the dialog box, enter the serial number to bind the user account with this serial number.

> **Note** When you add an authentication method to ASO, ASO calls the browser controls to automatically enter the serial number. If the serial number fails to be entered, you must enter it manually. The serial number of USB key authentication is stored within the USB key. Insert the USB key, install the drive and browser controls, and then read the serial number by using the browser controls.

e. Then, click **Enable Authentication** in the **Actions** column.

- PKI authentication

a. Set **Current Authentication Method** to **PKI Authentication**.

b. In the upper-right corner of the **Authentication Server Configuration** section, click **Add Server**. In the Add Server dialog box, specify the **IP Address** and **Port** parameters for the server. The added server is displayed in the server list. Click **Test** to test the connectivity of the authentication server.

c. In the **User List** section, click **Add User**. In the Add User dialog box, specify **Username**, **Full Name**, and **ID Card Number**, and then click OK. The added user is displayed in the user list.

d. (Optional)Find the user for whom you want to enable the PKI authentication, and then click **Bind** in the **Actions** column. Enter the full name and ID card number of the user to bind the user account with the name and ID card number.

e. Then, click **Enable Authentication** in the **Actions** column.

- No authentication

Set **Current Authentication Method** to **No Authentication**. Two-factor authentication is then disabled and all two-factor authentication methods become invalid.

3.4.9. Application whitelists

You can add, modify, or delete application whitelists as a system administrator.

Context

All access permissions on ASO services are managed by Operation Administrator Manager (OAM). Therefore, if an account does not have a corresponding role, it will not be allowed to access ASO services. The application whitelist feature allows you to access ASO services in scenarios where no permissions are granted. With the whitelist feature enabled, the application can be accessed by all users who have logged on. The valid application whitelist permissions are read-only and read/write. The configured value is the logon user permission.

The application whitelist is managed by the system administrator. You can access this page after you log on as a system administrator.

When you add a whitelist, enter the product name and service name. The current product name is aso, and the service name is the name of the backend service registered in ASO. The whitelist takes effect only if the configurations are valid.

Procedure

1. In the left-side navigation pane, choose **System Management > Application Whitelist**.

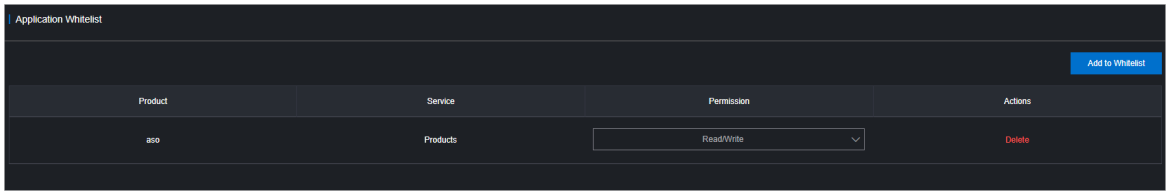

2. On the **Application Whitelist** page that appears, perform the following operations:
   - Add a whitelist

     In the upper-right corner, click **Add to Whitelist**. In the **Add to Whitelist** dialog box that appears, select the service and permission, and then click **OK**.
   - Modify permissions

     Set the service permission to **Read/Write** or **Read-only** in the **Permission** column.
   - Delete a whitelist

     Find the whitelist to be deleted, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

3.4.10. Server password management

The Server Password module allows you to configure and manage server passwords and search for history passwords in the Apsara Stack environment.

Context

Server password management covers passwords of all the servers in the Apsara Stack environment.

- The system automatically collects information of all the servers in the Apsara Stack environment.
- The server password is automatically updated periodically.
- You can configure the password expiration period and password length.
- You can manually update the password of one or more servers at a time.
- The system records the history of server password updates.
- You can search for server passwords by product, hostname, or IP address.

Procedure

1. In the left-side navigation pane, choose **System Management > Server Password**.

   The **Password Management** tab appears. The **Password Management** tab shows the passwords of all the servers in the current Apsara Stack environment.

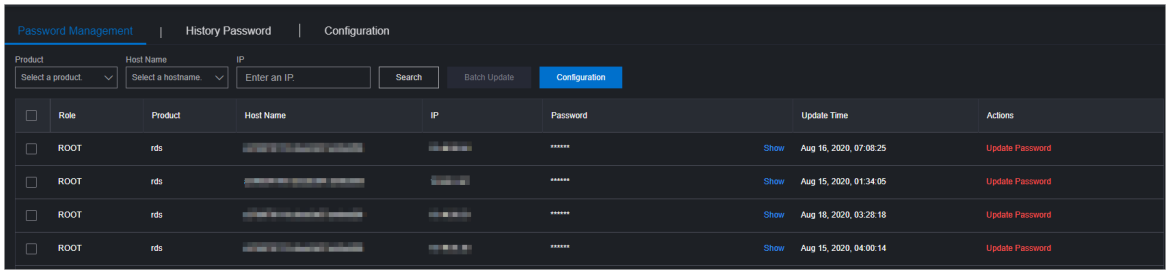

2. Perform the following operations:
   - Search for servers

     On the **Password Management** tab, select a product, server name, or IP address, and then click **query** to search for specific servers.
   - Show a password
     a. On the **Password Management** tab, find a server.
     b. Click **Show** in the **Password** column. The host password in plain text is displayed and turns into cipher text after 10 seconds. Alternatively, click **Hide** to show the cipher text.
   - Update a password
     a. On the **Password Management** tab, find a server.
     b. Click **Update Password** in the **Actions** column.
     c. In the **Update Password** dialog box, specify **Password** and **Confirm Password**, and then click **OK**.

        Then, the password of the corresponding server is updated.
   - Update multiple passwords
     a. On the **Password Management** tab, select multiple servers.
     b. Click **Batch Update** in the upper part of the tab.
     c. Specify **Password** and **Confirm Password**, and then click **OK**.

        The passwords of the selected servers are updated.
   - Configure the password expiration period
     a. On the **Password Management** tab, select one or more servers.
     b. Click **Configuration** in the upper part of the tab.
     c. In the **Configuration Item** dialog box, specify **Password Expiration Period** and **Unit**. Click **OK**.

        Server passwords are updated immediately after the configuration and will be updated again after an expiration period.
   - View the history of server password updates

     Click the **History Password** tab. Select a product, hostname, or IP address, and then click Search to view the history of server password updates in the search results.
   - Show historical passwords of servers
     a. On the **History Password** tab, find a server.

b. Click **Show** in the **Password** column. The host password in plain text is displayed and turns into the cipher text after 10 seconds. Alternatively, you can click **Hide** to show the cipher text.

- View and modify the password configuration policy

  Click the **Configuration** tab. On the **Configuration** tab, view the metadata of server password management, including the initial password, password length, and retry times. Notes:

  - The initial password is the one assigned when server password management is deployed in the Apsara Stack environment. This parameter is important, which is used to update the password of a server in the Apsara Stack environment.
  - The password length is the length of passwords automatically updated by the system.
  - Retry times is a limit of how many times a password can fail to be updated before the system stops trying to update it.

  To modify the configurations, click **Modify Configurations** in the **Actions** column. In the Modify Configurations dialog box, specify **Initial Password**, **Password Length**, and **Retry Times**. Click **OK**.

3.4.11. Operations logs

You can view logs to know the usage of all resources and the running status of all function modules on the platform in real time.

Context

The Operation Logs page allows you to view all the records of backend API calls, including audit operations. The auditor can filter logs by username and time period, and view the call details. You can also export the selected logs.

Procedure

1. In the left-side navigation pane, choose **System Management > Operations Logs**.
2. On the **Log Management** page, perform the following operations:
   - Query logs

     In the upper-left corner of the page, specify **User Name** and **Time Period**, and then click **Search**.
   - Delete logs

     Select one or more logs to be deleted, and then click **Delete** in the upper part of the page. In the message that appears, click **OK**.
   - Export logs

     Click the icon to export the displayed logs.

     > **Note** If the number of logs to be exported exceeds the threshold (10,000 by default), only the first 10,000 logs can be exported.

3.4.12. View authorization information

The Authorization page allows customers, field engineers, and operations engineers to query services that have authorization problems and troubleshoot the problems.

Prerequisites

Make sure that the current logon user has administrator permissions. Only after you are granted administrator permissions, you can view the trial authorization information or enter the authorization code to view the formal authorization information on the **Authorization Details** tab.

When you access this page if you are not granted administrator permissions, a message indicating that the user has insufficient permissions is displayed.

Procedure

1. In the left-side navigation pane, choose **System Management > Authorization**. The **Authorization Details** tab appears.

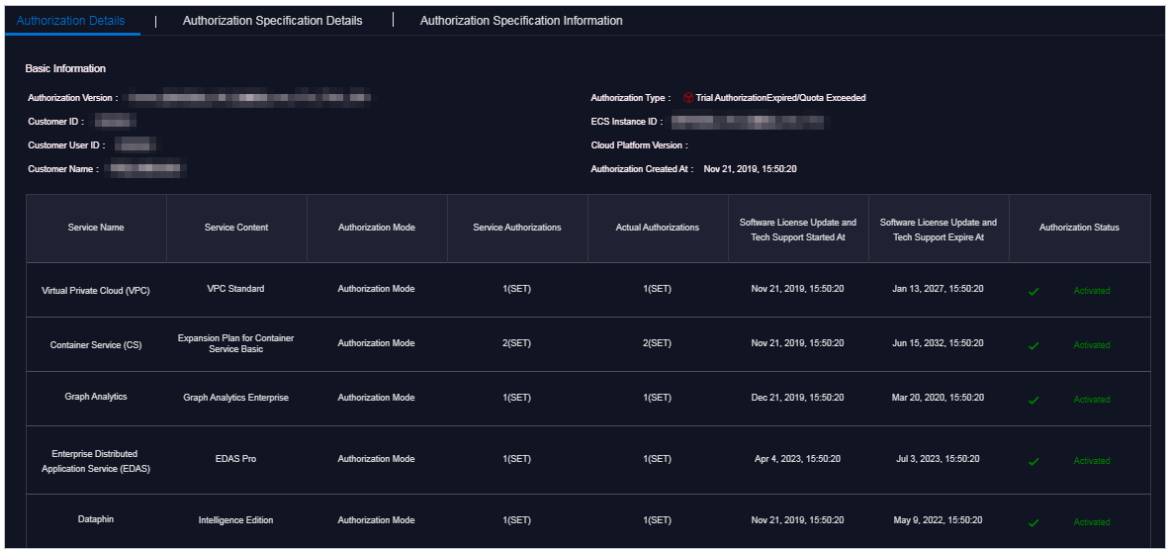

2. Perform the following operations to view the authorization information.

> **Note** For formal authorization, you must enter the authorization code to view the authorization information. Obtain the authorization code in the authorization letter attached by the project contract or contact the business manager (CBM) of your project to obtain the authorization code.

   - On the **Authorization Details** tab, view the basic authorization information.

     You can view authorization information, including authorization version, customer information, authorization type, Elastic Compute Service (ECS) instance ID, cloud platform version, the creation time of authorization, and the authorization information of all services within the current Apsara Stack environment.

     The following table describes the detailed authorization information.

Item	Description

Item	Description
Authorization Version	You can use the BP number in the version to associate with a project or contract. Note: ■ **TRIAL** in the version indicates that the authorization is trial authorization. The trial authorization is valid within 90 days from the date of deployment. ■ **FORMAL** in the version indicates that the authorization is a formal one. The authorization information of the service comes from the signed contract.
Authorization Type	Indicates the current authorization type and status.
Customer information	Includes the customer name, customer ID, and customer user ID.
ECS Instance ID	The ECS instance ID in the deployment planner of the field environment.
Cloud Platform Version	The Apsara Stack version of the current cloud platform.
Authorization Created At	The start time of the authorization.
Authorization information of a service	Includes the service name, service content, current authorization mode, service authorization quantity, actual authorization quantity, software license update and technical support start time, software license update and technical support end time, and real-time product authorization status. If the following information appears in the **Authorization Status** column of a service: ■ **RENEW Service Expired** Indicates that the customer must renew the subscription as soon as possible. Otherwise, field operations services (including ticket processing) will be terminated. ■ **Quota Exceeded** Indicates that the specifications deployed for a service have exceeded the contract quota, and the customer must scale up the service as soon as possible.

- Click the **Authorization Specification Details** tab to view the authorization specification information of a service.

  The following table describes the authorization specification information.

Item	Description
Service Name	The name of an authorized service.
Specification Name	The specification name of an authorized service.

Item	Description
Specifications	The total number of current authorizations of a specification for a service.
Specification Quota	The authorization quota of a specification for a service.
Specification Status	The current authorization status of a specification for a service.

- Click the **Authorization Specification Information** tab to view the authorization specification information and the authorization specification excess information of services.

  In the upper part of the tab, specify Licensing Specification Level as IDC Level, select IDC ID, service name, start time, and end time, and then click **Search**. You can view the authorization specification information of a service in the current environment, including the maximum and minimum number of specifications and their occurrence time points as well as the average number of specifications within the specified time range.

  In the **Authorization Specification Information** or **Authorization Specification Excess Information** section, click the + icon to the left of a service to view the specifications, specification quota, and recorded time of authorization specifications on the latest day of the specified time range for the specification of the service. Click **View More** to view the authorization specification information of the service within the specified time range by date.

3.4.13. Multi-cloud management

The Multi-cloud Management module provides the function of multi-cloud configurations. By using the multi-cloud configurations, you can perform Operations & Maintenance (O&M) operations on different data centers on an operations and maintenance platform.

3.4.13.1. Add multi-cloud configurations

If a multi-cloud environment is used, you can add multi-cloud configurations as a multi-cloud configuration administrator or super administrator. After that, you can switch to different data centers in the same console and then view or perform related operations.

Prerequisites

Before you add multi-cloud configurations, you must ensure that:

- Data centers are interconnected and share accounts with the same usernames and passwords.
- You are granted the permissions of a multi-cloud configuration administrator or super administrator.

Procedure

1. Log on to the ASO console as a multi-cloud configuration administrator or super administrator.
2. In the left-side navigation pane, choose **System Management > Multi-cloud Management**.
3. In the upper part of the page, click **Add**.

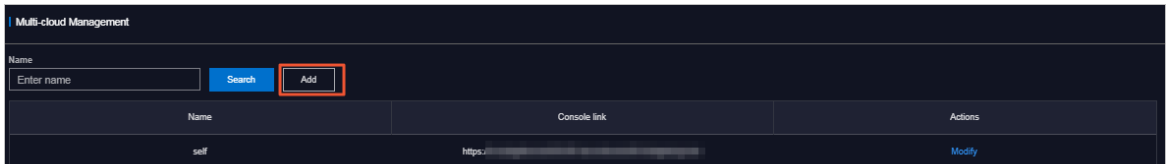

4. Add the console link of another data center, and then click **OK**.

Parameter	Description
Name	The name of another data center.
Console link	The console link of another data center. Ensure the console link is correct. Otherwise, an error message will be returned.

After that, you can log on to the ASO console with a shared account to switch to different data centers and then perform related operations.

3.4.13.2. Modify the name of a data center

After you add multi-cloud configurations, you can modify the name of a data center as a multi-cloud configuration administrator or super administrator.

Procedure

1. In the left-side navigation pane, choose **System Management > Multi-cloud Management**.
2. (Optional)Enter the target name in the Name search box and then click **Search**.
3. Find the target name and click **Modify** in the **Actions** column.
4. In the dialog box that appears, modify the name of the data center and click **OK**.

3.4.14. Menu settings

You can hide, add, modify, or delete a system menu based on business needs.

3.4.14.1. Add a level-1 menu

This topic describes how to add a level-1 menu.

Procedure

1. In the left-side navigation pane, choose **System Management > Menu Configuration**.
2. In the upper part of the page, click **Add**.
3. In the Add Level-1 Menu pane that appears, configure the menu parameters.

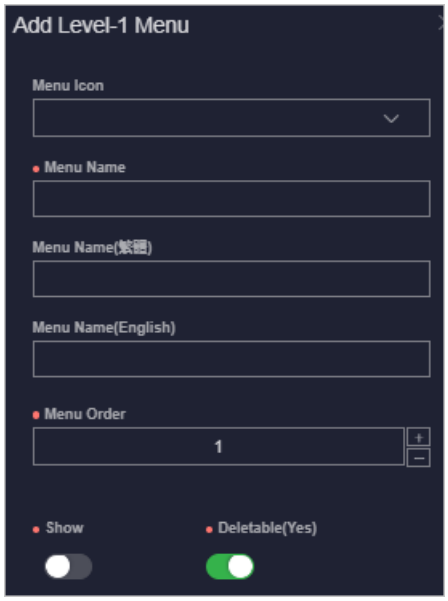

The following table describes the configuration of the parameters.

Parameter	Description
Menu Icon	Select the icon of the target level-1 menu from the drop-down list.
Menu Name	Specifies the name of the menu.
Menu Order	Specifies the order of items of this menu from top to bottom.
Show or Hide	Specifies whether to show the menu. Toggle the switch to hide or show the menu. By default, the menu is displayed.
Deletable	Specifies whether this menu can be deleted after being added. Toggle the switch to configure whether the menu can be deleted. By default, the menu can be deleted. This parameter cannot be modified after being specified.

4. Click **OK**.

Result

Then, you can view the added level-1 menu in the menu list and in the left-side navigation pane.

3.4.14.2. Add a submenu

This topic describes how to add a level-2 and a level-3 menu.

Procedure

1. In the left-side navigation pane, choose **System Management > Menu Configuration**.
2. Add a level-2 menu.
   i. Find the level-1 menu to which you want to add a level-2 menu, and then click **Add** in the **Actions** column.

ii. In the Add Submenu pane, configure the submenu parameters.

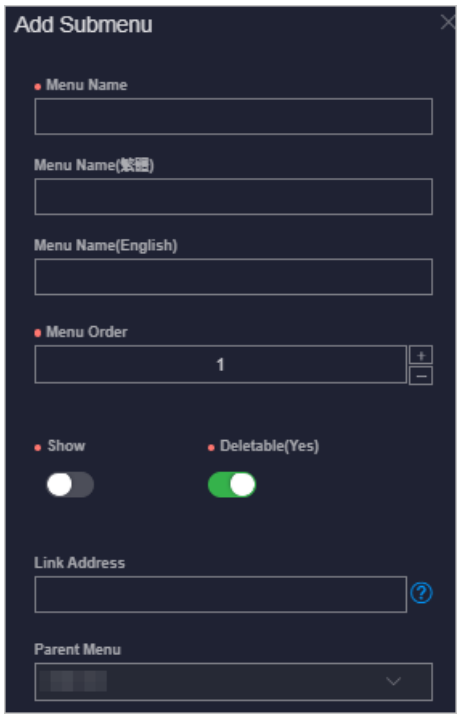

The following table describes the parameters.

Parameter	Description
Menu Name	Specifies the name of the level-2 menu.
Menu Order	Specifies the order of items of this level-2 menu from top to bottom.
Show or Hide	Specifies whether to hide this level-2 menu. Turn on or off the switch to hide or show the menu. By default, the menu is not hidden.
Deletable	Specifies whether this level-2 menu can be deleted after being added. Turn on or off the switch to configure whether the menu can be deleted. By default, the menu can be deleted. The settings cannot be modified after being configured.
Link Address	Specifies the menu path in the format of module name/path name. Example: /Dashboard/#/dashboardView.
Parent Menu	The parent menu of this menu.

iii. Click **OK**.
You can view the added level-2 menu under the corresponding level-1 menu in the menu list and the left-side navigation pane.

3. Click the fold button on the left side of the level-1 menu to expand the level-2 menus. Add a level-3 menu by following the preceding steps.

> **Note** The system only supports up to three levels of menus. You cannot add submenus for a level-3 menu.

After you add a level-3 menu, you can view it under the corresponding level-2 menu in the menu list and the left-side navigation pane.

3.4.14.3. Hide a menu

This topic describes how to hide a menu.

Prerequisites

> **Notice** You cannot hide the **System Management** menu and its submenus.

Procedure

1. In the left-side navigation pane, choose **System Management > Menu Configuration**.
2. Perform the following operations:
   - Hide a level-1 menu

     In the menu list, find the level-1 menu you are about to hide, and then click **Modify** in the **Actions** column. In the Modify Menu pane, turn on the switch to hide the menu, and then click **OK**.
   - Hide a level-2 or level-3 menu

     In the menu list, find the level-2 or level-3 menu you are about to hide, and then click **Modify** in the **Actions** column. In the Modify Menu pane, turn on the switch to hide the menu, and then click **OK**.

3.4.14.4. Modify a menu

This topic describes how to modify the icon, name, and order of a menu.

Procedure

1. In the left-side navigation pane, choose **System Management > Menu Management**.
2. In the menu list, find the menu or submenu to be modified. Click **Modify** in the **Actions** column.
3. In the Modify Menu pane, modify the icon, name, and order of a level-1 menu, or modify the name, order, and link address of a submenu.

3.4.14.5. Delete a menu

This topic describes how to delete a menu that is no longer needed.

Prerequisites

> **Notice** You can only delete a menu that had **Deletable** enabled when the menu was added.

Procedure

1. In the left-side navigation pane, choose **System Management > Menu Configuration**.
2. In the menu list, find the menu or submenu to be deleted. Click **Delete** in the **Actions** column.
3. In the message that appears, click **OK**.

4.Monitoring

4.1. Daily monitoring

4.1.1. Operations and maintenance

The Operations and Maintenance module displays the current usage and monitoring information of system resources by using graphs and a list and allows you to check the current operating conditions of the system.

In the left-side navigation pane, choose **Operations and Maintenance > Dashboard**.

The **Dashboard** page displays the product versions, inventory statistics, and alert statistics of the current console. By viewing the dashboard, operations engineers can know the overall operating conditions of Apsara Stack products in a timely manner. By viewing the dashboard, you can learn the overall operating conditions of Apsara Stack services.

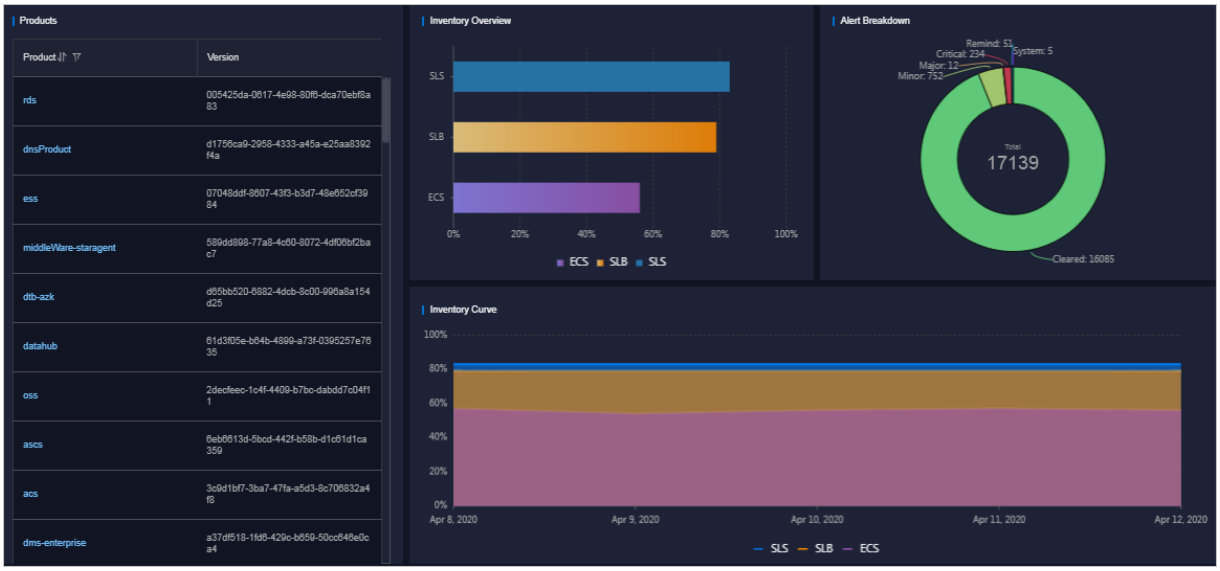

4.1.2. Alert Monitoring

The Alert Monitoring module allows operations engineers to quickly know the information of alerts generated by the system, locate the problems based on the alert information, track the problem processing, and configure the alerts.

4.1.2.1. Dashboard

The Alert Monitoring module allows you to view the overview information of alerts.

Context

You can configure filter conditions to filter alerts by adding a custom filter.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Dashboard**.

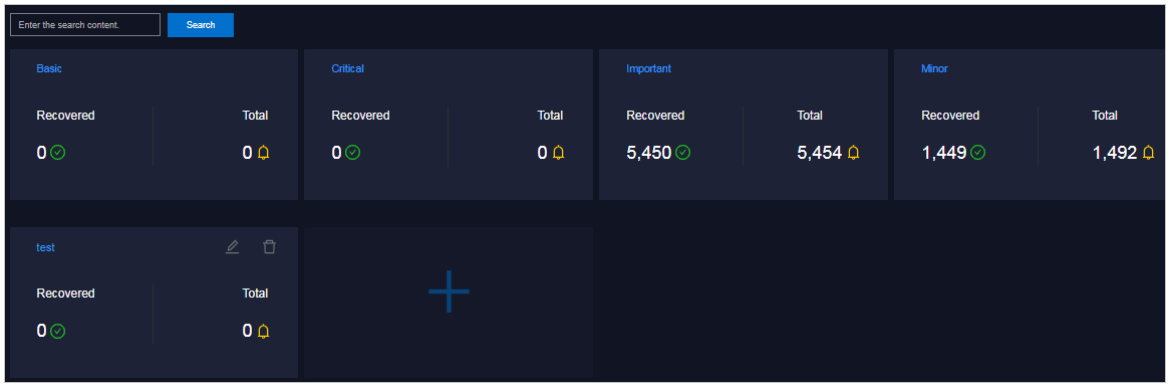

2. Perform the following operations:
   - View the total number of alerts and the number of recovered alerts in the basic, critical, important, and minor monitoring metrics, as well as custom filters.

     > **Note** Click a monitoring metric or custom filter to go to the corresponding **Alert Events** page.

   - Search for alerts

     Enter a keyword, such as cluster, product, service, severity, status, or monitoring metric name, in the search box. Click **Search** to search for the corresponding alert event.

   - Add a custom filter

     Click the icon. In the Add Filter pane, configure the parameters.

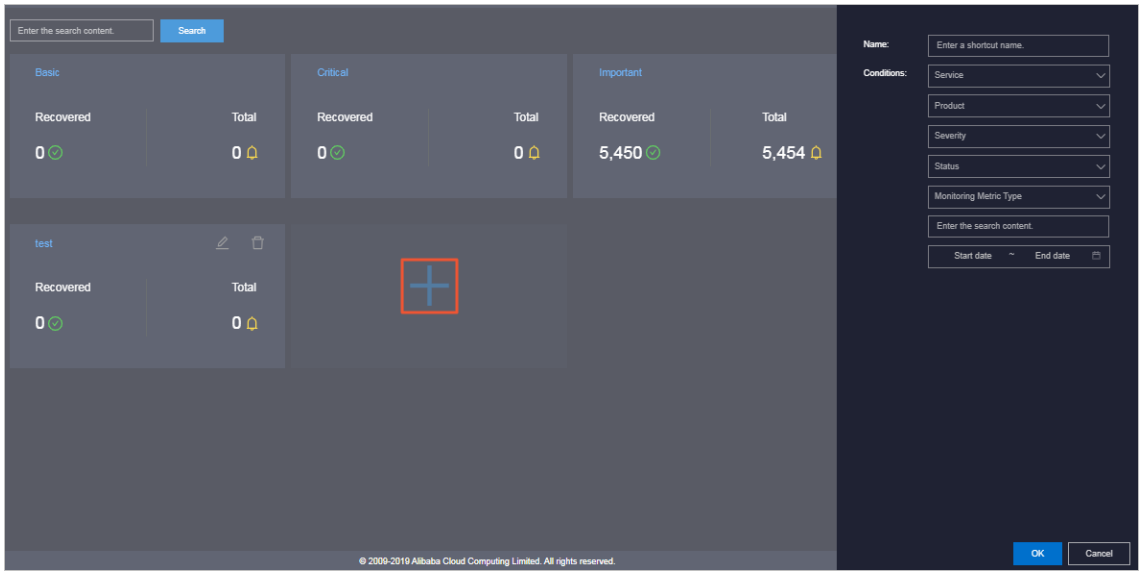

The following table describes the parameters for adding a filter.

Parameter	Description
Name	The filter name to be displayed on the **Dashboard** page.

Parameter	Description
Conditions	Configure the following filter conditions: ■ **Service**: the service to which the alerts to be filtered belong. ■ **Product**: the product to which the alerts to be filtered belong. ■ **Severity**: the severity of the alerts to be filtered. Alert levels are classified into the following types: ■ **P0**: indicates the cleared alerts, corresponding to alerts whose **Alert Level** is **Restored** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ■ **P1**: indicates the critical alerts, corresponding to alerts whose **Alert Level** is **P1** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ■ **P2**: indicates major alerts, corresponding to alerts whose **Alert Level** is **P2** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ■ **P3**: indicates the minor alerts, corresponding to alerts whose **Alert Level** is **P3** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ■ **P4**: indicates the alerts for notice, corresponding to alerts whose **Alert Level** is **P4** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ■ **P5**: indicates the system alerts. ■ **Status**: the current status of the alerts to be filtered. ■ **Monitoring Metric Type**: the type of the metric to which the alerts to be filtered belong. Valid values: ■ **Basic** ■ **Critical** ■ **Important** ■ **Minor** ■ **Enter the search content**: the information about the alerts to be filtered. ■ Select the start date and end date of the alerts to be filtered.

After you add a custom filter, you can view the overview information that meets the filter conditions on the **Dashboard** page.

- Modify a custom filter

  After you configure a custom filter, you can click the icon to modify the filter conditions and obtain the new filter results.

- Delete a custom filter

  After you add custom filters, you can click the icon to delete a filter that is no longer needed.

4.1.2.2. Alert events

The Alert Events module displays all alerts aggregated by metric and product name on different tabs. You can filter alerts with conditions such as monitoring metric type, product, service, severity, status, and time range. Then you can perform O&M on alerts.

Context

The Alert Events module contains the following tabs:

- **Hardware & System**: displays alert information related to the hardware or system in the Apsara stack environment.
- **Base Module**: displays the alert information related to base services such as baseserviceAll, webappAll, middlewareAll, https-proxy, dns, dnsProduct, and minirds.
- **Monitoring & Management**: displays the alert information related to the cloud management services except base and cloud services.
- **Cloud Product**: displays alert information related to cloud services such as OSS, ECS, SLB, VPC, RDS, DataWorks, DTS, NAS, MaxCompute, DataHub, Realtime Compute, Graph Analytics, AnalyticDB for MySQL, Apsara Stack Security (Advanced), Apsara Stack Security (Basic), EDAS, QuickBI, Graph Compute, Elasticsearch, and Tablestore.
- **Timeout Alert**: displays all timeout alerts.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Events**.

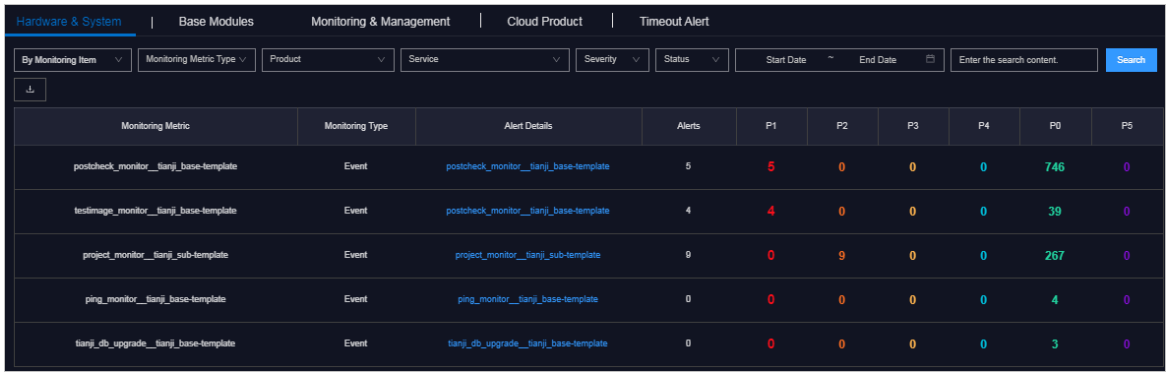

2. Click the **Hardware & System**, **Base Modules**, **Monitoring & Management**, **Cloud Product**, or **Timeout Alert** tab, and perform the following operations:
   - Search for an alert

     In the upper part of the tab, search for an alert by specifying **Monitoring Metric Type**, **Product**, **Service**, **Severity**, **Status**, **Start Date**, **End Date**, and search content.
   - View alert sources
     a. If the alert information is aggregated by **Product Name** on this tab, click + to the left of the product name to show the monitoring metrics. If the alert information is aggregated by **Monitoring Item** on this tab, skip this step.
     b. Find the monitoring metric and severity of the target alert, and then click the number in the specific severity column.

c. Move the pointer over the alert source information in blue in the **Alert Source** column to view the alert source details.

- View the details of a metric
  a. If the alert information is aggregated by **Product Name** on this tab, click + at the left of the product name to show the monitoring metrics. If the alert information is aggregated by **Monitoring Item** on this tab, skip this step.
  b. Find the monitoring metric and severity of the target alert, and then click the number in the specific severity column.
  c. Click the alert details in blue in the **Alert Details** column. On the **Alert Details** page that appears, view the alert information such as the alert description, reference, impact scope, and resolution.
- View the original alert information of an alert
  a. If the alert information is aggregated by **Product Name** on this tab, click + to the left of the product name to show the monitoring metrics. If the alert information is aggregated by **Monitoring Item**, skip this step.
  b. Find the monitoring metric and severity of the target alert, and then click the number in the specific severity column.
  c. Click the number in blue in the **Alerts** column. The **Alerts** pane appears.
  d. Click **Details** in the **Alert Information** column to view the original alert information.
- Handle alerts

  Find the monitoring metric and severity of the target alert, and then click the number in the specific severity column.

  > **Note** If the alert information is aggregated by **Product Name** on this tab, click + to the left of the product name to show the monitoring metrics.

  - If an alert is being processed by operations engineers, choose **Actions > Process** in the **Actions** column to set the alert status to **In process**.

    If multiple alerts are being processed by operations engineers, select these alerts and then click **Process** to process the alerts in batches.
  - If the alert has been processed, choose **Actions > Processed** in the **Actions** column to set the alert status to **Processed**.

    If multiple alerts have been processed, select these alerts and then click **Complete** to complete the alerts in batches.
  - To view the whole processing flow of an alert, choose **Actions > Alert Tracing** in the **Actions** column.
  - If an alert is considered as an incident when the alert is being processed, choose **Actions > Report to ITIL** in the **Actions** column. Then, an incident request is created in the Information Technology Infrastructure Library (ITIL) to track the issue. For more information, see Manage incidents.

    If multiple alerts are considered as incidents, select these alerts and then click **Report to ITIL** in the upper part of the page. Then, the system creates multiple incident requests in the ITIL to track the issues.

- View recent monitoring data

  Choose **Actions > Exploration** in the **Actions** column corresponding to an alert to view the trend chart of a monitoring metric of a product.

- Export reports

  Click the  icon in the upper part of the tab to download the alert list.

4.1.2.3. Alert history

The Alert History page shows all alerts generated by the system and their information in chronological order.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert History**.
2. On the **Alert History** page, perform the following operations:
   - Search for an alert

     In the upper part of the page, you can search for an alert by specifying **Monitoring Metric Type**, **Product**, **Service**, **Severity**, **Status**, **Start date**, **End date**, or search content.

   - Export the alert list

     Click the  icon in the upper part of the page to export a list of historical alerts.

   - View alert sources

     Move the pointer over an alert source name in blue in the **Alert Source** column to view the alert source details.

   - View the details of a metric

     Click an alert name in blue in the **Alert Details** column. On the **Alert Details** page, you can view the alert information such as the alert description, reference, impact scope, and resolution.

   - View the original alert information

     Click **Details** in the **Alert Information** column to view the original information of the alert.

   - View the alert duration

     The alert duration is the total duration of an alert from the start time to the time when the alert is terminated. You can view the duration of an alert in the **Duration** column. You can also move the pointer over a value in the **Duration** column to view the specific start time of the alert.

4.1.2.4. Alert configuration

The **Alert Configuration** module provides you with three functions: contacts, contact groups, and static parameter settings.

4.1.2.4.1. Alert contacts

You can query, add, modify, or delete an alert contact based on business needs.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Configuration**. The **Contacts** tab appears.
2. You can perform the following operations:
   - Search for alert contacts

     In the upper-lefter corner of the tab, specify the product name, contact name, and phone number and then click **Search**. The alert contacts that meet the search conditions are displayed in the list.
   - Add an alert contact

     In the upper-left corner of the tab, click **Add**. The **Add Contact** pane appears. Configure the parameters, and then click **OK**.
   - Modify an alert contact

     Find the alert contact to be modified and then click **Modify** in the **Actions** column. In the **Modify Contact** pane, modify the relevant information and then click **OK**.
   - Delete an alert contact

     Find the alert contact to be deleted and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

4.1.2.4.2. Alert contact groups

You can query, add, modify, or delete an alert contact group based on business needs.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Configuration**.
2. Click the **Contact Group** tab.
3. Perform the following operations:
   - Query an alert contact group

     Enter a group name in the search box and click **Search**. The information of the alert contact group that meets the search condition is displayed.
   - Add an alert contact group

     Click **Add** in the upper-left corner of the tab. In the **Add Contact Group** pane, enter a group name and select the contacts to be added to the contact group. Click **OK**.
   - Modify an alert contact group

     Find the contact group to be modified, and then click **Modify** in the **Actions** column. In the **Modify Contact Group** pane, modify the group name, description, contacts, and notification method. Click **OK**.
   - Delete one or more alert contact groups

     Find the contact group to be deleted, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

     Select one or more contact groups to be deleted and click **Delete All** in the upper part of the tab. In the message that appears, click **OK**.

4.1.2.4.3. Configure static parameters

You can configure alert-related static parameters based on your business needs. Only parameters related to timeout alerts can be configured.

Context

You cannot add new alert configurations in the current version. You can modify the default parameter configurations for timeout alerts.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Configuration**.
2. Click the **Static Parameter Settings** tab.
3. (Optional)Enter a parameter name in the search box and click **Search** to query the static parameter configurations.
4. Find the static parameter to be modified, and then click **Modify** in the **Actions** column.
5. In the **Modify Static Parameter** pane, modify the parameter name, parameter value, and description.

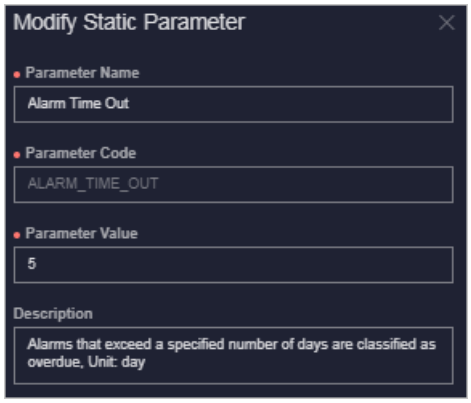

Parameter	Description
Parameter Name	Enter a parameter name related to the configuration.
Parameter Value	Enter the parameter value. The default value is 5, indicating five days. After you complete the configuration, you can choose **Alert Monitoring > Alert Events** and then click the **Timeout Alert** tab to view alert events that meet the condition specified by this parameter value. For example, if the parameter value is 5, you can choose **Alert Monitoring > Alert Events** and then click the **Timeout Alert** tab, alert events that are retained more than five days are displayed.
Description	Enter the description related to the configuration.

6. Click **OK**.

4.1.2.5. Alert overview

The Alert Overview module allows you to query the distribution of different levels of alerts for Apsara Stack services.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Overview** to go to the **Alert Overview** page.

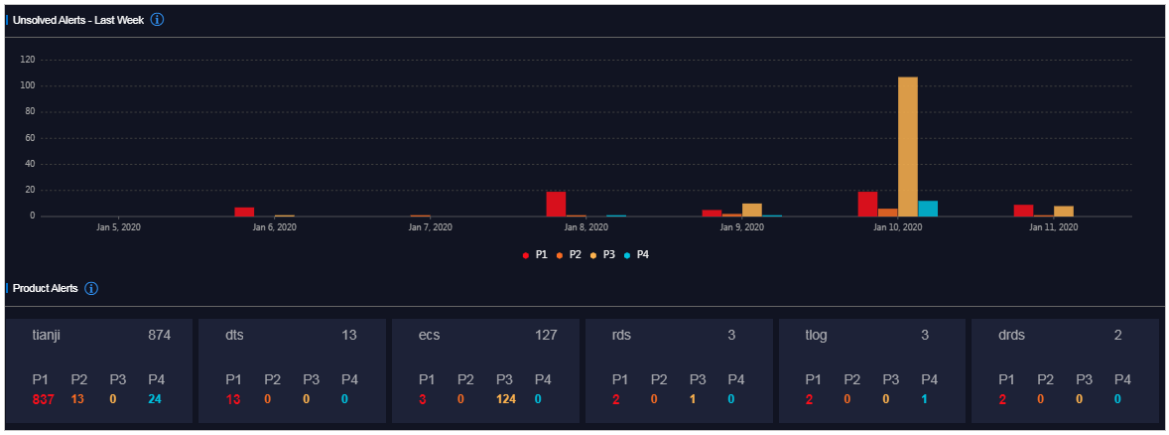

   - The column chart in the upper part of the page displays the number of unresolved alerts from the last seven days.
   - The section in the lower part of the page displays the alert statistics in the current system by service.

4.1.2.6. Alert subscription and push

The alert subscription and push feature allows you to configure alert notification channels and then push alerts to operations engineers.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Subscribe/Push**.

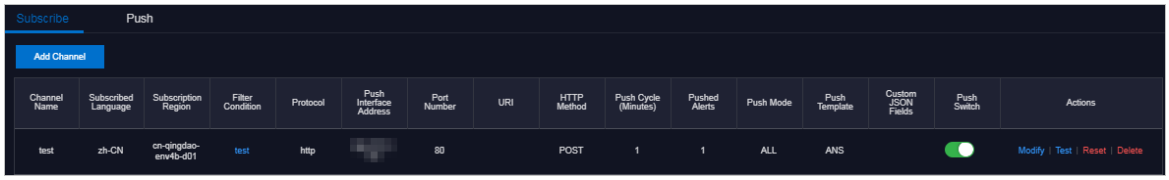

2. On the **Subscribe** tab, click **Add Channel**.
3. In the **Add Subscription** pane, configure the following parameters.

Parameter	Description
Channel Name	The name of the subscription channel.
Subscribed Language	The subscription language. Valid values: Chinese and English.

Parameter	Description
Subscription Region	The region where the subscription is located.
Filter Condition	The filter conditions used to filter alerts. Valid values: ○ **Basic** ○ **Critical** ○ **Important** ○ **Minor** ○ Custom filter
Protocol	The protocol used to push alerts. Only HTTP is supported.
Push Interface Address	The IP address of the push interface.
Port Number	The port number of the push interface.
URI	The URI of the push interface.
HTTP Method	The request method used to push alerts. Only the POST method is supported.
Push Cycle (Minutes)	The interval for pushing alerts. Unit: minutes.
Pushed Alerts	The number of alerts pushed each time.
Push Mode	The mode used to push alerts. Valid values: ○ **ALL**: All alerts are pushed each push cycle. ○ **TOP**: Only high priority alerts are pushed each push cycle.
Push Template	The template used to push alerts. Valid values: ○ ASO: the default template. ○ ANS: select this template to push alerts by DingTalk, short messages, or emails. You can only configure a single channel of this type. **Note** A preset ANS template exists if the system already connects with ANS. To restore the initial configurations of the template with one click, click **Reset** in the upper part of the page.

Parameter	Description
Custom JSON Fields	The person who receives the push can use this field to customize an identifier. The field must be in the JSON format.
Push Switch	Specifies whether to push alerts. If the switch in this pane is not turned on, after you configure the subscription channel, you can enable the push feature in the **Push Switch** column.

4. Click **OK**. To modify or delete a channel, click **Modify** or **Delete** in the **Actions** column corresponding to the channel.
5. (Optional)The newly added channel is displayed in the list. Click **Test** in the **Actions** column corresponding to the channel to test the connectivity of the push channel.

   > **Note** For the ANS push channel, after you click **Test** in the Actions column, you must enter the mobile phone number, email address, or DingTalk to which alerts are pushed.

6. After you configure the push channel and turn on the push switch, you can click the **Push** tab to view the push records.

4.1.2.7. Alert masking

The Alert Masking module allows you to mask a type of alerts and remove the masking as needed.

4.1.2.7.1. Add masking rules

Masking rules allow you to mask alerts that you no long need to pay attention to.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Masking**.
2. In the upper part of the page, click **Add**.
3. In the **Add** pane, configure parameters related to the alerts to be masked.

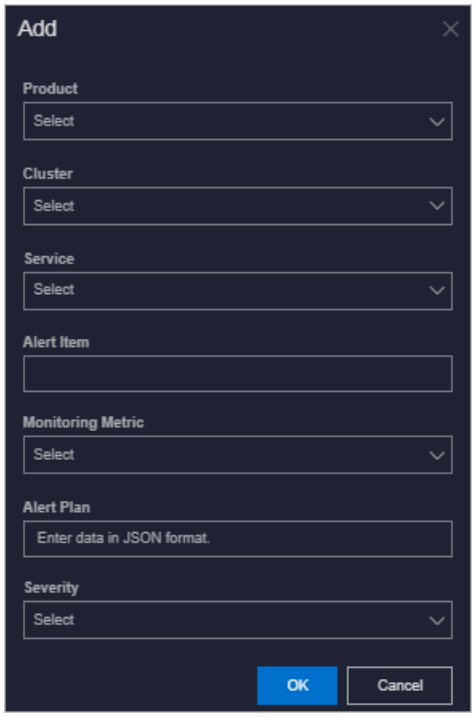

Parameter	Description
Product	Optional. The product to which the alerts to be masked belong.
Cluster	Optional. The cluster to which the alerts to be masked belong.
Service	Optional. The service to which the alerts to be masked belong.
Alert Item	Optional. The alert name to be masked. **Note** When you configure **Alert Item**, if the number of alerts is large, you may need to wait a few minutes.
Monitoring Metric	Optional. The monitoring metric to which the alerts to be masked belong.
Alert Plan	Optional. The alert details of the alerts to be masked. Example: {"serverrole":"ecs-yaochi.ServiceTest#","machine":"vm0100120****","level":"error"}

Parameter	Description
Severity	Optional. The severity levels of the alert. Valid values: ○ **P0**: indicates that the alert has been cleared, corresponding to alerts whose **Alert Level** is **Restored** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P1**: indicates critical alerts, corresponding to alerts whose **Alert Level** is **P1** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P2**: indicates major alerts, corresponding to alerts whose **Alert Level** is **P2** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P3**: indicates minor alerts, corresponding to alerts whose **Alert Level** is **P3** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P4**: indicates alerts for notice, corresponding to alerts whose **Alert Level** is **P4** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P5**: indicates system alerts.

4. Click **OK**.

Result

The added masking rule is displayed in the alert masking list.

After a masking rule is added, alerts that meet the conditions in the masking rule are not displayed in the **Alert Events** and **Alert History** tabs.

4.1.2.7.2. Remove the masking

You can remove the masking for masked alerts.

Procedure

1. In the left-side navigation pane, choose **Alert Monitoring > Alert Masking**.
2. (Optional)Specify a product, service, or an alert item. Click **Search**.
3. Find the alert masking rule to be removed, and then click **Delete** in the **Actions** column.

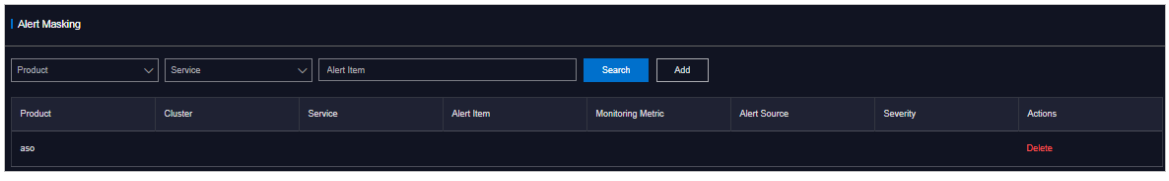

4. In the message that appears, click **OK**.

Result

After you remove the masking, alerts that were masked by the deleted masking rule are displayed in the **Alert Events** and **Alert History** tabs.

4.1.3. Physical servers

Operations personnel can monitor and view the physical servers where each product is located.

4.1.3.1. View the physical server information

This topic describes how to view the physical server list and the details of physical servers.

Product tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.

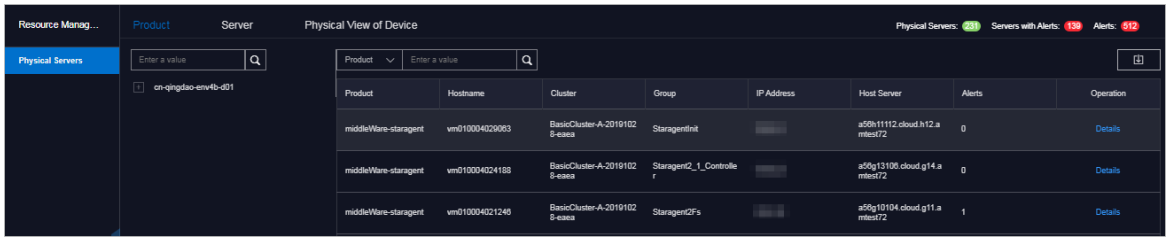

2. On the Product tab, perform the following operations to view the physical server information:
   - Expand the left-side navigation tree by selecting a region, product, and cluster in sequence to view the list of physical servers where a cluster of a service is located.
   - In the left-side search box, enter the product name, cluster name, group name, or hostname to search for the corresponding node.
   - In the right-side search box, search for physical servers by product, cluster, group, or hostname and view the details of a physical server.
   - Select a product and click **Details** in the **Actions** column. On the **Physical Server Details** page, you can view the basic information, monitoring details, and alert information of the physical server to which the product belongs.

     You can switch the tab to view the monitoring and alert information.

     Monitoring information includes the CPU usage, system load, disk usage, memory usage, network throughput, and disk IO. When you view the monitoring information, you can select a monitoring item in the upper-right corner of each monitoring graph and then select the time range to view the monitoring value in the specific time range.

     In the upper-right corner of the CPU usage, system load, disk usage, memory usage, network throughput, and disk IO sections, you can perform the following operations:

     - Click the icon to view the monitoring graph in full screen.
     - Click the icon to download the monitoring graph to your local computer.
     - Click the icon to manually refresh the monitoring data.
     - Click the icon. The icon will turn green. The system automatically refreshes the monitoring data every 10 seconds. To disable the auto refresh feature, click the icon again.

Server tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.

2. Click the **Server** tab.
3. On the Server tab, perform the following operations to view the physical server list:
   - Expand the left-side navigation tree by selecting an IDC and a rack in sequence to view the physical server list in a rack.
   - Enter the rack name in the left-side search box and press the Enter key to search for and view the list of all the physical servers in the rack.

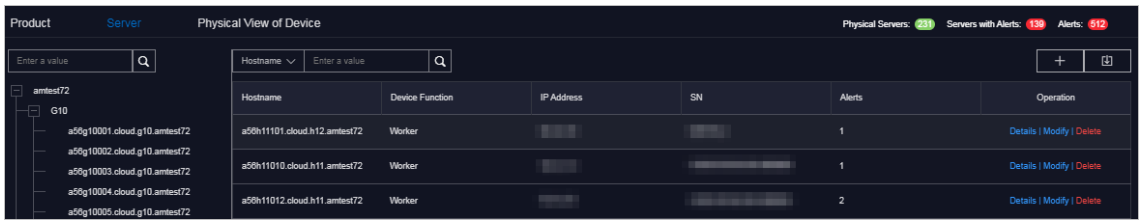

4. To view the details of a physical server, enter the hostname, IP address, device function, or serial number (SN) in the right-side search box and press the Enter key.
5. Find the physical server whose details you are about to view and then click **Details** in the **Actions** column. On the **Physical Machine Details** page, view the basic information, monitoring information, and alert information of the physical server.

   You can switch the tab to view the monitoring and alert information.

   Monitoring information includes the CPU usage, system load, disk usage, memory usage, network throughput, and disk IO. When you view the monitoring information, you can select a monitoring item in the upper-right corner of each monitoring graph and then select the time range to view the monitoring value in the specific time range.

   In the upper-right corner of the CPU usage, system load, disk usage, memory usage, network throughput, and disk IO sections, you can perform the following operations:

   - Click the icon to view the monitoring graph in full screen.
   - Click the icon to download the monitoring graph to your local computer.
   - Click the icon to manually refresh the monitoring data.
   - Click the icon. The icon will turn green. The system automatically refreshes the monitoring data every 10 seconds. To disable the auto refresh feature, click the icon again.

The Physical View of Device tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.

2. Click the **Physical View of Device** tab.

3. On the **Physical View of Device** tab, expand the left-side navigation tree by selecting an IDC and a rack in sequence to view the corresponding rack information on the right. In addition, the rack details pane appears on the right side of the tab and shows the server information of the rack.

   Racks and servers are displayed in different colors to indicate the alert condition of servers:

   - Red indicates a critical alert.
   - Orange indicates a moderate alert.
   - Blue indicates that the physical server is normal.

   In the upper-right corner, you can view the alert legend. By default, the check box at the left of the legend is selected, indicating that the information of racks or servers of this alert type is displayed on the rack graph or in the rack details pane. Clear the check box at the left of a legend to hide the information of racks or servers of this alert type on the rack graph or in the rack details pane.

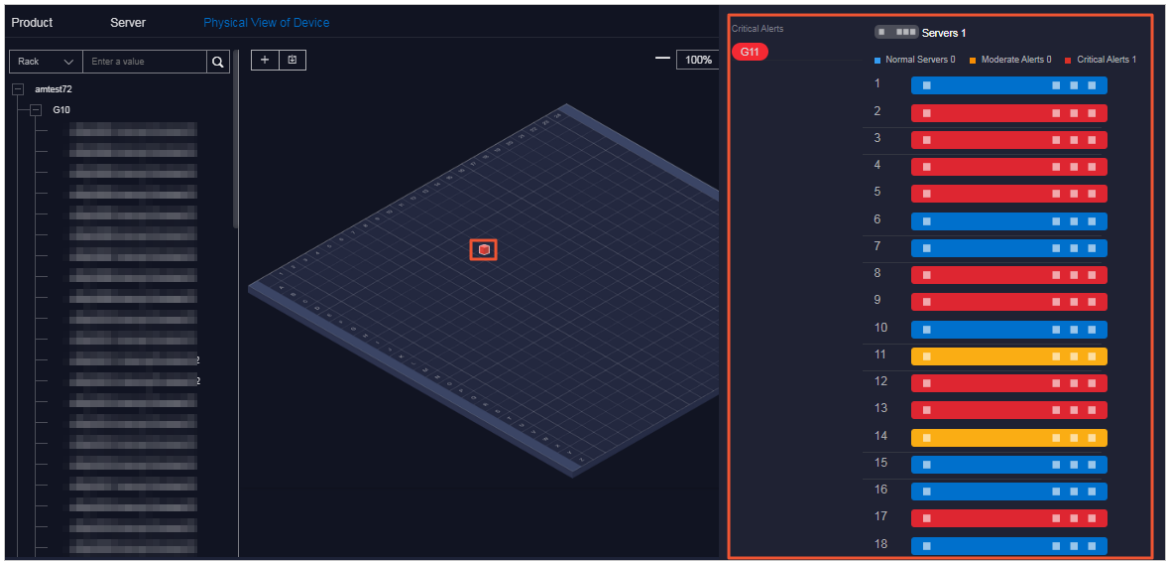

4. To view the details of a physical server, perform the following operations:

   i. Find the physical server whose details you are about to view in the left-side navigation tree or rack graph on the right side of the tab.

   ii. In the rack details pane that appears, click the color block of a server to view the basic information of the server.

   iii. Click **Details** in the **Operation** row of the basic information.

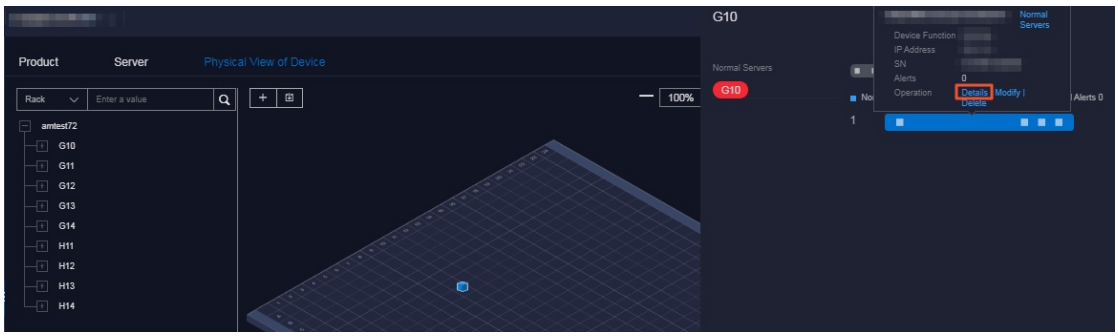

iv. On the **Physical Server Details** page, view the basic information, monitoring details, and alert information of the physical server.

You can switch the tab to view the monitoring information and alert information.

Monitoring information includes the CPU usage, system load, disk usage, memory usage, network throughput, and disk IO. When you view the monitoring information, you can select a monitoring item in the upper-right corner of each monitoring graph and then select the time range to view the monitoring value in the specific time range.

In the upper-right corner of the CPU usage, system load, disk usage, memory usage, network throughput, and disk IO sections, you can perform the following operations:

- Click the  icon to view the monitoring graph in full screen.
- Click the  icon to download the monitoring graph to your local computer.
- Click the  icon to manually refresh the monitoring data.
- Click the  icon. The icon will turn green. The system automatically refreshes the monitoring data every 10 seconds. To disable the auto refresh feature, click the icon again.

4.1.3.2. Add physical servers

Operations personnel can add the information of existing physical servers in the environment to the ASO console.

Procedure

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the page, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.
2. Click the **Server** or **Physical View of Device** tab.
3. In the upper-right corner of the **Server** tab or the upper-left corner of the **Physical View of Device** tab, click the  icon.
4. In the **Add Physical Server** pane, configure the parameters.

   The following table describes the parameters.

Parameter	Description
Zone	The zone where the target physical server is located.
Data Center	The data center where the target physical server is located.
Rack	The rack where the target physical server is located.
Room	The room where the target physical server is located.
Physical Server Name	The name of the target physical server.

Parameter	Description
Memory	The memory size of the target physical server.
Disk Size	The disk size of the target physical server.
CPU Cores	The CPU cores of the target physical server.
Rack Group	The rack group to which the target physical server belongs.
Server Type	The type of the target physical server.
Server Role	The function or purpose of the target physical server.
Serial Number	The serial number (SN) of the target physical server.
Operating System Template	The template used by the operating system of the target physical server.
IP Address	The IP address of the target physical server.

5. Click **OK**.

4.1.3.3. Modify a physical server

This topic describes how to modify the physical server information in the system when the information is changed in the Apsara Stack environment.

Server tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.

2. Click the **Server** tab.
3. (Optional)In the right-side search box, search for the physical server to be modified by hostname, IP address, device function, or serial number (SN).
4. Find the target physical server, and then click **Modify** in the **Actions** column.

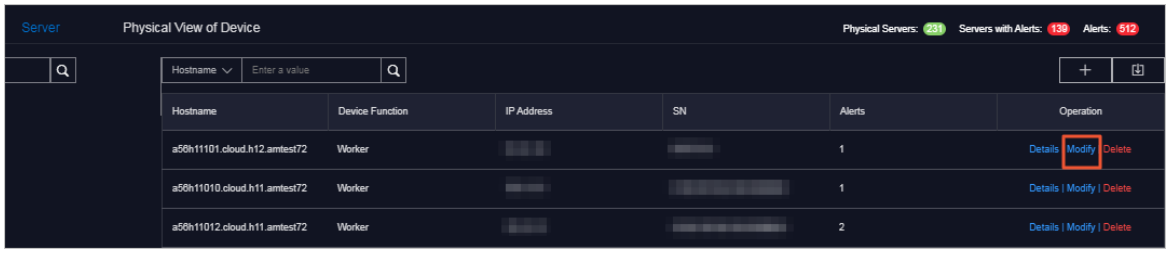

5. In the **Modify Physical Server** pane, modify the physical server information. You can modify the following physical server information: zone, data center, rack, room, physical server name, memory size, disk size, CPU cores, rack group, server type, server role, serial number, operating system template, and IP address.
6. Click **OK**.

Physical View of Device tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.

2. Click the **Physical View of Device** tab.
3. Expand the left-side navigation tree by selecting an IDC and a rack in sequence to find the physical server to be modified.

   > **Note** In the left-side search box, you can also search for the target physical server by rack, hostname, IP address, device function, SN, or IDC.

4. In the rack details pane that appears, click the color block of a server to view the basic information of the server.
5. Click **Modify** in the **Operation** row of the basic information.

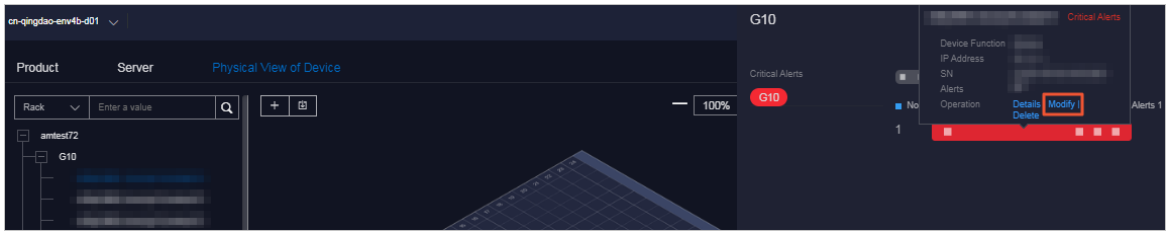

6. In the **Modify Physical Server** pane, modify the physical server information. You can modify the following physical server information: zone, data center, rack, room, physical server name, memory size, disk size, CPU cores, rack group, server type, server role, serial number, operating system template, and IP address.
7. Click **OK**.

4.1.3.4. Export server information

You can export the information of all physical servers within the system for offline viewing.

Product tab

The physical server information exported from the **Product** tab includes the zone, hostname, disk size, CPU cores, memory size, information about the data center (data center, rack, room, and rack group), model, device function, serial number, operating system template, IP address, out-of-band IP address, CPU architecture, host server, alerts, region, product, cluster, and service role group.

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the page, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.

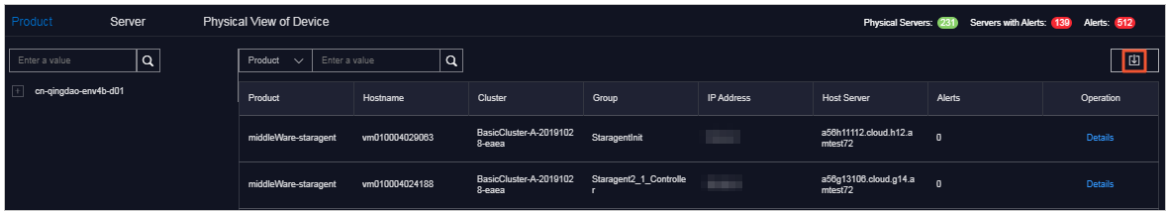

2. In the upper-right corner of the tab, click the icon to export the information of all the physical servers of all services to your local computer.

Server or Physical View of Device tab

The physical server information exported from the **Server** or the **Physical View of Device** tab includes the zone, hostname, disk size, CPU cores, memory size, information about the data center (data center, rack, room, and rack group), model, device function, serial number, operating system template, IP address, out-of-band IP address, CPU architecture, and alerts.

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the page, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.
2. Click the **Server** or the **Physical View of Device** tab.
3. In the upper-right corner of the **Server** tab or in the upper part of the **Physical View of Device** tab, click the icon to export all the information of physical servers to your local computer.

4.1.3.5. Delete a physical server

This topic describes how to delete a physical server that does not need to be monitored.

Server tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.
2. Click the **Server** tab.
3. (Optional)In the right-side search box, search for the physical server to be deleted by hostname, IP address, device function, or serial number (SN).
4. Find the target physical server, and then click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

Physical View of Device tab

1. In the left-side navigation pane, choose **Resource Management > Physical Servers**.

   The **Product** tab appears. In the upper-right corner of the tab, the number of current physical servers, the number of servers with alerts, and the number of alerts are displayed.
2. Click the **Physical View of Device** tab.
3. Expand the left-side navigation tree by selecting an IDC and a rack in sequence to find the physical server to be deleted.

   > **Note** In the left-side search box, you can also search for the physical server to be deleted by rack, hostname, IP address, device function, SN, or IDC.
4. In the rack details pane that appears, click the color block of a server to view the basic information of the server.

5. Click **Delete** in the **Operation** row of the basic information.

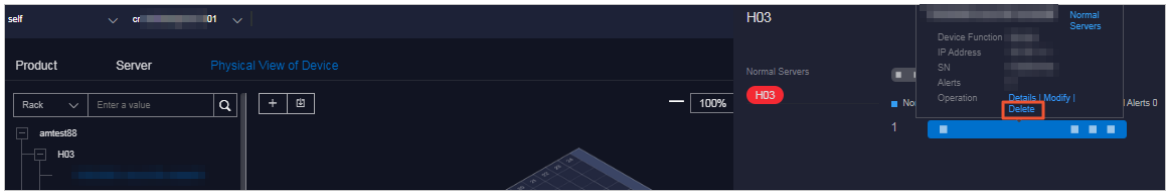

6. In the message that appears, click **OK**.

4.1.4. Inventory Management

The Inventory Management module allows you to view the current usage and inventory of various product resources, and manage resources in the system effectively.

4.1.4.1. View the ECS inventory

By viewing the Elastic Computing Service (ECS) inventory, you can query the usage and availability of ECS resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > ECS**.

> **Note** You can click the icon in the upper-right corner to configure the inventory thresholds.

2. Select a date in the upper part of the page and view the ECS inventory.

   The following information is displayed:

   - The **CPU Inventory Details (Core)** and **Memory Inventory Details (TB)** sections display the usage and availability of CPU (core) and memory (TB) of all ECS instance families for the last five days.
   - The **ECS Instances Inventory Details** section displays the inventory details of specified ECS instance tyeps at specified dates on multiple pages by **Region ID**, **Instance Type**, and **Date**.

3. (Optional)After you search for the data by specifying **Region ID**, **Instance Type**, and **Date** in the **ECS Instances Inventory Details** section, click **Export** to export the ECS inventory details to your local computer.

4.1.4.2. View the SLB inventory

By viewing Server Load Balancer (SLB) inventory, you can query the usage and availability of SLB resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > SLB**.

> **Note** You can click the icon in the upper-right corner to configure the inventory thresholds.

2. View the SLB inventory.

   The following information is displayed:

   - The Internal VIP Used Inventory and Public VIP Used Inventory sections display the amount and percentage of internal and public VIP inventory that are being used.
   - The Network Card Traffic section displays the inbound and outbound network card traffic.
   - The **SLB Inventory Details** section displays the SLB inventory details on multiple pages by **Type** and **Date**.

4.1.4.3. View the RDS inventory

By viewing the Relational Database Service (RDS) inventory, you can query the usage and availability of RDS resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > RDS**.

   > **Note** You can click the icon in the upper-right corner to configure the inventory thresholds of each engine.

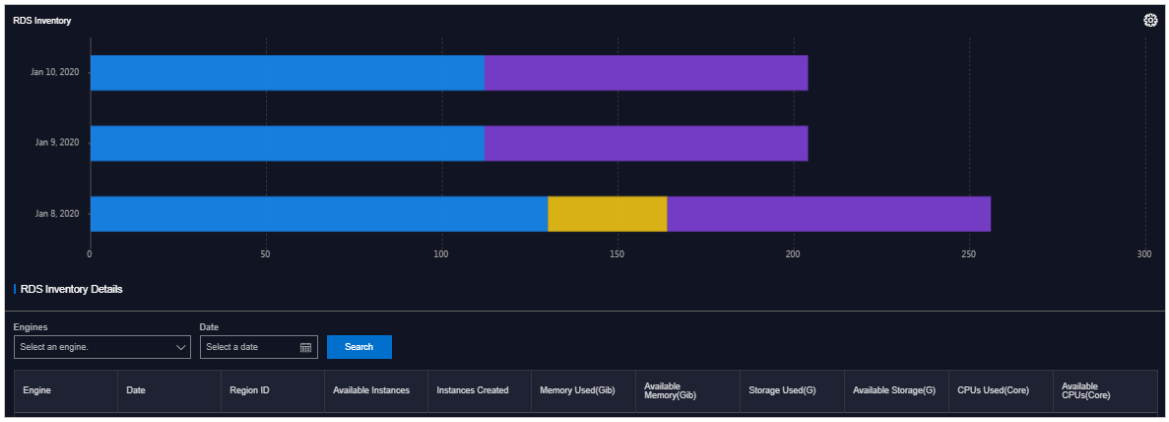

2. View the RDS inventory.

   The following information is displayed:

- The **RDS Inventory** section displays the inventories of different types of RDS instances for the last five days. Different colors represent different types of RDS instances.
- The **RDS Inventory Details** section shows the RDS inventory details on multiple pages by **Engine** and **Date**.

4.1.4.4. View the OSS inventory

By viewing the Object Storage Service (OSS) inventory, you can query the usage and availability of OSS resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > OSS**.

   > **Note** You can click the icon in the upper-right corner to configure the inventory thresholds.

   

2. View the OSS inventory.

   The following information is displayed:

   - The **Inventory Availability History (TB)** section shows the available OSS inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount and percentage of OSS inventory that are being used.
   - The **OSS Bucket Inventory Details** section shows the OSS inventory details on multiple pages by **Date**.

4.1.4.5. View the Tablestore inventory

By viewing the Tablestore inventory, you can query the usage and availability of Tablestore resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > OTS**.

> **Note** You can click the icon in the upper-right corner to configure the inventory thresholds.

2. View the Tablestore inventory.

   The following information is displayed:

   - The **Inventory Availability History (TB)** section shows the available Tablestore inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount and percentage of Tablestore inventory that are being used.
   - The **OTS Bucket Inventory Details** section shows the Tablestore inventory details on multiple pages by **Date**.

4.1.4.6. View the Log Service inventory

By viewing the Log Service inventory, you can query the usage and availability of Log Service resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > SLS**.

   > **Note** You can click the icon in the upper-right corner to configure the inventory thresholds and global quota.

2. Click the **sls-inner** tab to view the inventory of base Log Service instances.

   The following information is displayed:

   - The **Inventory Availability History (TB)** section shows the available and total Log Service inventory for the last five days.
   - The **Current Quota Details (G)** section shows the amount and percentage of Tablestore inventory that are being used.
   - The **Log Service Inventory Details** section shows the Log Service inventory details on multiple pages by **Date**.

3. Click the **sls-public** tab to view the inventory of Log Service instances that you have applied for.
   - The **Inventory Availability History (TB)** section shows the available Log Service inventory for

- The **Inventory Availability History (TB)** section shows the available Log Service inventory for the last five days.
- The **Current Inventory Usage (TB)** section shows the amount and percentage of Log Service inventory that are being used.
- The **SLS Bucket Inventory Details** section shows the Log Service inventory details in multiple pages by **Date**.

4.1.4.7. View the EBS inventory

By viewing the Elastic Block Storage (EBS) inventory, you can query the usage and availability of EBS resources to more efficiently perform O&M operations.

Context

> **Note** EBS is the Apsara Distributed File System storage allocated by the base and to ECS. The ECS IO cluster is the cluster for Apsara Distributed File System storage. Here you can view the EBS inventory in the ECS IO cluster.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > EBS**.

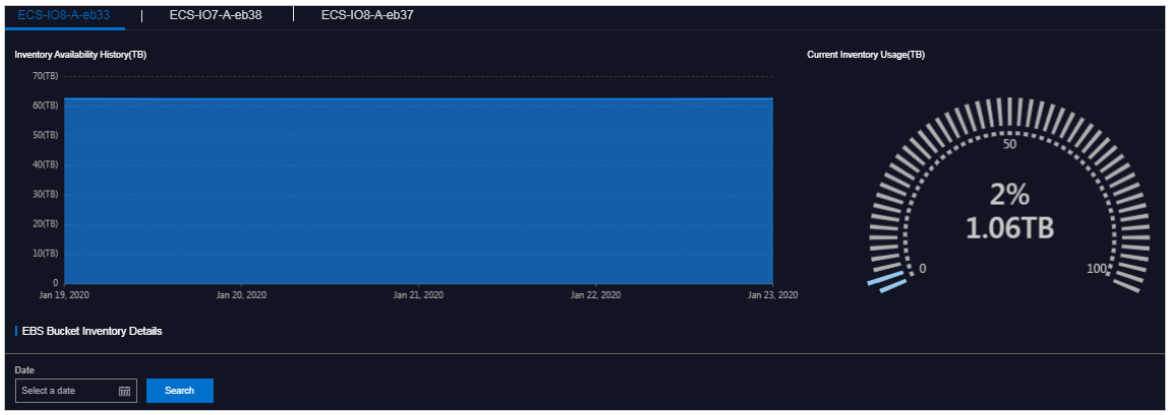

2. If multiple ECS IO clusters exist in the environment, click the tab of each ECS IO cluster to view the EBS inventory.

   The following information is displayed:

   - The **Inventory Availability History (TB)** section shows the available EBS inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount and percentage of EBS inventory that are being used.
   - The **EBS Bucket Inventory Details** section shows the EBS inventory details on multiple pages by Date.

4.1.4.8. View the Apsara File Storage NAS inventory

By viewing the Apsara File Storage NAS inventory, you can query the usage and availability of Apsara File Storage NAS resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > NAS**.

2. View the Apsara File Storage NAS inventory.

   The following information is displayed:

   - The **Inventory Availability History (TB)** section shows the available Apsara File Storage NAS inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount and percentage of Apsara File Storage NAS inventory that are being used.
   - The **NAS Bucket Inventory Details** section shows the Apsara File Storage NAS inventory details on multiple pages by Date.

4.1.4.9. View the HDFS inventory

By viewing the Hadoop Distributed File System (HDFS) inventory, you can query the usage and availability of HDFS resources to more efficiently perform O&M operations.

Procedure

1. In the left-side navigation pane, choose **Inventory Management > DFS**.

2. View the HDFS inventory.

   The following information is displayed:

   - The **Inventory Availability History (TB)** section shows the available HDFS inventory for the last five days.

- The **Current Inventory Usage (TB)** section shows the amount and percentage of HDFS inventory that are being used.
- The **DFS Bucket Inventory Details** section shows the HDFS inventory details on multiple pages by Date.

4.1.5. Full stack monitoring

The Full Stack Monitoring module allows you to perform aggregate queries for system alert events. You can query all end-to-end alert data by host IP address, instance ID, and time range, as well as view the end-to-end topology.

4.1.5.1. SLA

The **Standard Storage SLA** module allows you to view the current state, history data, instance availability, and product availability of each cloud product. You can view the current state and history data of products to obtain the SLA values and unavailable events of product instances within a time period.

4.1.5.1.1. View the current state of a cloud service

The Current State tab allows you to view the current state of a cloud service and the details of exception events.

Procedure

1. In the left-side navigation pane, choose **Full Stack Monitoring > Standard Storage SLA**.
2. The **Current Status** tab appears.

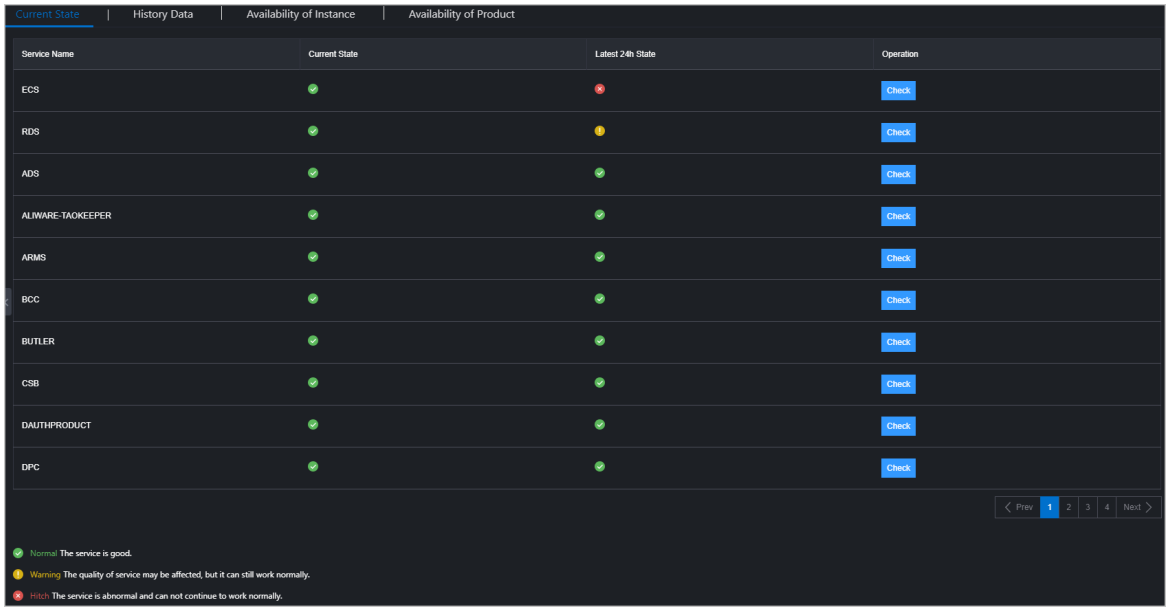

The current state and the status within the last 24 hours of each cloud service are displayed on this page. Services in different states are displayed in different colors:

- Green: normal. The service is running properly.
- Yellow: warning. The service has some latency, but can still work properly.

- Red: faulty. The service is temporarily interrupted and cannot work properly.

3. Find the service whose running status that you want to view, and then click **Check** in the Operation column.
   - The **Overall Availability** section shows the availability of a service. You can view the availability by hour, day, or minute.
   - The **Related Events** section shows the current exception event. Click **Show Details** corresponding to an event to view the event details.

4.1.5.1.2. View the history data of a cloud service

The History Data tab allows you to view the history status of a cloud service and the details of exception events.

Procedure

1. In the left-side navigation pane, choose **Full Stack Monitoring > Standard Storage SLA**.
2. Click the **History Data** tab.

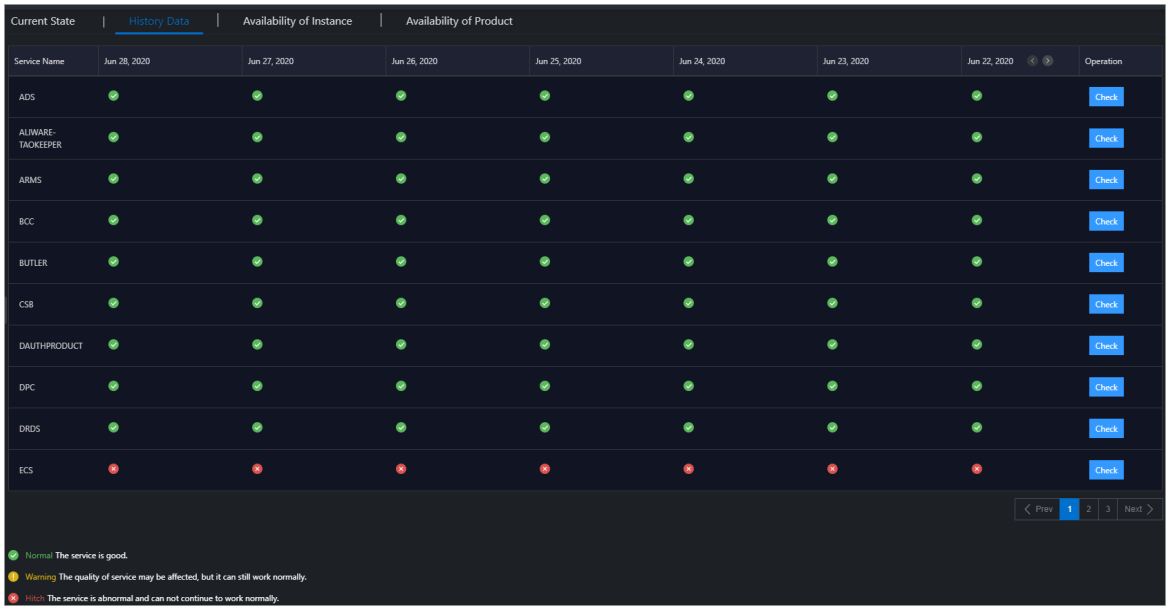

   The service availability of each service in the last two weeks is displayed on this page. Services in different states are displayed in different colors:
   - Green: normal. The service is running properly.
   - Yellow: warning. The service is experiencing some latency but otherwise working properly.
   - Red: faulty. The service has been temporarily interrupted and cannot function properly.
3. Find the service whose history status that you want to view. Click **Check** in the Operation column.
   - The **Overall Availability** section shows the historical availability of a service. You can view the availability by hour, day, or minute.
   - The **Related Events** section shows the historical exception events. Click **Show Details** corresponding to an event to view the event details.

4.1.5.1.3. View the availability of an instance

You can view the instance availability ratio of a cloud service to know the instance damages.

Procedure

1. In the left-side navigation pane, choose **Full Stack Monitoring > Standard Storage SLA**.
2. Click the **Availability of Instance** tab.

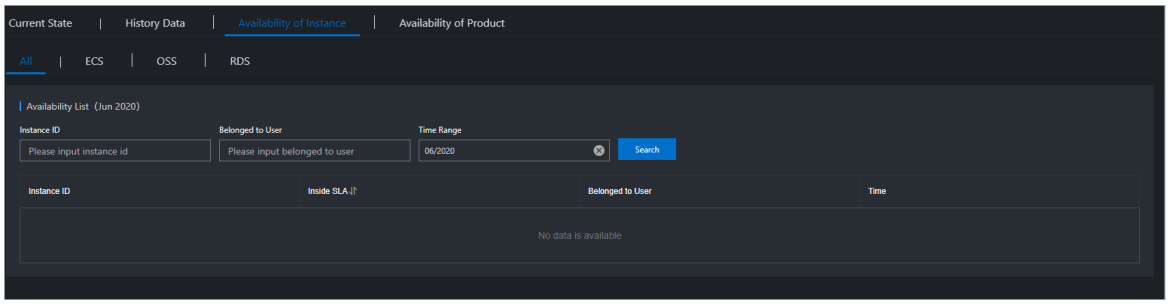

3. Specify **Instance ID**, **Belonged to User**, or **Time Range**. Then, click **Search**.
4. Click the **Instance ID** to view the following information of the instance:
   - **Basic Information**: the instance ID and the user to whom the instance belongs.
   - **Availability**: the availability ratio of the instance.
   - **Damage Event**: the exception event list.

4.1.5.1.4. View the availability of a service

You can view the availability ratio of a cloud service to determine its monthly availability index.

Procedure

1. In the left-side navigation pane, choose **Full Stack Monitoring > Standard Storage SLA**.
2. Click the **Availability of Product** tab.
3. Specify **Product** and **Time Range**, and then click **Search** to view the availability ratio of the product.For example, if the availability ratio of Elastic Compute Service (ECS) is 100.00%, it indicates that ECS runs properly this month, without any faults.

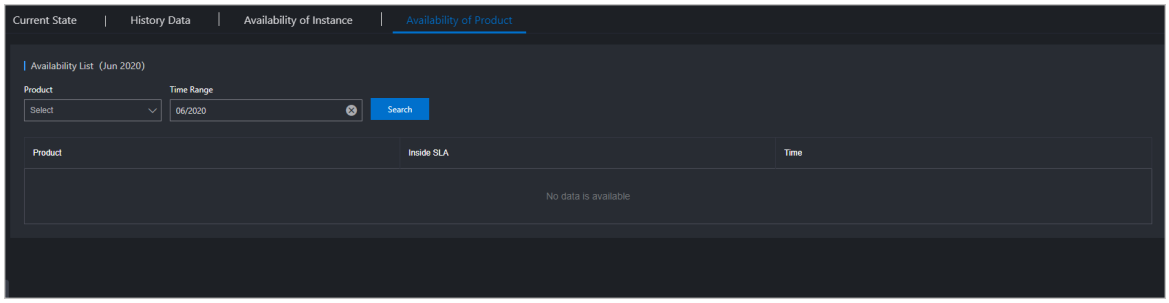

4.1.5.2. Full stack log monitoring

The Full Stack Log Monitoring module allows you to search for logs of ECS-, SLB-, and All in ECS-related applications.

Context

- You can search for the logs of a variety of product components on the ECS tab, such as pop,

openapi, pync, and opsapi.

- If each SLB service node enables the ilogtail reporting feature, you can search for logs of pop, slb-yaochi, and slb-control-master on the SLB tab.
- You can search for vm_adapter logs, all in ECS-Apsara Infrastructure Management Framework adaption layer logs, and all the other ECS operations logs on the All in ECS tab.

Procedure

1. In the left-side navigation pane, choose **Full Stock Monitoring > Full Stock Log Monitoring**.
2. Click the **ECS**, **SLB**, or **All in ECS** tab.
3. Enter a keyword in the **Query** field. Select the time range in the **Time** field. Click **Search**.

   > **Note** You can enter any string in the Query field as the search condition, such as the instance ID, request ID, or the keyword "error".

4. The search results are displayed. Click an application log.

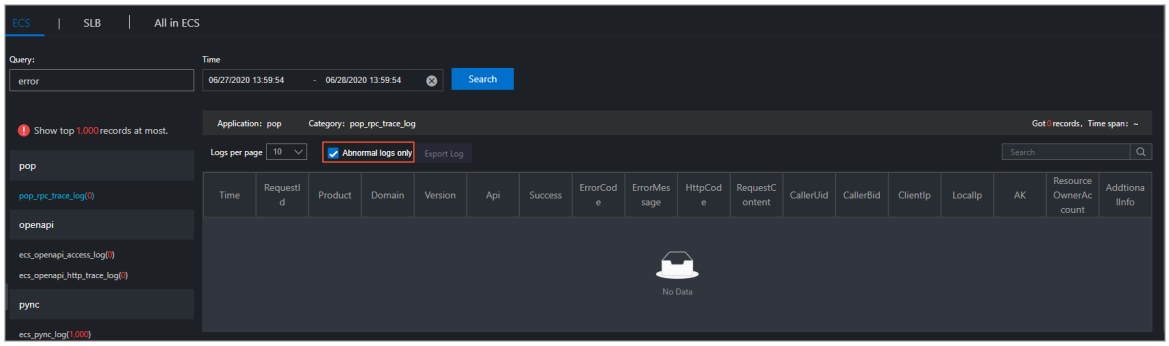

5. Select **Abnormal logs only** to view only the abnormal logs.

   If `code != 200`, `success=false`, or `error` exist in a log, the log is an abnormal log.

6. Enter a keyword in the search box to search for the related information in the search results.
7. (Optional)After the search is complete, click **Export Log** to export the search results to your local computer.

4.1.6. Storage Operation Center

The Storage Operation Center module consists of Apsara Distribute File System and EBS.

4.1.6.1. Apsara Distributed File System

The Apsara Distributed File System module shows the overview, cluster information, node information, and cluster status.

4.1.6.1.1. Overview

The Apsara Distributed File System module allows you to view the overview information, health heatmap, and data of the top five clusters.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > Apsara Distributed File System > Overview**.
2. Select the service that you want to view from the **Service** drop-down list.

   The Apsara Distributed File System module shows the overview information, health heatmap, and data of the top five clusters as of the current date.

   - **Overview**

     The Overview section shows the storage space, server information, and health information of the specified service. In the **Health** section, when the value of **Abnormal Disks**, **Abnormal Masters**, **Abnormal Chunk Servers**, or **Abnormal Water Levels** is greater than 0, the value is displayed in red.

     

   - **Heatmap of Health**

     The Heatmap of Health section shows the health information of all clusters within the specified service. Clusters in different health states are displayed in different colors:

     - Green indicates that the cluster works properly.
     - Yellow indicates that the cluster has a warning.
     - Red indicates that the cluster has an exception.
     - Dark red indicates that the cluster has a fatal error.
     - Grey indicates that the cluster is disabled.

     Click the name of an enabled cluster to go to the corresponding cluster information page.

     Move the pointer over the color block of each cluster to view the corresponding service name, server name, and IP address.

     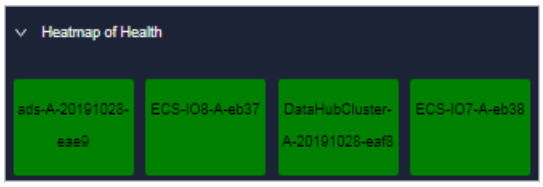

   - **Data of Top 5 Services**

     The Data of Top 5 Services section shows the data of the top five healthiest clusters of the specified service for the current date over the time range from 00:00 to the current time.

     This section shows the top five clusters in terms of abnormal water levels, abnormal masters, abnormal disks, and abnormal chunk servers. Click the cluster name to go to the corresponding cluster information page.

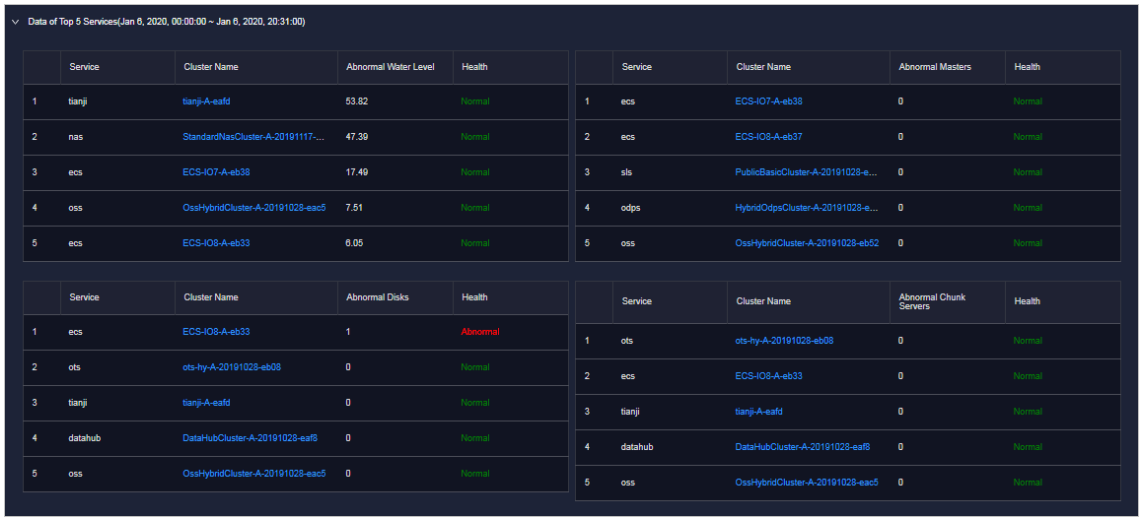

4.1.6.1.2. Cluster information

The Cluster Information module allows you to view the overview information and run charts of a cluster.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > Apsara Distributed File Storage > Cluster Information**.

   On the page that appears, the data of the first cluster in the **Cluster Name** drop-down list is displayed.

2. Select the cluster that you want to view from the **Cluster Name** drop-down list. The following information is displayed:

   > **Note** All the enabled clusters that are accessed within the current environment are displayed in the **Cluster Name** drop-down list.

   - **Overview**

     This section shows the storage space, server information, and health information of the specified cluster. In the **Health** section, when the value of **Abnormal Disks**, **Abnormal Masters**, **Abnormal Chunk Servers**, or **Abnormal Water Levels** is greater than 0, the value is displayed in red font.

     

   - **Alarm Monitor**

     This section shows the alert information of the specified cluster. You can query data by keyword.

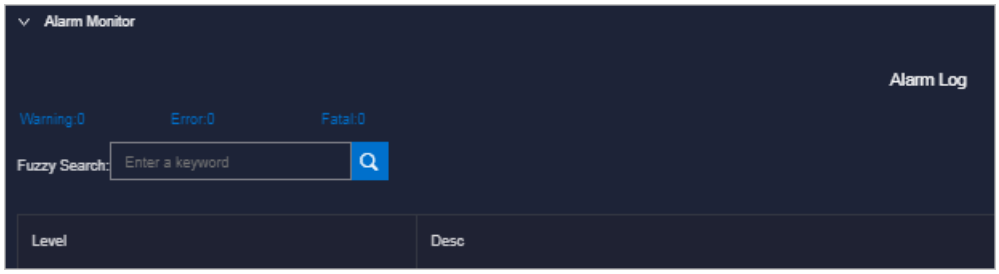

- **Replica**

  This section shows the replica information of the specified cluster.

- **Run Chart of Clusters**

  This section shows the charts of historical water levels, predicted water levels, number of files, number of chunk servers, and number of disks for the specified cluster.

  Predicted water levels predicts the run chart of the next seven days.

  > **Note** The water level can only be predicted if there is enough historical water level data. Some clusters may not have predicted water levels.

- **Rack Information**

  Rack information includes rack capacity and servers in rack.

  - **Servers in Rack** shows the number of machines in each rack of the specified cluster.

- **Storage** shows the total and used storage of each rack in the specified cluster.

4.1.6.1.3. Node information

The Node Information module allows you to view the master information and chunk server information in a cluster.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > Apsara Distributed File System > Node Information**.

   On the page that appears, the data of the first cluster in the **Cluster Name** drop-down list is displayed, including master information and chunk server information.

2. Select the name of the cluster that you want to view from the **Cluster Name** drop-down list. The following information is displayed:

   > **Note** All accessed clusters that are not disabled in the current environment are displayed in the **Cluster Name** drop-down list.

   - **Master Info**

     This section shows the master information of the specified cluster. You can click **Refresh** to refresh the master information of the specified cluster.

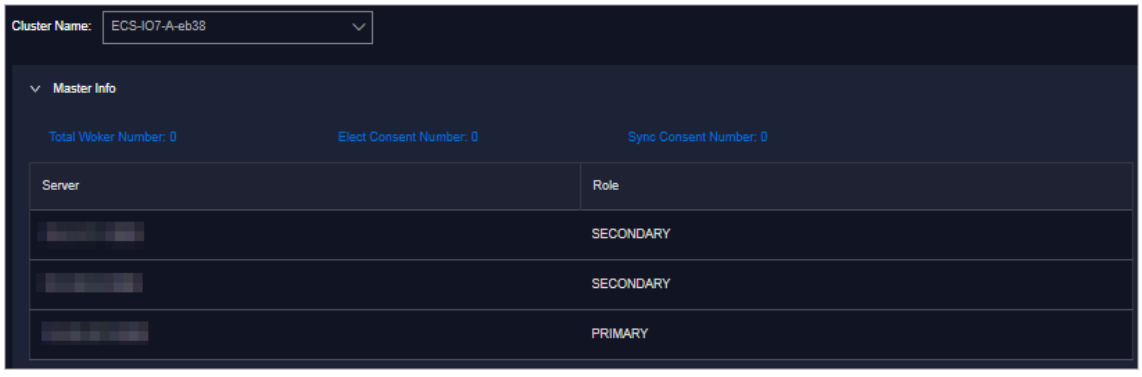

   - **Chunk Server Info**

     This section shows the chunk server information of the specified cluster. You can click **Refresh** to show the chunk server information of the cluster. Click the **+** icon in front of a server, the disk and SSD cache information of the server is displayed. Fuzzy search is supported in this section.

4.1.6.1.4. Operations and maintenance

The **Operations and Maintenance** module allows you to view the cluster status.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > Apsara Distributed File System > Operations and Maintenance**.
2. Select a service from the **Service** drop-down list to view the cluster status of the service. Clusters in different health status are displayed in different colors.
   - Green indicates that the cluster works properly.
   - Yellow indicates that the cluster has a warning.
   - Red indicates that the cluster has an exception.
   - Dark red indicates that the cluster has a fatal error.
   - Grey indicates that the cluster is disabled.

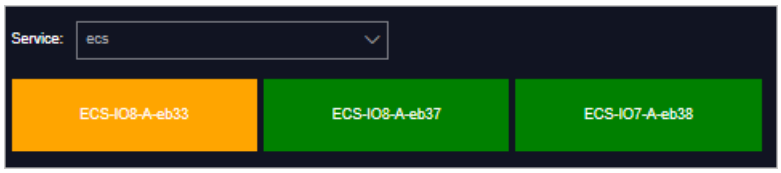

3. Move the pointer over a cluster name to view the service name, server name, and IP address of the cluster.

4.1.6.1.5. Product configuration

By default, the system configures thresholds for all clusters. You can modify the storage usage threshold, chunk server threshold, and disk threshold for each cluster.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > Apsara Distributed File System > Product Configuration**.
2. In the upper part of the page, select the cluster that you want to configure from the **Cluster Name** drop-down list.
3. In the lower part of the page, click **Modify** to modify the thresholds of the cluster.

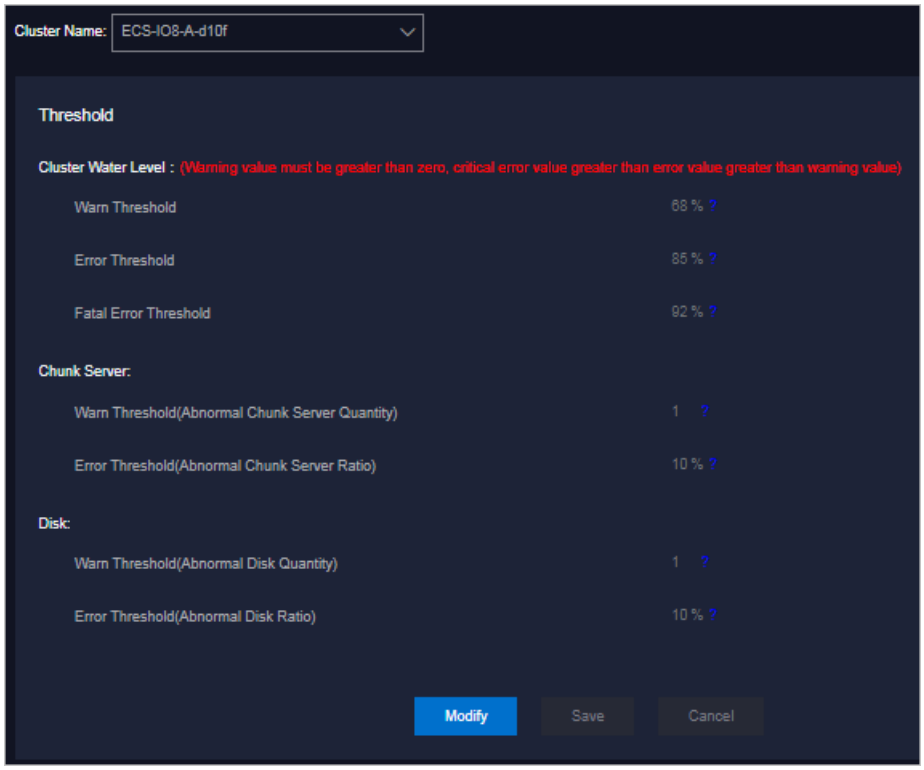

The following table describes the parameters.

Section		Description
Cluster Water Level	**Warn Threshold**	When the storage usage of the cluster is greater than or equal to this value, a warning alert is triggered and the health heatmap of the cluster is displayed in yellow. Value range: (0,100]. If this parameter is not specified, a warning alert is triggered by default when the storage usage of the cluster is greater than or equal to 65%.
	Error Threshold	When the storage usage of the cluster is greater than or equal to this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. Value range: (0,100]. If this parameter is not specified, an error alert is triggered by default when the storage usage of the cluster is greater than or equal to 85%.
	Fatal Error Threshold	When the storage usage of the cluster is greater than or equal to this value, a fatal-error alert is triggered and the health heatmap of the cluster is displayed in dark red. Value range: (0,100]. If this parameter is not specified, a fatal-error alert is triggered by default when the storage usage of the cluster is greater than or equal to 92%.

Section		Description
Chunk Server	**Warn Threshold (Abnormal Chunk Server Quantity)**	When the number of abnormal chunk servers is greater than or equal to this value, a warning alert is triggered and the health heatmap of the cluster is displayed in yellow. If this parameter is not specified, a warning alert is triggered by default when the number of abnormal chunk servers is greater than or equal to 1.
	Error Threshold (Chunk Server Ratio)	If the ratio of abnormal chunk servers to all chunk servers is greater than this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. If this parameter is not specified, an error alert is triggered by default when the ratio of abnormal chunk servers to all the chunk servers is greater than or equal to 10%.
Disk	**Warn Threshold (Abnormal Disk Quantity)**	When the number of abnormal disks is greater than or equal to this value, a warning alert is triggered and the health heatmap of the cluster is displayed in yellow. If this parameter is not specified, a warning alert is triggered by default when the number of abnormal disks is greater than or equal to 1.
	Error Threshold (Abnormal Disk Ratio)	When the ratio of abnormal disks to all disks is greater than this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. If this parameter is not specified, an error alert is triggered by default when the ratio of abnormal disks to all the disks is greater than or equal to 10%.

Note To reset the configurations during the modification, you can click **Cancel** to cancel the current configurations.

4. Click **Save**.

4.1.6.2. EBS

EBS provides the following features: EBS dashboard, block master O&M, block server O&M, snapshot server O&M, block gcworker O&M, device O&M, rebalance, IO HANG fault analysis, slow IO analysis, and inventory configuration.

4.1.6.2.1. EBS dashboard

The EBS Dashboard module shows the data overview and trend charts of cluster usage of EBS clusters.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > EBS Dashboard**.On the page that appears, cluster overview information and trend charts of cluster of all EBS clusters are displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. View the following information:
   - The **Overview** section shows data overview information of the selected cluster, including the storage space, server information, and health information.

     In the **Health** section, when the value of **Abnormal Cloud Disks**, **Abnormal Masters**, **Abnormal Block GcWorker**, or **Abnormal Block Servers** is greater than 0, it is displayed in red.
   - The **Trend Chart of Cluster Usage** section shows the storage usage curve of the cluster for the last 30 days.

4.1.6.2.2. Block master operations

The Block Master Operations module shows the block master node information of EBS clusters, including the IP address and role. The module also allows you to switch the role of a node to LEADER as well as query and configure flags.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Block Master Operations**.On the page that appears, the master node list and cluster information of the first cluster in the **Cluster Name** drop-down list are displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the master node list

     You can view the master node information of the selected cluster, including the IP address, role, log ID, and status.
   - Switch to LEADER

     A LEADER role for a master node has the same functions as a FOLLOWER role, including controlling and scheduling resources, as well as controlling deployment and service configurations.

     If a node in the master node list assumes a FOLLOWER role, you must switch its role to LEADER. Click **Switch to LEADER** in the **Actions** column. In the message that appears, click **OK**.
   - Query a flag

     In the master node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, set flag_key, and then click **Submit**. The deployment and service configurations of the block master node are displayed.

     Perform the following steps to query the flag_key value:

a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.
b. Enter EBS in the **Cluster** search box.
c. Find the EBS cluster and click the cluster name.
d. Click the **Configure** tab.
e. Find the *pangu_blockmaster_flag.json* file in */services/EbsBlockMaster/user/pangu_blockmaster*.

   The flag_key values of all block master nodes are stored in the *pangu_blockmaster_flag.json* file.

- Configure a flag

  If you want to modify the deployment and service configurations of a block master node, you can configure a flag and assign it to the node.

  In the master node list, click **Configure Flag** in the **Actions** column corresponding to a LEADER node. In the dialog box that appears, configure the parameters, and then click **OK**.

  The following table describes the parameters of a flag.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_blockmaster_flag.json* file.
flag_value	This value is customized.
flag_type	Select a flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Check the maser node status

  In the master node list, choose **More > Check Master Status** in the **Actions** column corresponding to a node.

- Query the version information

  In the master node list, choose **More > Query Version Information** in the **Actions** column corresponding to a node.

- Query the cluster overview information

  You can query the disk size, number of segments, total storage size, and storage usage of the cluster.

4.1.6.2.3. Block server operations

The Block Server Operations module shows the block server node information of EBS clusters, including the IP address, status, and real-time server load. The module also allows you to query and modify flags, configure server node status, as well as add nodes to and delete nodes from blacklists.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Block Server Operations**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the server node list

     You can view server node information of the cluster including the IP addresses, status, number of segments, and real-time load (read IOPS, write IOPS, read bandwidth, write bandwidth, read latency, and write latency).
   - Query a flag

     In the server node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, set flag_key, and then click **Submit**. The deployment and service configurations of the block server node are displayed.

     Perform the following steps to query the flag_key value:

     a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.

     b. Enter EBS in the **Cluster** search box.

     c. Find the EBS cluster and click the cluster name.

     d. Find the *pangu_blockserver_flag.json* file in */services/EbsBlockServer/user/pangu_blockserver*.

     The flag_key values of all block server nodes are stored in the *pangu_blockserver_flag.json* file.
   - Configure a flag

     In the server node list, click **Configure Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, flag_value, and flag_type, and then click **OK**.

     The following table describes the parameters of a flag.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_blockserver_flag.json* file.
flag_value	This value is customized.

Parameter	Description
flag_type	Select a flag type. Valid values: ■ int ■ bool ■ string ■ double

- Configure server node status

  In the server node list, choose **More > Set Server Status** in the **Actions** column corresponding to a node. In the dialog box that appears, specify server node status, and then click **OK**.

  The following table describes the server node status.

Status	Description
NORMAL	Indicates that the node is normal.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	The node has been disabled.
UPGRADE	The node has been upgraded.
RECOVERY	The node has been restored.

- Query the version information

  In the server node list, choose **More > Query Version Information** in the **Actions** column corresponding to a node. In the dialog box that appears, view the version information of the block server node.

- Add a block server node to the blacklist

  In the upper-right corner of the **Block Server Blacklist** section, click **Add**. In the dialog box that appears, select the IP address of the block server node that you want to add to the blacklist, and then click **OK**.

  The block server node that is added to the blacklist is disabled and will no longer provide services.

- View the block server blacklist

  You can view all block server nodes that are added to the blacklist in the **Block Server Blacklist** section.

- Remove a block server node from the blacklist

  In the **Block Server Blacklist** section, find the block server node that you want to remove from the blacklist, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

The block server node that is removed from the blacklist can continue to provide internal and external services.

4.1.6.2.4. SnapShotServer

The SnapShotServer module shows the snapshot server node information of EBS clusters, including the IP address, status, and other performance parameters. The module also allows you to query and modify flags and configure snapshot server node status.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > SnapShotServer**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the snapshot server node list

     You can view snapshot server node information of the cluster including node IP address, status, loading rate, and number of uploads, replicas, and delayed loadings.
   - Query a flag

     In the snapshot server node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, set flag_key, and then click **Submit**. The deployment and service configurations of the snapshot server node are displayed.

     Perform the following steps to query the flag_key value:

     a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.

     b. Enter EBS in the **Cluster** search box.

     c. Find the EBS cluster and click the cluster name.

     d. Find the *pangu_snapshotserver_flag.json* file in */services/EbsSnapshotServer/user/pangu_snapshotserver*.

     The flag_key values of all snapshot server nodes are stored in the *pangu_snapshotserver_flag.json* file.
   - Configure a flag

     In the snapshot server node list, click **Configure Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, flag_value, and flag_type, and then click **OK**.

     The following table describes the parameters of a flag.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_snapshotserver_flag.json* file.
flag_value	This value is customized.

Parameter	Description
flag_type	Select a flag type. Valid values: ■ int ■ bool ■ string ■ double

- Configure the snapshot server node status

  In the snapshot server node list, choose **More > Set snapshotserver Status** in the **Actions** column corresponding to a node. In the dialog box that appears, select the snapshot server node status, and then click **OK**.

  The following table describes the snapshot server node status.

Status	Description
NORMAL	Indicates that the node is normal.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	Indicates that the node has been disabled

- Query the version information

  In the snapshot server node list, choose **More > Version Information** in the **Actions** column corresponding to a node. In the dialog box that appears, view the version information of the node.

4.1.6.2.5. Block gcworker operations

The Block Gcworker Operations module allows you to view the IP addresses and statuses of block gcworker nodes in EBS clusters. You can also query and modify flags, configure the gcworker node status, and query version information.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Block GcWorker Operations**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the gcworker node list

     You can view the IP addresses and statuses of the block gcworker nodes in the selected cluster.
   - Query a flag

In the gcworker node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, and then click **Submit**. The deployment and service configurations of the block gcworker node are displayed.

Perform the following steps to query the flag_key value:

a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.
b. Enter EBS in the **Cluster** search box.
c. Find the EBS cluster and click the cluster name.
d. Click the **Configure** tab.
e. Find the *pangu_blockgcworker_flag.json* file in */services/EbsBlockGCWorker/user/pangu_blockgcworker*.

   The flag_key values of all block server nodes are stored in the *pangu_blockgcworker_flag.json* file.

- Configure a flag

  In the gcworker node list, click **Configure Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, flag_value, and flag_type, and then click **OK**.

  The following table describes the parameters.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_blockgcworker_flag.json* file.
flag_value	This value is customized.
flag_type	Select a flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the gcworker node status

  In the gcworker node list, choose **More > Configure gcworker Status** in the **Actions** column corresponding to a node. In the dialog box that appears, specify the gcworket node status and click **OK**.

  The following table describes the gcworker status.

Status	Description
NORMAL	Indicates that the node is running normally.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.

Status	Description
OFFLOADED	Indicates that the node has been disabled.

- Query the version information

  In the gcworker node list, choose **More > Query Version Information** in the **Actions** column corresponding to a node. In the dialog box that appears, view the version information of the block gcworker node.

4.1.6.2.6. Device operations

The Device Operations module allows you to view disk information in EBS clusters such as the disk ID, status, capacity, and type. You can also perform flush operations, modify disk configurations, query segment information, and open, close, delete, and restore devices.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Device Operations**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the device list

     You can view the total number of devices, the total logical space of devices, and information about each device in the cluster, including the device ID, status, logical capacity, number of segments, mode, and flags.
   - Global check segments

     In the upper-right corner of the **Device List** section, click **Global Check Segment**. You can view all the segments in the selected cluster and their indexes and statuses.
   - Check the status of disks

     In the upper-right corner of the **Device List** section, click **Check Cloud Disk Status**. You can view the number of invalid disks in the selected cluster.
   - Query device information

     In the device list, click **Query Device Information** in the **Actions** column corresponding to a device. In the dialog box that appears, view configuration information of the disk such as the disk ID, disk status, and disk capacity.
   - Delete a device

     In the device list, click **Delete** in the **Actions** column corresponding to a device.

     After you delete the disk, its status becomes **DELETING**, and the disk is unavailable. You are not allowed to perform operations such as enabling the device and modifying the configuration.
   - Restore a device

     In the device list, find the device whose status is **DELETING** and click **Restore** in the **Actions** column. In the dialog box that appears, click **OK**.

After you restore the disk, it becomes available. You can perform operations such as enabling the disk and modifying the configuration.

- Enable a device

  In the device list, choose **More > Turn On** in the **Actions** column corresponding to a device. In the dialog box that appears, configure the parameters and click **Submit**.

  > **Note** You can perform read and write operations on a disk only after the disk is enabled.

  The following table describes the parameters for enabling a device.

Parameter	Description
client_ip	Optional. Specifies the client where the disk is enabled. The client IP address is the IP address of the block server. If the client IP address is not specified, the IP address of the local server is used by default.
token	Specifies a string as a token to be used to disable the device.
mode	Specifies the disk mode. Valid values: ■ **ro**: read-only ■ **rw**: read/write Default value: **rw**

- Disable a device

  > **Notice** After a disk is disabled, data can no longer be read from or written to the disk. Proceed with caution.

  In the device list, choose **More > Turn Off** in the **Actions** column. In the dialog box that appears, configure the parameters and click **Submit**.

  The following table describes the parameters for disabling a device.

Parameter	Description
client_ip	Specifies the client IP address of the disk to be disabled. If the client IP address is not specified, the IP address of the local server is used by default.
token	Specifies the token for disabling the device, which is configured when the device is enabled. You can query the token by running the **dev -query** command.

Parameter	Description
open_ver	Specifies the current openversion of the device if the client IP address is not specified. If you specify a client IP address, you do not need to specify openversion. You can query openversion by running the **dev -query** command.

- Flush

  In the device list, choose **More > Flush** in the **Actions** column. In the dialog box that appears, configure the parameters and click **Submit** to clear the current disk or the segment transaction logs on the disk.

  The following table describes the parameters.

Parameter	Description
segment	Select the segment to be flushed. If you do not select any segments, all segments are flushed.
ifnsw	Valid values: ■ **0**: specifies that the index file is flushed during the flush. ■ **1**: specifies that the index file is not flushed during the flush.
dfnsw	Valid values: ■ **0**: specifies that data files are flushed during the flush. ■ **1**: specifies that data files are not flushed during the flush.

- Global flush

  You can perform the flush operation to clear disks or the transaction logs of segments.

  In the upper-right corner of the **Device List** section, click **Global Flush**. In the dialog box that appears, select ifnsw and dfnsw, and then click **OK** to clear all the disks or the transaction logs of segments in the selected cluster.

- Query configuration status

  In the device list, choose **More > Query Configuration Status** in the **Actions** column corresponding to a device. In the dialog box that appears, specify config_ver and click **OK**. You can determine whether the disk can be configured based on the query result.

  You can obtain the config_ver value from the device information.

- Modify device configurations

  You can modify the configurations of a disk, such as specifying whether to enable data compression, compression algorithms, and storage modes.

  In the device list, choose **More > Modify Device Configurations** in the **Actions** column corresponding to a device. In the dialog box that appears, modify the parameters and click **OK**.

  The following table describes the parameters.

Parameter	Description
compress	Select whether to enable data compression. Valid values: ■ **enable** ■ **disable**
algorithm	Select a data compression algorithm. Valid values: ■ **0**: indicates that no data compression algorithms are used. ■ **1**: indicates that the snappy data compression algorithm is used. ■ **2**: indicates that the lz4 data compression algorithm is used.
ec	Select whether to enable the ec storage mode. Default value: disable. Valid values: ■ **enable** ■ **disable**
data_chunks	Specifies the number of data chunks. Default value: 8.
parity_chunks	Specifies the number of parity chunks. Default value: 3.
packet_bits	Specifies the size of single data block in ec mode. Default value: 15.
copy	Specifies the number of data replicas. Default value: 3.
storage_mode	Specifies the storage mode of the disk.
cache	Select whether to enable the cache mode. Default value: 0. Valid values: ■ **0**: disabled ■ **1**: enabled
storage_app_name	Specifies the data storage name.

Parameter	Description
simsuppress	Select whether to enable the delay simulation feature. Default value: disable. Valid values: ■ enable ■ disable
baselatency	Specifies the basic latency. Default value: 300.
consumespeed	Specifies the processing speed. Default value: 256 bit/μs.
lat80th	Specifies the quantile jitter control of the latency as 80%.
lat90th	Specifies the quantile jitter control of the latency as 90%.
lat99th	Specifies the quantile jitter control of the latency as 99%.

- Query segment information

  In the device list, choose **More > Segment Information** in the **Actions** column corresponding to a device. In the dialog box that appears, view the segment information such as index and status.

- Query segments

  In the device list, choose **More > Check Segment** in the **Actions** column corresponding to a device. In the dialog box that appears, select the segment to be checked and click **Submit**. You can view the information of the selected segment such as index and status.

4.1.6.2.7. Enable or disable Rebalance

When segments are unevenly distributed among a block server, you can enable the Rebalance feature to redistribute the segments. After the redistribution, you can disable Rebalance.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Rebalance**.
2. Click **Enable Rebalance** or **Disable Rebalance**.

   After you click **Enable Rebalance**, the status of Rebalance changes to **running**.

   After you click **Disable Rebalance**, the status of Rebalance changes to **stopped**.

4.1.6.2.8. IO hang fault analysis

The IO HANG Fault Analysis module allows you to view the affected virtual machine (VM) list, VM cluster statistics, and device cluster statistics.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **Storage Operation Center > EBS > IO HANG**. By

default, the system displays the affected VM list, VM cluster statistics, and device cluster statistics in the last 24 hours.

3. Select the time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or customize the time range) that you are about to view and then click **Search**. View the following information:
   - **Affected VM List**

     The **Affected VM List** section displays the IO HANG start time and recovery time of all the VMs, and the cluster name and user ID to which these VMs belong.

     To view the information of a cluster, user, or VM, enter the cluster name, user ID, or VM name in the search box to perform a fuzzy search.

     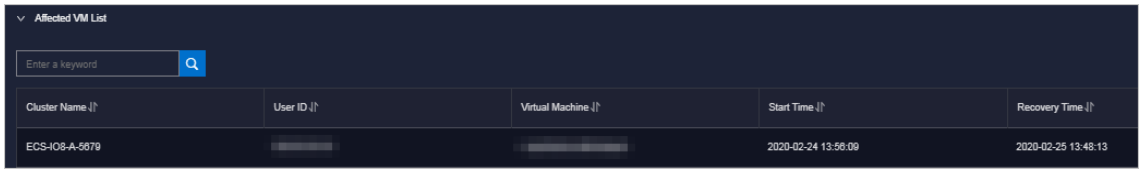

   - **VM Cluster Statistics**

     The **VM Cluster Statistics** section displays the number of affected VMs in a cluster.

     To view the VM statistics of a cluster, enter the cluster name in the search box to perform a fuzzy search.

     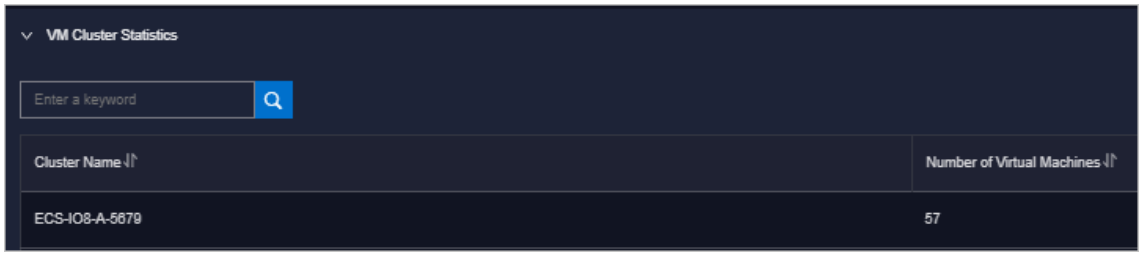

   - **Device Cluster Statistics**

     The **Device Cluster Statistics** section displays the number of affected devices in a cluster.

     To view the device statistics of a cluster, enter the cluster name in the search box to perform a fuzzy search.

     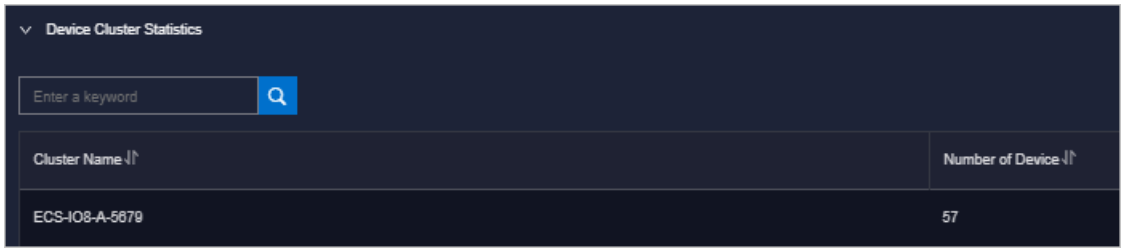

4.1.6.2.9. Slow IO analysis

The Slow IO Analysis module allows you to view the Slow IO list, top 10 NCs, cluster statistics, top 5 cluster statistics, and reasons.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **Storage Operation Center > EBS > Slow IO**. By default,

the system displays the Slow IO list, top 10 NCs, cluster statistics, top 5 cluster statistics, and reasons in the last 24 hours.

3. Select the time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or customize the time range) that you are about to view and then click **Search**. View the following information:
   - **Slow IO List**

     The **Slow IO List** section displays the following Slow IO-related data: cluster name, NC IP address, virtual machine, device ID, storage type, start time, recovery time, number of Slow IO, and reason.

     To view the information of a cluster, NC, or block device, you can enter the cluster name, NC IP address, or device ID in the search box to perform a fuzzy search.

     You can also sort by Cluster Name, NC IP, Virtual Machine, Device ID, Storage Type, Start Time, Recovery Time, Number of Slow IO, and Reason as needed.

     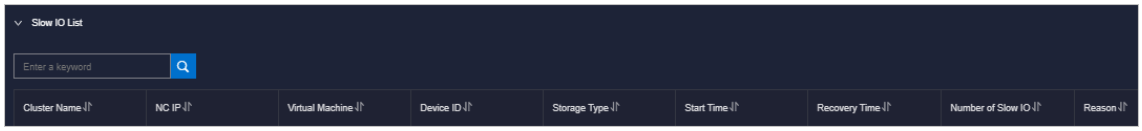

   - **Top Ten NC**

     The system displays the information of top 10 NCs by using a graph and a list.

     Where,

     - The **Graphic Analysis** section displays the proportion for the number of Slow IO in each cluster of the top 10 NCs by using a pie chart.
     - The **Top Ten NC** section displays the NC IP address, cluster name, Slow IO, percentage, and major reason of the top 10 NCs with the most Slow IO by using a list.

       To view the information of a cluster or NC, you can enter the NC IP address or cluster name in the search box to perform a fuzzy search.

       You can also sort by NC IP, Cluster Name, Slow IO, and Major Reason as needed.

       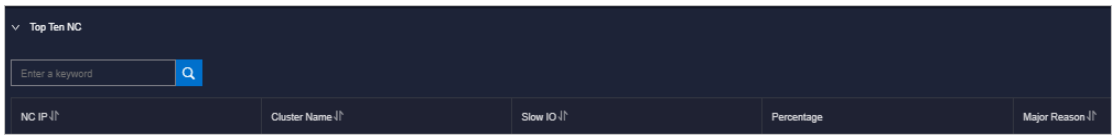

   - **Cluster Statistics**

     The **Cluster Statistics** section displays the cluster name, number of devices, number of Slow IO, percentage, and major reason of a cluster with Slow IO.

     To view the information of a cluster, you can enter the cluster name in the search box to perform a fuzzy search.

     You can also sort by Cluster Name, Number of Device, Number of Slow IO, and Major Reason as needed.

   - **Top Five Cluster Statistics**

     The system displays the statistics of top 5 clusters by using a graph and a list.

     Where,

- The **Top Five Cluster Statistics** section displays the cluster name, number of devices, number of Slow IO, percentage, and major problem of the top 5 clusters with the most Slow IO by using a list.

  To view the information of a cluster, you can enter the cluster name in the search box to perform a fuzzy search.

  You can also sort by Top Five Cluster, Number of Device, Number of Slow IO, and Major Problem as needed.

  

- The **Graphic Analysis** section displays the proportion for the number of Slow IO in each of the top 5 clusters by using a pie chart.

- **Reason**

  The system displays the reason statistics by using a graph and a list.

  Where,

  - The **Reason** section displays the number of Slow IO from the dimension of reasons.

    To view the information of a reason, you can enter the reason information in the search box to perform a fuzzy search.

    You can also sort by Reason and Number of Slow IO as needed.

    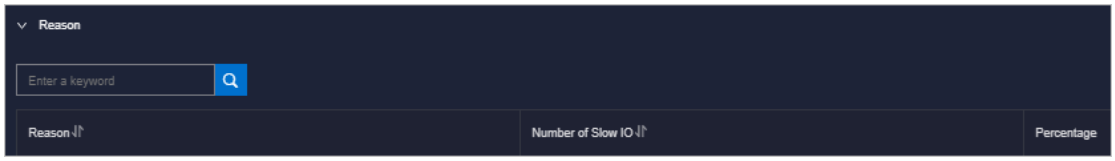

  - The **Graphic Analysis** section displays the proportion of reasons by using a pie chart.

4.1.6.2.10. Inventory configuration

The **Inventory Settings** module allows you to view the sales status of a cluster, configure the oversold ratio of a cluster, and configure whether a cluster is on sale.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **Storage Operation Center > EBS > Inventory Settings**. By default, the system displays the data, namely the cluster name, oversold ratio, and sales status, of all the clusters in the current environment.

3. Complete the following configurations:
   - Select a cluster. Enter a number in the **Adjust Setting Oversell Ratio(%)** field, and then click **Confirm** to configure the oversold ratio of the cluster.
   - Select a cluster. Turn on or off the **Adjustment of sales status** switch to configure whether the cluster is on sale.

5.Operations tools

5.1. Offline Backup

The Offline Backup module is used to back up the key metadata of Apsara Stack. Currently, you can only back up the pangu metadata. The backed up metadata is used for the fast recovery of Apsara Stack faults.

5.1.1. Service configuration

The Service Configuration module consists of Backup Server Configuration and Product Management.

5.1.1.1. Configure the backup server

You can configure the backup server for the subsequent storage of backup files.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Service Configuration > Backup Server Configuration**.The **Backup Service Configuration** page appears.

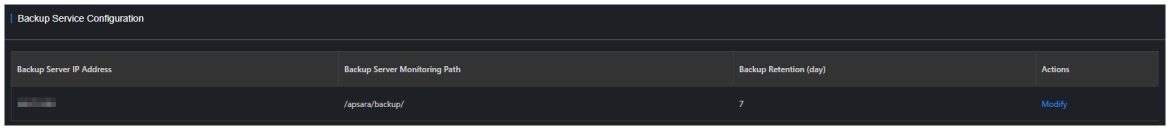

2. Click **Modify** in the **Actions** column corresponding to the backup server to configure the parameters.

Parameter	Description
Backup Server IP Address	The IP address of the backup server. The backup server must meet the following requirements: ○ The backup server is an independent physical server. ○ The backup server is managed and controlled by Apsara Infrastructure Management Framework. ○ The backup server has its network connected with other servers in Apsara Stack. ○ Apsara Distributed File System cannot be deployed on the server or on the disk where backup metadata is stored.
Backup Service Monitoring Path	The storage path of backup files on the backup server. The backup service detects new backup files by monitoring the specified folder on the backup server and determines whether the backup is successful by comparing the MD5 value of the backup file with that of the original file.
Backup Retention	The actual time period (in days) that backup files are stored. Backup files that exceeds the specified time period are deleted.

5.1.1.2. Add a backup product

The Product Management module allows you to add product backup information. Currently, only the metadata of Apsara Distributed File System can be backed up.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Service Configuration > Product Management**.
2. In the upper-right corner of the page, click **Add**.

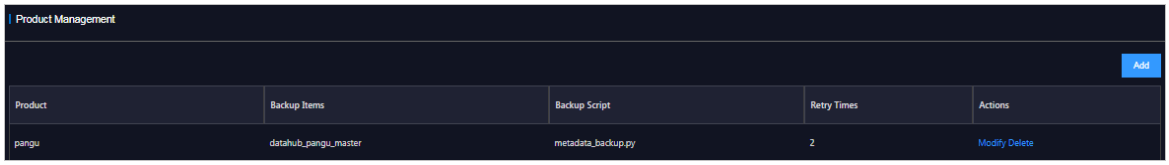

3. In the **Add Product** dialog box that appears, configure the parameters based on the following table, and then click **OK**.

Parameter	Description
Product	Set this parameter to pangu to back up the Apsara Distributed File System data.
Backup Items	Specify this parameter based on the Apsara Distributed File System information of the cloud product to be backed up in the format of product name_pangu. Example: ecs_pangu.
Backup Script	The backup script name. Example: metadata_backup.py.
Retry Times	The number of times to retry after an error occurs. Typically, set this parameter to 3.

   You can view the added product on the Backup Configuration page by choosing **Backup Service > Backup Configuration**.
4. To add more product backup items, follow the preceding steps again.

   You can click **Modify** or **Delete** in the **Actions** column to modify or delete a product backup item.

5.1.2. Backup service

The **Backup Service** module consists of the backup configuration, backup details, and service status.

5.1.2.1. Backup configuration

After you add a product backup item, you must configure backup parameters in the ASO console.

Prerequisites

A product backup item is added. For more information about how to add a product backup item, see Add a backup product.

Context

The backup item is the minimum unit for backup. You can back up the metadata of Apsara Distributed File System for different products such as ecs pangu and ots pangu.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Backup Service > Backup Configuration**. The left part of the Backup Configuration page shows the current backup configurations in a hierarchical tree structure. The root node is a product list and shows the products whose data can be backed up in the current backup system. Currently, only the metadata of Apsara Distributed File System can be backed up.
2. Click a product backup item on the left and then configure the parameters on the right.

Parameter	Description
Product Cluster Location	The IP address of the actual transfer server.
Backup File Folder	A folder on the transfer server. You are only required to enter the value in the field, instead of manually creating the folder. Example: /apsarapangu/disk8/pangu_master_bak /*product name*_pangu/bin
Script Execution Folder	A folder on the transfer server. You are only required to enter the value in the field, instead of manually creating the folder. Example: /apsarapangu/disk8/pangu_master_bak /*product name*_pangu/bin
Script Parameters	You must enter the value in the format of `--ip=xxx.xxx.xxx.xxx`, in which the IP address is any IP address of the pangu master.
Backup Schedule	In this example, a value of 1 is entered to specify that the backup is performed only once.
Backup Schedule Unit	Select **Day**, **Hour**, or **Minute**. In this example, **Hour** is selected to specify that the backup is performed by the hour.
Time-out	Select the timeout period. In this example, **3600** is entered.

3. Click **Modify** to complete the configurations and trigger the backup.
4. Follow the preceding steps to configure all the product backup items.

5.1.2.2. View the backup details

You can view the backup details of each backup item in the ASO console during the backup.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Backup Service > Backup Details**.
2. On the **Backup Details** page, specify the product and backup item, select the start date and end

date, and then click **Search**.

3. View the backup details of a backup item, including the product, backup item, the name of the file to back up, start time, and status. The backup status includes **Not started**, **In transmission**, **Time-out** and **Failed**.

5.1.2.3. View the backup server status

You can view the memory and disk usage, as well as CPU utilization of the current backup server before and after the backup.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Backup Service > Service Status**.
2. On the **Service Status** page, view the memory and disk usage, as well as CPU utilization of the current backup server.

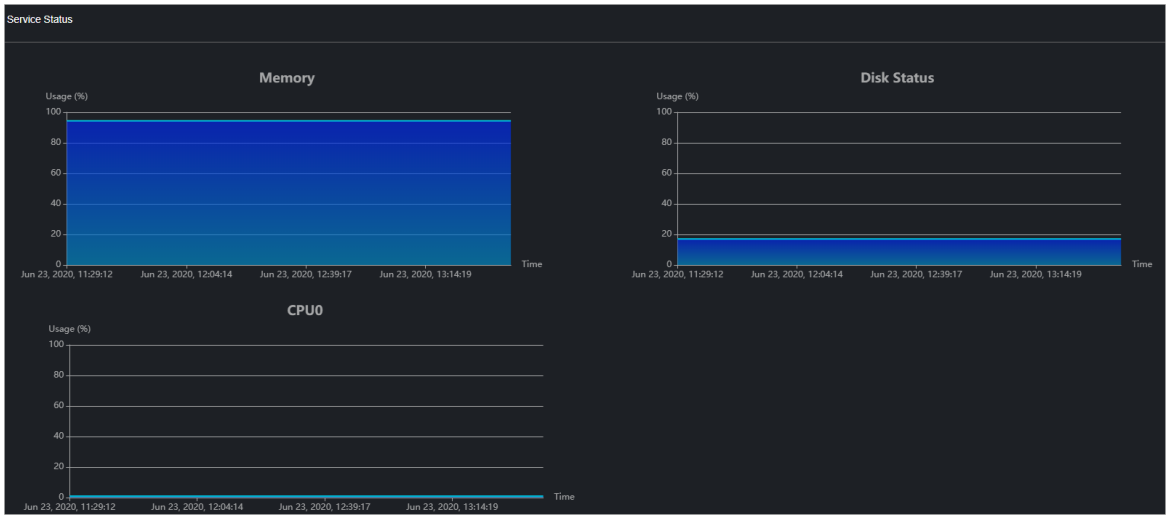

5.1.3. View the backup status

The Service Status module allows you to view the status of the current backup server and the status of the current backup service, including the backup product, completed backup items, timeout backup items, and failed backup items.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Service Status > Status**.

2. On the **Status** page, view the current backup status.
   - View the number of backup items that are in the **In Process**, **Complete**, **Time-out**, and **Failed** state.
   - View the status of the latest backup items of the current service

     The backup status contains the following types: **Success**, **Not Started**, **In Process**, **Timeout**, and **Failure**.
   - On the right side of the page, view the status of the current backup server, including the memory and disk usage, as well as CPU utilization.

5.1.4. Use cases

5.1.4.1. Offline backup of metadata

To guarantee the availability of cloud platforms, you must back up the pangu data of each product.

5.1.4.1.1. Preparations before the backup

This topic describes the preparations before the backup.

Before the backup, prepare the following machines as required:

- Prepare an independent buffer machine as the backup server.

  If no buffer machine exists in the environment, select the physical machine with large disk space and good network performance in the environment. Otherwise, the security of the backup data cannot be guaranteed.

  Offline backup files cannot be stored on backed up objects. Therefore, if the on-site environment does not have extra physical machines or sufficient disk capacity, you must increase physical machines or disk space before the offline backup.
- A transfer machine is required for backup products to store one-time backup data and backup scripts of each product.

  No other requirements are for the transfer machine.
- The network of the backup server must be connected with the network of the Docker container

where the offline backup service is located to make sure that the backup container in the ASO cluster can log on to the transfer machine and backup server by using SSH, without providing the username and password.

5.1.4.1.2. Collect Apsara Distributed File System information of each product

You can collect the Apsara Distributed File System information of products to be backed up to add the product backup information to ASO.

Context

In this topic, product names are customized as oss, ecs, ads, and ots, and the information of these products are collected. The products whose Apsara Distributed File System information you are about to collect are subject to the on-site environment.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.

   > **Note** In this topic, operations are performed in the new Apsara Infrastructure Management Framework console.

2. In the left-side navigation pane, choose **Operations > Service Operations**.
3. Enter **pangu** in the search box to search for the Apsara Distributed File System service.
4. Click **Operations** in the **Actions** column to go to the service details page.
5. Click the **Clustes** tab.
6. Click the name of the target cluster.
7. On the **Services** tab, click **pangu.PanguMaster#**.

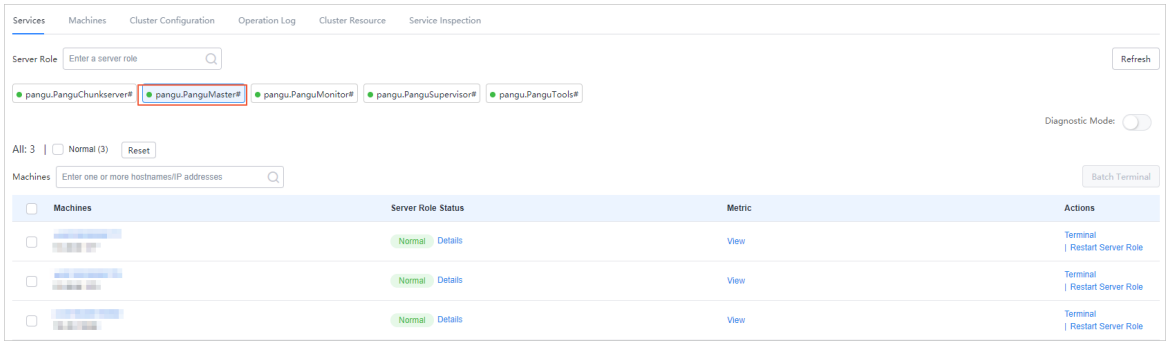

8. View and record the IP address of Apsara Distributed File System master in the server list.Record any one IP address of the three **PanguMaster#**.
9. Repeat Steps 6 to 8 to view and record the Apsara Distributed File System information of each product. The recorded results are similar to those in the following table.

Cluster name	pangumaster IP	Product name
AdvanceOssCluster-A-xx	10.10.10.1	oss

Cluster name	pangumaster IP	Product name
ECS-IO7-A-xx	10.10.10.2	ecs
ads-A-xx	10.10.10.3	ads
otsv3_p-A-xx	10.10.10.4	ots

Note You can customize the product name. Make sure that the product name is unique and recognizable.

5.1.4.1.3. Configure the backup server

This topic describes how to configure the backup server in the ASO console.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Service Configuration > Backup Server Configuration**.
2. Click **Modify** in the **Actions** column corresponding to the backup server to configure the parameters.

Parameter	Description
Backup Server IP Address	The IP address of the actual backup server. The backup server must be an independent physical server that is managed by Apsara Infrastructure Management Framework and has its network connected with other servers in Apsara Stack. Apsara Distributed File System cannot be deployed on the server or on the disk where backup metadata is stored.
Backup Service Monitoring Path	The storage path of backup files on the actual backup server. The backup service detects new backup files by monitoring the specified folder on the backup server and determines whether the backup is successful by comparing the MD5 value of the backup file with that of the original file.
Backup Retention	The actual time period (in days) that backup files are stored. Backup files that exceeds the specified time period are deleted.

5.1.4.1.4. Add a backup product

You can add backup product information in the ASO console.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Service Configuration > Product Management**.
2. In the upper-right corner of the page, click **Add**.
3. In the **Add Product** dialog box, configure the parameters based on the following table, and then click **OK**.

Parameter	Description
Product	Set this parameter to pangu to back up the Apsara Distributed File System data.
Backup Items	Specifies this parameter based on the product information described in the Collect Apsara Distributed File System information of each product topic in the format of *product name*_pangu. Example: oss_pangu.
Backup Script	The backup script name. Example: metadata_backup.py.
Retry Times	The number of times to retry after an error occurs. Typically, set this parameter to 3.

4. Repeat Steps 2 and 3 to add all the product backup items. The following figure shows the result after the backup items are added.

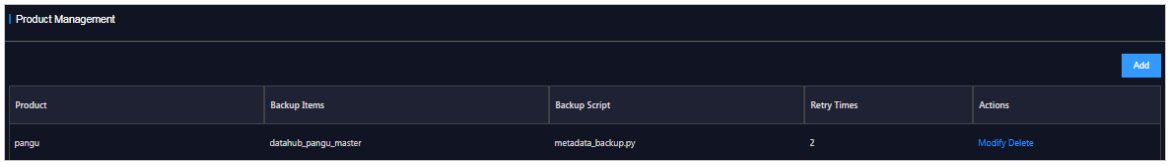

5.1.4.1.5. Configure backup parameters

After you add a product backup item, you must configure backup parameters in the ASO console.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Backup Service > Backup Configuration**.
2. Click a product backup item on the left and then configure the parameters on the right.

Parameter	Description
Product Cluster Location	The IP address of the actual transfer server.
Backup File Folder	A folder on the transfer server. You are only required to enter the value in the field, instead of manually creating the folder. Example: /apsarapangu/disk8/pangu_master_bak /*product name*_pangu/bin.

Parameter	Description
Script Execution Folder	A folder on the transfer server. You are only required to enter the value in the field, instead of manually creating the folder. Example: /apsarapangu/disk8/pangu_master_bak /*product name*_pangu/bin.
Script Parameters	You must enter the value in the format of `--ip=xxx.xxx.xxx.xxx`, in which the IP address is any IP address of the pangu master recorded in the Collect Apsara Distributed File System information of each product topic.
Backup Schedule	In this example, a value of 1 is entered to specify that the backup is performed only once.
Backup Schedule Unit	Select **Day**, **Hour**, or **Minute**. In this example, **Hour** is selected to specify that the backup is performed by the hour.
Time-out	Select the timeout period. In this example, **3600** is entered.

3. Click **Modify** to complete the configurations and trigger the backup.
4. Repeat Steps 2 and 3 to configure all the product backup items.

5.1.4.1.6. View the backup details

After you configure the backup items, you can check whether the backup items work properly in the ASO console.

Procedure

1. In the left-side navigation pane, choose **Offline Backup > Backup Service > Backup Details**.
2. On the Backup Details page, specify the product and backup item, select the start date and end date, and then click Search.

   If the status of a backup item is **Complete**, the backup item is working properly.

   If the backup is finished, view the MD5 value of the backup file to check whether the MD5 value of the offline backup service is the same as that of the backup server. If yes, the backup was successful.

5.2. NOC

5.2.1. Overview

The Network Operation Center (NOC) module is an all-round operations tool platform that covers the whole network (virtual network and physical network).

NOC provides the operations capabilities such as the visualization of network-wide monitoring, automated implementation, automated fault location, and network traffic analysis, which enhances the operations efficiency of network operations engineers, reduces the operations risk, and greatly improves the quality of Apsara Stack network services.

5.2.2. Dashboard

5.2.2.1. Dashboard

The Dashboard tab allows you to monitor the current devices, network, and traffic.

Procedure

1. In the left-side navigation pane, choose **NOC > Dashboard**.
2. On the **Dashboard** tab, view the dashboard information.

Item		Description
Device Management	Device Overview	The model distribution of used network devices.
	Ports Usage	○ **Ports Utilization**: the proportion of ports in use to the total ports in the network devices. ○ **Error Packets by Port (Top 5)**: the total number of error packets generated by device ports within a certain time range, of which the top 5 are displayed.
	Configuration Management	○ **Automatic Backup**: the backup of startup configurations for all network devices. ○ **Configuration Sync**: the synchronization of running configurations and startup configurations for all network devices.
Network Monitoring	Alerts	The total number of alerts generated by network devices.
	Alerting Devices	The number of network devices that generate alerts and the total number of network devices.
	Alarm Details	The details of the alert.

Item		Description
Traffic Dashboard	SLB Overview	The bandwidth utilization of SLB clusters.
	XGW Overview	The bandwidth utilization of XGW clusters.

5.2.2.2. Network topology

The **Network Topology** tab allows you to view the physical network topology.

Procedure

1. In the left-side navigation pane, choose **NOC > Dashboard**.
2. Click the **Network Topology** tab.
3. On the **Network Topology** tab, view the physical network topology of a physical data center.

   You can select **Standard Topology** or **Dynamic Topology** as **Topology Type**.

   If an offset exists between the dynamic topology and the standard topology, when you go to the **Network Topology** tab, a message appears in the upper-right corner of the tab and disappears after a few seconds. You can click **Update Topology** to update the standard topology.

   > **Note**
   >
   > The colors of connections between network devices indicate the connectivity between the network devices:
   >
   > - Green: The connection works properly.
   > - Red: The connection has an error.
   > - Grey: The connection is not enabled.

   By default, if you select **Standard Topology** as **Topology Type**, the **Refresh Alert** switch is turned on. You can turn off **Refresh Alert**, and then devices or connection statuses within the topology are not updated after new alerts are triggered.

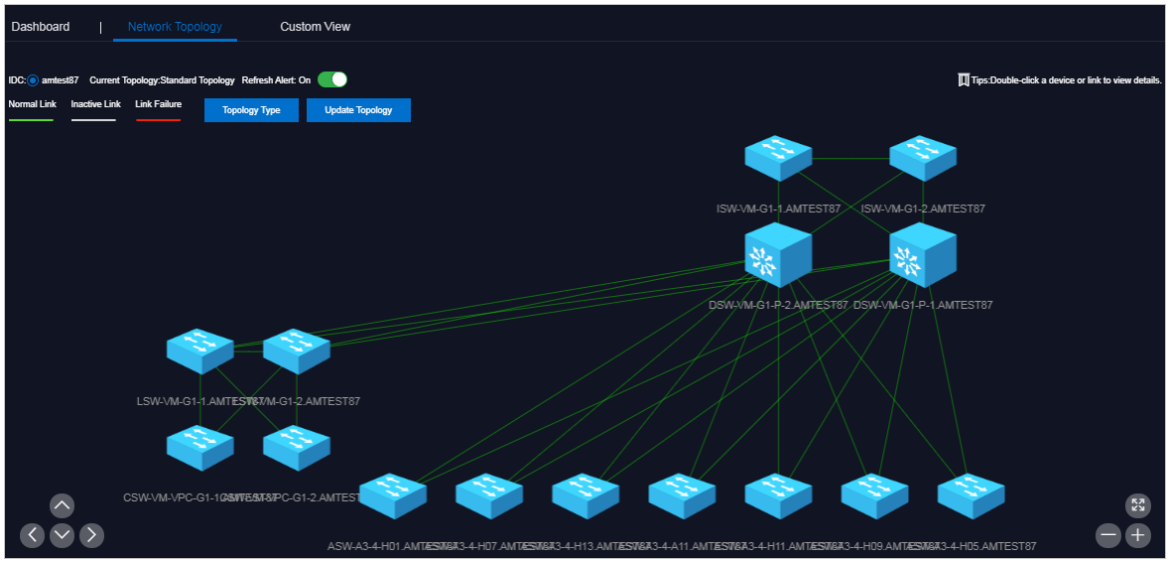

4. In the topology, double-click a connection between two devices to view the connections and alerts between the two devices.
5. In the topology, double-click a physical network device to view the basic information and node alerts of the device on the right.

5.2.2.3. Manage custom views

You can create a custom view to configure how to show the independent monitoring data collection. By configuring the content and rules to show in the view, you can summarize and show the monitoring data and graph information you are interested in.

Create a view

1. In the left-side navigation pane, choose **NOC > Dashboard**.
2. Click the **Custom View** tab.
3. Create a view.
   i. In the upper part of the tab, click **Create View**.

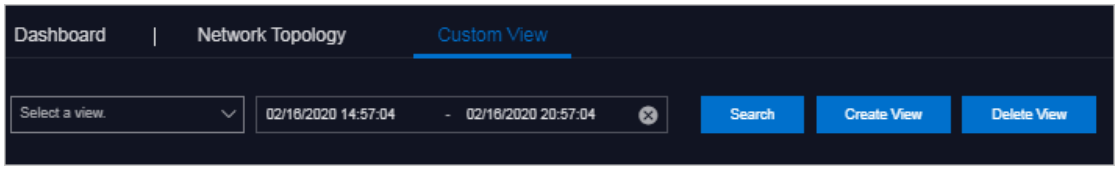

   ii. In the dialog box that appears, enter the view name and description, and then click **OK**.

   The view name must be unique. If the message **A view with the same name already exists** appears, you must change the view name to something unique and then click **OK**.
4. Add a subview. By default, no subviews exist in a view after you create the view.

i. In the search box, select the view and then click **Search**.

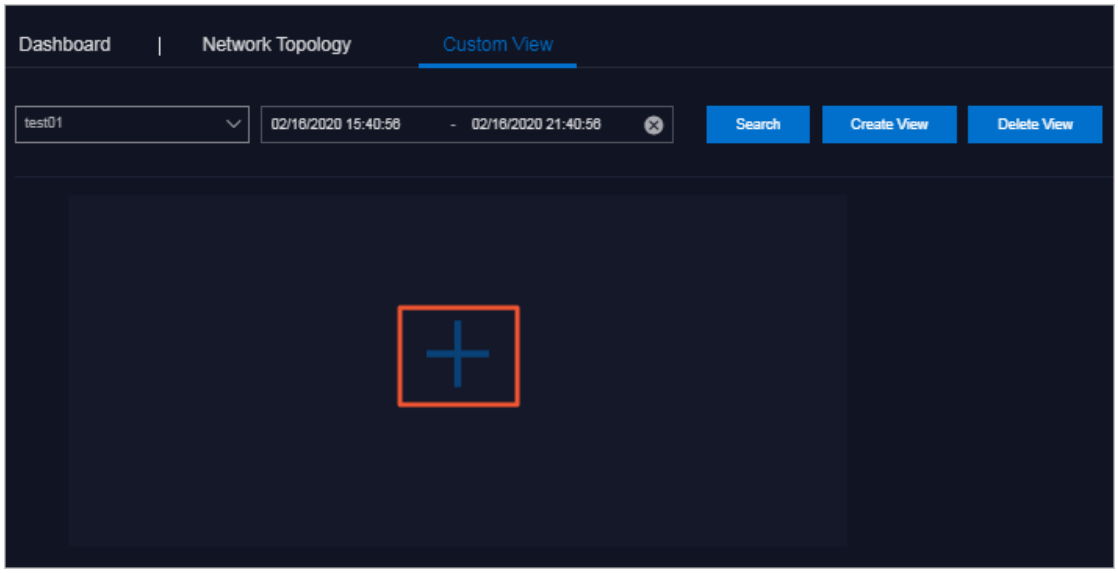

ii. Click the + icon.

iii. In the pane that appears, select the device, monitoring object, monitoring metric, and monitoring submetric.

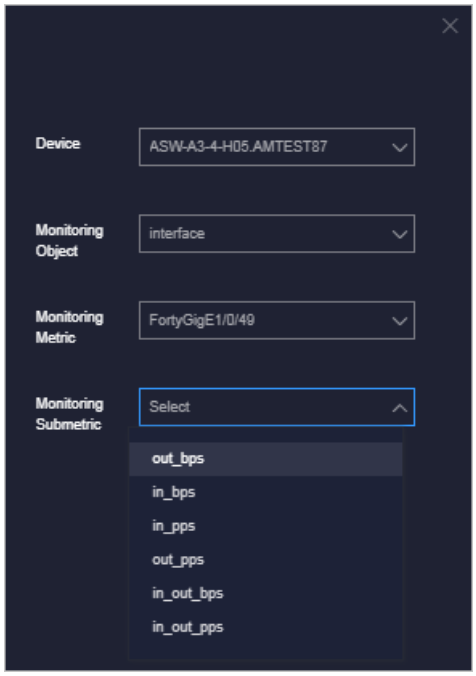

Parameter	Description
Device	Required. Select the device to be monitored from the drop-down list.
Monitoring Object	Required. Select the monitoring object from the drop-down list. ■ **interface**: the switch interface, including the watermark, packet error, and packet loss of the interface. ■ **hardware**: the switch hardware, including the memory usage and CPU utilization. ■ **capacity**: others, which is not supported currently.
Monitoring Metric	Required. Select the corresponding monitoring metric from the drop-down list based on the selected monitoring object.
Monitoring Submetric	Optional. Select the corresponding monitoring submetric from the drop-down list based on the selected monitoring metric.

iv. Click **OK**.

After the subview is added, the system automatically shows the subview on the view to which the subview belongs.

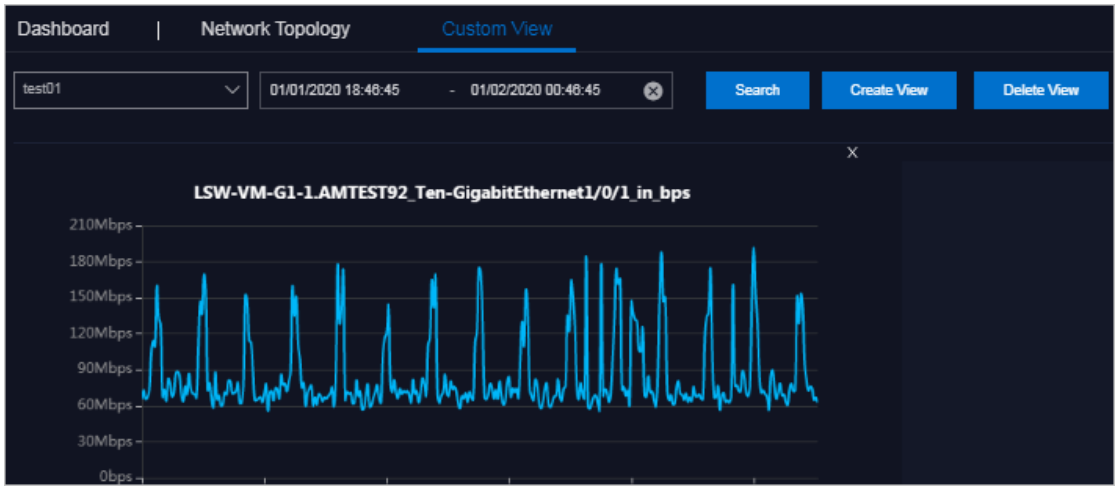

v. Add other subviews.

Delete a subview

1. In the left-side navigation pane, choose **NOC > Dashboard**.
2. Click the **Custom View** tab.
3. In the search box, select the view to which the subview to be deleted belongs and then click **Search**.
4. Click the X icon in the upper-right corner of the subview to be deleted.

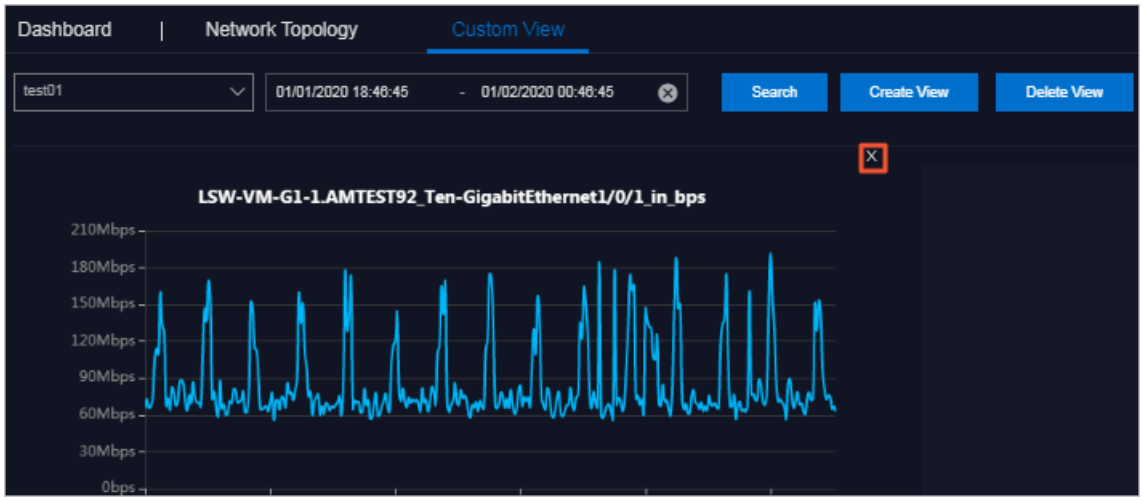

5. In the message that appears, click **OK**.

Delete a view

> **Notice** If you delete a view, its subviews are also deleted. Proceed with caution.

1. In the left-side navigation pane, choose **NOC > Dashboard**.
2. Click the **Custom View** tab.

3. In the search box, select the view to be deleted and then click **Search**.
4. Click **Delete View** in the upper part of the tab.
5. In the message that appears, click **OK**.

5.2.3. Network Service Provider

5.2.3.1. View access gateway instances

The Instance Management tab allows you to view information such as the access gateway name, IBGP role, and creation time of access gateway instances.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **Instance Management** tab.
3. Enter the region ID in the upper-left corner.

> **Note** To view the instances in other regions, click **Reset** in the upper-right corner and then enter the ID of another region.

4. Click **Display Device List** to view the access gateway device list in the current environment.

> **Note** If you add a device, click **Scan for New Devices** and then click **Display Device List**.

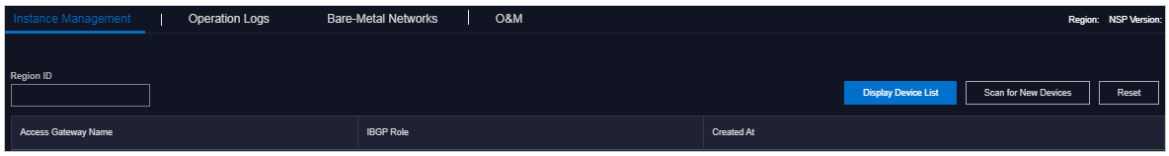

Column name	Description
Access Gateway Name	The access gateway name in the current system.
IBGP Role	The role of the access gateway in the environment. Note: ○ **RR-Active**: indicates that the role of the current gateway device is RR active device. ○ **Client**: indicates that the role of the current gateway is not RR active device.
Created At	Indicates the time when the current VSwitch began to act as an access gateway instance.

5.2.3.2. View operation logs

You can view API operation logs of bare metals based on O&M needs.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **Operation Logs** tab.
3. Configure filter conditions such as VSwitch ID, bare metal name, access gateway name, and time range, and then click Search to search for operation logs that meet the filter conditions.

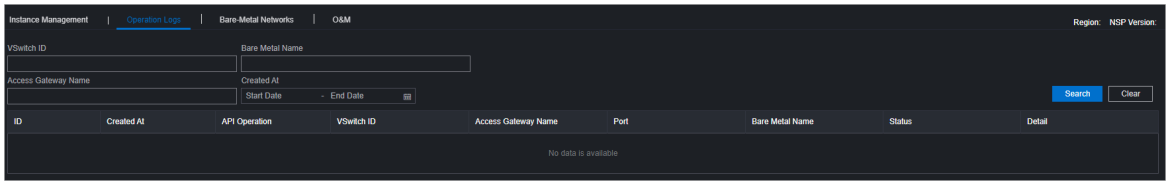

The following table describes some of the filter conditions.

Filter condition	Description
Vswitch ID	The ID of the VSwitch when the bare metal is applied for or released in the VPC.
Bare Metal Name	The name of the bare metal in the VPC that was applied for or released. To identify the bare metal as a unique one in the region, the serial number of the bare metal is used.
Access Gateway Name	The name of the access gateway to query.
Created At	The time range of the API operation to query.

> **Note** To modify the filter conditions, click **Clear** in the upper-right corner of the tab and configure the filter conditions again.

The following table describes the fields in the query result.

Column name	Description
ID	The index of the operation log.
Created At	The time when the operation was performed.
API Operation	The category of the API operation, such as applying for or releasing a bare metal in the VPC. Note: ○ **add** indicates that a bare metal is applied for in the VPC. ○ **del** indicates that a bare metal is released in the VPC. ○ **del_pc** indicates that a physical connection is deleted. ○ **del_vbr** indicates that a Virtual Border Router (VBR) is deleted. ○ **del_router_intf** indicates that a router interface is deleted. ○ **del_route_entry** indicates that a route table entry is deleted.

Column name	Description
Vswitch ID	The ID of the VSwitch when the bare metal is applied for or released in the VPC.
Access Gateway Name	The name of the access gateway involved with the current operation.
Port	The port to which the bare metal belongs.
Bare Metal Name	The name of the bare metal that is applied for or released in the VPC. To identify the bare metal as a unique one in the region, the serial number of the bare metal is displayed.
Status	The status of the API operation. **success** indicates that the operation was successful. If the API operation is in progress, the value indicates the real-time status of the API operation. If the API operation is complete but the value is not **success**, you can view the failure information in this column.

4. Find an operation log in the search results, and then click **View Details** in the **Detail** column to view the details of the API operation.

5.2.3.3. View network information of bare metals in the VPC

The **Bare-Metal Network** tab allows you to view the information of bare metals that are added to the VPC.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **Bare-Metal Networks** tab.By default, the network information of bare metals in the current system are displayed by page.
3. Configure the filter conditions such as bare metal name, VPC ID, VSwitch ID, VBR ID, BD ID, access gateway name, and time range, and then click Search to search for the information of bare metals that meet the filter conditions.

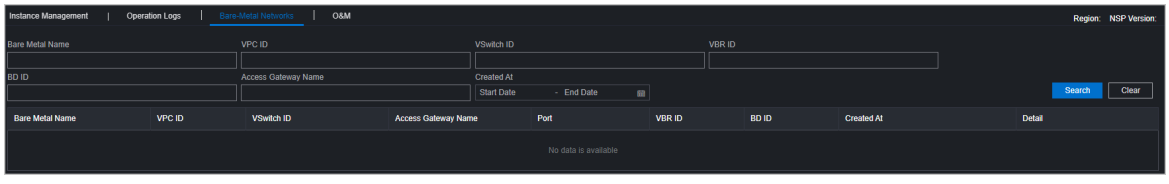

Filter condition	Description
Bare Metal Name	The name of the bare metal in the VPC that was applied for or released. To identify the bare metal as a unique one in the region, the serial number of the bare metal is used.

Filter condition	Description
VPC ID	The ID of the VPC to which the target bare metal belongs.
Vswitch ID	The ID of the VSwitch to which the target bare metal belongs.
VBR ID	The VBR ID of the physical connection created on HSW by the VPC to which the target bare metal belongs.
BD ID	The value of the hardware bridge-domain (BD) to which the target bare metal is added.
Access Gateway Name	The name of the access gateway to which the target bare metal belongs.
Created At	The time range when the current bare metal is allocated to the VPC.

> **Note** To modify the filter conditions, click **Clear** in the upper-right corner of the tab and configure the filter conditions again.

4. Find a bare metal in the search result, and then click **View Details** in the **Detail** column to view the details of the bare metal.

5.2.3.4. O&M configurations

5.2.3.4.1. Apply for a bare metal in the VPC

In O&M emergency scenarios, you can use this feature to add the physical port of the access gateway to which a bare metal belongs to the VPC.

Prerequisites

> **Notice** This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Before you use this feature, note that:

- Typically, you cannot use this feature to apply for a bare metal in the VPC. You can use the bare metal controller to call an API operation to activate the bare metal network.
- This feature can only be used to connect the bare metal to the access gateway port but cannot be used to perform operations on the bare metal. To configure the network port IP address and routing information of the bare metal, contact the corresponding product team for guidance.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Apply for Bare Metal in VPC** from the drop-down list in the upper-left corner of the tab.

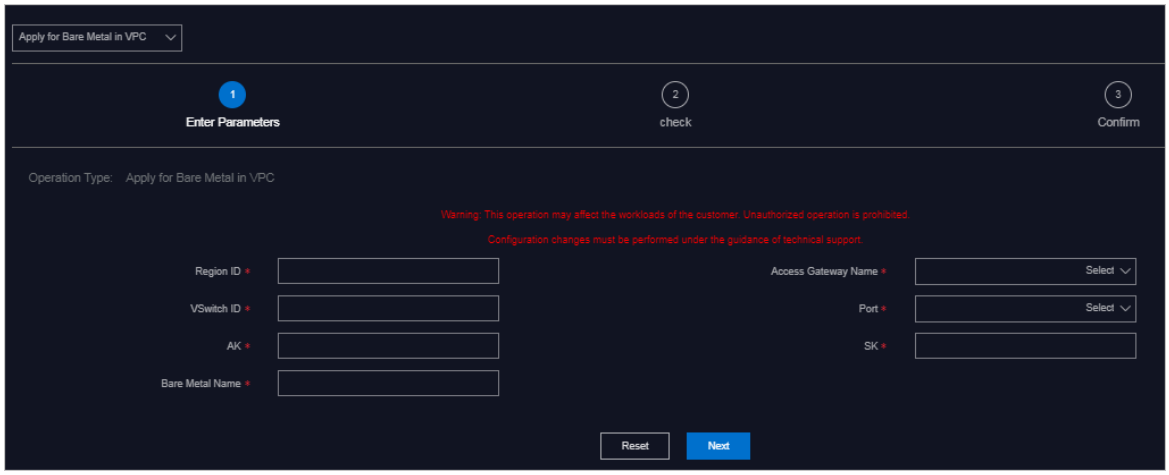

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Access Gateway Name	Select the name of the access gateway to which the bare metal is connected.
Vswitch ID	Enter the ID of the Vswitch to which the bare metal is to be added. You can obtain the Vswitch ID from the VPC console.
Port	Select the port of the access gateway to which the bare metal is connected.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VPC belongs.
Bare Metal Name	The name of the bare metal. Set this value to the serial number of the bare metal.

> **Note** If the specified values are not correct, click **Rest** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, you can search for the bare metal based on the bare metal name on the **Bare-Metal Networks** tab. If you can view the bare metal. it is added to the VPC .

5.2.3.4.2. Release a bare metal in the VPC

In O&M emergency scenarios, you can use this feature to disconnect the physical port of the access gateway to which a bare metal belongs from the VPC.

Prerequisites

> **Notice** This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Before you use this feature, note that:

- Typically, you cannot use this feature to release a bare metal in the VPC. You can use the bare metal controller to call an API operation to delete the bare metal network.
- This feature can only be used to connect the bare metal to the access gateway port but cannot be used to perform operations on the bare metal. To configure the network port IP address and routing information of the bare metal, contact the corresponding product team for guidance.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Release Bare Metal in VPC** from the drop-down list in the upper-left corner of the tab.

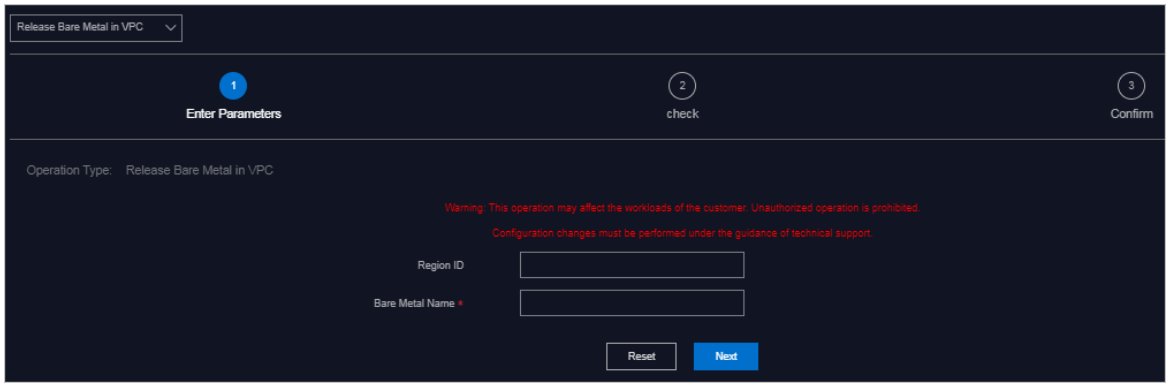

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Bare Metal Name	The name of the bare metal to be released. Set the value to the serial number of the bare metal.

> **Note** If the specified values are not correct, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, you can search for the bare metal based on the bare metal name on the **Bare-Metal Networks** tab. If the bare metal does not exist in the VPC, it is released.

5.2.3.4.3. Delete a VPC route table entry

In O&M emergency scenarios, you can use this feature to delete route table entries that point to the bare metal subnet in VPC.

Prerequisites

Notice This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Before you use this feature, note that:

- In most cases, you cannot use this feature to delete a VPC route table entry. This operation is only for emergency situations.
- You can perform this operation to delete only one route table entry at a time. To delete multiple route table entries, you must perform this operation multiple times.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Delete Route Table Entry** from the drop-down list in the upper-left corner of the tab.

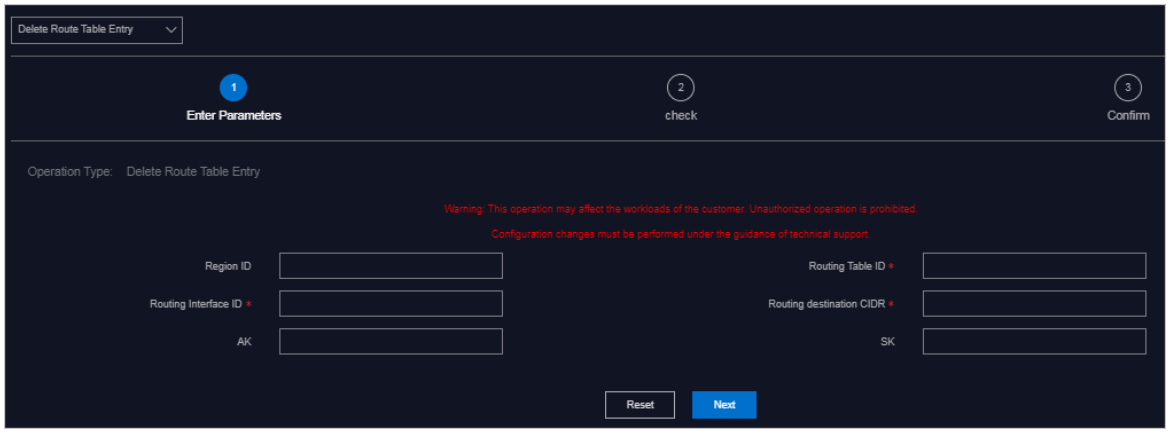

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Routing Table ID	The VPC route table ID, which can be obtained from the VPC console. For more information about how to obtain the route table ID, see the *VPC* User Guide.
Routing Interface ID	The VPC router interface ID, which can be obtained from the VPC console. For more information about how to obtain the router interface ID, see the *VPC* User Guide.

Parameter	Description
Routing destination CIDR	The destination CIDR block to which the VPC points, which can be obtained from the VPC console. For more information about how to obtain the routing destination CIDR block, see the *VPC* User Guide.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VPC belongs.

> **Note** If the specified values are not correct, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, you can log on to the VPC console and view the route table entry of the entered destination CIDR block. If the route table entry does not exist, it is deleted.
7. (Optional)In actual fault scenarios, if multiple route table entries exist in the VPC route table, repeat Step 3 to Step 6 to delete other route table entries.

5.2.3.4.4. Delete a VBR route table entry

In O&M emergency scenarios, you can use this feature to delete the default route table entry of a VBR.

Prerequisites

> **Notice** This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Delete Route Table Entry** from the drop-down list in the upper-left corner of the tab.
4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.

Parameter	Description
Routing Table ID	The VBR route table ID. If the bare metal involved with the VBR has already been added to the VPC, you can search for the bare metal on the **Bare-Metal Networks** tab based on the bare metal name, and then click View Details. The VBR Route Table ID in the details is the value of this parameter. If the bare metal involved with the VBR is not added to the VPC, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add** on the **Operation Logs** tab. Click **View Details** and the VBR Route Table ID in the details is the value of this parameter.
Routing Interface ID	The VBR router interface ID. If the bare metal involved with the VBR has already been added to the VPC, you can search for the bare metal on the **Bare-Metal Networks** tab based on the bare metal name, and then click View Details. The VBR RI in the details is the value of this parameter. If the bare metal involved with the VBR is not added to VPC, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add** on the **Operation Logs** tab. Click View Details and the VBR RI in the details is the value of this parameter.
Routing destination CIDR	Set the value to 0.0.0.0/0.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VBR belongs.

> **Note** If the specified values are not correct, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, you can click **Apsara Network Intelligence** of the **Product** module. On the homepage of Apsara Network Intelligence, enter the VBR route table ID and search for the VBR route table. If the route table entry 0.0.0.0/0 does not exist in the VBR route table, the route table entry is deleted.

5.2.3.4.5. Delete a VPC router interface

In O&M emergency scenarios, you can use this feature to delete a VPC router interface.

Prerequisites

.

Notice This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Delete Router Interface** from the drop-down list in the upper-left corner of the tab.

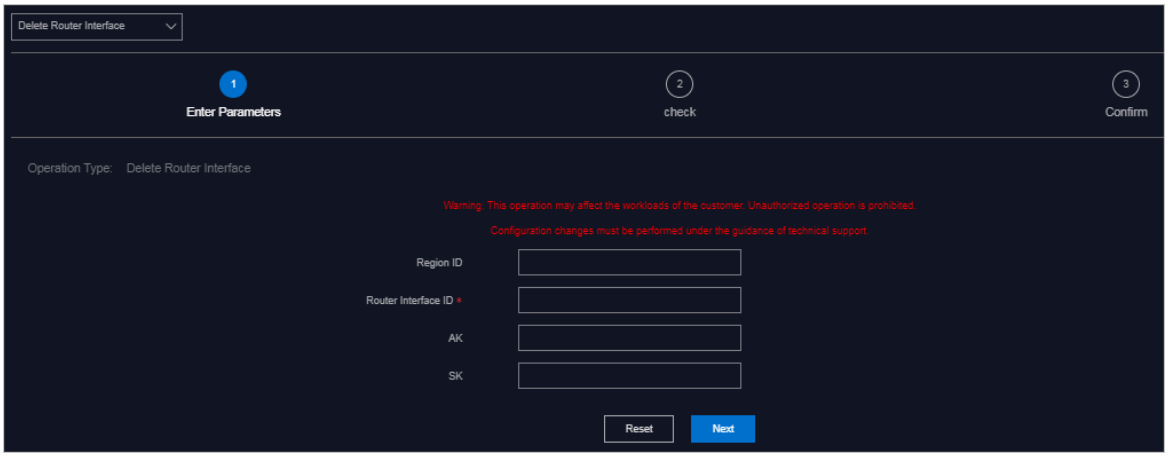

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Router Interface ID	The VPC router interface ID. On the **Operation Logs** tab, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add**. Click **View Details** and the VPC RI in the details is the value of this parameter.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VPC belongs.

Note If the specified values are incorrect, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, go to the **Products** module, and then click **Apsara Network Intelligence**. On the homepage of Apsara Network Intelligence, enter the VPC router interface ID to search for the router interface. If no search result appears, the router interface is deleted.

5.2.3.4.6. Delete a VBR router interface

In O&M emergency scenarios, you can use this feature to delete a VBR router interface.

Prerequisites

Notice This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > VPC**.
2. Click the **O&M** tab.
3. Select **Delete Router Interface** from the drop-down list in the upper-left corner of the tab.
4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Router Interface ID	The VBR router interface ID. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the **Bare-Metal Networks** tab based on the bare metal name, and then click View Details. The VBR RI in the details is the VBR router interface ID. If the bare metal involved with the VBR is not added to VPC, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add** on the **Operation Logs** tab. Click View Details and the VBR RI in the details is the value of this parameter.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VPC belongs.

Note If the specified values are incorrect, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, go to the **Products** module, and then click **Apsara Network Intelligence**. On the homepage of Apsara Network Intelligence, enter the VBR router interface ID to search for the router interface. If no search result appears, the router interface is deleted.

5.2.3.4.7. Delete a VBR

In O&M emergency scenarios, you can use this feature to delete a VBR.

Prerequisites

Notice This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > VPC**.
2. Click the **O&M** tab.
3. Select **Delete VBR** from the drop-down list in the upper-left corner of the tab.

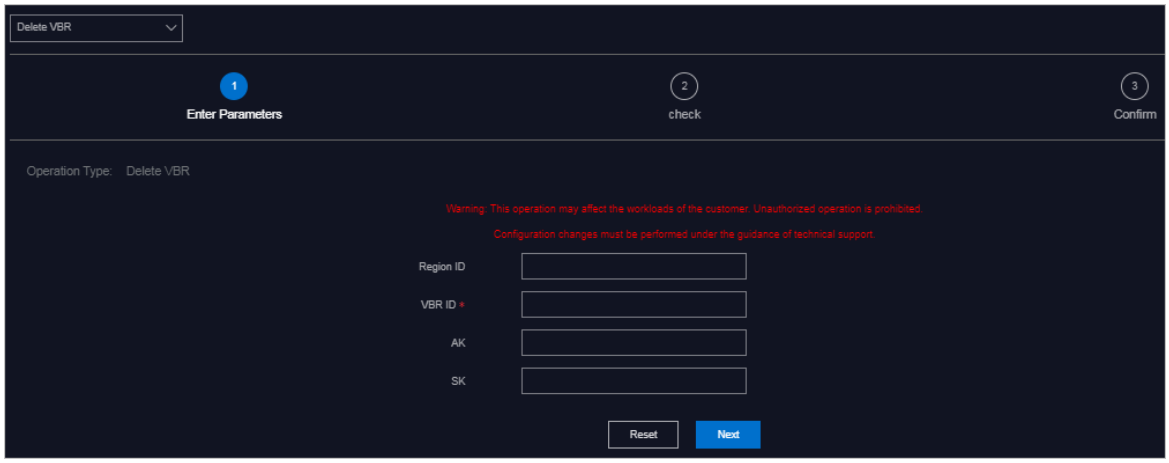

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
VBR ID	The ID of the VBR to be deleted. On the **Operation Logs** tab, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add**. Click View Details and the VBR ID in the details is the value of this parameter.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VBR belongs.

Note If the specified values are not correct, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.

6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message `Result: Successful` appears.

   After the configurations are pushed, go to the **Products** module, and then click **Apsara Network Intelligence**. On the homepage of Apsara Network Intelligence, enter the VBR ID to search for the VBR. If no search result appears, the VBR is deleted.

5.2.3.4.8. Delete a physical connection

In O&M emergency scenarios, you can use this feature to delete a physical connection.

Prerequisites

Notice This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > VPC**.
2. Click the **O&M** tab.
3. Select **Delete Express Connect Circuit** from the drop-down list in the upper-left corner of the tab.

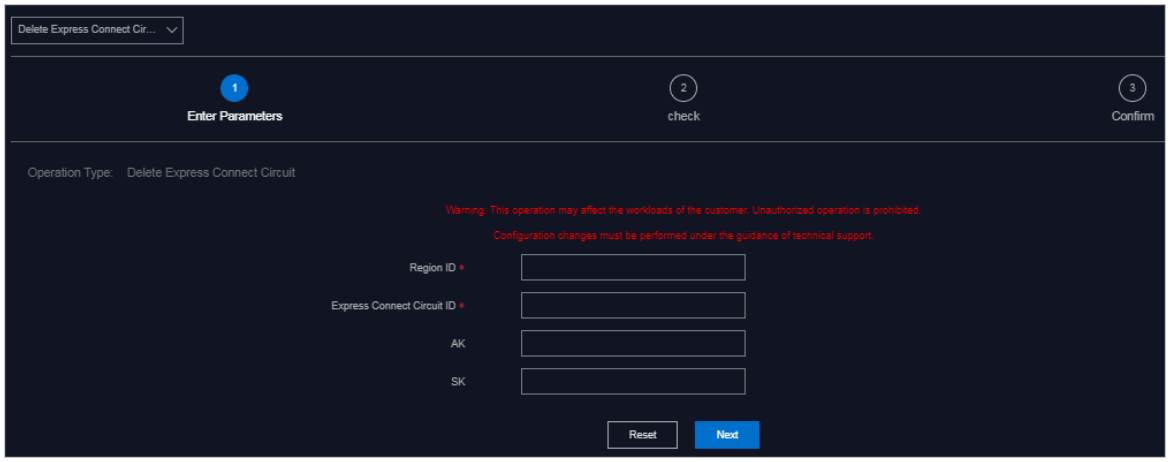

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Express Connect Circuit ID	The ID of the physical connection to be deleted. On the **Operation Logs** tab, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add**. Click **View Details** and the Express Connect Circuit ID in the details is the value of this parameter.

Parameter	Description
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VBR belongs.

> **Note** If the specified values are incorrect, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, go to the **Products** module, and then click **Apsara Network Intelligence**. On the homepage of Apsara Network Intelligence, enter the physical connection ID to search for the physical connection. No search result appears, indicating the physical connection is deleted.

5.2.3.4.9. Delete all resources with one click

In O&M emergency scenarios, you can use this feature to delete all resources, including the Virtual Private Cloud (VPC) route table entries, Virtual Border Router (VBR) route table entries, VPC router interfaces, VBR router interfaces, VBRs, and physical connections.

Prerequisites

> **Notice** This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > VPC**.
2. Click the **O&M** tab.
3. Select **Delete All Resources** from the drop-down list in the upper-left corner of the tab.

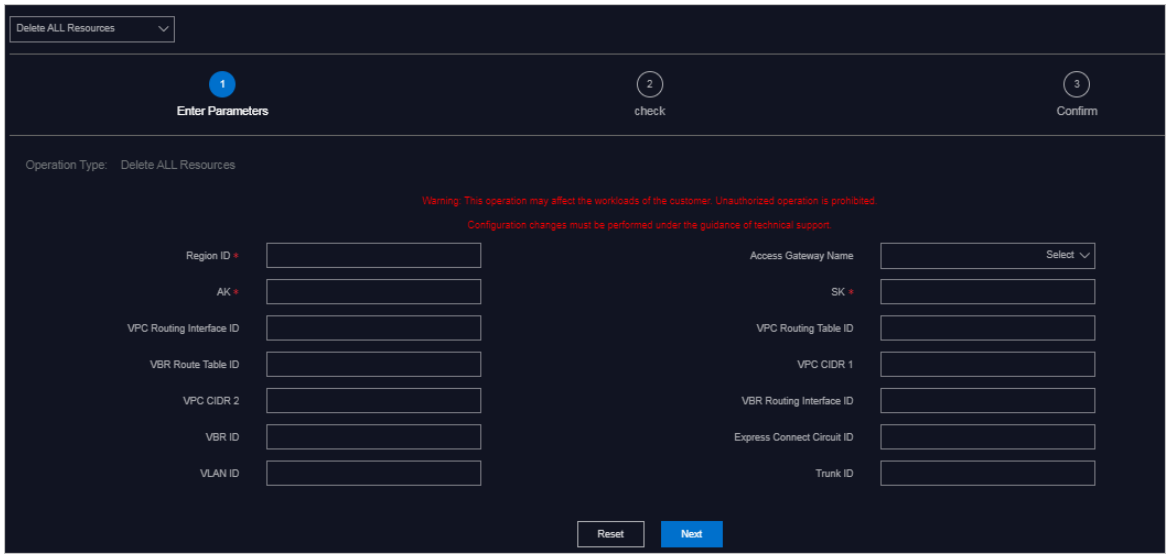

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Access Gateway Name	Select the name of the access gateway to which the bare metal is connected.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VBR belongs.
VPC Routing Interface ID	The router interface ID of VPC, which can be obtained from the VPC console. For more information about how to obtain the router interface ID, see User Guide in the *VPC* documentation.
VPC Routing Table ID	The route table ID of VPC, which can be obtained from the VPC console. For more information about how to obtain the route table ID, see User Guide in the *VPC* documentation.

Parameter	Description
VBR Route Table ID	The route table ID of VBR. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the **Bare-Metal Networks** tab based on the bare metal name, and then click **View Details**. The VBR Route Table ID in the details is the value of this parameter. If the bare metal involved with the VBR is not added to VPC, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add** on the **Operation Logs** tab. Click **View Details** and the VBR Route Table ID in the details is the value of this parameter.
vpc cidr1	The destination CIDR block 1 to which the VPC points, which can be obtained from the VPC console. For more information bout how to obtain the VPC CIDR block 1, see User Guide in the *VPC* documentation.
vpc cidr2	The destination CIDR block 2 to which the VPC points, which can be obtained from the VPC console. For more information bout how to obtain the VPC CIDR block 2, see User Guide in the *VPC* documentation.
VBR Routing Interface ID	The router interface ID of VBR. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the **Bare-Metal Networks** tab based on the bare metal name, and then click **View Details**. The VBR RI in the details is the value of this parameter. If the bare metal involved with the VBR is not added to VPC, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add** on the **Operation Logs** tab. Click **View Details** and the VBR ID in the details is the value of this parameter.

Parameter	Description
VBR ID	The ID of the VBR to be deleted. On the **Operation Logs** tab, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add**. Click **View Details** and the VBR ID in the details is the value of this parameter.
Express Connect Circuit ID	The ID of the physical connection to be deleted. On the **Operation Logs** tab, specify the bare metal name and creation time to search for operation logs and find an operation log whose **API Operation** is **add**. Click **View Details** and the Express Connect Circuit ID in the details is the value of this parameter.
Vlan ID	Set the value to 10.
Trunk ID	You do not need to specify this parameter.

> **Note** If the specified values are not correct, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**.The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, use the methods provided in Delete a VPC route table entry, Delete a VBR route table entry, Delete a VPC router interface, Delete a VBR router interface, Delete a VBR, and Delete a physical connection to check whether the VPC route table entries, VBR route table entries, VPC router interfaces, VBR router interfaces, VBRs, and physical connections are deleted.

5.2.3.4.10. View physical connection bandwidth

You can view the physical connection bandwidth when the access gateway instance is connected to a VPC in the system based on O&M needs.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **View Express Connect Bandwidth** from the drop-down list in the upper-left corner.

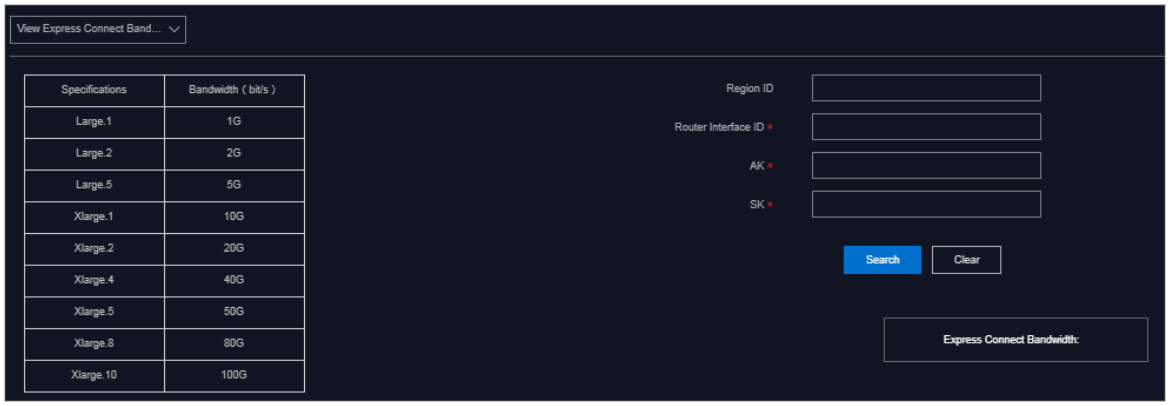

4. Configure the filter conditions and then click **Search**.

Parameter	Description
Region ID	The name of the region in the current environment.
Router Interface ID	The VBR router interface ID. On the **Bare-Metal Networks** tab, specify the VPC ID and access gateway name, and then click View Details. The VBR RI in the details is the value of this parameter.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VBR belongs.

The system shows physical connection bandwidth information that meets the filter conditions.

The obtained bandwidth information describes the specifications of the physical connection bandwidth on HSW of the current VPC. View the table on the left and obtain the bandwidth (bit/s) based on the specification.

5.2.3.4.11. Modify the physical connection bandwidth

In O&M emergency scenarios, you can use this feature to modify the physical connection bandwidth.

Prerequisites

Notice This operation is for emergency situations only and must be performed under the guidance of technical personnel. Otherwise, normal business operations will be affected.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Change Express Connect** from the drop-down list in the upper-left corner.

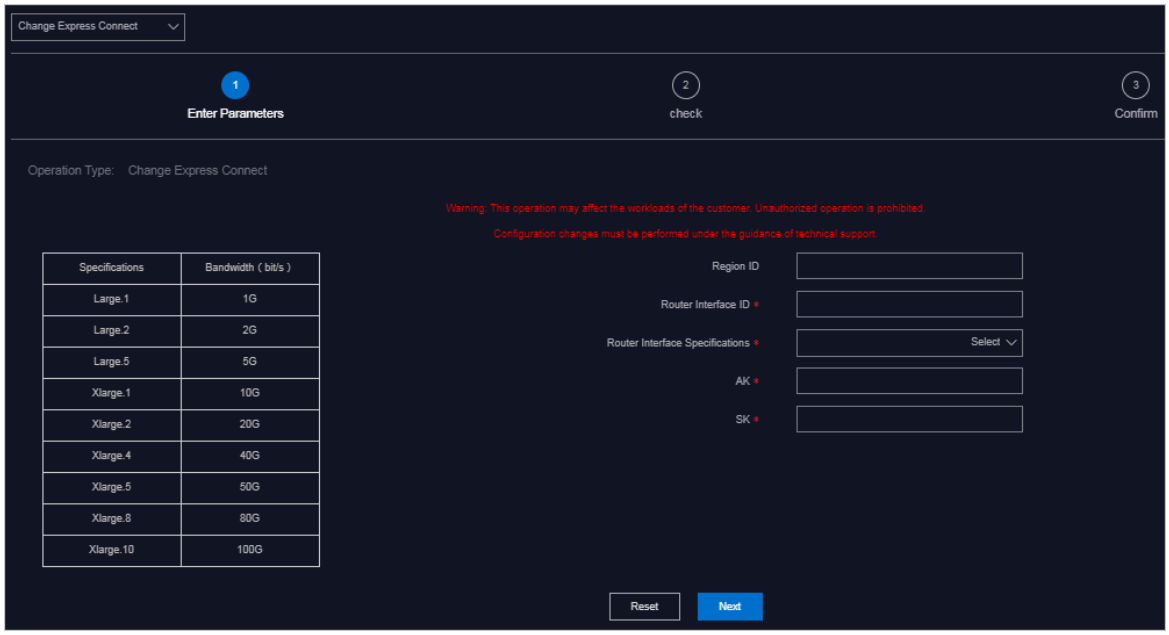

4. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region.
Router Interface ID	The ID of the router interface to which the physical connection bandwidth to be modified corresponds. On the **Bare-Metal Networks** tab, specify the VPC ID and access gateway name to search for the bare metal. Click **View Details** and the VBR RI in the details is the value of this parameter.
Router Interface Specifications	Select a new specification of the physical connection bandwidth.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Stack Cloud Management (ASCM) console based on the organization to which the VPC belongs.

> **Note** If the specified values are incorrect, click **Reset** in the lower part of the tab and configure the parameters again.

5. Click **Next**.
6. Check the information. If the information is correct, click **Confirm**. The system will begin to push the configurations. After the configurations have been pushed, the message **Result: Successful** appears.

   After the configurations are pushed, check whether the bandwidth connection has been modified. Form more information, see View physical connection bandwidth.

5.2.3.4.12. View the BD usage

You can view the BD usage to learn about BD configuration distribution in a timely manner.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **Show BD usage** from the drop-down list in the upper-left corner.
4. Configure filter conditions and then click **Search**.

5.2.3.4.13. View the trunk usage

You can view the trunk usage to learn about the usage of hardware ports in a timely manner.

Procedure

1. In the left-side navigation pane, choose **NOC > NSP**.
2. Click the **O&M** tab.
3. Select **View Trunk Usage** from the drop-down list in the upper-left corner.

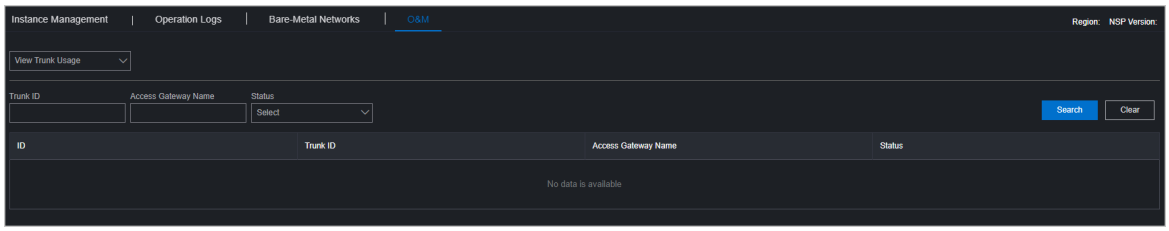

4. Specify the trunk ID and access gateway name, select the trunk status, and then click **Search**. Trunk status includes Idle, Used, Creating, and Deleting.

   The value of **Trunk ID** can be set to the last integer of the **Port** value that is obtained from the **Bare-Metal Networks** tab. For example, if the port number is 10GE1/0/40, set **Trunk ID** to 40.

   > **Note** To modify search conditions, you can click **Clear** in the upper-right corner and configure the parameters again.

5.2.4. Resource management

The **Resource Management** module is used to manage network-related resources, including the information of physical network element devices, virtual network products, and IP addresses.

5.2.4.1. Network elements

Network elements are network devices, including switches and routers. The **Network Elements** module displays the basic information and running status of physical network devices, and allows you to configure and manage physical network devices, including device management, password management, and configuration comparison.

5.2.4.1.1. Device management

The **Device Management** tab displays the basic information, running status, traffic monitoring, and logs of physical network element devices, and allows you to configure the collection settings of network devices.

5.2.4.1.1.1. View the network monitoring information

The Network Monitoring tab allows you to view the basic information, running status, and traffic monitoring of Apsara Stack physical network devices and check the health status of network devices in a timely manner.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. On the **Device Management** tab, click the **Network Monitoring** tab.
3. In the upper part of the tab, select an IDC and perform the following operations:
   - View the basic information, ping status, and SNMP status of Apsara Stack physical network devices.

     > **Note** You can also click **Export to CSV** to export network device information to your local computer.

     If a device has a business connectivity or gateway connectivity problem, the value in the Ping Status or SNMP Status column turns from green to red. The operations personnel are required to troubleshoot the problem.
   - In the search box in the upper-right corner, enter the device name or IP address to search for the monitoring information of a specific device.
   - View the port information, CPU utilization, memory usage, aggregation port information, and alert information of a device.
     a. Click a device name, or click **View** in the **Details** column corresponding to a device.
     b. On the **Port** tab, view the ports, port operation status, and other link information of the device.
     c. On the **CPU Utilization** tab, view all the CPU utilization information of the device.
     d. On the **Memory Usage** tab, view all the memory usage information of the device.
     e. On the **Aggregation Port Management** tab, view all the aggregation port information of the device. You can click **View** in the **Operation** column corresponding to a port to view the usage of the aggregation port.

f. On the **Alert Info** tab, view the alert information of the device.

During routine O&M, take note of the alert list of the device. Typically, if no data is displayed on the **Alert Info** tab, the device is operating normally.

If alert events occur, unrecovered alert events are displayed in the list. You must handle these exceptions in a timely manner. When exceptions are handled, their corresponding alerts are automatically cleared from the list.

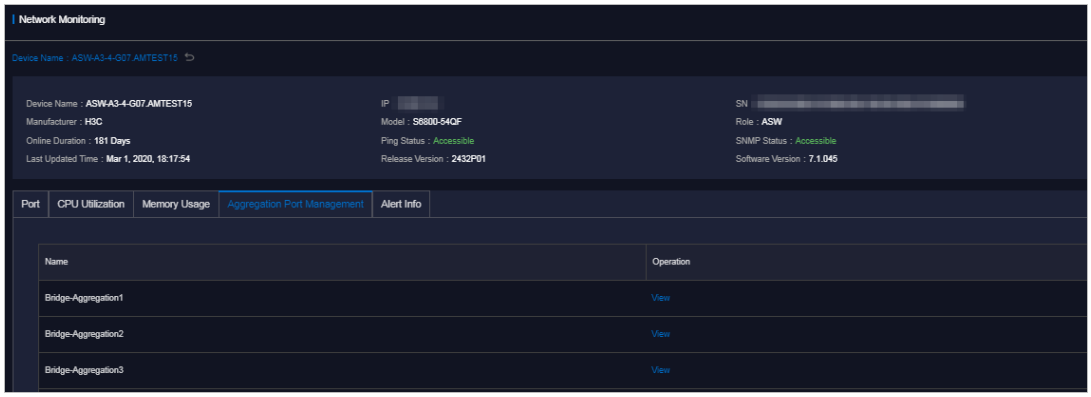

- View the traffic information of a device for a specified port and time range.
  a. Click a device name, or click **View** in the **Details** column corresponding to a device.
  b. Search for the port that you are about to view by using the search box in the upper-right corner of the **Port** tab. Click **View** in the **Details** column corresponding to the port.
  c. Select a time range on the right, and then click **Search** to view the traffic in the selected time range.

  You can select 5MIN, 30MIN, 1H, 6H in the **Quick Query** section to view the traffic within the last 5 minutes, 30 minutes, 1 hour, or 6 hours.

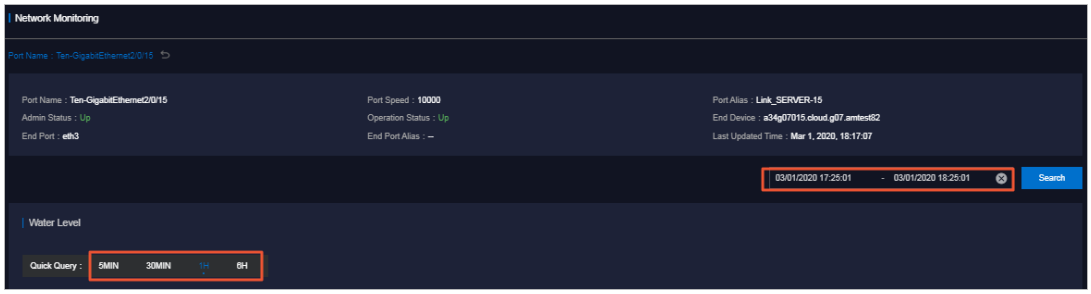

5.2.4.1.1.2. View logs

The Syslogs tab allows you to view logs of physical network element devices, providing necessary data for fault location and diagnosis information collection if a fault occurs.

Context

During the daily inspection, you can search for logs generated by a specific network device during a specific time range on the **Syslogs** tab.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. On the **Device Management** tab, click the **Syslogs** tab.

3. In the upper-right corner of the tab, select a device name from the drop-down list, select a time range, and then click **Search** to check whether the device has generated system logs in the specified time range.

   If the device has a configuration exception or does not have any generated logs for the specified time range, no search results will be returned.

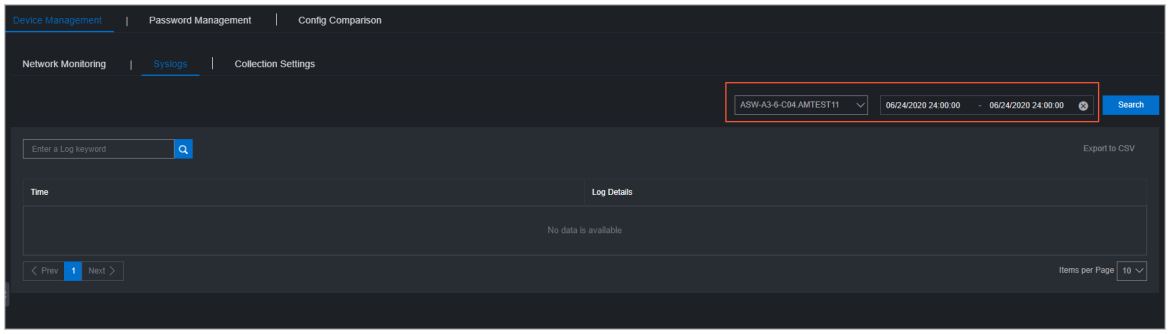

4. (Optional)You can filter the search results based on the log keyword.
5. (Optional)Click **Export to CSV** in the upper-right corner to export the search results to your local computer.

5.2.4.1.1.3. Collection settings

The **Collection Settings** tab allows you to configure the collection interval of physical network element devices and manage OOB network segments.

5.2.4.1.2. Modify the device password

You can modify the passwords of physical network devices.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. Click the **Password Management** tab.
3. (Optional)Enter the name of the device whose password is to be modified in the search box of the **Devices on Live Network** section and then click **Search**. To search for another device, you can click **Reset** to reset the configured search condition.
4. Select one or more devices and then click **Add**. The selected devices are displayed in the **Target Devices** section on the right.

   > **Note** To remove a device from the **Target Devices** section, choose **Manage > Delete** in the **Actions** column corresponding to the device. You can also click **Clear** in the upper-right corner to remove all the devices from the Target Devices section.

5. The system must verify the old password before you modify it. Specify the **Username** and **Old Password** in the lower-right corner and then click **Verify**. You must verify the old password for all the devices in the **Target Devices** section.
6. After the verification is passed, modify the password for one or more devices.
   - Modify the password of a device

Choose **Manage > Set Username and Password** in the **Actions** column corresponding to a device. In the dialog box that appears, enter the username and password, and then click **OK**.

- Modify the passwords of all devices

  Click **Modify** below the **Target Devices** section to modify the passwords of all the devices added to the **Target Devices** section.

5.2.4.1.3. Compare device configurations

For a device, you can compare its current configuration with its configuration at startup and check whether they are consistent.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. Click the **Config Comparison** tab.
3. (Optional)Enter the name of the device whose configurations you are about to compare in the **Device Name** search box and then click **Search**. To search for another device, you can click **Reset** to reset the configured search condition.
4. Select the device and then click **Compare Configuration**. After the comparison, click **Refresh** and then click **Export Results** to export the results.

5.2.4.2. Server Load Balancers

Server Load Balancers displays the basic information, running status, and water level of network product Server Load Balancer by using cluster monitoring and instance monitoring.

5.2.4.2.1. View the cluster monitoring information

The **Cluster Monitoring** tab allows you to view the basic information, inbound limit (bit/s), outbound limit (bit/s), inbound limit (PPS), outbound limit (PPS), active connection limit, inactive connection limit, and usage of a single device node in a cluster.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Server Load Balancers**.
2. The **Cluster Monitoring** tab appears.
3. Select the cluster that you want to view from the drop-down list and then click **Search**.The information of all the device nodes in the cluster is displayed.

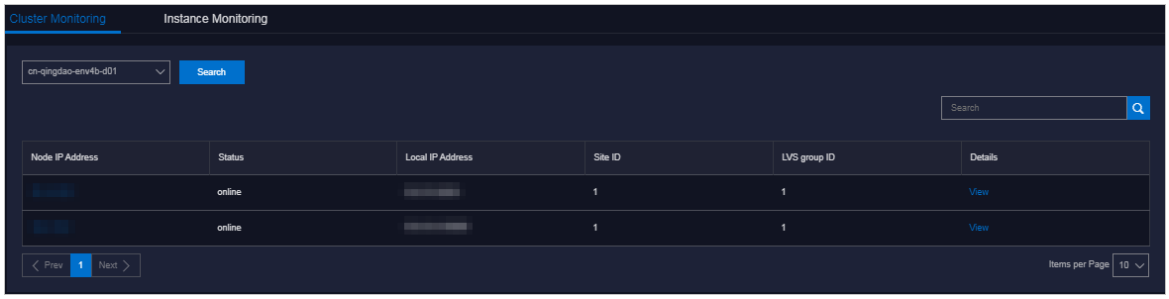

4. Find a device node and click **View** in the **Details** column.

5. On the **Node Message** page, view the basic information, inbound limit (bit/s), outbound limit (bit/s), inbound limit (PPS), outbound limit (PPS), active connection limit, and inactive connection limit of the device node.

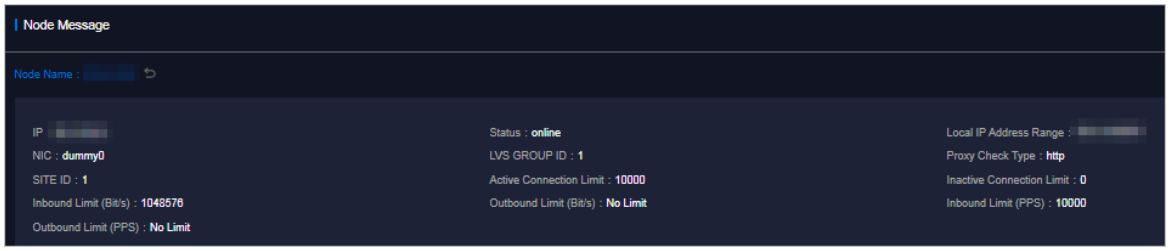

5.2.4.2.2. View the instance monitoring information

The Instance Monitoring tab allows you to view the basic information and usage of an instance, including the BPS and PPS.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Server Load Balancers**.
2. Click the **Instance Monitoring** tab.
3. Select the cluster where the instance that you want to view is located from the Cluster drop-down list. Enter the lb-id or VIP address that you want to search for in the field and then click **Search**.
4. In the search result, view the monitoring information of the instance.

   The following information is displayed:

   - The first section shows the basic information of the SLB instance, which allows operations engineers to troubleshoot problems and confirm the owner of a device.
   - The second section shows the operating graph of the instance. Select a time range and then click **Search** or select 5 MIM, 30 MIN, 1H, or 6H in the Quick Query section to view the operating graph of the instance in a specific time range, including the detailed BPS and PPS.

5.2.4.3. Collect IP addresses

The system regularly collects the IP addresses of all physical networks within the current Apsara Stack environment based on a configured collection interval. You can search for the information of devices and ports to which a CIDR block or an IP address belongs based on the CIDR block or IP address, and subnet mask.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > IP Address Collection**.
2. Enter the CIDR block or IP address, and subnet mask in the corresponding search boxes, and then click **Search**.If the CIDR block you are searching for belongs to a CIDR block in the current Apsara Stack environment, the system shows the information of devices and ports to which the specified CIDR block belongs.

   > **Note** If you enter an IP address in the search box and then click Search, the system calculates the corresponding CIDR block based on the IP address and subnet mask.

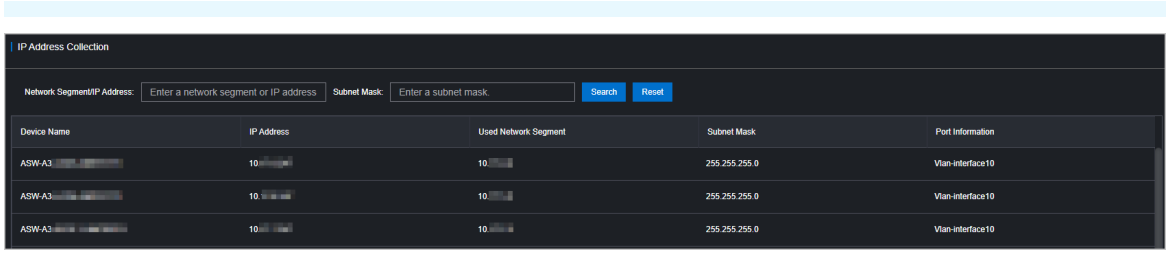

5.2.4.4. IP address ranges

The **IP Address Ranges** module is used to manage the planning information in the Apsara Stack environment, including the network architecture and IP address planning. You can modify, import, and export the planning information.

5.2.4.4.1. Import the planning file

No data is imported when the system is initialized. You must import the planning file to obtain the IP address allocation information of the current Apsara Stack environment. You can also import a new planning file for a change in the environment.

Prerequisites

The IP address allocation list is obtained from the Apsara Stack deployment planner.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > IP Address Ranges**.
2. Click **Import** in the upper-right corner.
3. In the dialog box that appears, click **Browse** and then select the IP address allocation list.
4. Click **Import**.

5.2.4.4.2. Manually add the IP address pool information

You can manually add new IP address pool information to the ASO console for centralized management.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > IP Address Ranges**.
2. Click **Add**.
3. In the dialog box that appears, enter the IP address pool information.
4. Click **Add**.

5.2.4.4.3. Modify the IP address pool information

If an IP address range is changed, you can modify the IP address pool information.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > IP Address Ranges**.

2. (Optional)On the **IP Address Ranges** page, configure the search conditions and then click **Search**.

   > **Note** To reset the search conditions, you can click **Reset** to clear your configurations with one click.

3. Find the IP address pool whose information you are about to modify and then choose **Manage > Modify** in the **Actions** column.
4. In the dialog box that appears, modify the network architecture and IP address planning.
5. Click **Edit**.

5.2.4.4.4. Export the IP address pool information

You can export the IP address pool information to your local computer and then view the information offline.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > IP Address Ranges**.
2. Select the IP address pool that you want to export and then click **Export**.

5.2.4.4.5. Delete the IP address pool information

You can delete IP address pool information that is no longer needed.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > IP Address Ranges**.
2. Find the IP address pool whose information you are about to delete and then choose **Manage > Delete** in the **Actions** column.

5.2.4.5. View Anytunnel information

You can view the Anytunnel information to know the Anytunnel resources registered by projects within the current environment or whether a project registers Anytunnel. The system allows you to search for the registration information of Anytunnel resources based on project, cluster, service instance, and server role. You can use the global query feature to search for the usage of all the Anytunnel resources in the current environment.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Anytunnel Management**.
2. Perform the following operations:
   - In the upper part of the page, click **Query Details** to view all the Anytunnel information in the environment.

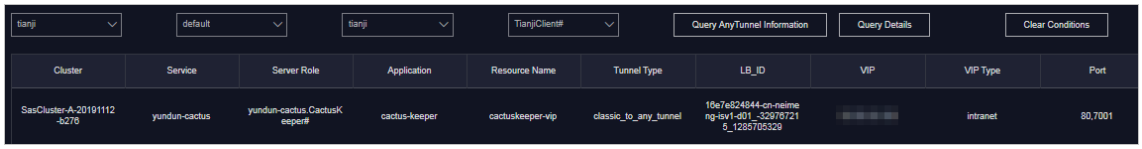

- In the upper part of the page, select the project, cluster, service instance, or service role, and then click **Query AnyTunnel Information** to view the AnyTunnel information that meets the search conditions.

  > **Note** To modify the search conditions, click **Clear Conditions** and configure the search conditions again.

5.2.4.6. XGW management

The XGW Management module allows you to view the basic information, running status, aggregated traffic, and usage of each device node of XGW network products.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > XGW Management**.
2. Select the cluster that you want to view from the drop-down list and then click **Search**. The system shows the basic information and usage information of the aggregated traffic of all device nodes in the selected cluster. By default, the usage information of the last one hour is displayed. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range to search for the usage information.

3. Find a device node and click **View** in the **Details** column.
4. On the page that appears, view the traffic usage information of the device node.

5.2.4.7. Fire wall

If the cloud firewall is deployed in your environment, you can use the Fire Wall feature to isolate or restore the firewall.

Prerequisites

Notice Confirm with the administrator that the cloud firewall is deployed in your environment. Otherwise, you cannot use the Fire Wall feature to isolate or restore the firewall.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Fire Wall**.
2. Select the operation type, firewall type, and data center from the drop-down lists and then click **Confirm**.

   You can select one of the following operation types:

   - **Isolate Firewall**: physically isolates the firewall from the network structure. If the cloud firewall service has an exception, the system removes the firewall device from the network forwarding path, making sure that the normal business traffic forwarding is not affected by faults.
   - **Restore Firewall**: restores the firewall from the network isolated status to the normal status. After the exception of the cloud firewall is recovered, the system adds the firewall device back to the network forwarding path, making sure that the firewall is restored to the initial online status.

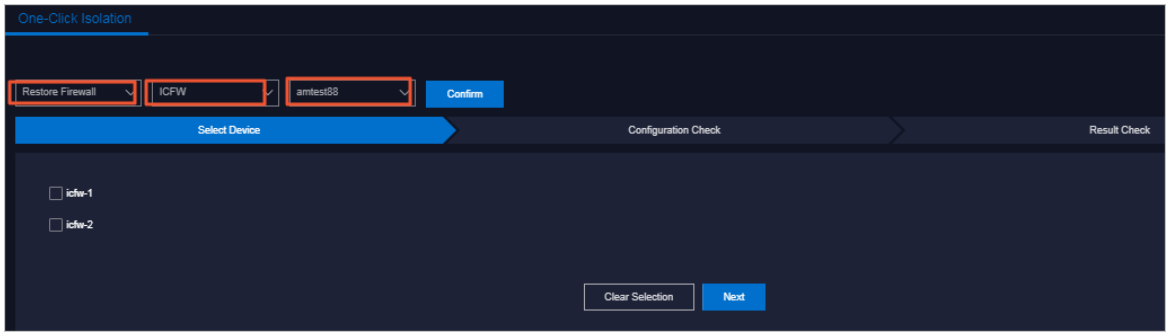

3. In the **Select Device** step, select devices and then click **Next**.

4. In the **Configuration Check** step, check the selected devices and template information. If the information is correct, click **Confirm**.

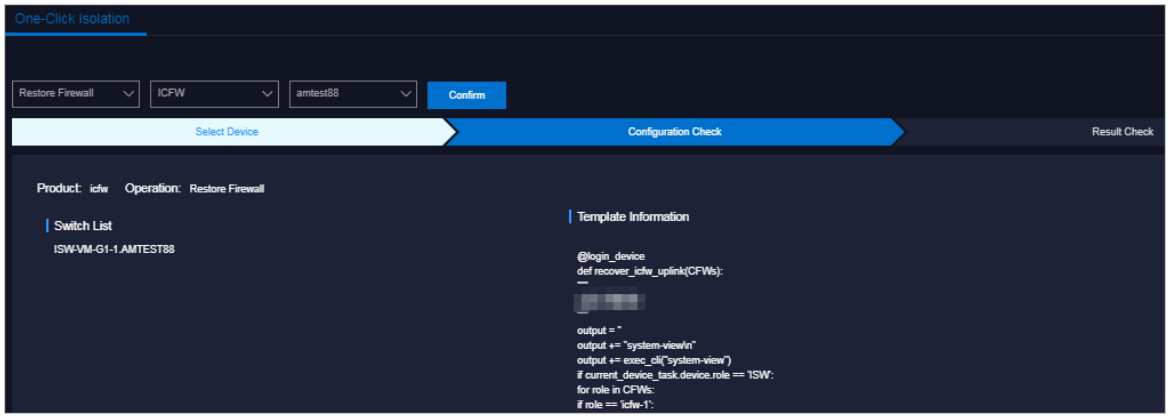

5. In the message that appears, click **OK**.

   Then, the system automatically isolates or restores the firewall in the selected devices based on the configuration template.

   The results are automatically displayed in the **Result Check** step.

6. In the **Result Check** step, click **Details** in the **Details** column corresponding to each device to view the corresponding result.

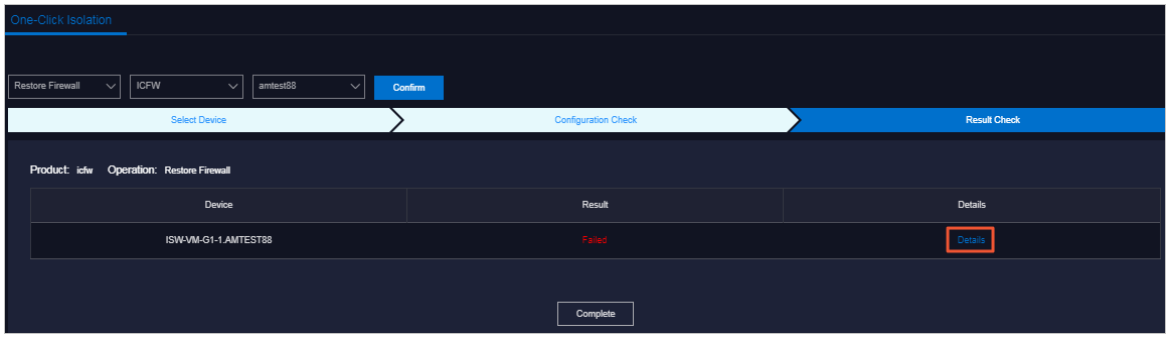

7. Click **Complete**.

5.2.5. Alert management

The **Alert Management** module provides you with the real-time alert dashboard, history alert dashboard, and the alert settings function.

5.2.5.1. View and process current alerts

You can view and process current alerts on the Current Alerts tab.

Procedure

1. In the left-side navigation pane, choose **NOC > Alert Management > Alert Dashboard**.
2. Click the **Current Alerts** tab.
3. Enter a keyword in the search box in the upper-right corner, and then click **Search**.Alerts that meet the search conditions are displayed.
4. (Optional)You can filter the search results by device name, device IP address, or alert name.

5. Click **Details** in the **Details** column corresponding to an alert to view detailed information about the alert.
6. Find the reason why the alert is triggered and then process the alert.
   - If the alert does not affect the normal operation of the system, you can click **Ignore** in the **Actions** column to ignore the alert.
   - If the alert is no longer significant, you can click **Delete** in the **Actions** column to delete the alert.

     After the alert is processed, you can search for it on the **History Alerts** tab.
7. (Optional)Click **Export to SCV** to export the alert information to your local computer.

5.2.5.2. View historical alerts

You can view historical alerts on the History Alerts tab.

Procedure

1. In the left-side navigation pane, choose **NOC > Alert Management > Alert Dashboard**.
2. Click the **History Alerts** tab.
3. Select Alert Source, Alerting IP Address, Alerting Device, Alert Name, Alert Item, or Alerting Instance from the drop-down list, and then enter a keyword in the field. Select a time range, and then click **Search**.Alerts that meet the search conditions are displayed.
4. Click **Details** in the **Details** column corresponding to an alert to view detailed information about the alert.
5. (Optional)Click **Export to SCV** to export the alert information to your local computer.

5.2.5.3. Add a trap

If the initially configured trap subscription does not meet the monitoring requirements, you can add a trap for monitoring match.

Context

The trap in this topic is the Simple Network Management Protocol (SNMP) trap. SNMP trap is a part of SNMP and a mechanism that devices being managed (here refers to network devices such as switches and routers) send SNMP messages to the NOC monitoring server. If an exception occurs on the side being monitored, namely the switch monitoring metrics have an exception, the SNMP agent running in a switch sends an alert event to the NOC monitoring server.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Alert Configuration**.
2. On the **Alert Settings** page, click **Configure Trap**.
3. In the **Configure Trap** dialog box, configure the parameters.

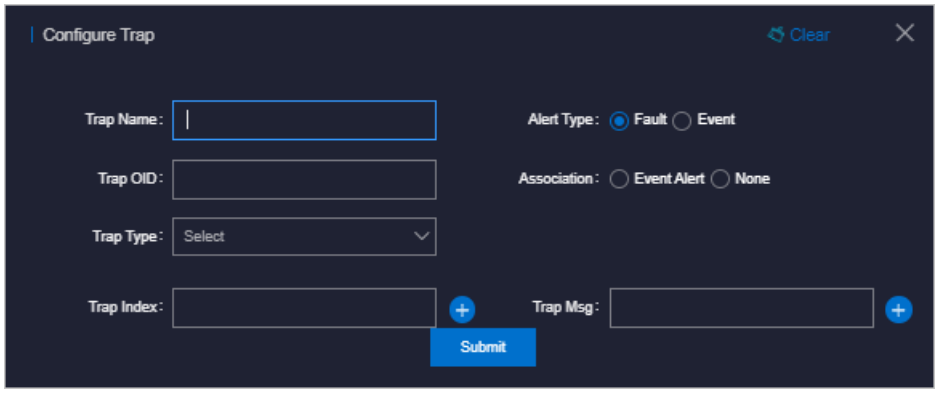

The following table describes the parameters.

Parameter	Description	Example
Trap Name	The name of the alert event.	linkdown or BGPneighbor down. You can customize the value.
Trap OID	The OID of the alert event.	.1.3.6.1.4.1.25506.8.35.12.1.12 Configure the value based on the device document. You cannot customize the value.
Trap Type	The type of the alert event.	None
Trap Index	The index ID of the alert item. This value is the KV information in the trap message, which is used to identify the alert object. Typically, this value can be an API name, protocol ID, or index ID. Configure the value based on the device document. You cannot customize the value.	None
Trap Msg	The message of the alert item. This value is the KV information in the trap message, which is used to identify the alert data. Typically, this value can be the additional information of the alert item, such as a system message or a message indicating the location of the state machine or the current status. Configure the value based on the device document. You cannot customize the value.	None

Parameter	Description	Example
Alert Type	Specifies whether the alert is of the fault type or the event type.	None
Association	Specifies whether the alert has an event alert. If **Alert Type** is set to **Fault** and the alert has an associated alert, set Association to Event Alert and then add the trap of the associated alert.	None

4. Click **Submit**. After the configuration is submitted, the system checks whether the values of Trap OID and Trap Name are the same as the existing ones. If not, the configuration of the trap is complete.

   After the trap is added, the alert events of the configured Trap OID are monitored and displayed on the **Current Alerts** and **Alert History** tabs in the **Alert Management** module.

5.2.5.4. View traps

You can view traps configured in the current system.

Procedure

1. In the left-side navigation pane, choose **NOC > Alert Management > Alert Configuration**.
2. Enter a keyword in the search box in the upper-right corner, and then click **Search**.

> **Note** After the search results are displayed, you can click **Export to CSV** in the upper-right corner to export the trap information to your local computer.

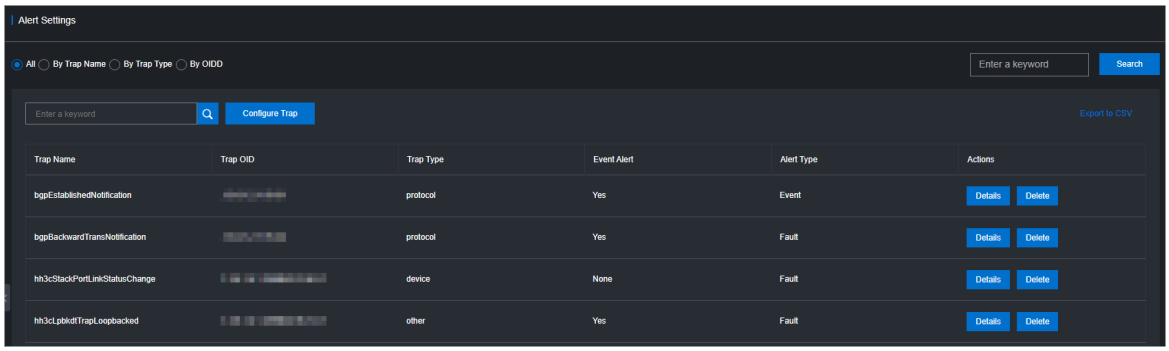

3. (Optional)You can filter the search results by trap name, trap type, or OID.
4. Move the pointer over **Details** in the **Actions** column corresponding to a trap to view detailed information about the trap.

> **Note** If a trap is no longer needed, you can click **Delete** in the **Actions** column corresponding to the trap.

5.2.6. Network reconfiguration

The **Network Reconfiguration** module allows you to automatically reconfigure the network of the data center in Apsara Stack Operations (ASO).

5.2.6.1. Physical network integration

The Physical Network Integration module allows network operations engineers to perform automated integration of physical networks in the ASO console by entering the integration parameters. Network Operation Center (NOC) automatically generates and issues the configurations to specific devices and then automatically performs the network integration test.Enter the project name and then click Create to create a project.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Reconfiguration > Physical Network Integration**.
2. Enter the project name and then click **Create** to create a project.You must create a project file for this change to store the parameters related to the change. You can click **Manage > Import** in the **History** section to import the project information for later usage.
3. In the upper-right corner, click **Save Project** to save the project details.
4. Click **Next**.
5. Select a device.
   i. In the **Select Device** step, enter a device name in the search box of the **Devices on Live Network** section and then click **Search**.After you add a device, you can click **Reset** to clear the search condition and then search for and add another device.

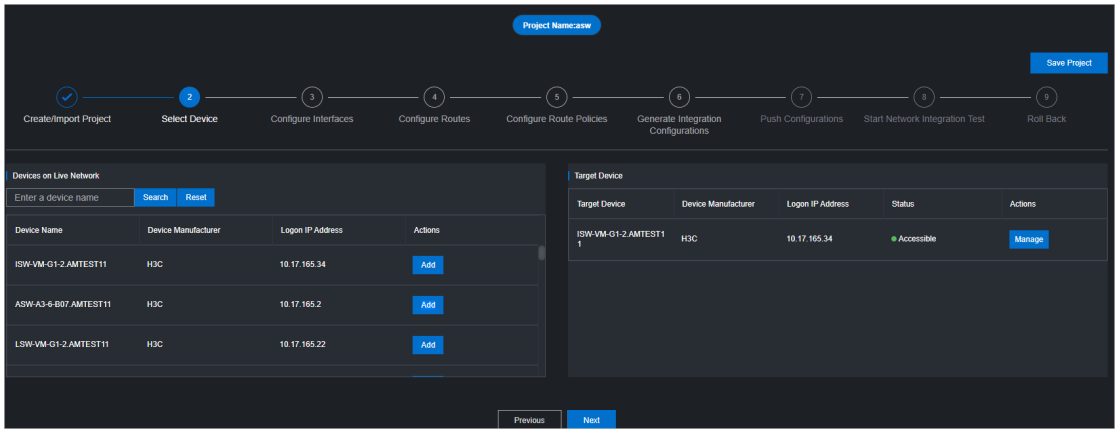

   ii. Find the target device and click **Add** in the Actions column to add it to the Target Device section on the right.To remove the device from the Target Device section, choose **Manage > Delete** in the Actions column. You can also choose **Manage > Set the username and password** to modify the logon username and password of the device.
   iii. In the upper-right corner, click **Save Project** to save the information of devices added to the Target Device section.
6. Click **Next**.
7. Configure the interface parameters.

i. In the **Configure Interfaces** step, click **Edit**. The **Configure Interfaces** section appears.

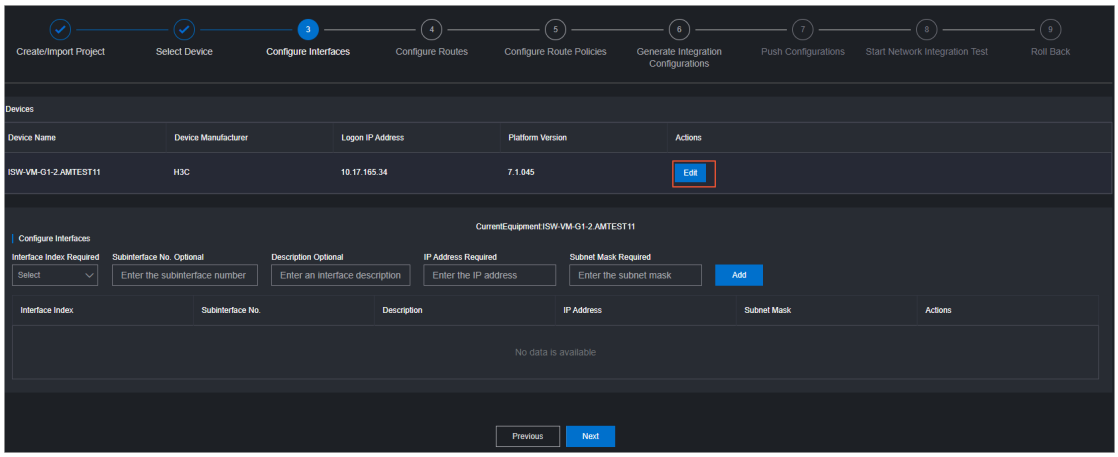

ii. In the **Configure Interfaces** dialog box that appears, configure the parameters and then click **Add** to add the interface to the list.You can choose **Manage > Edit** or **Manage > Delete** in the list to modify or delete the interface.

iii. In the upper-right corner, click **Save Project** to save the information.

8. Click **Next**.
9. Configure the route parameters.

i. In the **Configure Routes** step, click **Edit**. The **Configure Routes** section appears.

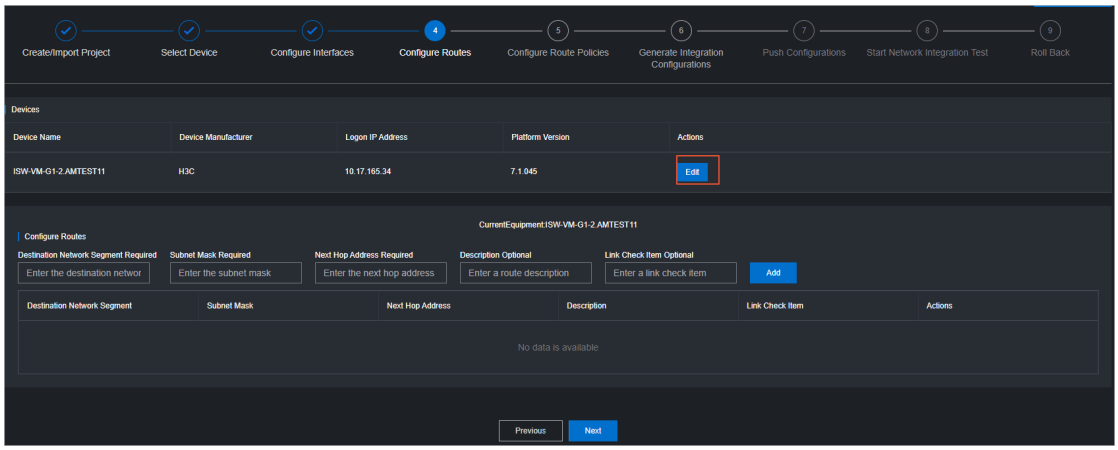

ii. In the **Configure Routes** dialog box that appears, configure the parameters and then click **Add** to add the interface to the list.You can choose **Manage > Edit** or **Manage > Delete** in the list to modify or delete the route.

iii. In the upper-right corner, click **Save Project** to save the information.

10. Click **Next**.
11. Configure the route policies.

i. In the **Configure Route Policies** step, click **Edit**. The **Configure Route Policies** section appears.

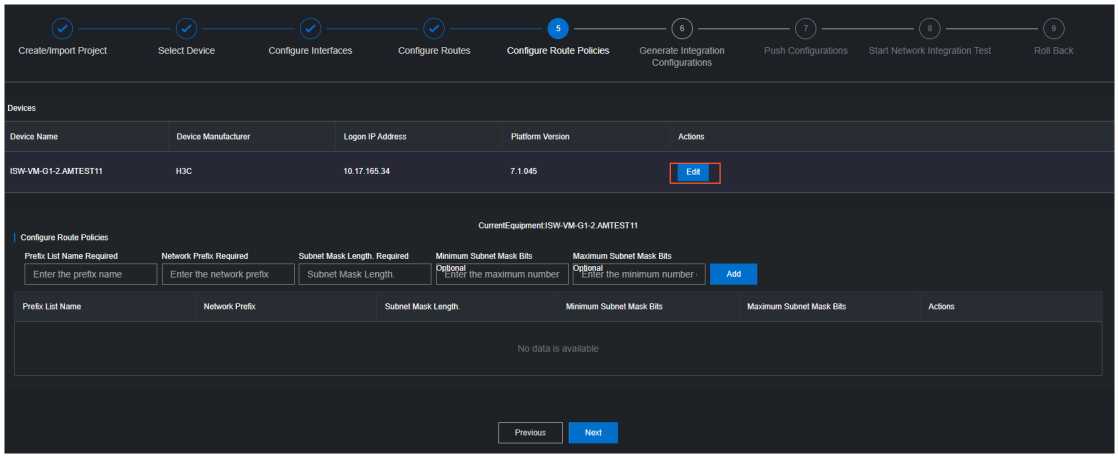

ii. In the Configure Route Policies dialog box that appears, configure the parameters and then click **Add** to add the interface to the list.You can choose **Manage > Edit** or **Manage > Delete** in the list to modify or delete the route policy.

iii. In the upper-right corner, click **Save Project** to save the information.

12. Click **Next**.
13. In the **Generate Integration Configurations** step, click **Generate** to generate the integration configurations.

    The system generates the integration configuration commands and rollback commands of all the devices with parameters configured.

    You can generate the configurations of each device based on the configured parameters. After the generation, click **View** in the **Actions** column to view the corresponding commands on the left.

    You can also click **Export** to export the file, which contains the configuration commands and rollback commands of detection devices, to your local computer.

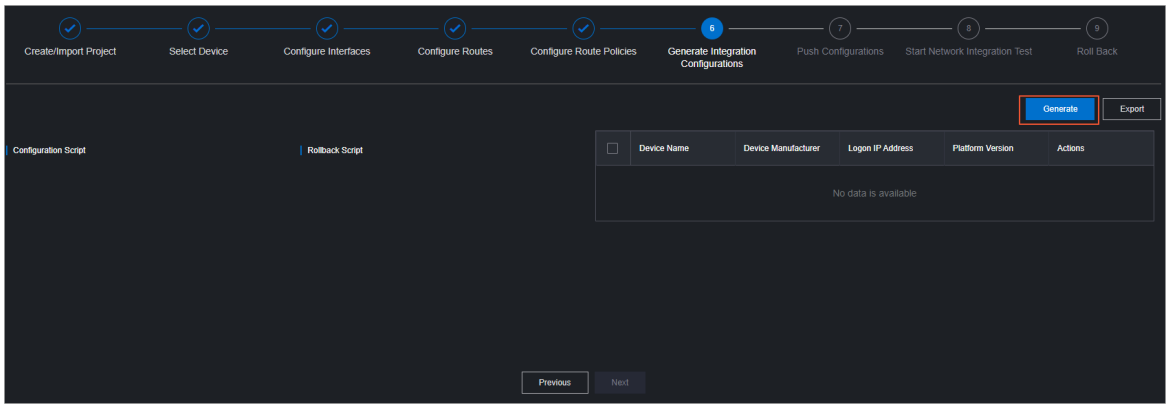

5.2.6.2. ASW scale-up

The ASW Scale-up module allows network operations engineers to automatically scale up ASW devices in the ASO console. After network operations engineers configure the scale-up parameters, NOC automatically generates the configuration and pushes the configuration to a specific device for automatic scale-up.

Prerequisites

Before you scale up ASW devices in the ASO console, you must plan the IP addresses and configure the ASW.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Reconfiguration > ASW Scale-up**.
2. Select target devices.
   i. In the **Select Device** step, enter a device name in the search box of the **Devices on Live Network** section and then click **Search**.After you add a device, you can click **Reset** to clear the search condition and then search for and add another device.
   ii. Find the target device and click **Add** in the Actions column to add it to the Target Device section on the right.To remove the device from the Target Device section, choose **Manage > Delete** in the Actions column. You can also choose **Manage > Set the username and password** to modify the logon username and password of the device.
3. Click **Next**.
4. Disable the DSW ports.
   i. In the **Disable DSW Port** step, find the target device and then click **Port Settings** in the Actions column.
   ii. Disable the port and then click **Implement** in the Actions column.
   iii. In the message that appears, click **OK** to run the script command.
5. Click **Next**.
6. Configure the DSW ports.
   i. In the **Configure DSW Port** step, find the target device and then click **Edit** in the Actions column. The **Interface Parameter Configuration** section appears.
   ii. In the **Interface Parameter Configuration** section that appears, set **Display Ports**, **Port Description**, **IP Address**, and **Subnet Mask**, and then click **Add** to add the interface parameters to the list.You can choose **Manage > Edit** or **Manage > Delete** in the list to modify or delete the interface.
   iii. After adding the interface parameters, click **Implement** in the Actions column.
   iv. In the message that appears, click **OK** to run the script command.If an exception occurs after the implementation, you can click **Back** to roll back to the version before the implementation.
7. Click **Next**.
8. Configure BGP.
   i. In the **Configure BGP** step, find the target device and then click **Edit** in the Actions column. The **Interface Parameter Configuration** section appears.
   ii. In the **Interface Parameter Configuration** section that appears, set **Group Name**, **Peer ASN**, **Peer IP Address**, and **Local Port Name**, and then click **Add** to add the interface parameters to the list.You can choose **Manage > Edit** or **Manage > Delete** in the list to modify or delete the interface.
   iii. After adding the interface parameters, click **Implement** in the Actions column.
   iv. In the message that appears, click **OK** to run the script command.If an exception occurs after the implementation, you can click **Back** to roll back to the version before the implementation.

9. Click **Next**.
10. In the **Upload ASW Configurations** step, upload the new ASW configurations.
11. Click **Next**.
12. Enable the DSW ports.
    i. In the **Enable DSW Port** step, find the target device and click **Port Settings** in the Actions column.
    ii. Enable the port and then click **Implement** in the Actions column.
    iii. In the message that appears, click **OK** to run the script command.
13. Click **Next**.
14. Perform the scale-up test.
    i. In the **Test Scale-up** step, find the target device and click **Select**. The route table is displayed.
    ii. In the **ASW IP Address** search box, enter the IP address to be tested and then click **Add** to add it to the ASW Connectivity Test list.
    iii. Click **Test** and then the system returns the test results.

5.2.6.3. Push IPv6 configurations

The system can automatically push IPv6 configurations. After network operations engineers configure the IPv6 parameters in the IPV6 Configuration module, the system automatically generates the IPv6 configurations and pushes the configurations to specified devices.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Reconfiguration > IPV6 Configuration**.
2. In the **Initialize Device** step, select the target devices and configure the parameters to complete the initialization.
    i. Find the target device in the device list and click **Check** in the **Check** column to check if the device is accessible.You can check multiple devices.
    ii. Select one or more devices whose **Status** is **Accessible** after you click Check.

iii. Configure the parameters on the right.

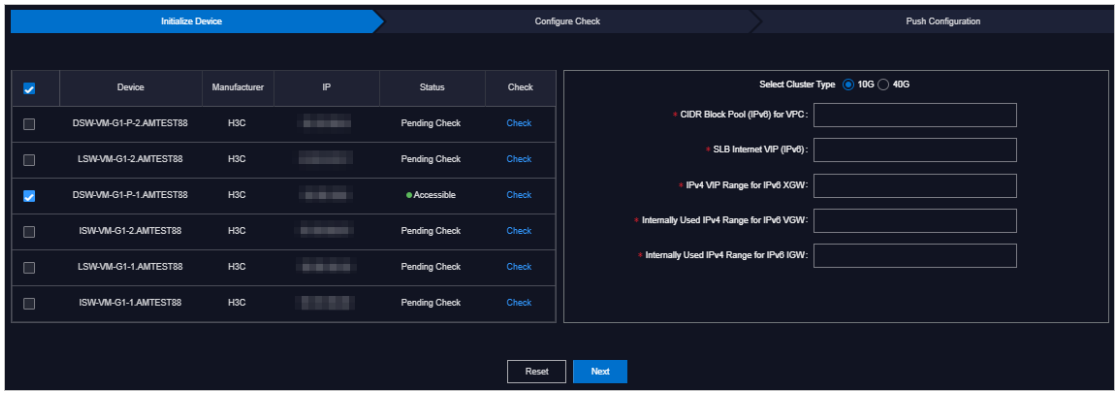

Parameter	Description
Select Cluster Type	The type of the cluster. Valid values: **10G** and **40G**. Select a value based on the planned cluster type.
CIDR Block Pool (IPv6) for VPC	The VPC CIDR block pool in the format of IPv6.
SLB Internet VIP (IPv6)	The SLB public VIP address in the format of IPv6.
IPv4 VIP Range for IPv6 XGW	The XGW VIP CIDR block in the format of IPv4.
Internally Used IPv4 Range for IPv6 VGW	The internally used CIDR block for VGW in the format of IPv4.
Internally Used IPv4 Range for IPv6 IGW	The internally used CIDR block for IGW in the format of IPv4.

3. Click **Next**.
4. In the **Configure Check** step, check the configurations.

   During the configuration check, the system automatically checks the current configurations of the selected devices and generates the IPv6 configuration script based on the check results. Click **View** at the right of the script file to view the generated configuration script, or click **Download** to download the configuration script to your local computer.

   > **Note** If you select multiple devices in the **Initialize Device** step, click **Batch Download** to download multiple configuration scripts to your local computer at a time.

   Generally, the following results occurs during the configuration check:

   - **The configuration is generated. Pending Pushing**
   - **Failed to check the configuration. No BGP processes have been found.**
   - **Failed to check the configuration. Failed to generate the configuration.**
   - **Failed to check the configuration. The IPv6 configuration already exists.**
5. Click **Next**.

The system checks if the configuration pushing function is enabled. If not, the message "**Contact the onsite manager to enable the function before you continue**" appears. If yes, check if the pushing condition is met based on the configuration check results and generation conditions of IPv6 configuration scripts.

- If the configuration check is successful and the IPv6 configuration scripts are generated in the previous step, a dialog box appears. Click **Continue** to automatically push the configuration scripts to the selected devices.
- If the configuration check result is **Failed to check the configuration. No BGP processes have been found.**, **Failed to check the configuration. Failed to generate the configuration.**, or **Failed to check the configuration. The IPv6 configuration already exists.** in the previous step, a dialog box appears, and the system does not automatically push the configurations.

6. After the configurations are pushed, view the pushing results in the **Push Configuration** step.

   If the system indicates that it is pushing the configurations, click **Refresh** to refresh the pushing results.

   After the configurations are pushed, click **View** to view the current running configurations of the selected devices to check if the IPv6 configurations are pushed.

5.2.7. Fault check

The **Fault Check** module consists of IP address conflict check, leased line discovery, and network inspection.

5.2.7.1. Check IP address conflicts

The IP Address Conflict Check module allows you to check whether conflicting IP addresses exist in the current Apsara Stack environment.

Procedure

1. In the left-side navigation pane, choose **NOC > Fault Check > IP Address Conflict Check**. After you access the **IP Address Conflict Check** page, the system automatically checks if conflicting IP addresses exist in the current Apsara Stack environment. If yes, the conflicting IP addresses are displayed in the list. You can also view the port information, device name, and corresponding logon IP address of each conflicting IP address.
2. On the **IP Address Conflict Check** page, view the information of conflicting IP addresses.

5.2.7.2. Leased line discovery

You can configure the leased line discovery of devices in the ASO console and implement it automatically. After network operations engineers configure the discovery parameters, network operation center (NOC) automatically generates the discovery configuration, pushes the configuration to a specific device, and then automatically performs the discovery test.

Procedure

1. In the left-side navigation pane, choose **NOC > Fault Check > Leased Line Discovery**.
2. Select a discovery source.

i. In the **Select Source** step, enter a device name in the search box of the **Devices on Live Network** section, and then click **Search**.After you add a device, you can click **Reset** to clear the search condition, and then search for and add another device.

ii. Click **Add for Discovery** in the Actions column corresponding to the target device to add the device on the live network to the Devices for Discovery list on the right.To remove a device from the Devices for Discovery list, choose **Manage > Delete** in the Actions column corresponding to the device. You can also modify the logon username and password of the device by choosing **Manage > Set the username and password** in the Actions column corresponding to the device.

3. Click **Next**.
4. Configure the discovery parameters.
   i. In the **Configure Parameters** step, click **Edit**. The **Configure Parameters** list is displayed.
   ii. Specify **Link Name**, **Destination IP Address**, **Source IP**, **Discovery Interval**, **Discoveries**, and **Discovery Timeout**, and then click **Add** to add the information to the list.You can choose **Manage > Edit** or **Manage > Delete** to modify or delete the discovery parameters.
5. Click **Next**.
6. In the **Generate Discovery Configuration** step, click **Generate** to generate the discovery configuration and rollback commands of all devices with discovery parameters configured.Click **View** in the **Actions** column. The corresponding commands are displayed on the left.

   You can also select one or more devices and click **Export** to export the files that contain configuration and rollback commands of discovery devices to your local computer.
7. Click **Next**.
8. In the **Push Configuration** step, click **Push Configuration**.
9. In the message that appears, click **Continue** to push the discovery configuration commands to the corresponding device.After the configuration is pushed, you can click **View Logs** to view detailed pushed logs.
10. Click **Next**.
11. In the **Start Discovery** step, click **Started** in the Actions column corresponding to a device to perform the leased line discovery test.
12. After the test is complete, click **Next**.
13. In the **Roll Back Discovery** step, click **Roll Back** in the Actions column corresponding to the device on which you have performed the leased line test to roll back the corresponding NQA configurations in the device.After the rollback is complete, you can click **View Logs** to view detailed rollback logs.

5.2.7.3. Configuration baseline audit

The Configuration Baseline Audit module allows you to compare the baseline configurations of devices with the current running configurations.

Procedure

1. In the left-side navigation pane, choose **NOC > Fault Check > Configuration Baseline Audit**.
2. Select one or more devices in the device list and then click **Audit**. The system starts to audit the baseline configurations of the selected devices.

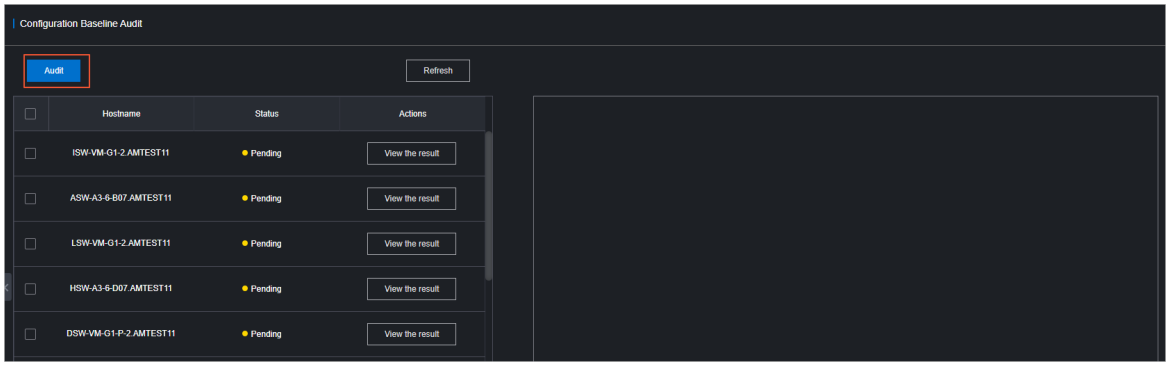

The following table describes status during the audit process and the corresponding descriptions.

Status	Description
Pending	The initial status.
Auditing	The baseline configurations of the device are being audited in the background.
Pass	The current configuration is consistent with the baseline configuration.
Fail	The current configuration is not consistent with the baseline configuration.
Disconnected	The system cannot connect to the device.
No Data	The system cannot obtain the baseline configurations of the device.

3. After the audit is complete, click **Refresh** to update the audit results.
4. Click **View the result** in the **Actions** column of the device. The audit result is displayed on the right.

5.2.8. Network inspection

The Network Inspection module allows you to configure inspection tasks for network devices in the ASO console and implement the tasks automatically for daily fault checking of network devices. You can configure tasks to be executed once or schedule tasks to be executed by the hour, day, week, or month. The Network Inspection module consists of Inspection Dashboard, Inspection History, Inspection Tasks, and Inspection Template.

5.2.8.1. Inspection dashboard

The Inspection Dashboard page shows the inspection data and the last 10 records.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Dashboard**.
2. Perform the following operations:
   - View the inspection statistics of the current day (number of successful tasks, failed tasks, and

tasks scheduled for today, as well as the progress) and the last 10 inspection records.

- View inspection records

  In the **Recent Inspection Tasks** section, click **Details** in the **Result** column corresponding to a task. The following information about the task is displayed: inspection time, inspection template, execution progress, inspection health status, task type, task status, task name, and inspection details of each subtask.

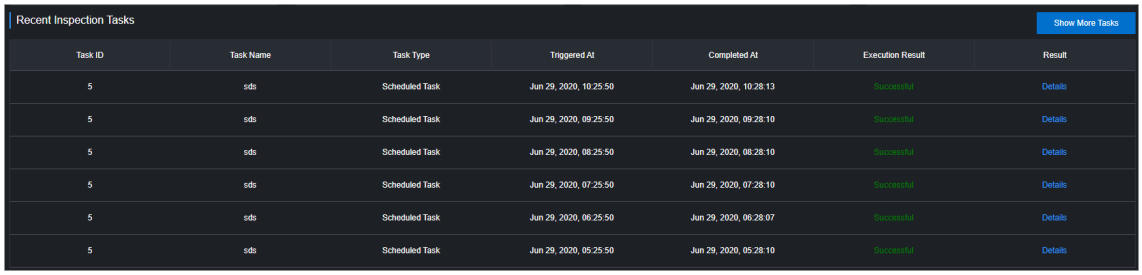

  In the **Inspection Details** section, click **Details** in the **Rollback** column corresponding to a subtask. The inspection result of the inspection subtask is displayed.

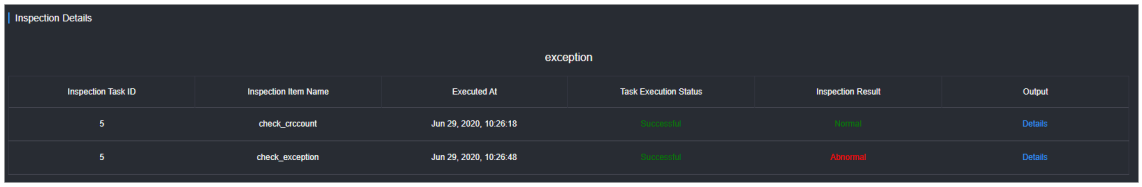

- Click **Show More Tasks** to go to the **Inspection History** page to view the inspection history.

5.2.8.2. Inspection history

You can query the inspection history and view the details by task type and time range.

Context

Inspection tasks can be divided into one-time tasks and scheduled tasks. A one-time task can be executed only once. You can set an execution interval for a scheduled task.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection History**. By default, all inspection records in the last 24 hours are displayed.
2. Select the inspection type (**All**, **One-time Task**, or **Scheduled Task**), specify the time range, and then click **Search**.
3. View inspection history that meets the conditions.
4. Click **Details** in the **Result** column corresponding to an inspection record. The following information is displayed: inspection time, inspection template, execution progress, inspection health status, task type, task status, task name, and inspection details of each subtask.
5. In the **Inspection Details** section, click **Details** in the **Rollback** column corresponding to a subtask. The inspection result of the inspection subtask is displayed.

5.2.8.3. Inspection tasks

The Inspection Tasks module allows you to create, view, modify, start, suspend, and delete inspection tasks.

5.2.8.3.1. Create a one-time task

This topic describes how to create a one-time task.

Context

By default, a one-time task can be executed only once after it is created. After a one-time task is executed once, the task automatically enters the **Suspended** state. The task can be manually started and then executed again.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Tasks**.The **New Inspection Tasks** tab appears.
2. Click the **One-time Task** icon.The wizard for configuring an inspection task appears.
3. In the Specify Inspection Task Name step, enter the inspection task name, and then click **Next**.
4. In the **Add Device for Inspection** step, select one or more devices from the drop-down list, and then click **Next**.

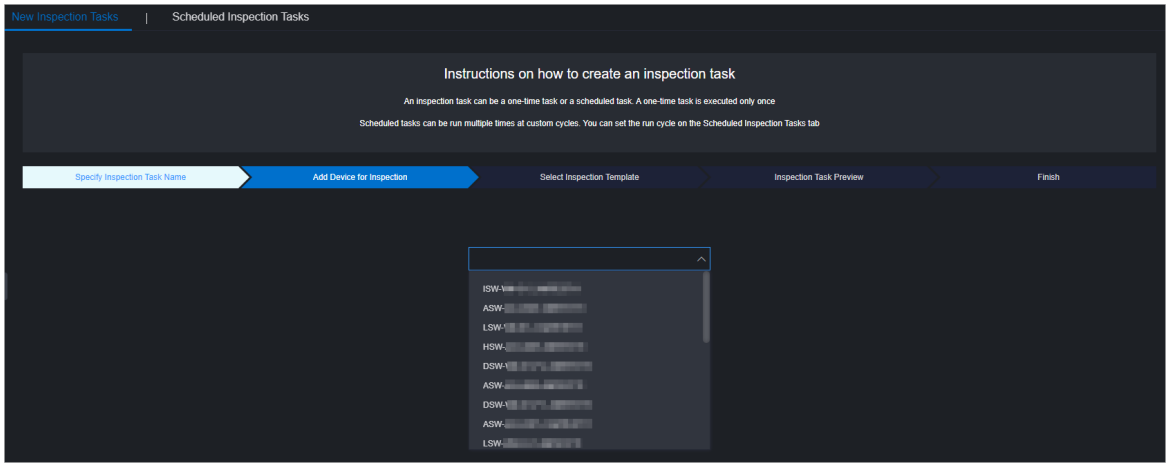

5. In the **Select Inspection Template** step, select an existing template from the drop-down list or click **Create Temporary Inspection Template**.To create a temporary inspection template, click **Create Temporary Inspection Template**. In the dialog box that appears, select the inspection items that you want to associate with the temporary inspection template, and then click **OK**.

> **Note** Some inspection items are provisioned by manufacturers. You must select proper inspection templates or inspection items based on devices. For more information about inspection templates and items in the system, see View template details and View inspection items.

6. Click **Next**.
7. In the **Inspection Task Preview** step, confirm the inspection task information, and then click **Next**.

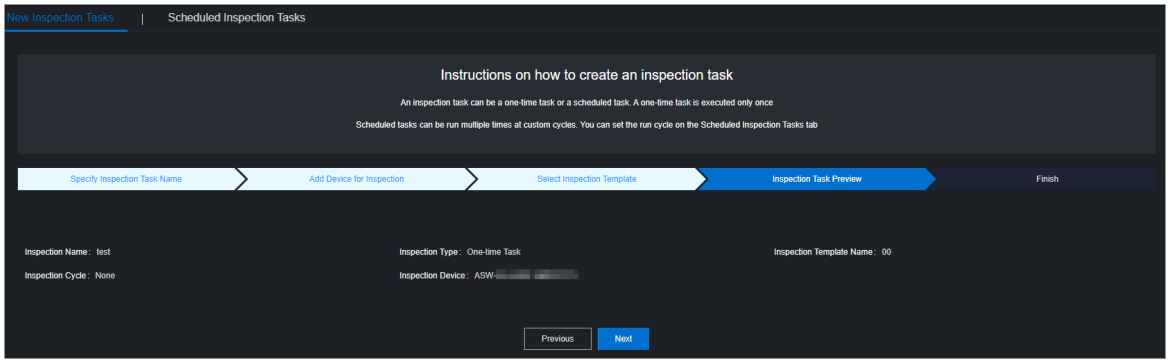

8. Click **Finish**.

   The message **Created** is displayed. You can choose **NOC > Network Inspection > Inspection Tasks** to view the newly created one-time inspection task on the **Scheduled Inspection Tasks** tab.

5.2.8.3.2. Create a scheduled task

This topic describes how to create a scheduled task based on routine inspection requirements. You can set an execution interval for the scheduled task.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Tasks**.The **New Inspection Task** tab appears.
2. Click the **Scheduled Task** icon.The wizard for configuring an inspection task appears.
3. In the Specify Inspection Task Name step, enter the inspection task name, and then click **Next**.
4. In the **Add Device for Inspection** step, select one or more devices from the drop-down list, and then click **Next**.
5. In the **Select Inspection Template** step, select an existing template from the drop-down list or click **Create Temporary Inspection Template**.To create a temporary inspection template, click **Create Temporary Inspection Template**. In the dialog box that appears, select the inspection items that you want to associate with the temporary inspection template, and then click **OK**.
6. Click **Next**.
7. Specify the inspection cycle and the time point when the task is triggered, and then click **Next**.

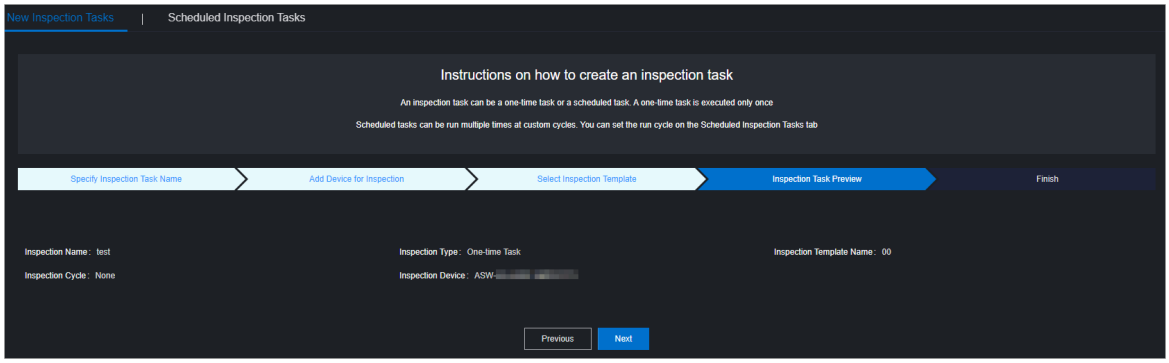

8. In the **Inspection Task Preview** step, confirm the inspection task information, and then click **Next**.
9. Click **Finish**.

The message **Created** is displayed. You can choose **NOC > Network Inspection > Inspection Tasks** to view the newly created scheduled task on the **Scheduled Inspection Tasks** tab.

5.2.8.3.3. Manage scheduled inspection tasks

After an inspection task is created, you can view, modify, start, suspend, or delete the task.

View tasks

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Tasks**.
2. Click the **Scheduled Inspection Tasks** tab.
3. View the information of all the created inspection tasks in the system, including the task ID, task name, task type, associated template, creation time, and running status.

Modify task parameters

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Tasks**.
2. Click the **Scheduled Inspection Tasks** tab.
3. In the task list, find the task that you want to modify and click **Modify** in the **Actions** column.
4. In the dialog box that appears, modify the task parameters.

   Note:

   - For a one-time task, you can modify the inspection name, inspection type, inspection template, and inspection device.
   - For a scheduled task, you can modify the inspection name, inspection type, inspection template, inspection device, and inspection cycle.
5. Click **OK**.

Start or suspend a task

You can start a suspended task or suspend a running task based on O&M requirements.

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Tasks**.
2. Click the **Scheduled Inspection Tasks** tab.
3. In the task list, find the target task and click **Start** or **Suspend** in the **Actions** column.

   > **Note** After a one-time task is executed, it automatically enters the **Suspended** state. Click **Start**, and it can be executed again.
4. In the message that appears, click **OK**.

Delete a task

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Tasks**.
2. Click the **Scheduled Inspection Tasks** tab.
3. In the task list, find the task that you want to delete and click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

5.2.8.4. Template management

The Template Management module allows you to manage inspection templates, including creating, viewing, modifying, and deleting inspection templates.

5.2.8.4.1. Create a template

You can create a common inspection template to facilitate routine inspection task creation.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Templates**.The **Inspection Templates** tab appears.
2. Click **Create Template**.
3. In the dialog box that appears, enter the template name and template tag, and select a manufacturer and a template inspection item collection for the device.

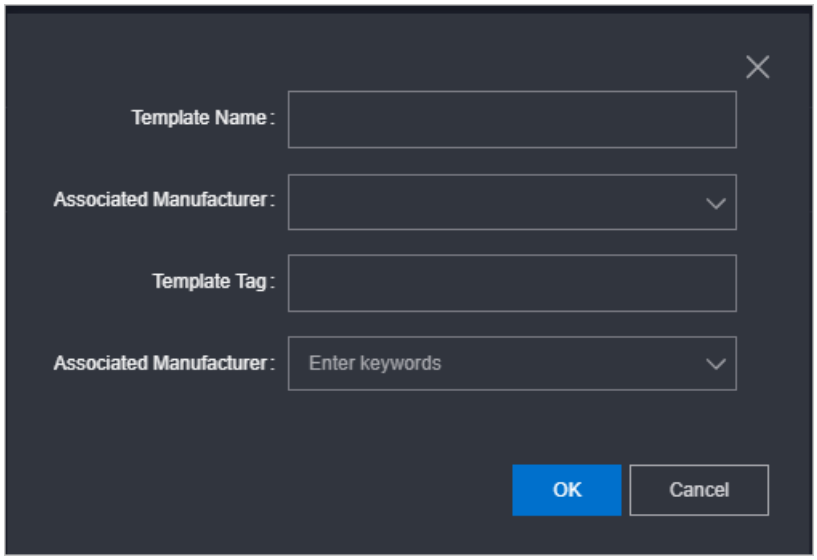

Parameter	Description
Template Name	The name of the inspection template. The name must be unique.
Associated Manufacturers	The manufacturer of the device.
Template Tag	Add a tag to the template to differentiate the template.
Template inspection item collection	The collection of inspection items associated with the template.

4. Click **OK**.After the template is created, you can view the new template in the template list.

5.2.8.4.2. View template details

Before you use an inspection template, you can view the template details to check whether it meets your requirements.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Templates**.The **Inspection Templates** tab appears.
2. In the template list, find the template that you want to view and click **Details** in the **Details** column.

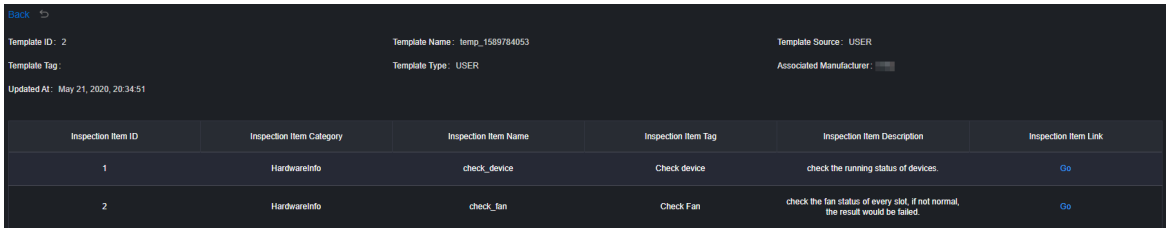

3. View the basic information about the template and the inspection items associated with the template.
4. (Optional)To manage an inspection item in the template, click **Go** in the **Inspection Item Link** column to go to the inspection item management page.For more information about how to manage inspection items, see the following topic in *Apsara Stack Enterprise Operations and Maintenance Guide*:**Configure templates > Network operations > Network management and operations > Network automation**.

> **Note** Typically, no other management operations are required for inspection items.

5.2.8.4.3. Modify a template

After you create a template, you can modify its information.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Templates**. The **Inspection Templates** tab appears.
2. In the template list, find the template that you want to modify and click **Modify** in the **Actions** column.
3. In the dialog box that appears, modify the template name, associated manufacturer, template tag, and template inspection item collection.
4. Click **OK**.

5.2.8.4.4. Delete a template

You can delete an inspection template that is no longer needed based on routine O&M requirements.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Templates**. The **Inspection Template** tab appears.
2. In the template list, find the template that you want to delete and click **Delete** in the **Actions**

column.

> **Notice** If you delete a template, the inspection tasks and inspection records associated with the template are also deleted. Exercise caution when you delete templates.

3. In the message that appears, click **OK**.

5.2.8.4.5. View inspection items

You can view the details of all inspection items in the system, including the item ID, category, name, tag, and description.

Procedure

1. In the left-side navigation pane, choose **NOC > Network Inspection > Inspection Templates**. The **Inspection Templates** tab appears.
2. Click the **Inspection Items** tab.
3. View the information of all inspection items in the system.
4. To perform other management operations on an inspection item, click **Go** in the **Inspection Item Link** column to go to the inspection item management page. For more information about how to manage inspection items, see the following topic in *Apsara Stack Enterprise Operations and Maintenance Guide:* **Configure templates > Network operations > Network management and operations > Network automation**.

> **Note** Typically, no other management operations are required for inspection items.

5.2.9. Use case

5.2.9.1. Troubleshoot network failures

This topic uses a common case to describe how to use the NOC module to troubleshoot network failures.

Scenarios

If a cloud service has latency of visits and the times of retransmission are increased, you must make sure whether this is caused by network failures.

Procedure

1. In the left-side navigation pane, choose **NOC > Dashboard**.
2. Click the **Network Topology** tab.
3. On the tab that appears, click **Topology Type** and select **Standard Topology**.

   Wait five seconds. After the page finishes loading, the system shows the network-wide topology and device connections of the AZ in the current environment.

If device alerts are not triggered in the network, the device icon is blue, the link between devices is green, and the device name is white in the topology. If device alerts are triggered in the network, the topology updates the alert information in the current network every five seconds and shows the updated alert information.

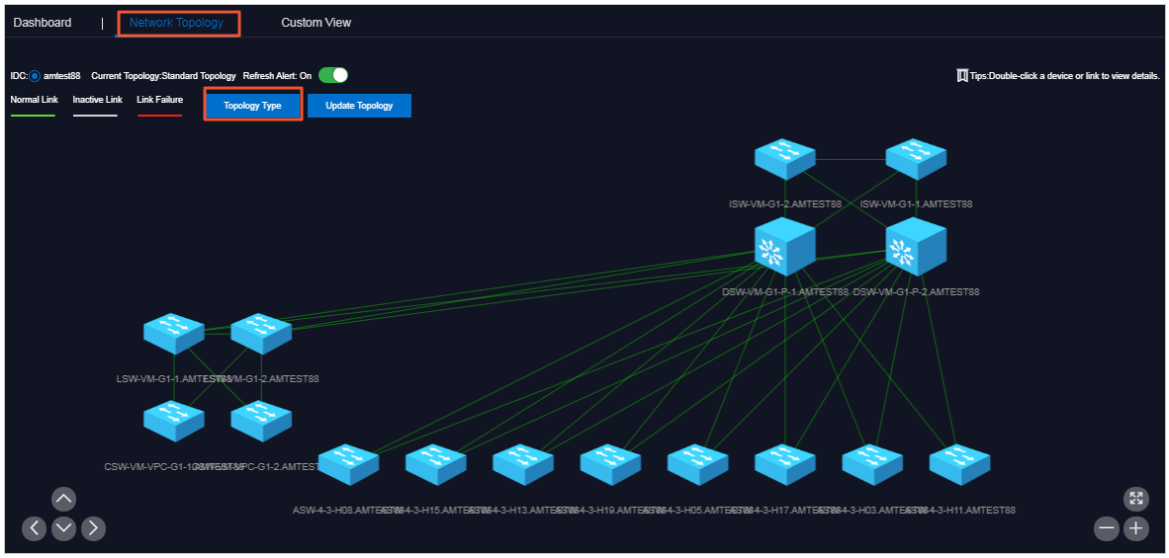

4. If the device name or the link becomes red in the topology, it indicates that alerts are detected in the network device or link port. Double-click the icon of the red device name. In the pane that appears, you can view the basic information of this device and the network alert information related to the device.

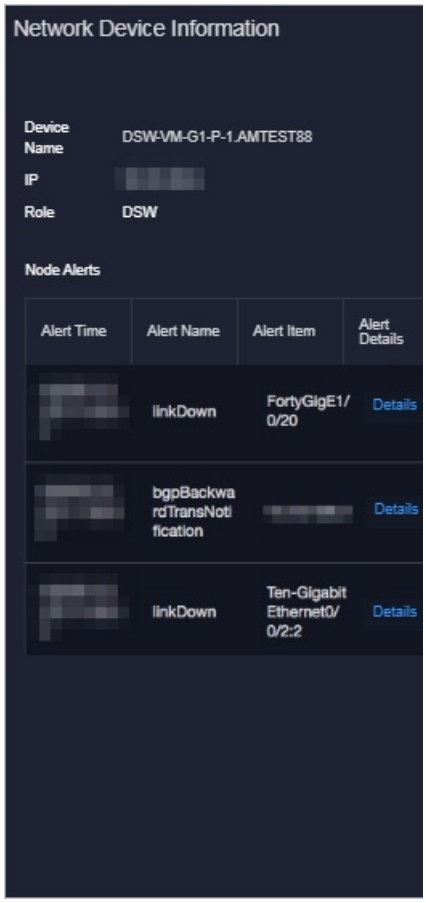

In this example, the port that is connected to this DSW has a **linkDown** alert and a bgp peer alert. An ASW is identified based on the IP address of the BGP peer. Therefore, you can determine that a link between DSW and ASW has a problem, which causes the port down and triggers the alerts.

5. Click the red link in the topology. In the pane that appears, you can view one or more actual physical links contained in the logical link and the alert information of the link between devices.

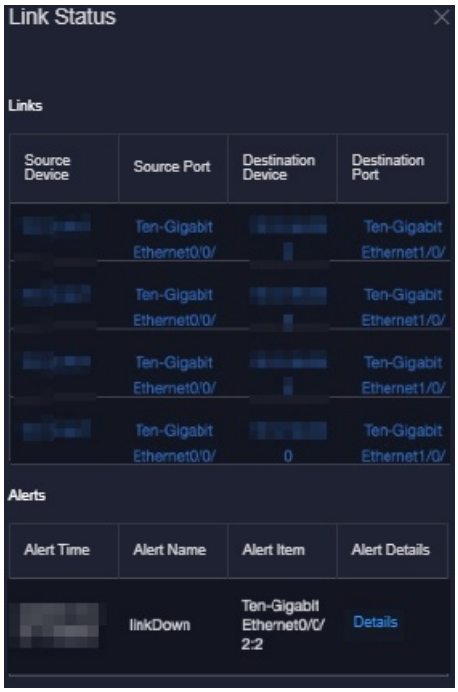

In this example, the logical link connected to the two devices contains four actual end-to-end links. The port 0/0/2:2 has a port **linkDown** alert. Then, you can proceed to log on to the device and check whether this is caused by the low optical power or damaged module.

6. When the problem corresponding to the previous alerts is solved, the system automatically updates the alert information. If the fault is repaired, the alert automatically disappears and the topology is restored to the normal status. As a result, no device names or links are red.

Use the Alert Management module as a supplement to troubleshoot the problem

If a device name or link in the topology becomes red, namely a problem exists in the network device or link, you can choose **NOC > Alert Management > Alert Dashboard** and view the current alerts that are not recovered in the network on the **Current Alerts** tab.

The Current Alerts tab shows more detailed alert information.

If an alert is for test or generated because of a cutover, you can click **Ignore** or **Delete** in the **Actions** column corresponding to the alert to ignore or delete the alert.

Use the syslog log query tool as a supplement to troubleshoot the problem

If a device name or link within a topology becomes red and you have confirmed that the device alert is not caused by expected changes or because of a cutover by using the Alert Management module, you must view the detailed exception logs. You can use the syslog log query tool of the switch to search for logs.

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**. The **Device Management** tab appears.
2. Click the **Syslogs** tab.
3. In the upper-right cornet, select the device and time range, and then click **Search** to search for logs within the selected time range.By default, you can search for a maximum of 1000 logs.
4. In the upper-left corner, enter the keyword in the search box and then click the **Search** icon to search for specific logs in the search results.
5. After the search, if you want to export logs to submit a ticket or submit logs to the device manufacturer for location, click **Export to CSV** in the upper-right corner to save logs as a .csv file to your local computer.

5.3. Task Management

The system allows you to run operations scripts on the cloud platform, which reduces your actions by using command lines, lowers misoperations, and improves the security and stability of the cloud platform.

5.3.1. Overview

The Task Management module has the following functions:

- Supports viewing task overview and creating tasks quickly.
- Supports the following four methods to run tasks: manual execution, scheduled execution, regular execution, and advanced mode.
- Supports the breakpoint function, which allows a task to stop between its two scripts and wait for manual intervention.
- Supports searching for tasks by name, status, and created time.
- Supports uploading the .tar package as the script.

5.3.2. View the task overview

The Task Overview page shows the overall running conditions of tasks in the system. You can also create a task on this page.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Overview**.

   The **Task Overview** page appears.

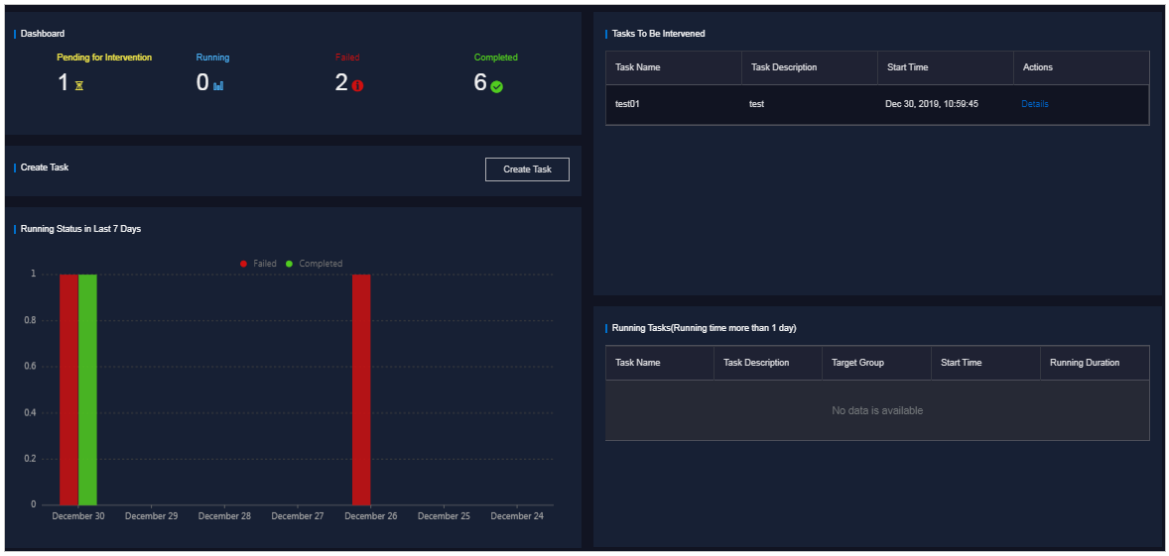

2. You can perform the following operations:
   - In the **Dashboard** section, view the number of tasks that are in the **Pending for Intervention**, **Running**, **Failed**, or **Completed** state in the system.

     Click a state or number to view the task list of the corresponding state.
   - In the **Create Task** section, click **Create Task** to create an operations task.

     For more information about how to create a task, see Create a task.
   - If a task has a breakpoint and reaches the breakpoint, the task stops and waits for manual confirmation. You can view and process tasks that require manual intervention in the **Tasks To Be Intervened** section.
   - In the **Running Status in Last 7 Days** section, view the running trend of tasks and whether tasks are successful within the last seven days.
   - In the **Running Tasks** section, view tasks running within the last 24 hours.

5.3.3. Create a task

You can make regular modifications as tasks to run in the ASO console.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. Click **Create**.
3. In the dialog box that appears, configure the parameters.

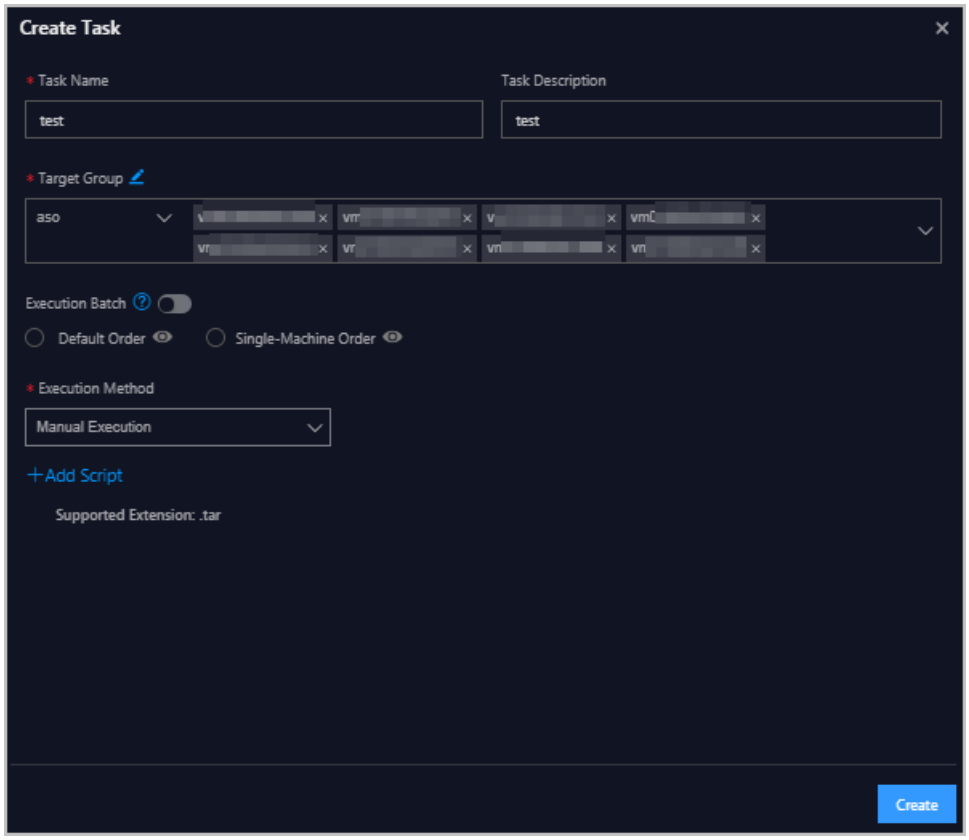

Parameter	Description
Task Name	The name of the operations task.
Task Description	The description of the operations task.
Target Group	The task target. You can use one of the following methods to configure the target group: ○ Select the **product**, **cluster**, **service**, **server role**, and **virtual machine (VM) or physical machine** in sequence. ○ Select a product. Enter the VM or physical machine in the field and then press the Enter key. You can enter multiple VMs or physical machines in sequence. ○ Click the icon next to **Target Group**. In the dialog box that appears, enter the target group, with one VM or physical machine in one line. Click **OK**.

Parameter	Description
Execution Batch	Optional. This option appears after you specify the target group. If **Execution Batch** is not specified, **Target Group** is displayed in the **Target Group** column, which can be viewed by choosing **Task Management > Task Management**. If you specify **Execution Batch**, **Batch Execution Policy** is displayed in the **Target Group** column. You can set **Execution Batch** to one of the following values: ○ **Default Order** By default, if the number of machines is less than or equal to 10, the machines are allocated to different batches, with one machine in batch 1, one machine in batch 2, two machines in batch 3, three machines in batch 4, and the other machines in batch 5. You can change the number of machines in each batch. By default, if the number of machines is greater than 10, the machines are allocated to different batches, with one machine in batch 1, three machines in batch 2, five machines in batch 3, N/3-1 (an integer) machines in batch 4, N/3-1 (an integer) machines in batch 5, until all of the machines are allocated. N is the total number of servers in the cluster. You can change the number of machines in each batch. ○ **Single-Machine Order**: By default, each batch has one machine. You can change the number of machines in each batch.
Execution Method	If **Execution Batch** is specified, **Execution Method** can only be set to **Manual Execution**. If **Execution Batch** is not specified, you can select one of the following execution methods: ○ **Manual Execution**: You must manually start the task. With **Manual Execution** specified, you must click **Start** in the **Actions** column to run the task after the task is created. ○ **Scheduled Execution**: Select the execution time. The task automatically starts when the execution time is reached. ○ **Regular Execution**: Select the time interval and times to run the task. The task starts again if the execution condition is met. ○ **Advanced**: Configure the command to run the task periodically.
Add Script	Click **Add Script**. Select one or more .tar packages to upload the script file. After the upload, you can delete and re-upload the script. After you upload the script, if **Execution Method** is set to **Manual Execution**, you must specify whether to enable **Intervention Required**. If manual intervention is enabled, the task will stop and wait for manual intervention after you run the script.

4. Click **Create**.

Result

The created task is displayed in the task list.

5.3.4. View the execution status of a task

After a task starts, you can view the execution status of the task.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Query** to search for tasks.
3. Find the task that you want to view, and then click **Target Group** or **Batch Execution Policy** in the **Target Group** column.

> **Note** If **Execution Batch** is not selected when you create a task, **Target Group** is displayed in the **Target Group** column. If you select **Execution Batch** when you create a task, **Batch Execution Policy** is displayed in the **Target Group** column.

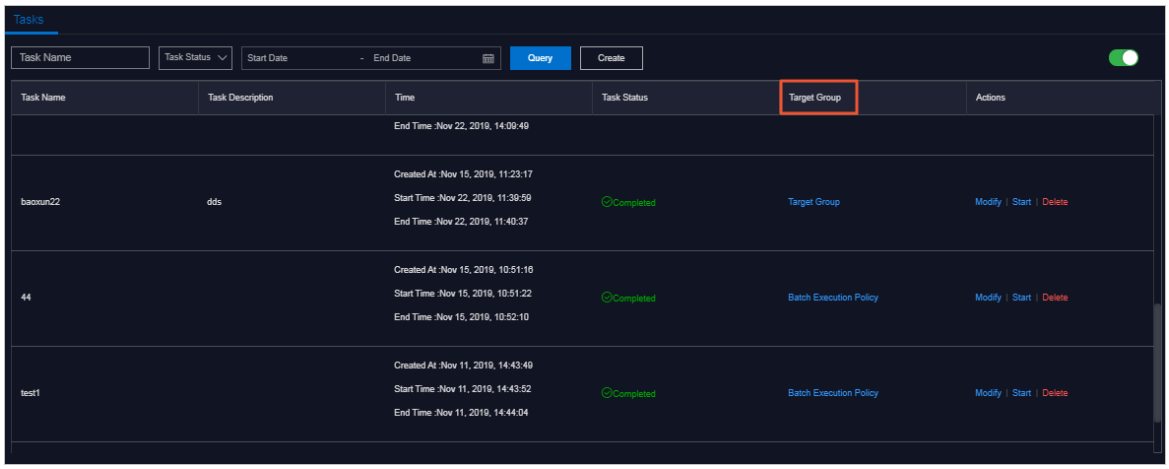

4. In the dialog box that appears, view the task execution status based on the machine color. Click a machine to view the execution results of the task.

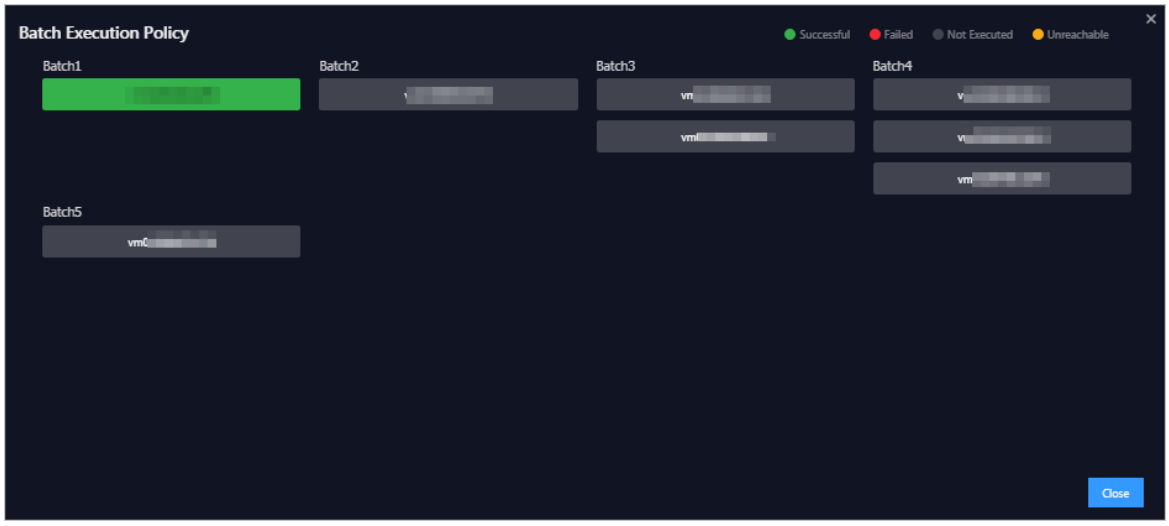

5.3.5. Start a task

If you select **Manual Execution** when you create a task, you must manually start the task after it is created.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Query** to search for tasks.
3. Find the task that you are about to start, and then click **Start** in the **Actions** column.
4. In the dialog box that appears, select the batches to start, and then click **Start**.

   For a new task, after you click **Start** for the first time, the system will indicate that the task is started. The virtual machines (VMs) or physical machines in batch 1 start to run the task. Click **Start** again and you can select VMs or physical machines in one or more batches to run the task.

   If the task has enabled Intervention Required, you must intervene the script after you click **Start**. The **Task Status** turns to **Pending for Intervention**, and you can continue to run the task only by clicking **Continue** in the **Actions** column.

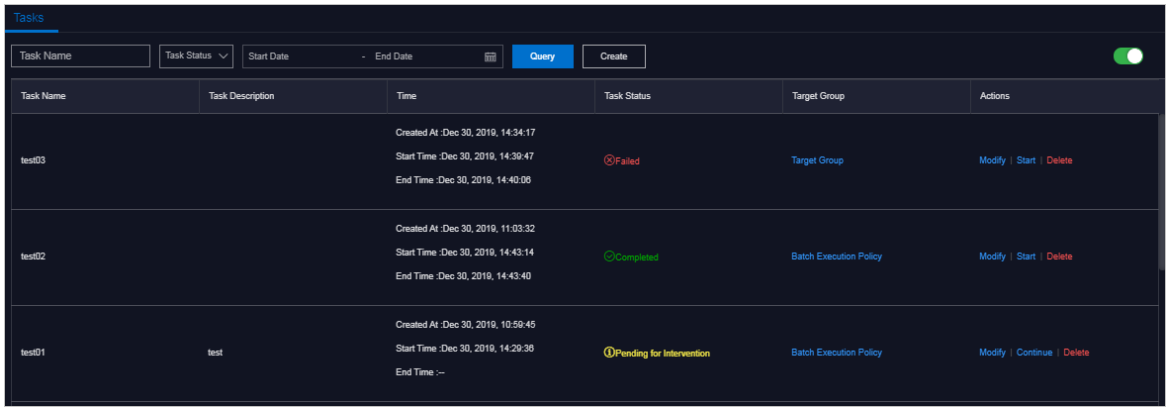

5.3.6. Delete a task

You can delete tasks that are no longer needed.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Query** to search for the task.
3. Find the task to be deleted, and then click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

5.3.7. Process tasks to be intervened

If a task reaches a breakpoint, the task will stop and wait for manual confirmation. The task will continue only after receiving manual confirmation.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Overview**.
2. In the **Tasks To Be Intervened** section, find the task to be intervened, and then click **Details** in the **Actions** column.

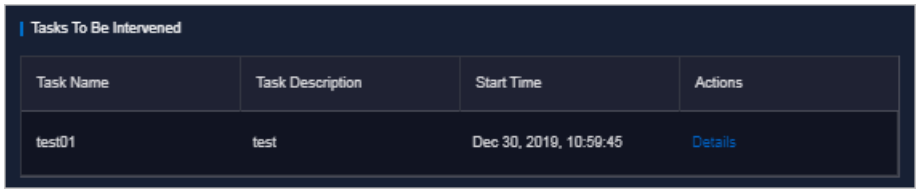

3. On the **Task Details** tab, check the information and then click **Continue** to continue to run the task.

5.4. Apsara Stack Doctor (ASD)

5.4.1. Apsara Stack Doctor introduction

Apsara Stack Doctor (ASD) checks the health of services for Apsara Stack Management Console and troubleshoots faulty services. Data in Apsara Stack Doctor comes from Apsara Infrastructure Management Framework SDK. The data includes the raw data of deployed Apsara Stack products, network topology metadata, and monitoring data.

Basic features

- Provides data filtering, analysis, and processing for O&M data consumers.
- Provides encapsulation, orchestration, and rights management of O&M operations.
- Provides O&M experience accumulation and archiving capabilities.
- Provides troubleshooting, pre-diagnosis, health check, and early warning capabilities.
- Records O&M experience, prescriptions, monitoring data, and log data to support intelligent O&M.

Benefits

- Provides unified management of Apsara Stack O&M data.
- Complements on-site O&M tools.
- Provides a unified tool for automated inspection of Apsara Stack.
- Allows you to perform O&M through Web interfaces, eliminating highly risky black screen operations.
- Allows you to have a periodic offline backup of Apsara Stack metadata, providing out-of-band support for metadata recovery.

Terms

Apsara Stack has five levels of release granularity, as shown in Levels of release granularity.

Levels of release granularity

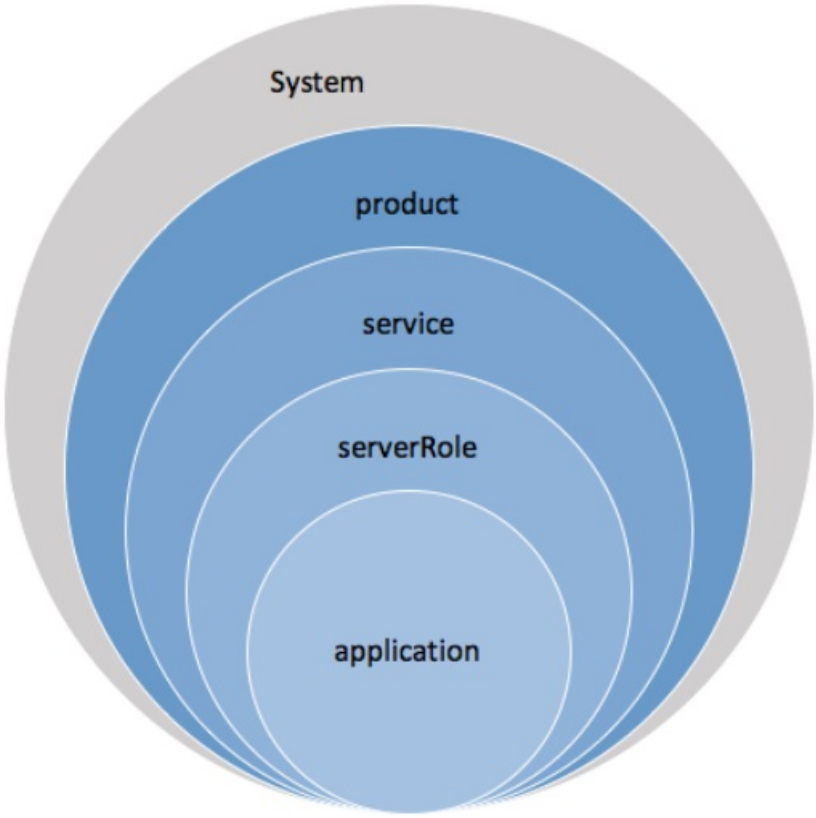

- system

  The greatest granularity at which Apsara Stack is available to external users. It is a collection of one or more Apsara Stack products.

- product

  A category of product visible to users in Apsara Stack. It provides users with a kind of relatively independent features. For example, both ECS and SLB are products. Each product provides one or more features. Each product feature may be provided by one or more types of clusters.

- service

  A type of software that provides independent features. It represents a product module or component. Each service can be managed separately or combined with other services into a product. If a service provides a complete set of features, it can also serve as a separate product alone.

- server role (sr)

  A service component. A service can contain multiple server roles, each of which serves as a submodule of the service and provides a separate feature. Server role is also the smallest granularity monitored during Apsara Infrastructure Management Framework deployment and O&M. Some examples of server roles include PanguMaster and PanguChunkserver. Server roles are mapped to servers. Applications can be deployed to servers by their server role. A server role can contain multiple applications. Multiple applications belonging to a server role are packaged together for deployment. Different applications in a single server role can only be deployed to the same server. Multiple server roles are combined into a server role group (srg) for software deployment purposes. Only one server role group can be deployed to a server.

- application (app)

An independent process. Applications are one component of a server role, the other two being docker and file. All applications are built from source code.

- docker: a Docker image that is built from source code.
- file: a file that is placed on a server.
- application: a piece of software that is built from source code files and can be started directly from a start executable.

5.4.2. Log on to Apsara Stack Doctor

This topic describes how to log on to Apsara Stack Doctor.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   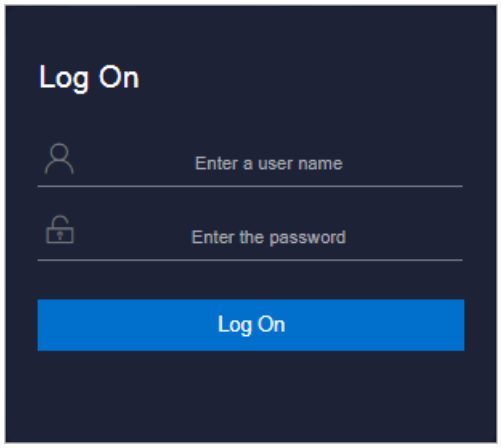

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.

- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, click **Products**.
6. In the **Basic O&M** region, click **ASD**.

5.4.3. ASA

ASA is a tool provided to help you improve the efficiency in testing, operating, maintaining, and releasing cloud products in Apsara Stack while ensuring the stability of version qualities. ASA retains the features of Apsara Stack V2, including inspection, scanning, and version tracking. This continues and precipitates all the long-term experience of Apsara Stack.

5.4.3.1. RPM Check

The RPM Check module allows you to check whether the RPM service is available on all machines, including Docker virtual machines and NCs.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > RPM Check**.

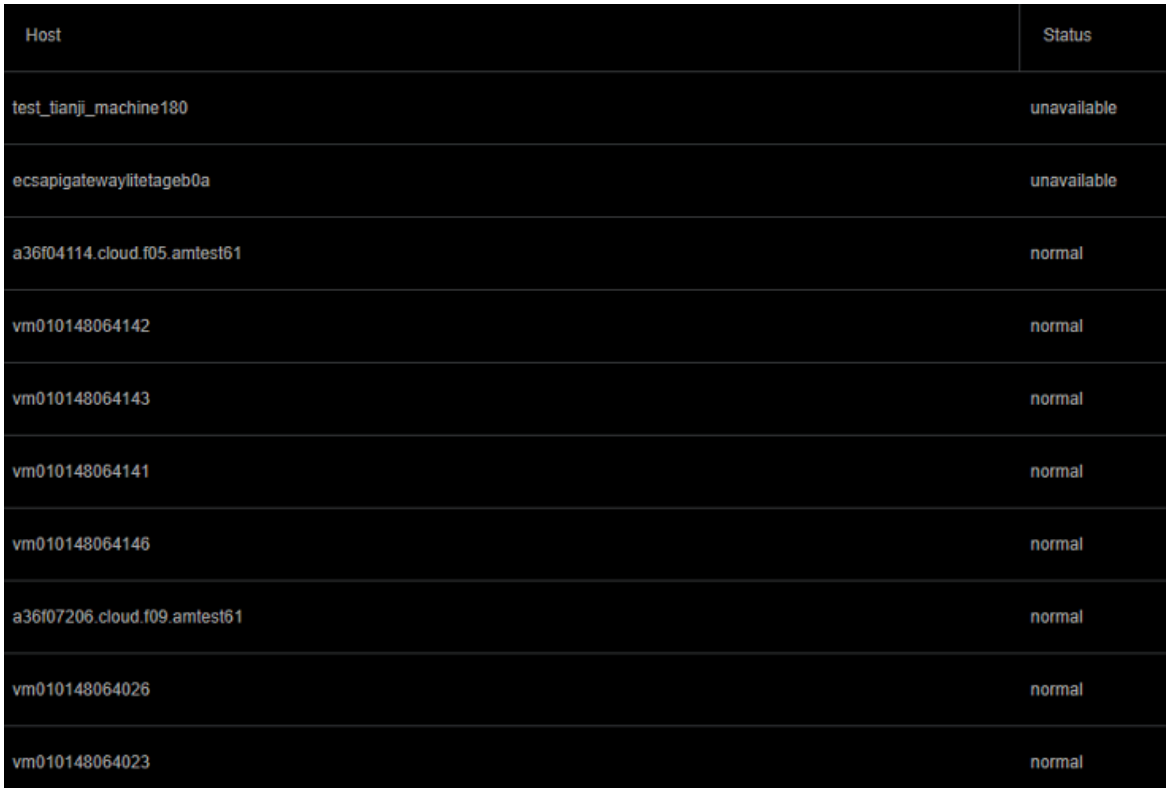

Host	Status
test_tianji_machine180	unavailable
ecsapigatewaylitetageb0a	unavailable
a36f04114.cloud.f05.amtest61	normal
vm010148064142	normal
vm010148064143	normal
vm010148064141	normal
vm010148064146	normal
a36f07206.cloud.f09.amtest61	normal
vm010148064026	normal
vm010148064023	normal

Description of parameters on the RPM Check page

Parameter	Description
Host	The name of a host.
Status	The status of a machine. Valid values: ○ **normal**: indicates that the machine is operating normally. ○ **unavailable**: indicates that the machine is not operating normally or unavailable.

5.4.3.2. Virtual IP Check

The Virtual IP Check module allows you to obtain the virtual IP addresses that are incorrectly bound to IP addresses of backend services.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > Virtual IP Check**.

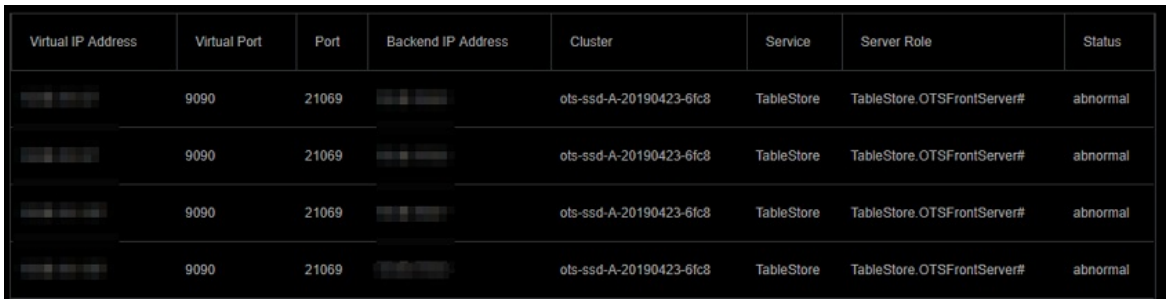

Parameters on the Virtual IP Check page

Parameter	Description
Virtual IP Address	The virtual IP address.
Virtual Port	The port corresponding to a virtual IP address.
Port	The port corresponding to the IP address of a backend service.
Backend IP Address	The IP address of a backend service.
Cluster	The cluster to which the IP address of a backend service belongs.
Service	The service to which the IP address of a backend service belongs.
Server Role	The server role to which the IP address of a backend service belongs.

Parameter	Description
Status	The health status, indicating whether the binding between the virtual IP address and the IP address of the backend service is normal. ○ **normal**: indicates that the virtual IP address is correctly bound to the IP address of the backend service. ○ **abnormal**: indicates that the virtual IP address is not bound to the backend IP address properly.

5.4.3.3. Volume Check

The Volume Check module allows you to view the volume details of Docker hosts.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > Volume Check**.

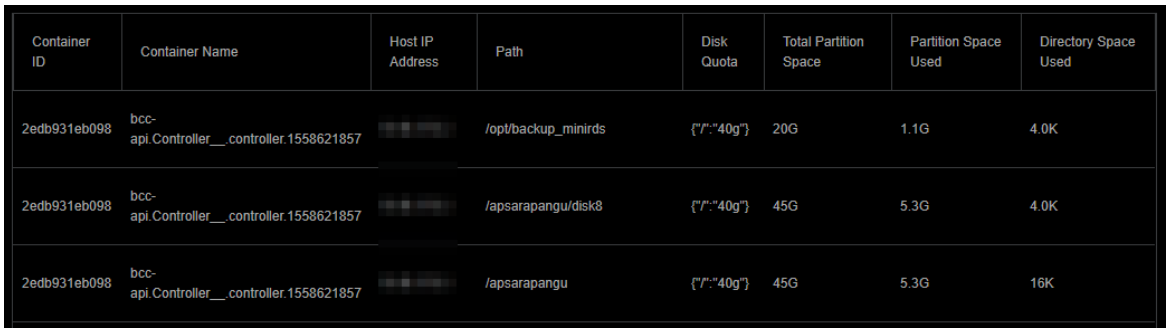

Container ID	Container Name	Host IP Address	Path	Disk Quota	Total Partition Space	Partition Space Used	Directory Space Used
2edb931eb098	bcc-api.Controller__.controller.1558621857	[redacted]	/opt/backup_minirds	{"/":"40g"}	20G	1.1G	4.0K
2edb931eb098	bcc-api.Controller__.controller.1558621857	[redacted]	/apsarapangu/disk8	{"/":"40g"}	45G	5.3G	4.0K
2edb931eb098	bcc-api.Controller__.controller.1558621857	[redacted]	/apsarapangu	{"/":"40g"}	45G	5.3G	16K

Parameters on the Volume Check page

Parameter	Description
Container ID	The unique ID of a Docker container.
Container Name	The name of a Docker container.
Host IP Address	The IP address of a Docker host. Typically, a Docker virtual machine can be either a physical host or virtual host.
Path	The disk partition mount point of a Docker volume.
Disk Quota	The quota of a disk.
Total Partition Space	The total space of a mount point calculated by running the **df** command.

Parameter	Description
Partition Space Used	The space used by a mount point directory.
Directory Space Used	The total space of a mount point calculated by running the **du** command.

5.4.3.4. NTP Check

The NTP Check module allows you to check whether the system time of all machines, including Docker virtual machines and physical machines, is synchronized with the NTP time. If not, the time offset is reported in milliseconds.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > NTP Check**.

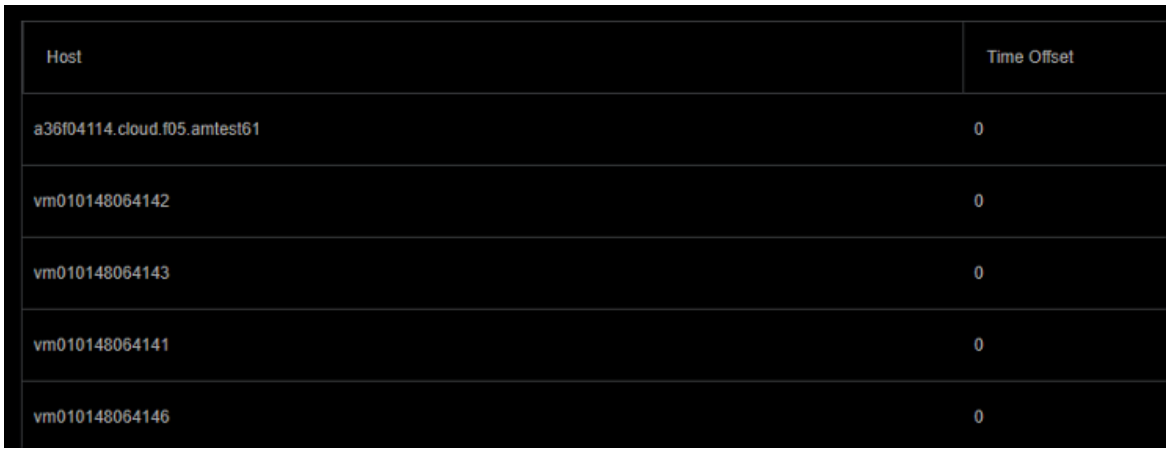

Parameters on the NTP Check page

Parameter	Description
Host	The name of a host.
Time Offset	The time offset. Unit: milliseconds.

5.4.3.5. IP Conflict Check

The IP Conflict Check module allows you to check for IP address conflicts in the current environment.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > IP Conflict Check**.

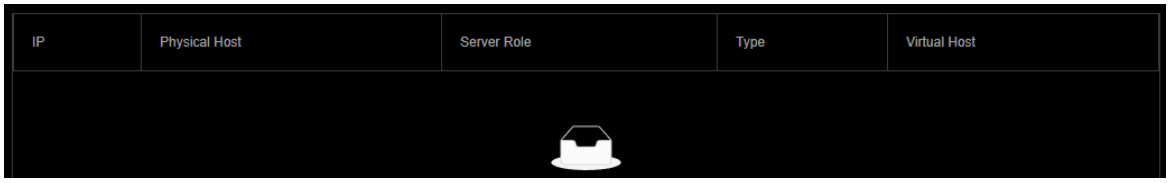

Parameters on the IP Conflict Check page

Parameter	Description
IP	A conflicting IP address.
Physical Host	The name of the physical host with the conflicting IP address.
Server Role	The server role that requests the resource.
Type	The IP address type. Valid values: docker, vm, and physical.
Virtual Host	The hostname of the Docker virtual machine.

5.4.3.6. DNS Check

The DNS Check module allows you to check whether the IP address bound to a domain name is the same as the requested IP address.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > DNS Check**.

Parameters on the DNS Check page

Parameter	Description
Domain	The domain name requested by Apsara Infrastructure Management Framework.
Virtual IP Address	The IP address that is bound to the domain name requested by Apsara Infrastructure Management Framework.
Owner	The application that requests the DNS resource.
IP	The physical IP address that is bound to the domain name.

5.4.3.7. IP Details

The IP Details module allows you to check the details of all IP addresses in the current environment, including the IP addresses of physical machines, Docker machines, and virtual machines, as well as virtual IP addresses.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > IP Details**.

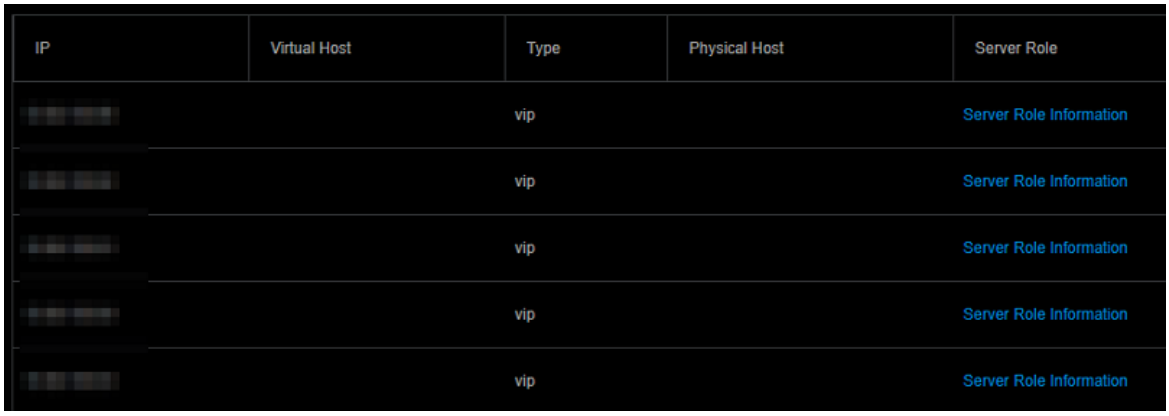

Parameters on the IP Details page

Parameter	Description
IP	The IP address of a resource.
Virtual Host	The name of a virtual machine.
Type	The resource type. Valid values: ○ **physical** ○ **docker** ○ **vm**
Physical Host	The name of a physical host.
Server Role	The server role that requests the resource.

3. Move the pointer over **Server Role Information** in the **Server Role** column to view server role details.

5.4.3.8. Quota Check

The Quota Check module allows you to check the memory, CPU, and disk quotas of containers.

Procedure

1. Log on to Apsara Stack Doctor.

2. In the left-side navigation pane, choose **ASA > Quota Check**.

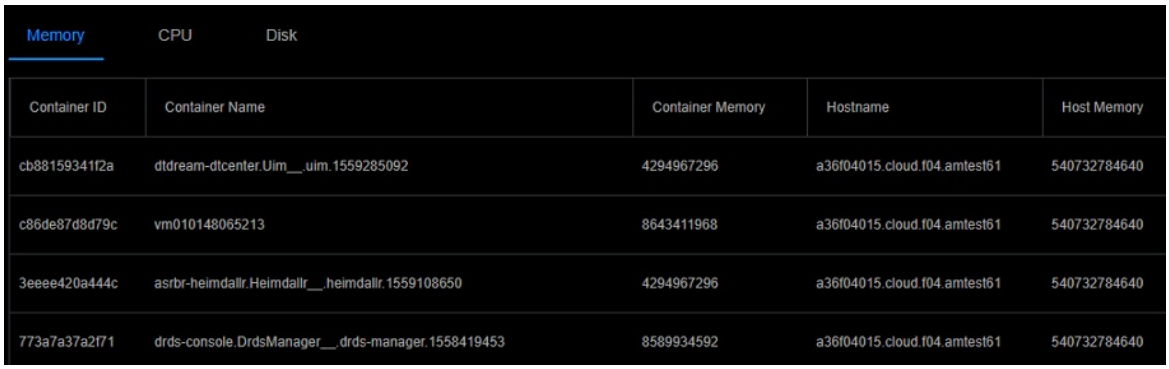

3. On the Quota Check page, you can view memory, CPU, and disk quota information.
   - Memory quota check

     Click the **Memory** tab to view the memory allocation of specified machines.
   - CPU quota check

     Click the **CPU** tab to view the CPU allocation of specified machines.
   - Disk quota check

     Click the **Disk** tab to view the disk allocation of specified machines.

5.4.3.9. Error Diagnostics

Context

The Error Diagnostics page consists of the following tabs:

- Resource Errors: displays resource errors.
- Error with Self: displays internal errors.
- Error with Dependency: displays dependency errors.
- Normal: displays resources with no errors.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > Error Diagnostics**.
3. Switch between tabs to view the corresponding information.

5.4.3.10. Versions

The Versions module allows you to obtain version information and upgrade information of all services in the current environment.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **ASA > Versions**.
3. You can perform the following operations:
   - Click the **Product Versions** tab to view information related to service versions, such as the IDC

- Click the **Product Versions** tab to view information related to service versions, such as the IDC, service, and version.
- Click the **Server Role Versions** tab to view information related to server role versions, such as the IDC, service, version, server role, and type.
- Click the **Version Tree** tab to view information related to version trees.

5.4.4. Support tools

5.4.4.1. Diagnose with the OS tool

The OS tool allows you to perform OS diagnostics on physical machines in Apsara Stack.

Context

The OS tool allows you to diagnose the following metrics: disk file metadata usage, memory usage, process statuses, time synchronization, kernel errors, high-risk operations, system loads, fstab files, read-only file systems, kdump services, kdump configurations, conman configurations, domain name resolution, disk I/O loads, file deletion exceptions, system errors, RPM databases, fgc, tair, route_curing, default routes, unusual network packets, TCP connection status exceptions, TCP queue exceptions, network packet loss, bonding exception, NIC exception, SN retrieval exceptions, OOB IP retrieval exceptions, sensor exceptions, sensor record exceptions, SEL record exceptions, Docker status exceptions, and RAID exceptions.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Support Tools > OS Tool**.

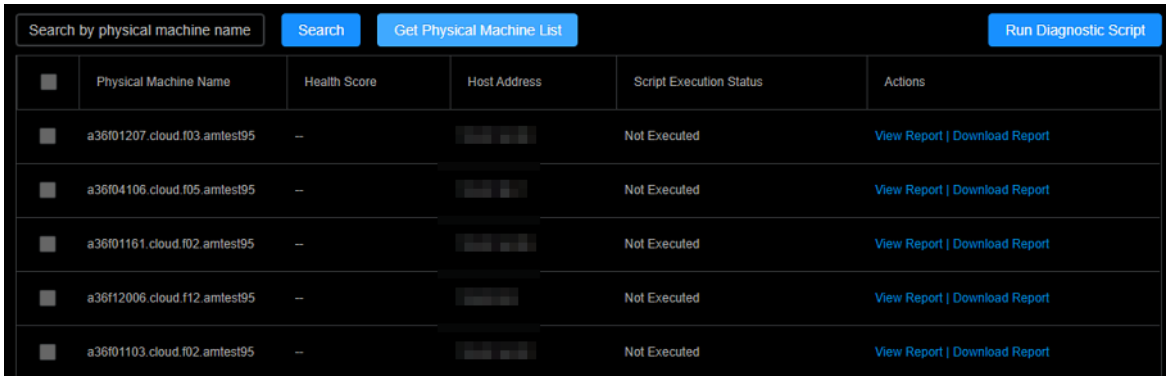

3. Click **Get Physical Machine List** to obtain a list of all the physical machines in the system.
4. (Optional)In the search bar, enter the name of a physical machine and click **Search**. The section below the search bar displays the physical machines.
5. Select the physical machine and click **Run Diagnostic Script** in the upper-right corner.
6. When **Script Execution Status** changes from **Not Executed** to **Diagnostic Result Decompression Finished**, you can view the health score of the physical machine in the **Health Score** column.
7. After the diagnostics are completed, click **View Report** in the **Actions** column to view the diagnostic result.
8. (Optional)For more information, click **View Result** or **Download Report** in the **Actions** column.

5.4.4.2. Use Support Tools

Support Tools allows you to diagnose some services and export diagnostic reports.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Support Tools > Support Tools**.
3. (Optional)Select the target service, enter the host name or IP address, and click **Search**. The search results appear in the section below. The following table lists the supported diagnostic items.

Diagnostic item	Description
Apsara Distributed File System Diagnostics	Collects and analyzes the running status of Apsara Distributed File System and its dependent services and environments, and provides diagnostic reports in case of exceptions.
ecs_vmdisk_usage_V3	Checks the ECS disk usage.
oss_used_summary	Checks the usage of OSS resources.
ots_examine	Checks the following information: ○ NTP ○ Consistency of the Table Store versions ○ Chunkserver status of Apsara Distributed File System ○ Status of Apsara Name Service and Distributed Lock Synchronization System ○ SQL status ○ SQL partition and distribution ○ Service availability of DNS ○ Service availability of SLB ○ Service availability of RDS ○ Service availability of OTS Cluster Management (OCM) ○ Service availability of Red Hat Package Manager (RPM) databases
ecs_error_log	Collects ECS logs.
ots_used_summary	Checks the usage of Table Store resources.
docker	Collects and analyzes data from Docker hosts, and generates reports based on the data.
ecs_diagnostor_v3	Collect the logs of end-to-end ECS links.

Diagnostic item	Description
os	Collects and analyzes system logs, including the following operations: ○ Collects information about the OS, network, disk, and hardware. ○ Diagnoses and analyze system logs. ○ Generates reports.
oss_examine	Diagnoses OSS.

4. Find the row that contains the target machine and click **Run Diagnostics** in the **Actions** column corresponding to the target machine.

   > **Note** Alternatively, you can select the target service and click **Search**. In the search results, select multiple machines and click **Run Diagnostics** for batch diagnostics.

   When **Diagnostics Execution Status** changes from **Running** to **Succeeded**, the diagnostics are completed.

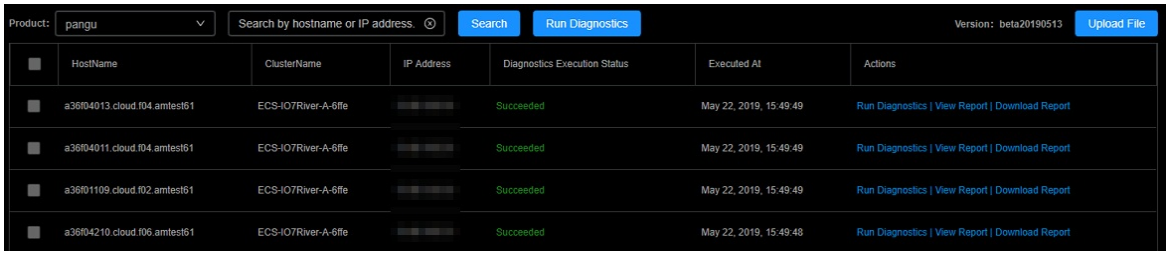

5. After the diagnostics are complete, click **View Report** in the **Actions** column to view the diagnostic result.
6. (Optional)After the diagnostics are complete, click **Download Report** in the **Actions** column to download the diagnostic results to your local machine.

5.4.4.3. Update Support Tools

When the Support Tools toolkit has updates, you can update it to the latest version by uploading files.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Support Tools > Support Tools**.
3. In the upper-right corner of the page, click **Upload File**.
4. Select the toolkit file to upload, enter the verification code, and click **Upload File**. Contact level-2 support engineers to obtain the verification code.

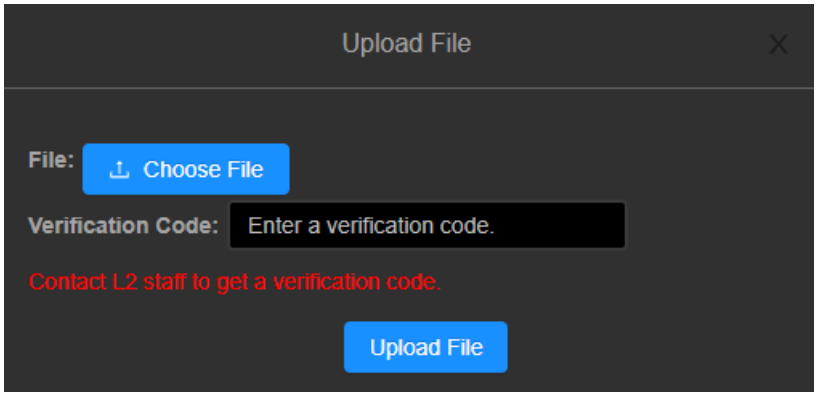

5.4.4.4. Diagnose with inspection tools

You can use inspection tools to diagnose and inspect services, such as Apsara File Storage NAS (NAS), Block Storage, and Apsara Name Service and Distributed Lock Synchronization System.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Support Tools > Inspection Tool**.
3. Select the target service from the Product drop-down list and click **Search**. The search result appears in the section below. Apsara Stack Doctor (ASD) supports diagnostics for services, including NAS, Block Storage, and Apsara Name Service and Distributed Lock Synchronization System.
   - NAS diagnostics

     It allows you to collect NAS information, including disk status, KV (key-value) status, KV server spacing, version, recycle bin, memory, and TCP.
   - EBS diagnostics

     It allows you to collect the utilization information about storage clusters.
   - Diagnostics of Apsara Name Service and Distributed Lock Synchronization System

     It allows you to check the following information about this service:

     - The health status of the E2E service link.
     - The disk space of the service.
     - Whether the nuwazk log is properly stored.
     - Whether the nuwaproxy log is properly stored.
4. You can select multiple machines and click **Run Diagnostics** to perform batch diagnosis. Alternatively, you can select only one machine and click **Run Diagnostics** in the **Actions** column corresponding to the machine.

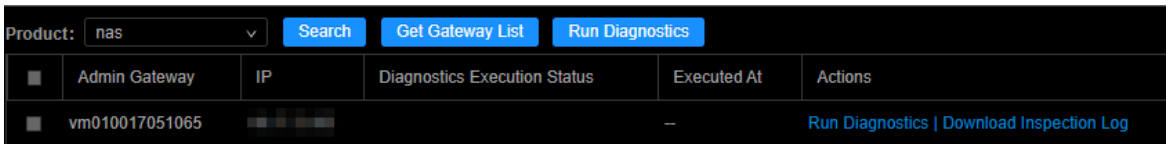

5. After the diagnostics is complete, you can click **Download Report** in the **Actions** column corresponding to the machine to download the diagnostic results to your local machine.

5.4.4.5. Upload script files for EDAS diagnostics

Before the diagnostics, you can unload script files to be executed for server roles.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Support Tools > EDAS Diagnostics**.
3. In the upper-right corner of the page, click **Upload Diagnostic Script**.
4. Select the product, service, and server role.

   If the server role has script files, the script files will be displayed in the **Existing Scripts** field. You can click the name of a script file to view details.

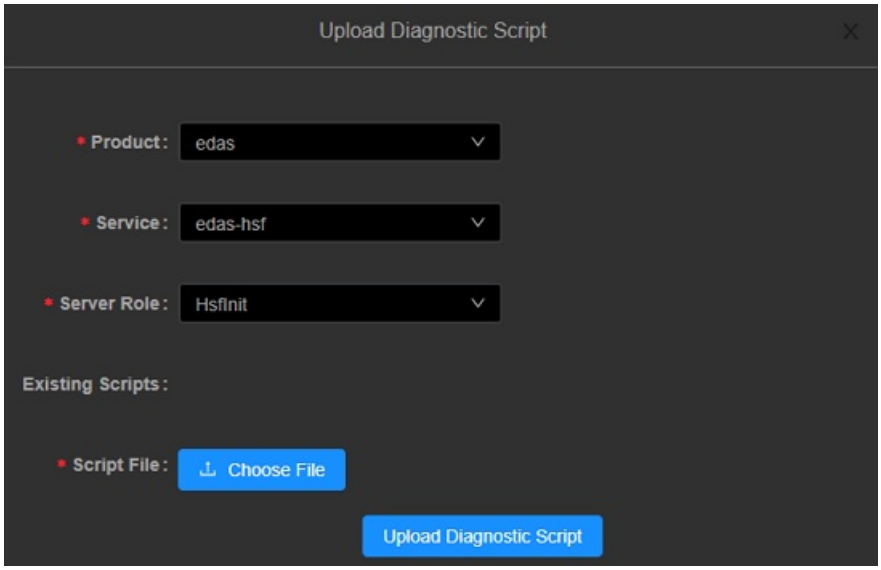

5. Click **Choose File**. In the dialog box that appears, select the script file to be uploaded. Click **Open** to add the script file to be uploaded.
6. Click **Upload Diagnostic Script**.

5.4.4.6. EDAS diagnostics

The EDAS diagnostics tool allows you to inspect EDAS.

Prerequisites

Before the diagnosis, make sure that the server role to be diagnosed has an executable script file. If not, you need to upload the script file to be executed for the server role. For more information about how to upload the script file, see Upload script files for EDAS diagnostics.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Support Tools > EDAS Diagnostics**.
3. (Optional)Select one or more services from the Service drop-down list and click **Refresh**. The filtered services appear in the section below.

4. Find the server role to diagnose, and click **Run Diagnostics** in the **Actions** column corresponding to the server role.

   > **Note** You can select multiple server roles at a time from the filtered services and click **Run Diagnostics**. In the dialog box that appears, click **OK** to run diagnostics.

   When **Diagnostic Status** changes from **Diagnosing** to **Diagnostics Succeeded**, the tasks are completed.

5. After the tasks are completed, you can click **Download Report** in the **Actions** column corresponding to the server role to download the original diagnostic information.

5.4.5. Service Availability

5.4.5.1. View Service Availability

Service Availability allows you to view the availability statuses of cloud services in Apsara Stack.

Context

It is used to verify the continuity of these cloud services.

During the hot upgrade of a service, you can use Service Availability to check whether the upgrade causes a service interruption, helping you detect and solve problems in a timely manner.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Service Availability > Service Availability**.
3. In the search bar, select the service you want to view and click **Search** to view its service status. The following table describes the service statuses.

Service status	Description
Pending	The service availability inspection is not enabled for this service.
UNKNOW	The service availability status of the service is unknown.
ERROR	The service availability status of the service is abnormal.
OK	The service availability status of the service is normal.

5.4.5.2. View Control Service Availability

The Control Service Availability page displays the statistics of the global environment, product response times, and product QPS.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Service Availability > Control Service Availability**.
3. View the following information:
   - Global statistics

     Global Statistics displays the environment information of all control gateways, including global queries per second (QPS), global response time statistics, and error details.

     On the **Global Statistics** tab, select **Last 1 Hour**, **Last 2 Hours**, **Last 24 Hours**, or **Select Time** from the Time drop-down list and select HTTP status code. Click **Update** to view the information of the global environment within the specified time range.

     The following table describes the HTTP status codes.

HTTP status code	Description
200	The request is successful. It is generally used for GET and POST requests.
400	The syntax of the request from the client is incorrect, which cannot be understood by the server.
403	The server understands the request from the client but refuses to execute it.
404	The server cannot find the resource based on the request from the client.
500	The request cannot be completed because the server has an internal error.
503	The server is temporarily unable to process the request from the client.
201	Created. The request is successful, and a new resource is created.
204	No content. The server has processed the request but does not return any content.

HTTP status code	Description
409	A conflict occurs when the server processes the request.
202	Accepted. The request has been accepted but has not been processed.
405	The method specified in the request from the client is forbidden.

- Product response time statistics

  Product Response Time Statistics displays the latency of each service from a specified period of time. You can view product response time statistics to identify whether exceptions have occurred in a service API based on the number of responses within a specified period of time.

  On the **Product Response Time Statistics** tab, select **Last 1 Hour**, **Last 2 Hours**, **Last 24 Hours**, or **Select Time** from the Time drop-down list, Product to be queried, and HTTP Status Code. Click **Update** to view the average latency of a service within a specified period of time.

- Product QPS statistics

  Product QPS statistics displays the requests of each service within a specified period of time. You can view product QPS statistics to identify whether exceptions have occurred in the service status based on the number of requests within a specified period of time.

  On the **Product QPS Statistics** tab, select **Last 1 Hour**, **Last 2 Hours**, **Last 24 Hours**, or **Select Time** from the Time drop-down list, Product to be queried, and HTTP Status Code. Click **Update** to view the latency of a service from a specified period of time.

5.4.6. Monitoring

The Monitoring module allows you to view alert templates, alerts, and alert status in the system.

5.4.6.1. View alert templates

Alert templates are used to configure alert monitoring settings. You can filter alert template content by service and template.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Monitoring > Monitoring Templates**.
3. In the search bar, select a service and a template, and click **Search**.
4. View the alert template content in the search result.

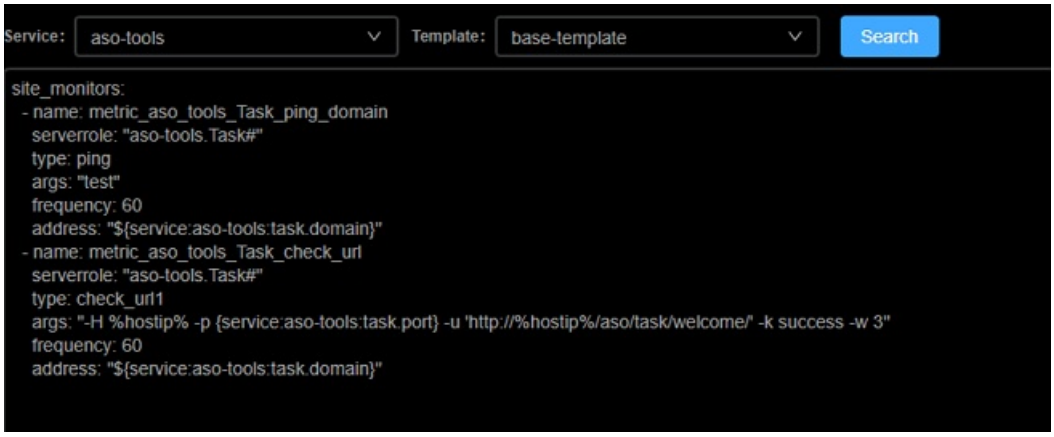

5.4.6.2. View alert information

During routine O&M, you can view alert information to obtain up-to-date information about services. When a service fails, you can filter out the alert information that you need based on the service, cluster name, and alert name to quickly resolve the failure.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Monitoring > Alerts**.
3. Perform the following operations:
   - To view all alerts in the system, click **Search** without selecting any filters.

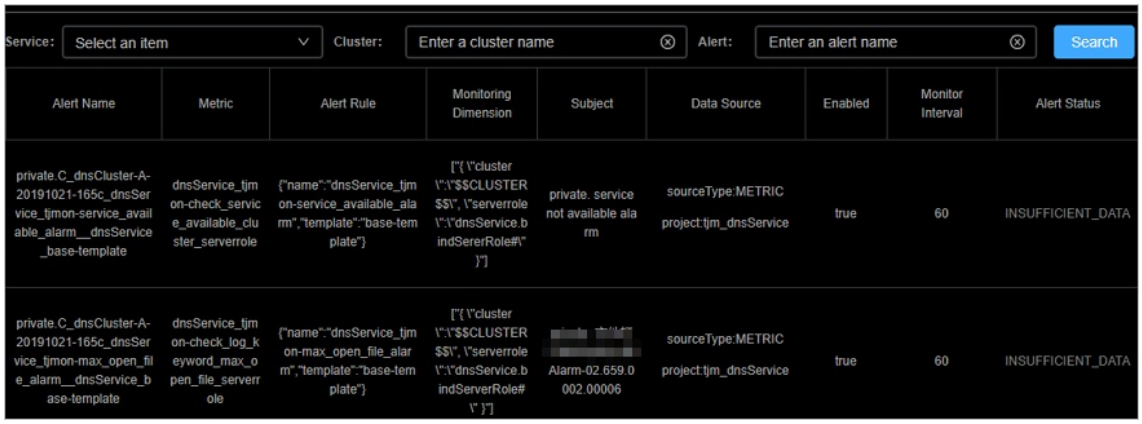

   - In the search bar, select a service, enter a cluster name and an alert name, and then click **Search** to view information about an alert.

5.4.6.3. View the alert status

After alerts are triggered, you can view the status of all alerts in the system.

Procedure

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **Monitoring > Alert Status**.
3. Perform the following operations:
   - To view the status of all alerts in the system, click **Search** without selecting any filters.

- To view the status of all alerts in the system, click **Search** without selecting any filters.

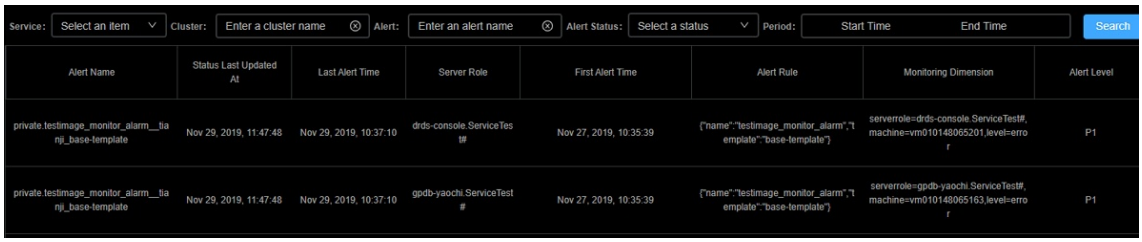

- In the search bar, select a service, enter a cluster name and an alert name, and select a status and a time range. Then, click **Search** to view the status of an alert.

5.5. Apsara Infrastructure Management Framework

5.5.1. Old version

5.5.1.1. What is Apsara Infrastructure Management Framework?

This topic describes what Apsara Infrastructure Management Framework is from the aspects of core functions and basic concepts.

5.5.1.1.1. Overview

Apsara Infrastructure Management Framework is a distributed data center management system, which manages applications on clusters containing multiple machines and provides the basic functions such as deployment, upgrade, expansion, contraction, and configuration changes.

Apsara Infrastructure Management Framework also supports monitoring data and analyzing reports, which facilitates users to perform one-stop Operation & Maintenance (O&M) control. Based on the Apsara Infrastructure Management Framework services, automated O&M is implemented in the large-scale distributed environment, which greatly improves the operations efficiency and enhances the system availability.

Apsara Infrastructure Management Framework is mainly composed of TianjiMaster and TianjiClient. Apsara Infrastructure Management Framework installs TianjiClient as the agent on machines it manages. Then, TianjiMaster accepts and issues the upper-layer instructions to TianjiClient for execution. In the upper layer, Apsara Infrastructure Management Framework is divided into different components based on different functions, and then provides API server and portal for external use.

Core functions

- Network initialization in data centers
- Server installation and maintenance process management
- Deployment, expansion, and upgrade of cloud products
- Configuration management of cloud products
- Automatic application for cloud product resources

- Automatic repair of software and hardware faults
- Basic monitoring and business monitoring of software and hardware

5.5.1.1.2. Basic concepts

Before using Apsara Infrastructure Management Framework, you must know the following basic concepts for a better understanding.

project

A collection of clusters, which provides service capabilities for external entities.

cluster

A collection of physical machines, which logically provides services and is used to deploy project software.

- A cluster can only belong to one project.
- Multiple services can be deployed on a cluster.

service

A set of software, which provides relatively independent functions. A service is composed of one or more server roles and can be deployed on multiple clusters to form multiple sets of services and provide the corresponding service capabilities. For example, Apsara Distributed File System, Job Scheduler, and Apsara Name Service and Distributed Lock Synchronization System are all services.

service instance

A service that is deployed on a cluster.

server role

One or more indivisible deployment units into which a service can be divided based on functions. A server role is composed of one or more specific applications. If a service is deployed on a cluster, all the server roles of the service must be deployed to machines of this cluster. Multiple server roles, such as PanguMaster and TianjiClient, can be deployed on the same server.

server role instance

A server role that is deployed on a machine. A server role can be deployed on multiple machines.

application

Applications correspond to each process-level service component in a server role and each application works independently. The application is the minimum unit for deployment and upgrade in Apsara Infrastructure Management Framework, and can be deployed to each machine. Generally, an application is an executable software or Docker container.

If a server role is deployed on a machine, all applications in the server role must be deployed to this machine.

rolling

Each time when a user updates configurations, Apsara Infrastructure Management Framework upgrades services and modifies the cluster configurations based on the updated configurations. This process is called rolling.

service configuration template

Some configurations are the same when services are deployed on clusters. A service configuration template can be created to quickly write the same configurations to different clusters. The service configuration template is basically used for large-scale deployment and upgrade.

associated service template

A *template.conf* file that exists in the configurations. This file declares the service configuration template and its version, of which the configuration is used by the service instance.

final status

If a cluster is in this status, all hardware and software on each of its machines are normal and all software are in the target version.

dependency

The dependency between server roles in a service defines that server roles with dependencies run tasks or are upgraded based on the dependency order. For example, if A depends on B, B is upgraded first. A starts to be downloaded after B is downloaded successfully, and upgraded after B is successfully upgraded. (By default, the dependency does not take effect for configuration upgrade.)

upgrade

A way of aligning the current status with the final status of a service. After a user submits the version change, Apsara Infrastructure Management Framework upgrades the service version to the target version. With the server role as the processing unit, upgrade aims to update the versions of all machines to the target version.

At the beginning, the final status and current status of the cluster are the same. When a user submits the change, the final status is changed, whereas the current status is not. A rolling task is generated and has the final status as the target version. During the upgrade, the current status is continuously approximating to the final status. Finally, the final status and the current status are the same when the upgrade is finished.

5.5.1.2. Log on to the Apsara Infrastructure Management Framework console

This topic describes how to log on to the Apsara Infrastructure Management Framework console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

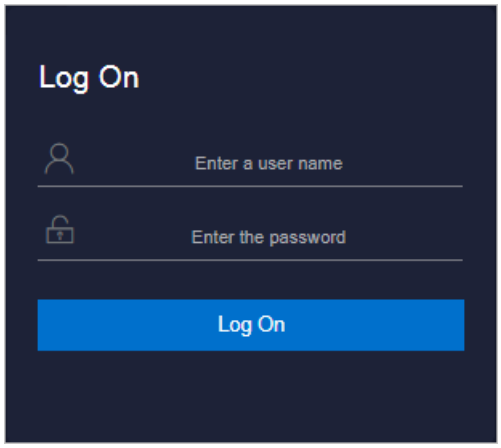

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products** > Product List.
6. In the **Apsara Stack O&M** section, choose Basic O&M > Apsara Infrastructure Management Framework.

5.5.1.3. Web page introduction

Before performing Operation & Maintenance (O&M) operations on Apsara Infrastructure Management Framework, you must have a general understanding of the Apsara Infrastructure Management Framework page.

5.5.1.3.1. Instructions for the homepage

After you log on to the Apsara Infrastructure Management Framework console, the homepage appears. This topic describes the basic operations and functions on the homepage.

Log on to Apsara Infrastructure Management Framework. The homepage appears, as shown in Homepage of the Apsara Infrastructure Management Framework console.

Homepage of the Apsara Infrastructure Management Framework console

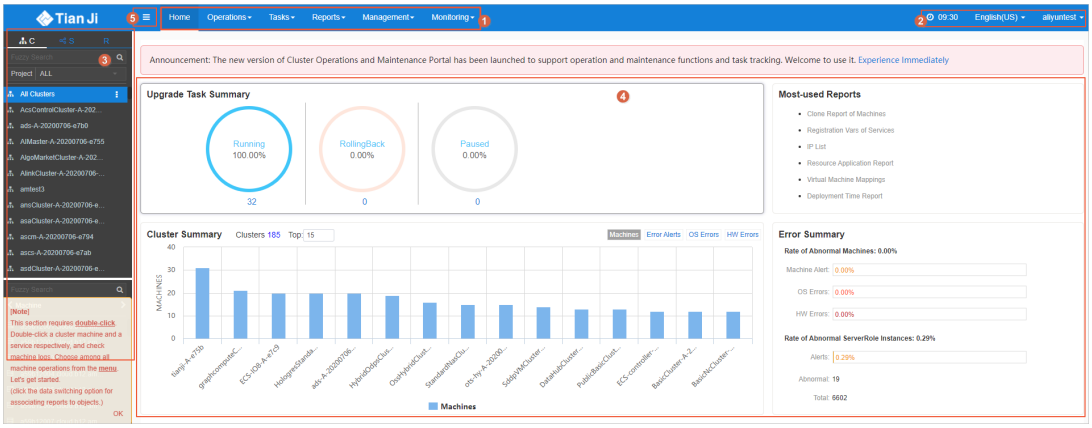

Description of functional sections describes the functional sections on the homepage.

Description of functional sections

Section		Description

①	Top navigation bar	• **Operations**: the quick entrance to operations and maintenance (O&M) operations and their objects. This menu consists of the following submenus:
		  ◦ **Cluster Operations**: allows you to use the project permissions to perform O&M and management operations on clusters. For example, you can view the cluster status.
		  ◦ **Service Operations**: allows you to use the service permissions to manage services. For example, you can view the service list.
		  ◦ **Machine Operations**: allows you to perform O&M and management operations on machines. For example, you can view the machine status.
		• **Tasks**: Rolling tasks are generated when you modify the configurations in the system. This menu allows you to view the running tasks, task history, and deployment of clusters, services, and server roles in all projects.
		• **Reports**: allows you to view monitoring data in tables and find specific reports by using fuzzy search.
		• **Monitoring**: monitors metrics during system operations and sends alert notifications for abnormal conditions. This menu allows you to view the alert status, modify alert rules, and search alert history.

Section		Description
②	Upper-right buttons	•  : ◦ **TJDB Synchronization Time**: the time when the data on the current page is generated. ◦ **Final Status Computing Time**: the time when the desired-state data on the current page is calculated. The system processes data as fast as it can after the data is generated. Latency exists because Apsara Infrastructure Management Framework is an asynchronous system. Time information helps explain why data on the current page is generated and determine whether the system experiences an error. • 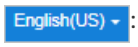 : the current display language of the console. You can select another language from the drop-down list. •  : your logon account. You can select **Logout** from the drop-down list to log out of your account.
③	Left-side navigation pane	In the left-side navigation pane, you can view the logical architecture of Apsara Infrastructure Management Framework. The tabs allow you to view details and perform operations. For more information, see Introduction on the left-side navigation pane.
④	Workspace	The workspace shows a summary of tasks and other information. • **Upgrade Task Summary**: shows the numbers and proportions of running, rolling back, and suspended upgrade tasks. • **Cluster Summary**: shows the numbers of machines, error alerts, operating system errors, and hardware errors in each cluster. • **Error Summary**: shows metric values about the rate of abnormal machines and the rate of abnormal server role instances. • **Most-used Reports**: shows links of common statistical reports.
⑤	Show/hide button	If you do not need to use the left-side navigation pane, click this button to hide the pane and enlarge the workspace.

5.5.1.3.2. Instructions for the left-side navigation pane

The left-side navigation pane contains three tabs: **C** (cluster), **S** (service), and **R** (report). This topic describes how to use the tabs to view information.

Cluster

You can search for clusters in a project and their information such as the cluster status, cluster operations and maintenance (O&M), service desired state, and logs by fuzzy match.

On the **C** tab of the left-side navigation pane, you can perform the following operations:

- Enter a cluster name or a part of a cluster name in the search box to filter clusters.
- Select a project from the **Project** drop-down list to view all clusters in the project.
- Move the pointer over the icon next to a cluster and select menu items to perform corresponding operations on the cluster.

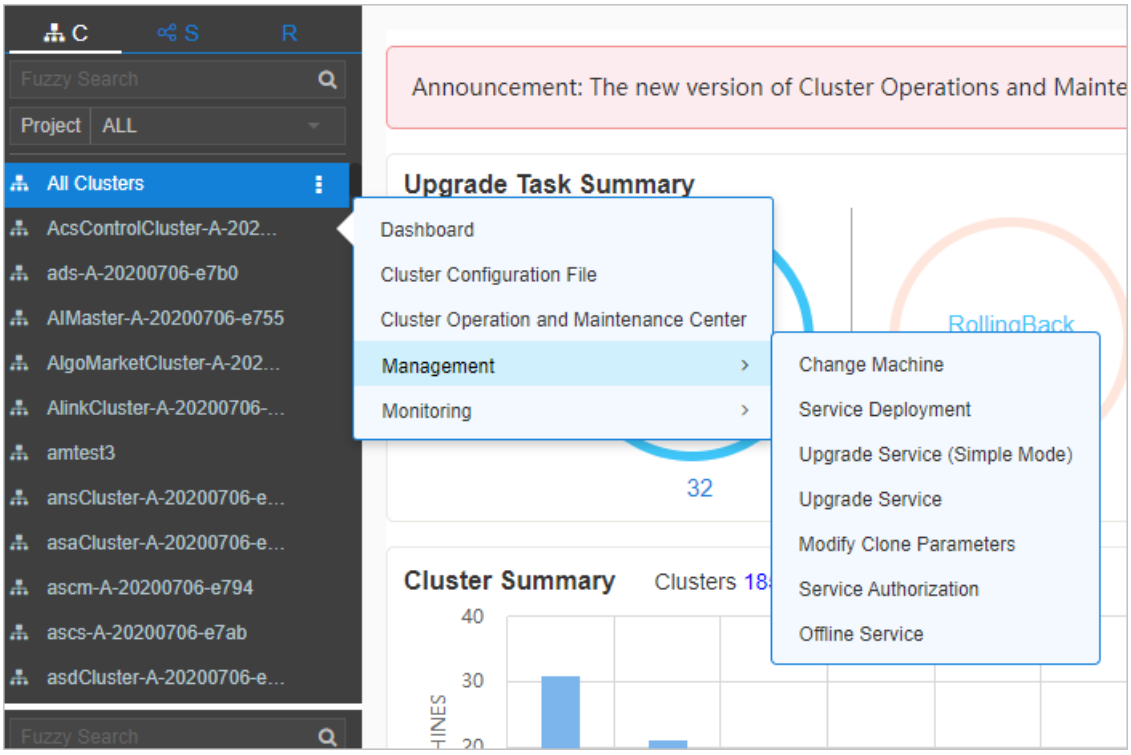

- Click a cluster. All machines and services within the cluster are displayed in the lower part of the left-side navigation pane. Move the pointer over the icon next to a machine or service on the **Machine** or **Service** tab and select menu items to perform corresponding operations on the machine or service.

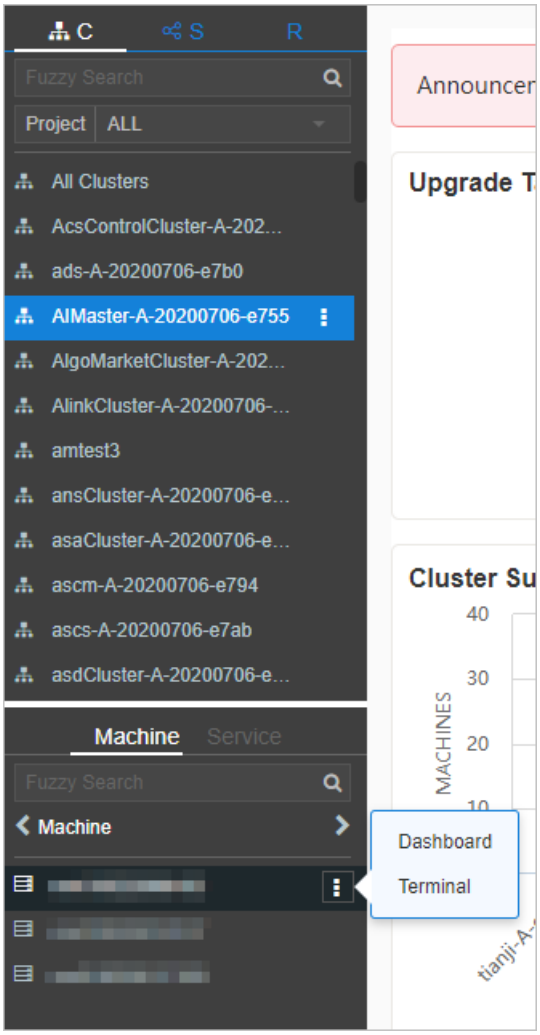

- Click the **Machine** tab. Double-click a machine to view information about all server roles on the machine. Double-click a server role to view applications, and then double-click an application to view log files.
- Click the **Service** tab. Double-click a machine to view information about all server roles on the machine. Double-click a server role to view machines, double-click a machine to view applications, and then double-click an application to view log files.
- Double-click a log file. Move the pointer over the log file, click the icon next to the log file, and then click **Download** to download the log file.

  Alternatively, move the pointer over a log file and click **View** next to the log file. The time-ordered log details are displayed on the **Log Viewer** page. You can search for log details by keyword.

Service

You can search for services and view information about services and service instances by fuzzy match.

On the **S** tab of the left-side navigation pane, you can perform the following operations:

- Enter a service name or a part of a service name in the search box to filter services.
- Move the pointer over the icon next to a service and select menu items to perform corresponding operations on the service.

- Click a service. All service instances within the service are displayed in the lower part of the left-side navigation pane. Move the pointer over the icon next to a service instance and select menu items to perform corresponding operations on the service instance.

Report

You can search for reports by fuzzy match and view report details.

On the **R** tab of the left-side navigation pane, you can perform the following operations:

- Enter a report name or a part of a report name in the search box to filter reports.
- Click **All Reports** or **Favorites**. Corresponding groups are displayed in the lower part of the left-side navigation pane. Double-click a group to view all reports in the group. Double-click a report to view details of the report.

5.5.1.4. Cluster operations

This topic describes the actions about cluster operations.

5.5.1.4.1. View configuration information of a cluster

This topic describes how to view the basic information, deployment plan, and configuration information of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Cluster Operations**.The **Cluster Operations** page contains the following information:
   - Cluster

     The name of a cluster. Click a cluster name to go to the Cluster Dashboard page. For more information, see View dashboard information of a cluster.
   - Scale-Out/Scale-In

     The numbers of machines and server roles that are scaled in and out. Click a number to go to the Cluster Operation and Maintenance Center page. For more information, see View information of the cluster O&M center.
   - Abnormal Machine Count

     The number of machines that are not in the Good state within a cluster. Click the number to go to the Cluster Operation and Maintenance Center page. For more information, see View information of the cluster O&M center.
   - Final Status of Normal Machines

     Specifies whether a cluster has reached the desired state. Select **Clusters not Final** above the cluster list to view all clusters that have not reached the desired state. Click a link in the column to view desired state information. For more information, see View the desired state of a service.
   - Rolling

Specifies whether rolling tasks are running within a cluster. Select **Rolling Tasks** above the cluster list to view all clusters that have rolling tasks. Click rolling in the column to view rolling tasks. For more information see View rolling tasks.

3. (Optional)Select a project from the drop-down list or enter a cluster name to search for the cluster.
4. Click the cluster name or click **Cluster Configuration** in the **Actions** column to go to the **Cluster Configuration** page.

   Cluster configuration description describes the parameters on the **Cluster Configuration** page.

Cluster configuration description

Section	Parameter	Description
Basic Information	**Cluster**	The name of the cluster.
	Project	The project to which the cluster belongs.
	Clone Switch	○ **Pseudo-clone**: The system is not cloned when a machine is added to the cluster. ○ **Real Clone**: The system is cloned when a machine is added to the cluster.
	Machines	The number of machines included in the cluster. Click View Clustering Machines to view the list of machines.
	Security Verification	The access control among processes. By default, security verification is disabled in non-production environments. You can enable or disable security verification based on your business requirements.
	Cluster Type	○ RDS ○ NETFRAME ○ T4: a type of cluster that renders special configurations for the mixed deployment of e-commerce ○ Default
Deployment Plan	**Service**	The service that is deployed within the cluster.
	Dependency Service	The service on which the current service depends.
	Service Information	The service that you want to view. Select a service from the drop-down list to view its configuration information.
	Service Template	The template that is used by the service.

Section	Parameter	Description
Service Information	Monitoring Template	The monitoring template that is used by the service.
	Machine Mappings	The machines where server roles of the service are deployed.
	Software Version	The version of the software that is included in server roles of the service.
	Availability Configuration	The percentage of availability configuration for server roles of the service.
	Deployment Plan	The deployment plan of server roles of the service.
	Configuration Information	The configuration file that is used for the service.
	Role Attribute	The server roles and their parameter information.

5. Click **Operation Logs** in the upper-right corner to view version differences. For more information about operation logs, see View operation logs.

5.5.1.4.2. View dashboard information of a cluster

This topic describes how to view the basic information and related statistics of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the **Cluster Dashboard** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and select **Dashboard**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click the name of the target cluster.
3. View all information about the cluster on the **Cluster Dashboard** page. The following table describes the information that you can view, such as basic information, desired state information, rolling tasks, dependencies, resources, virtual machine (VM) mappings, and monitoring status.

Parameter	Description

Parameter	Description
Basic Cluster Information	The basic information about the cluster. ○ **Project Name**: the name of the project. ○ **Cluster Name**: the name of the cluster. ○ **IDC**: the data center to which the cluster belongs. ○ **Final Status Version**: the latest version of the cluster. ○ **Cluster in Final Status**: specifies whether the cluster has reached the desired state. ○ **Machines Not In Final Status**: the number of machines that have not reached the desired state. ○ **Real/Pseudo Clone**: specifies whether the system is cloned when a machine is added to the cluster. ○ **Expected Machines**: the number of machines that are expected within the cluster. ○ **Actual Machines**: the number of machines that are deployed in the current environment. ○ **Machines Not Good**: the number of machines that are not in the Good state within the cluster. ○ **Actual Services**: the number of services that are deployed within the cluster. ○ **Actual Server Roles**: the number of server roles that are deployed within the cluster. ○ **Cluster Status**: specifies whether the cluster is starting or shutting down machines.
Machine Status Overview	The status of machines within the cluster.
Machines In Final State	The distribution of machines where services are deployed, based on whether the machines have reached the desired state.
Load-System	The statistics chart of the cluster system load.
CPU-System	The statistics chart of the CPU load.
Mem-System	The statistics chart of the memory load.
Disk_Usage-System	The statistics chart of the disk usage.
Traffic-System	The statistics chart of the system traffic.
TCP State-System	The statistics chart of the CPU request status.
TCP Retrans-System	The statistics chart of the CPU retransmission traffic.
Disk_IO-System	The statistics chart of the disk I/O information.

Parameter	Description
Service Instances	The service instances that are deployed within the cluster and their desired state information. ○ **Service Instance**: the service instance that is deployed within the cluster. ○ **Final Status**: specifies whether the service instance has reached the desired state. ○ **Expected Server Roles**: the number of server roles that are expected to deploy in the service instance. ○ **Server Roles in Final Status**: the number of server roles that have reached the desired state in the service instance. ○ **Server Roles Going Offline**: the number of server roles that are being unpublished from the service instance. ○ Actions: Click **Details** to go to the **Service Instance Information Dashboard** page. For more information about the service instance dashboard, see View the service instance dashboard.
Upgrade Tasks	The upgrade tasks within the cluster. ○ **Cluster Name**: the name of the cluster. ○ **Type**: the type of the upgrade task. Valid values: app and config. app indicates version upgrade, and config indicates configuration change. ○ **Git Version**: the change version of the upgrade task. ○ **Description**: the description of the change. ○ **Rolling Result**: the result of the upgrade task. ○ **Submitted By**: the user who submits the change. ○ **Submitted At**: the time when the change is submitted. ○ **Start Time**: the time when rolling starts. ○ **End Time**: the time when the upgrade task ends. ○ **Time Used**: the time consumed for the upgrade. ○ Actions: Click **Details** to go to the **Rolling Task** page. For more information about rolling tasks, see View rolling tasks.
Cluster Resource Request Status	○ **Version**: the version of the resource request. ○ **Msg**: the error message. ○ **Begintime**: the time when the resource request analysis starts. ○ **Endtime**: the time when the resource request analysis ends. ○ **Build Status**: the build status of resources. ○ **Resource Process Status**: the resource request status of the version.

Parameter	Description
Cluster Resource	○ **Service**: the name of the service. ○ **Service Role**: the name of the server role. ○ **App**: the name of the application of the server role. ○ **Name**: the name of the resource. ○ **Type**: the type of the resource. ○ **Status**: the status of the resource request. ○ **Error_Msg**: the error message. ○ **Parameters**: the parameters of the resource. ○ **Result**: the result of the resource request. ○ **Res**: the ID of the resource. ○ **Reprocess Status**: the request status of AnyTunnel VIP addresses. ○ **Reprocess Msg**: the error message reported when AnyTunnel VIP addresses are requested. ○ **Reprocess Result**: the request result of AnyTunnel VIP addresses. ○ **Refer Version List**: the version that uses the resource.
VM Mappings	The VMs within the cluster. VM information is displayed only when VMs are deployed within the cluster. ○ **VM**: the hostname of the VM. ○ **Currently Deployed On**: the hostname of the physical machine where the VM is deployed. ○ **Target Deployed On**: the hostname of the physical machine where you expect to deploy the VM.
Service Dependencies	The dependency configuration of service instances and server roles within the cluster, and the desired state information of dependency services or server roles. ○ **Service**: the name of the service. ○ **Server Role**: the name of the server role. ○ **Dependent Service**: the service on which the server role depends. ○ **Dependent Server Role**: the server role on which the server role depends. ○ **Dependent Cluster**: the cluster where the dependency server role is deployed. ○ **Dependency in Final Status**: specifies whether the dependency server role has reached the desired state.

5.5.1.4.3. View information of the cluster O&M center

This topic describes how to view the status and statistics of services and machines within a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the **Cluster Operation and Maintenance Center** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and select **Cluster Operation and Maintenance Center**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and choose **Monitoring > Cluster Operation and Maintenance Center** in the Actions column.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click the name of the target cluster. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**.
3. View information on the **Cluster Operation and Maintenance Center** page.

Parameter	Description
SR not in Final Status	All server roles that have not reached the desired state within the cluster. Click the number to view the list of server roles. Click a server role to view information of machines where the server role is deployed.
Running Tasks	Specifies whether rolling tasks are running within the cluster. Click **Rolling** to go to the **Rolling Task** page. For more information about rolling tasks, see View rolling tasks.
Head Version Submitted At	The time when the HEAD version is submitted. Click the time to view details.
Head Version Analysis	The status of desired state analysis. During desired state analysis, Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to specific change contents. Desired state analysis can be in one of the following states: ○ **Preparing**: No new version is detected. ○ **Waiting**: The latest version has been detected, but the analysis module has not started. ○ **Doing**: The application to be changed is being analyzed. ○ **done**: The desired state analysis succeeds. ○ **Failed**: The desired state analysis fails to parse change contents. Apsara Infrastructure Management Framework can obtain change contents of server roles in the latest version only when the desired state analysis is in the **done** state. Click a state to view related information.
Service	The service deployed within the cluster. Select a service from the drop-down list.

Parameter	Description
Server Role	The server role of a service within the cluster. Select a server role from the drop-down list. **Note** After you select a service and a server role, machines that are related to the service or the server role are displayed.
Total Machines	The total number of machines within the cluster or machines where the selected server roles are deployed.
Scale-in Scale-out	The numbers of machines and server roles that are scaled in and out.
Abnormal Machines	The numbers of machines in an abnormal state for the following reasons: ◦ **Ping Failed**: the number of machines that experience ping_monitor errors because TianjiMaster cannot ping the machines. ◦ **No Heartbeat**: the number of machines that experience TianjiClient or network errors because TianjiClient does not report data on a regular basis. ◦ **Status Error**: the number of machines that experience critical or fatal errors. Resolve problems based on alert information.
Abnormal Services	The number of machines that have abnormal services. The following rules are used to check whether a service has reached the desired state: ◦ Each server role on the machine is in the GOOD state. ◦ The actual version of each application of each server role on the machine is consistent with the HEAD version. ◦ Before the Image Builder builds an application of the HEAD version, Apsara Infrastructure Management Framework cannot obtain the value of the HEAD version, and the desired state of the service is unknown. This process is called change preparation. The desired state of the service cannot be obtained when the preparation process is in progress or if the preparation fails.

Parameter	Description
Machines	All machines within the cluster or machines where the selected server roles are deployed. ○ Click the Machine Search search box. In the dialog box that appears, enter one or more machines. Fuzzy match and batch search are supported. ○ Click the name of a machine to view its physical information in the Machine Information dialog box. Click **DashBoard** to go to the **Machine Details** page. For more information about machine details, see View the machine dashboard. ○ Move the pointer over the **Final Status** or **Final SR Status** column and click **Details** to view the machine status and system service information, as well as status information and error messages of server roles on the machine. ○ Before you filter machines by service and service role, move the pointer over the **Running Status** column and click **Details** to view status information and error messages of the machine. After you filter machines by service and service role, move the pointer over the **SR Running Status** column and click **Details** to view status information and error messages of server roles on the machine. ○ Click **Error**, **Warning**, or **Good** in the **Monitoring Statistics** column to view machine and server role metrics. ○ Click **Terminal** in the **Actions** column to log on to the machine and perform operations. ○ Click **Machine Operation** in the **Actions** column to perform reboot, out-of-band reboot, or reclone operations on the machine.

5.5.1.4.4. View the desired state of a service

This topic describes how to check whether a service within a cluster has reached the desired state and how to view desired state details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the **Service Final Status Query** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and choose **Monitoring > Service Final Status Query**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and choose **Monitoring > Service Final Status Query** in the Actions column.
3. View information on the **Service Final Status Query** page.

Parameter	Description
Project Name	The project to which the cluster belongs.

Parameter	Description
Cluster Name	The name of the cluster.
Head Version Submitted At	The time when the HEAD version is submitted.
Head Version Analysis	The status of desired state analysis. During desired state analysis, Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to specific change contents. Desired state analysis can be in one of the following states: ○ **Preparing**: No new version is detected. ○ **Waiting**: The latest version has been detected, but the analysis module has not started. ○ **Doing**: The application to be changed is being analyzed. ○ **done**: The desired state analysis succeeds. ○ **Failed**: The desired state analysis fails to parse change contents. Apsara Infrastructure Management Framework can obtain change contents of server roles in the latest version only when the desired state analysis is in the **done** state.
Cluster Rolling Status	Specifies whether the cluster has reached the desired state. If a rolling task is running, its task information is displayed.
Cluster Machine Final Status Statistics	The status of all machines within the cluster. Click **View Details** to go to the **Cluster Operation and Maintenance Center** page and view machine details. For more information about the operations and maintenance (O&M) center, see View the cluster operation and maintenance center.
Final Status of Cluster SR Version	The desired state of services within the cluster. **Note** This section includes only the services that have not reached the desired state due to version inconsistency or status exceptions. For other services that fail to reach the desired state due to machine errors, see desired state information of machines within the cluster.
Final Status of SR Version	The number of machines that have not reached the desired state. The number is displayed if server roles have rolling tasks.

5.5.1.4.5. View operations logs

This topic describes how to view differences between Git versions from operation logs.

Procedure

1. Log on to Apsara Infrastructure Management Framework.

2. Use one of the following methods to go to the **Cluster Operation Logs** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and choose **Monitoring > Operation Logs**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and choose **Monitoring > Operation Logs** in the Actions column.
3. On the **Cluster Operation Logs** page, click **Refresh** in the upper-right corner to view the Git version, description, submission information, and task status.
4. (Optional)On the **Cluster Operation Logs** page, view differences between versions.
   i. Find the target operation log and click **View Release Changes** in the **Actions** column.
   ii. On the **Version Difference** page, configure the following parameters:
      - **Select Base Version**: Select a basic version.
      - **Configuration Type**: Select **Extended Configuration** or **Cluster Configuration**. **Extended Configuration** allows you to view differences between the merging results of cluster and template configurations. **Cluster Configuration** allows you to view differences between cluster configurations.
   iii. Click **Obtain Difference**.

   Difference files are displayed.
   iv. Click each difference file to view its difference details.

5.5.1.5. Service operations

This topic describes the actions about service operations.

5.5.1.5.1. View the service list

The service list allows you to view the list of all services and the related information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Service Operations**.
3. View the information on the **Service Operations** page.

Item	Description
Service	The service name.
Service Instances	The number of service instances in the service.
Service Configuration Templates	The number of service configuration templates.

Item	Description
Monitoring Templates	The number of monitoring templates.
Service Schemas	The number of service configuration validation templates.
Actions	Click **Management** to view the service instances, service templates, monitoring templates, monitoring instances, service schemas, and detection scripts.

5.5.1.5.2. View dashboard information of a service instance

This topic describes how to view the basic information and related statistics of a service instance.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, click the **S** tab.
3. (Optional)Enter a service name in the search box to search for the service.
4. Click the service name to view service instances of the service.
5. Move the pointer over the icon next to the target service instance and select **Dashboard**.
6. View information on the **Service Instance Information Dashboard** page.

Parameter	Description

Parameter	Description
Service Instance Summary	The basic information about the service instance. ◦ **Cluster Name**: the name of the cluster where the service instance is deployed. ◦ **Service Name**: the name of the service to which the service instance belongs. ◦ **Actual Machines**: the number of machines that are deployed in the current environment. ◦ **Expected Machines**: the number of machines that are expected for the service instance. ◦ **Target Total Server Roles**: the number of server roles that are expected for the service instance. ◦ **Actual Server Roles**: the number of server roles that are deployed in the current environment. ◦ **Template Name**: the name of the service template that is used by the service instance. ◦ **Template Version**: the version of the service template that is used by the service instance. ◦ **Schema**: the name of the service schema that is used by the service instance. ◦ **Monitoring System Template**: the name of the Monitoring System template that is used by the service instance.
Server Role Statuses	The status of server roles in the service instance.
Machine Statuses for Server Roles	The status of machines where server roles are deployed.
Service Monitoring Information	◦ **Monitored Item**: the name of the metric. ◦ **Level**: the level of the metric. ◦ **Description**: the description of the metric. ◦ **Updated At**: the time when the data is updated.
Service Alert Status	◦ **Alert Name** ◦ **Instance Information** ◦ **Alert Start** ◦ **Alert End** ◦ **Alert Duration** ◦ **Severity Level** ◦ **Occurrences**: the number of occurrences of the alert.

Parameter	Description
Server Role List	○ Server Role ○ Current Status ○ Expected Machines ○ Machines In Final Status ○ Machines Going Offline ○ Rolling Task Status ○ Time Used: the time that is used for the execution of rolling tasks. ○ Actions: Click Details to go to the View the server role dashboard page.
Service Alert History	○ Alert Name ○ Alert Time ○ Instance Information ○ Severity Level ○ Contact Group
Service Dependencies	The dependency configuration of service instances and server roles, and the desired state information of dependency services or server roles. ○ Server Role: the name of the server role. ○ Dependent Service: the service on which the server role depends. ○ Dependent Server Role: the server role on which the server role depends. ○ Dependent Cluster: the cluster where the dependency server role is deployed. ○ Dependency in Final Status: specifies whether the dependency server role has reached the desired state.

5.5.1.5.3. View the server role dashboard

The server role dashboard allows you to view the statistics of a server role.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, click the S tab.
3. (Optional)Enter the service name in the search box. Services that meet the search condition are displayed.
4. Click a service name and then service instances in the service are displayed in the lower-left corner.
5. Move the pointer over at the right of a service instance and then select **Dashboard**.
6. In the **Server Role List** section of the **Service Instance Information Dashboard** page, click **Details** in the **Actions** column.

7. View the information on the **Server Role Dashboard** page.

Item	Description
Server Role Summary	Displays the basic information of the server role as follows: ○ **Project Name**: the name of the project to which the server role belongs. ○ **Cluster Name**: the name of the cluster to which the server role belongs. ○ **Service Instance**: the name of the service instance to which the server role belongs. ○ **Server Role**: the server role name. ○ **In Final Status**: whether the server role reaches the final status. ○ **Expected Machines**: the number of expected machines. ○ **Actual Machines**: the number of actual machines. ○ **Machines Not Good**: the number of machines whose status is not Good. ○ **Machines with Role Status Not Good**: the number of server roles whose status is not Good. ○ **Machines Going Offline**: the number of machines that are going offline. ○ **Rolling**: whether a running rolling task exists. ○ **Rolling Task Status**: the current status of the rolling task. ○ **Time Used**: the time used for running the rolling task.
Machine Final Status Overview	The statistical chart of the current status of the server role.
Server Role Monitoring Information	○ **Updated At**: the time when the data is updated. ○ **Monitored Item**: the name of the monitored item. ○ **Level**: the level of the monitored item. ○ **Description**: the description of the monitored item.

Item	Description
Machine Information	○ **Machine Name**: the hostname of the machine. ○ **IP**: the IP address of the machine. ○ **Machine Status**: the machine status. ○ **Machine Action**: the action that the machine is performing. ○ **Server Role Status**: the status of the server role. ○ **Server Role Action**: the action that the server role is performing. ○ **Current Version**: the current version of the server role on the machine. ○ **Target Version**: the expected version of the server role on the machine. ○ **Error Message**: the exception message. ○ **Actions**: ■ Click **Terminal** to log on to the machine and perform operations. ■ Click **Restart** to restart the server roles on the machine. ■ Click **Details** to go to the **Machine Details** page. For more information about the machine details, see View the machine dashboard. ■ Click **Machine System View** to go to the **Machine Info Report** page. For more information about the machine info report, see Machine info report. ■ Click **Machine Operation** to restart, out of band restart, or clone the machine again.
Server Role Monitoring Information of Machines	○ **Updated At**: the time when the data is updated. ○ **Machine Name**: the machine name. ○ **Monitored Item**: the name of the monitored item. ○ **Level**: the level of the monitored item. ○ **Description**: the description of the monitored item.
VM Mappings	The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster. ○ **VM**: the hostname of the virtual machine. ○ **Currently Deployed On**: the hostname of the physical machine where the virtual machine is currently deployed. ○ **Target Deployed On**: the hostname of the physical machine where the virtual machine is expected to be deployed.

Item	Description
Service Dependencies	The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role. ○ **Dependent Service**: the service on which the server role depends. ○ **Dependent Server Role**: the server role on which the server role depends. ○ **Dependent Cluster**: the cluster to which the dependent server role belongs. ○ **Dependency in Final Status**: whether the dependent server role reaches the final status.

5.5.1.6. Machine operations

This topic describes the actions about machine operations.

5.5.1.6.1. View the machine dashboard

The machine dashboard allows you to view the statistics of a machine.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, click the **C** tab.
3. (Optional)On the **Machine** tab in the lower-left corner, enter the machine name in the search box. Machines that meet the search condition are displayed.
4. Move the pointer over  at the right of a machine and then select **Dashboard**.
5. On the **Machine Details** page, view all the information of this machine. For more information, see the following table.

Item	Description
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
DISK Usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-System	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
DISK IO-System	The statistical table of the disk input and output.

Item	Description
Machine Summary	◦ **Project Name**: the name of the project to which the machine belongs. ◦ **Cluster Name**: the name of the cluster to which the machine belongs. ◦ **Machine Name**: the machine name. ◦ **SN**: the serial number of the machine. ◦ **IP**: the IP address of the machine. ◦ **IDC**: the data center of the machine. ◦ **Room**: the room in the data center where the machine is located. ◦ **Rack**: the rack where the machine is located. ◦ **Unit in Rack**: the location of the rack. ◦ **Warranty**: the warranty of the machine. ◦ **Purchase Date**: the date when the machine is purchased. ◦ **Machine Status**: the running status of the machine. ◦ **Status**: the hardware status of the machine. ◦ **CPUs**: the number of CPUs for the machine. ◦ **Disks**: the disk size. ◦ **Memory**: the memory size. ◦ **Manufacturer**: the machine manufacturer. ◦ **Model**: the machine model. ◦ **os**: the operating system of the machine. ◦ **part**: the disk partition.
Server Role Status of Machine	The distribution of the current status of all server roles on the machine.
Machine Monitoring Information	◦ **Monitored Item**: the name of the monitored item. ◦ **Level**: the level of the monitored item. ◦ **Description**: the description of the monitored contents. ◦ **Updated At**: the time when the monitoring information is updated.

Item	Description
Machine Server Role Status	○ Service Instance ○ Server Role ○ Server Role Status ○ Server Role Action ○ Error Message ○ Target Version ○ Current Version ○ Actual Version Update Time ○ Actions: ■ Click Details to go to the Server Role Dashboard page. For more information about the server role dashboard, see View the server role dashboard. ■ Click Restart to restart the server roles on the machine.
Application Status in Server Roles	○ Application Name: the application name. ○ Process Number ○ Status: the application status. ○ Current Build ID: the ID of the current package version. ○ Target Build ID: the ID of the expected package version. ○ Git Version ○ Start Time ○ End Time ○ Interval: the interval between the time when Apsara Infrastructure Management Framework detects that the process exits and the time when Apsara Infrastructure Management Framework repairs the process. ○ Information Message: the normal output logs. ○ Error Message: the abnormal logs.

5.5.1.7. Monitoring center

You can view the alert status, alert rules, and alert history in the monitoring center.

5.5.1.7.1. Modify an alert rule

You can modify an alert rule based on the actual business requirements.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Service Operations**.

3. (Optional)Enter the service name in the search box.
4. Find the service and then click **Management** in the **Actions** column.
5. Click the **Monitoring Template** tab.
6. Find the monitoring template that you are about to edit and then click **Edit** in the **Actions** column.
7. Configure the monitoring parameters based on actual conditions.
8. Click **Save Change**.

   Wait about 10 minutes. The monitoring instance is automatically deployed. If the status becomes Successful and the deployment time is later than the modified time of the template, the changes are successfully deployed.

5.5.1.7.2. View the status of a monitoring instance

After a monitoring instance is deployed, you can view the status of the monitoring instance.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Service Operations**.
3. (Optional)Enter the service name in the search box.
4. Find the service and then click **Management** in the **Actions** column.
5. Click the **Monitoring Instance** tab. In the **Status** column, view the current status of the monitoring instance.

5.5.1.7.3. View the alert status

This topic describes how to view the alerts related to different services and the alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert Status**.
3. (Optional)Search for an alert by service name, cluster name, alert name, or alert time range.
4. View alert details on the **Alert Status** page. The following table describes the related parameters.

Parameter	Description
Service	The name of the service.
Cluster	The name of the cluster where the service is deployed.
Instance	The name of the monitored instance. Click the name of an instance to view the alert history of the instance.
Alert Status	Two alert states are available, which are Normal and Alerting.

Parameter	Description
Alert Level	Alerts are divided into five levels in descending order of severity: ○ P0: an alert that has been cleared ○ P1: an urgent alert ○ P2: a major alert ○ P3: a minor alert ○ P4: a reminder alert
Alert Name	The name of the alert. Click the name of an alert to view alert rule details.
Alert Time	The time when the alert is triggered and how long the alert lasts.
Actions	Click **Show** to view the data before and after the alert time.

5.5.1.7.4. View alert rules

This topic describes how to view alert rules.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert Rules**.
3. (Optional)Search for alert rules by service name, cluster name, or alert name.
4. View alert rules on the **Alert Rules** page. The following table describes the related parameters.

Parameter	Description
Service	The name of the service.
Cluster	The name of the cluster where the service is deployed.
Alert Name	The name of the alert.
Alert Conditions	The conditions that trigger the alert.
Periods	The frequency at which the alert rule is executed.
Alert Contact	The groups and members to notify when the alert is triggered.
Status	The status of the alert rule. ○ **Running**: Click it to stop the alert rule. ○ **Stopped**: Click it to execute the alert rule.

5.5.1.7.5. View the alert history

This topic describes how to view the historical alerts related to different services and the alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert History**.
3. (Optional)Search for an alert by service name, cluster name, or alert time range.
4. View the alert history on the **Alert History** page. The following table describes the related parameters.

Parameter	Description
Service	The name of the service to which the alert belongs.
Cluster	The name of the cluster where the service is deployed.
Alert Instance	The name of the instance where the alert is triggered.
Status	Two alert states are available, which are Normal and Alerting.
Alert Level	Alerts are divided into five levels in descending order of severity: ○ P0: an alert that has been cleared ○ P1: an urgent alert ○ P2: a major alert ○ P3: a minor alert ○ P4: a reminder alert
Alert Name	The name of the alert. Click the name of an alert to view alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members to notify when the alert is triggered.
Actions	Click **Show** to view the data before and after the alert time.

5.5.1.8. Tasks and deployment summary

This topic describes how to view rolling tasks, running tasks, history tasks, and deployment summary on Apsara Infrastructure Management Framework.

5.5.1.8.1. View rolling tasks

This topic describes how to view rolling tasks and their status.

Procedure

1. Log on to Apsara Infrastructure Management Framework.

2. In the top navigation bar, choose **Operations > Cluster Operations**.
3. Select **Rolling Tasks** to view all clusters that have rolling tasks.
4. Click **rolling** in the **Rolling** column.
5. On the **Rolling Task** page. view the change task information and change details.

Change task parameters

Parameter	Description
Change Version	The source version of the rolling task.
Description	The description of the change.
Head Version Analysis	The status of desired state analysis. During desired state analysis, Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to specific change contents. Desired state analysis can be in one of the following states: ○ **Preparing**: No new version is detected. ○ **Waiting**: The latest version has been detected, but the analysis module has not started. ○ **Doing**: The application to be changed is being analyzed. ○ **done**: The desired state analysis succeeds. ○ **Failed**: The desired state analysis fails to parse change contents. Apsara Infrastructure Management Framework can obtain change contents of server roles in the latest version only when the desired state analysis is in the **done** state.
Blocked Server Role	The server role that is blocked by dependencies in the rolling task.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.
Actions	Click **View Difference** to go to the **Version Difference** page. For more information, see View operation logs. Click **Stop** to terminate the rolling task. Click **Pause** to suspend the rolling task.

Change details parameters

Parameter	Description
Service Name	The name of the service that has changes.

Parameter	Description
Status	The current status of the service. The rolling status of a service is an aggregation result of rolling statuses of multiple server roles. Services can be in one of the following states: ○ **succeeded**: A task succeeds. ○ **blocked**: A task is blocked. ○ **failed**: A task fails.
Server Role Status	The status of the server role. Click **>** to the left of a service name to view the rolling task status of each server role in the service. Server roles can be in one of the following states: ○ **Downloading**: A task is being downloaded. ○ **Rolling**: A rolling task is in progress. ○ **RollingBack**: A rolling task fails and is performing rollback.
Depend On	The services on which the service depends, or the server roles on which the server role depends.
Actions	Click **Stop** to terminate the change of the server role. Click **Pause** to suspend the change of the server role.

5.5.1.8.2. View running tasks

This topic describes how to view running tasks.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > Running Tasks**.
3. (Optional)Search for running tasks by cluster name, server role name, task status, task submitter, Git version, or time range.
4. Find the target task, move the pointer over the **Rolling Task Status** column, and then click **View Tasks** to go to the **Rolling Task** page. For more information about rolling task details, see View rolling tasks.

5.5.1.8.3. View historical tasks

This topic describes how to view historical tasks.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > History Tasks**.
3. (Optional)Search for historical tasks by cluster name, Git version, submitter, or time range.
4. Find the target task and click **Details** in the **Actions** column to go to the **Rolling Task** page. For more information about rolling task details, see View rolling tasks.

5.5.1.8.4. View the deployment summary

On the **Deployment Summary** page, you can view the deployment conditions of clusters, services, and server roles in all projects on Apsara Infrastructure Management Framework.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > Deployment Summary**.
   - View the deployment status and the duration of a certain status for each project.
     - Gray: wait to be deployed. It indicates that some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed.
     - Blue: being deployed. It indicates that the project has not reached the final status for one time yet.
     - Green: has reached the final status. It indicates that all clusters in the project have reached the final status.
     - Orange: not reaches the final status. It indicates that a server role does not reach the final status for some reason after the project reaches the final status for the first time.
   - Configure the global clone switch.
     - **normal**: Clone is allowed.
     - **block**: Clone is forbidden.
   - Configure the global dependency switch.
     - **normal**: All configured dependencies are checked.
     - **ignore**: The dependency is not checked.
     - **ignore_service**: None of the service-level dependencies, including the server role dependencies across services, are checked, and only the server role-level dependencies are checked.
3. Click the **Deployment Details** tab to view the deployment details.

   For more information, see the following table.

Item	Description

Item	Description
Status Statistics	The general statistics of deployment conditions, including the total number of projects that are currently available. Click each status to display the projects in the corresponding status in the list. The projects have five deployment statuses: ○ **Final**: All the clusters in the project have reached the final status. ○ **Deploying**: The project has not reached the final status for one time yet. ○ **Waiting**: Some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed. ○ **Non-final**: A server role does not reach the final status for some reason after the project reaches the final status for the first time. ○ **Inspector Warning**: An error is detected on service instances in the project during the inspection.
Start Time	The time when Apsara Infrastructure Management Framework starts the deployment.
Progress	The proportion of server roles that reach the final status to all the server roles in the current environment.
Deployment Status	The time indicates the deployment duration for the following statuses: **Final**, **Deploying**, **Waiting**, and **Inspector Warning**. The time indicates the duration before the final status is reached for the **Non-final** status. Click the time to view the details.
Deployment Progress	The proportion of clusters, services, and server roles that reach the final status to the total clusters, services, and server roles in the project. Move the pointer over the blank area at the right of the data of roles and then click **Details** to view the deployment statuses of clusters, services, and server roles. The deployment statuses are indicated by icons, which are the same as those used for status statistics.
Resource Application Progress	**Total** indicates the total number of resources related to the project. ○ **Done**: the number of resources that have been successfully applied for. ○ **Doing**: the number of resources that are being applied for and retried. The number of retries (if any) is displayed next to the number of resources. ○ **Block**: the number of resources whose applications are blocked by other resources. ○ **Failed**: the number of resources whose applications failed.
Inspector Error	The number of inspection alerts for the current project.

Item	Description
Monitoring Information	The number of alerts generated for the machine monitor and the machine server role monitor in the current project.
Dependency	Click the icon to view the project services that depend on other services, and the current deployment status of the services that are depended on.

5.5.1.9. Reports

The system allows you to search for and view reports based on your business needs, and add commonly used reports to your favorites.

5.5.1.9.1. View reports

The **Reports** menu allows you to view the statistical data.

Context

You can view the following reports on Apsara Infrastructure Management Framework.

- System reports: default and common reports in the system.
- All reports: includes the system reports and custom reports.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You can go to the report list in the following three ways:
   - In the top navigation bar, choose **Reports > System Reports**.
   - In the top navigation bar, choose **Reports > All Reports**.
   - In the left-side navigation pane, click the **R** tab. Move the pointer over at the right of **All Reports** and then select **View**.

   See the following table for the report descriptions.

Item	Description
Report	The report name. Move the pointer over next to **Report** to search for reports by report name.
Group	The group to which the report belongs. Move the pointer over next to **Group** to filter reports by group name.
Status	Indicates whether the report is published.
Public	Indicates whether the report is public.
Created By	The person who creates the report.

Item	Description
Published At	The published time and created time of the report.
Actions	Click **Add to Favorites** to add this report to your favorites. Then, you can view the report by choosing **Reports > Favorites** in the top navigation bar or moving the pointer over at the right of **Favorites** on the **R** tab in the left-side navigation pane and then selecting **View**.

3. (Optional)Enter the name of the report that you are about to view in the search box.
4. Click the report name to go to the corresponding report details page. For more information about the reports, see Appendix.

5.5.1.9.2. Add a report to favorites

You can add common reports to favorites. Then, find them quickly on the **Favorites** page.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You can go to the report list in the following three ways:
   - In the top navigation bar, choose **Reports > System Reports**.
   - In the top navigation bar, choose **Reports > All Reports**.
   - In the left-side navigation pane, click the **R** tab. Move the pointer over at the right of **All Reports** and then select **View**.
3. (Optional)Enter the name of the report that you are about to add to favorites in the search box.
4. At the right of the report, click **Add to Favorites** in the **Actions** column.
5. In the displayed **Add to Favorites** dialog box, enter tags for the report.
6. Click **Add to Favorites**.

5.5.1.10. Appendix

5.5.1.10.1. Project component info report

This report displays the name and status for each type of project components, including services, server roles, and machines.

Item	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.

Item	Description
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

5.5.1.10.2. IP list

This report displays the IP addresses of physical machines and Docker applications.

IP List of Physical Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.

IP List of Docker Applications

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.
Docker IP	The Docker IP address.

5.5.1.10.3. Machine info report

This report displays the statuses of machines and server roles on the machines.

Machine Status

Displays all the machines currently managed by Apsara Infrastructure Management Framework and their corresponding statuses. In the **Global Filter** section at the top of the page, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists, and then click **Filter** on the right to filter the data.

Item	Description
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The machine status.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status.
Status Description	The description about the machine status.

Expected Server Role List

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The name of the expected server role on the machine.

Abnormal Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

Server Role Version and Status on Machine

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.
Server Role Status	The status of the server role.
Target Version	The expected version of the server role on the machine.
Current Version	The current version of the server role on the machine.
Status Description	The description about the status.
Error Message	The exception message of the server role.

Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

5.5.1.10.4. Rolling info report

This report displays the running and completed rolling tasks and the task-related status.

Choose a rolling action

This list only displays the running rolling tasks. If no rolling task is running, no data is available in the list.

Item	Description
Cluster	The cluster name.
Git Version	The version of change that triggers the rolling task.
Description	The description about the change entered by a user when the user submits the change.
Start Time	The start time of the rolling task.

Item	Description
End Time	The end time of the rolling task.
Submitted By	The ID of the user who submits the change.
Rolling Task Status	The current status of the rolling task.
Submitted At	The time when the change is submitted.

Server Role in Job

Select a rolling task in the **Choose a rolling action** section to display the rolling status of server roles related to the selected task. If no rolling task is selected, the server role statuses of all historical rolling tasks are displayed.

Item	Description
Server Role	The server role name.
Server Role Status	The rolling status of the server role.
Error Message	The exception message of the rolling task.
Git Version	The version of change to which the rolling task belongs.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Approve Rate	The proportion of machines that have the rolling task approved by the decider.
Failure Rate	The proportion of machines that have the rolling task failed.
Success Rate	The proportion of machines that have the rolling task succeeded.

Server Role Rolling Build Information

The source version and target version of each application under the server role in the rolling process.

Item	Description
App	The name of the application that requires rolling in the server role.
Server Role	The server role to which the application belongs.
From Build	The version before the upgrade.

Item	Description
To Build	The version after the upgrade.

Server Role Statuses on Machines

Select a server role in the **Server Role in Job** section to display the deployment status of this server role on the machine.

Item	Description
Machine Name	The name of the machine on which the server role is deployed.
Expected Version	The target version of the rolling.
Actual Version	The current version.
State	The status of the server role.
Action Name	The Apsara Infrastructure Management Framework action currently performed by the server role.
Action Status	The action status.

5.5.1.10.5. Machine RMA approval pending list

Some Apsara Infrastructure Management Framework actions (such as restart) on machines and server roles can be triggered by users, but this type of actions must be reviewed and approved. This report is used to process the actions that must be reviewed and approved.

Machine

Displays the basic information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
State	The running status of the machine.
Action Name	The action on the machine.
Action Status	The status of the action on the machine.

Item	Description
Actions	The approval button.

Machine Serverrole

Displays the information of server roles on the pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
Serverrole	The server role name.
State	The running status of the server role.
Action Name	The action on the server role.
Action Status	The status of the action on the server role.
Actions	The approval button.

Machine Component

Displays the hard disk information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
Component	The hard disk on the machine.
State	The running status of the hard disk.
Action Name	The action on the hard disk.
Action Status	The status of the action on the hard disk.
Actions	The approval button.

5.5.1.10.6. Registration vars of services

This report displays values of all service registration variables.

Item	Description
Service	The service name.
Service Registration	The service registration variable.
Cluster	The cluster name.
Update Time	The updated time.

5.5.1.10.7. Virtual machine mappings

Use the global filter to display the virtual machines of a specific cluster.

Displays the information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.

Item	Description
Project	The project name.
Cluster	The cluster name.
VM	The hostname of the virtual machine.
Currently Deployed On	The hostname of the physical machine on which the virtual machine is currently deployed.
Target Deployed On	The hostname of the physical machine on which the virtual machine is expected to be deployed.

5.5.1.10.8. Service inspector report

Use the global filter to display the service inspection reports of a specific cluster.

Service Inspector: Data is available only for services with inspection configured.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Description	The contents of the inspection report.
Level	The level of the inspection report.

5.5.1.10.9. Resource application report

In the **Global Filter** section, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists and then click **Filter** on the right to display the corresponding resource application data.

Change Mappings

Item	Description
Project	The project name.
Cluster	The cluster name.
Version	The version where the change occurs.
Resource Process Status	The resource application status in the version.
Msg	The exception message.
Begintime	The start time of the change analysis.
Endtime	The end time of the change analysis.

Changed Resource List

Item	Description
Res	The resource ID.
Type	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.
Parameters	The resource parameters.
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.

Item	Description
APP	The application of the server role.
Name	The resource name.
Type	The resource type.
Status	The resource application status.
Parameters	The resource parameters.
Result	The resource application result.
Res	The resource ID.
Reprocess Status	The status of the interaction with Business Foundation System during the VIP resource application.
Reprocess Msg	The error message of the interaction with Business Foundation System during the VIP resource application.
Reprocess Result	The result of the interaction with Business Foundation System during the VIP resource application.
Refer Version List	The version that uses the resource.
Error Msg	The exception message.

5.5.1.10.10. Statuses of project components

This report displays the status of all server roles in an abnormal status on machines of the project, and the monitoring information (alert information reported by the server role to Apsara Infrastructure Management Framework monitor) of server roles and machines.

Error State Component Table

Only displays the information of server roles that are not in GOOD status and server roles to be upgraded.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.

Item	Description
Need Upgrade	Whether the current version reaches the final status.
Server Role Status	The current status of the server role.
Machine Status	The current status of the machine.

Server Role Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Machine Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Service Inspector Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

5.5.1.10.11. Relationship of service dependency

This report displays the dependencies among server roles. Use the global filter to display the data of a specific cluster in the list.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Dependent Service	The service on which the server role depends.
Dependent Server Role	The server role on which the server role depends.
Dependent Cluster	The cluster to which the dependent server role belongs.
Dependency in Final Status	Whether the dependent server role reaches the final status.

5.5.1.10.12. Check report of network topology

This report checks if network devices and machines have wirecheck alerts.

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Network Instance	The name of the network device.
Level	The alert level.
Description	The description about the alert information.

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Machine Name	The server (machine) name.
Level	The alert level.
Description	The description about the alert information.

5.5.1.10.13. Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Status	The running status of the machine.
Clone Progress	The progress of the current clone process.

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.

Item	Description
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

5.5.1.10.14. Auto healing/install approval pending report

The list structure is the same as the machine RMA approval pending list, whereas this view is used for the approval during the installation. For more information, see Machine RMA approval pending list.

5.5.1.10.15. Machine power on or off statuses of clusters

After a cluster starts or shuts down machines, you can view the related information in this report.

Cluster Running Statuses

If a cluster is starting or shutting down machines, the corresponding data is available in this list. No data indicates that no cluster has machines shut down.

Item	Description
Project	The project name.
Cluster	The cluster name.
Action Name	The startup or shutdown action that is being performed by the cluster.
Action Status	The status of the action.

Server Role Power On or Off Statuses

Displays the power on or off statuses of server roles in the cluster selected in the **Cluster Running Statuses** section.

Select a row in the **Cluster Running Statuses** section to display the information of the corresponding cluster in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.

Item	Description
Action Name	The startup or shutdown action that is being performed by the server role.
Action Status	The status of the action.

Statuses on Machines

Displays the running status of the selected server role on machines.

Select a row in the **Server Role Power On or Off Statuses** section to display the information of the corresponding server role in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
Machine Name	The machine name.
Server Role Status	The running status of the server role.
Server Role Action	The action currently performed by the server role.
Server Role Action Status	The status of the action.
Error Message	The exception message.

Machine Statuses

Displays the running statuses of machines in the selected cluster.

Select a row in the **Statuses on Machines** section to display the information of the corresponding machine in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status of the machine.
Error Message	The exception message.

5.5.2. New version

5.5.2.1. What is Apsara Infrastructure Management Framework?

This topic describes what Apsara Infrastructure Management Framework is from the aspects of core functions and basic concepts.

5.5.2.1.1. Introduction

Apsara Infrastructure Management Framework is a distributed data center management system, which manages applications on clusters containing multiple machines and provides the basic functions such as deployment, upgrade, expansion, contraction, and configuration changes.

Overview

Apsara Infrastructure Management Framework also supports monitoring data and analyzing reports, which facilitates users to perform one-stop Operation & Maintenance (O&M) control. Based on the Apsara Infrastructure Management Framework services, automated O&M is implemented in the large-scale distributed environment, which greatly improves the operations efficiency and enhances the system availability.

Apsara Infrastructure Management Framework is mainly composed of TianjiMaster and TianjiClient. Apsara Infrastructure Management Framework installs TianjiClient as the agent on machines it manages. Then, TianjiMaster accepts and issues the upper-layer instructions to TianjiClient for execution. In the upper layer, Apsara Infrastructure Management Framework is divided into different components based on different functions, and then provides API server and portal for external use.

Core functions

- Network initialization in data centers
- Server installation and maintenance process management
- Deployment, expansion, and upgrade of cloud products
- Configuration management of cloud products
- Automatic application for cloud product resources
- Automatic repair of software and hardware faults
- Basic monitoring and business monitoring of software and hardware

5.5.2.1.2. Basic concepts

Before using Apsara Infrastructure Management Framework, you must know the following basic concepts for a better understanding.

project

A collection of clusters, which provides service capabilities for external entities.

cluster

A collection of physical machines, which logically provides services and is used to deploy project software.

- A cluster can only belong to one project.
- Multiple services can be deployed on a cluster.

service

A set of software, which provides relatively independent functions. A service is composed of one or more server roles and can be deployed on multiple clusters to form multiple sets of services and provide the corresponding service capabilities. For example, Apsara Distributed File System, Job Scheduler, and Apsara Name Service and Distributed Lock Synchronization System are all services.

service instance

A service that is deployed on a cluster.

server role

One or more indivisible deployment units into which a service can be divided based on functions. A server role is composed of one or more specific applications. If a service is deployed on a cluster, all the server roles of the service must be deployed to machines of this cluster. Multiple server roles, such as PanguMaster and TianjiClient, can be deployed on the same server.

server role instance

A server role that is deployed on a machine. A server role can be deployed on multiple machines.

application

Applications correspond to each process-level service component in a server role and each application works independently. The application is the minimum unit for deployment and upgrade in Apsara Infrastructure Management Framework, and can be deployed to each machine. Generally, an application is an executable software or Docker container.

If a server role is deployed on a machine, all applications in the server role must be deployed to this machine.

rolling

Each time when a user updates configurations, Apsara Infrastructure Management Framework upgrades services and modifies the cluster configurations based on the updated configurations. This process is called rolling.

service configuration template

Some configurations are the same when services are deployed on clusters. A service configuration template can be created to quickly write the same configurations to different clusters. The service configuration template is basically used for large-scale deployment and upgrade.

associated service template

A template.conf file that exists in the configurations. This file declares the service configuration template and its version, of which the configuration is used by the service instance.

desired state

If a cluster is in this state, all hardware and software on each of its machines are normal and all software are in the target version.

dependency

The dependency between server roles in a service defines that server roles with dependencies run tasks or are upgraded based on the dependency order. For example, if A depends on B, B is upgraded first. A starts to be downloaded after B is downloaded successfully, and upgraded after B is successfully upgraded. (By default, the dependency does not take effect for configuration upgrade.)

upgrade

A way of aligning the current state with the desired state of a service. After a user submits the version change, Apsara Infrastructure Management Framework upgrades the service version to the target version. With the server role as the processing unit, upgrade aims to update the versions of all machines to the target version.

At the beginning, the desired state and current state of the cluster are the same. When a user submits the change, the desired state is changed, whereas the current state is not. A rolling task is generated and has the desired state as the target version. During the upgrade, the current state is continuously approximating to the desired state. Finally, the desired state and the current state are the same when the upgrade is finished.

5.5.2.2. Log on to Apsara Infrastructure Management Framework

This topic describes how to log on to Apsara Infrastructure Management Framework.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

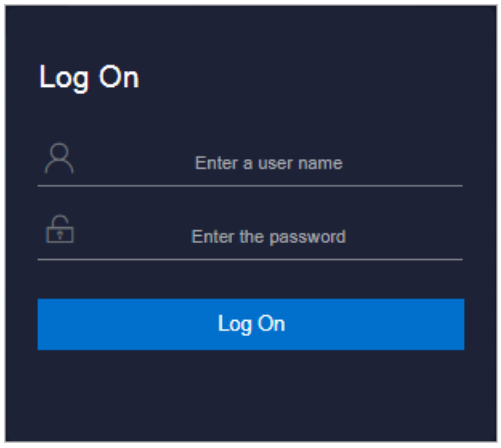

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, select **Products**.
6. In the Product List, select **Apsara Infrastructure Management Framework**.

5.5.2.3. Homepage introduction

After you log on to Apsara Infrastructure Management Framework, the homepage appears. This topic allows you to get a general understanding of the basic operations and functions of Apsara Infrastructure Management Framework.

Log on to Apsara Infrastructure Management Framework. The homepage appears, as shown in the following figure.

Homepage of Apsara Infrastructure Management Framework

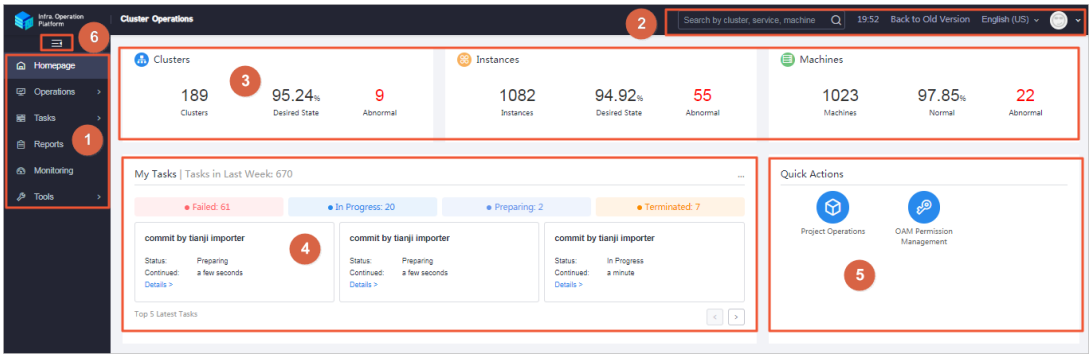

For more information about the descriptions of functional areas on the homepage, see the following table.

Descriptions of functional areas

Area		Description

1	Left-side navigation pane	• **Operations**: the quick entrance of Operation & Maintenance (O&M) operations, which allows operations engineers to quickly find the corresponding operations and operation objects. This menu consists of the following sections:
		  ◦ **Project Operations**: manages projects with the project permissions.
		  ◦ **Cluster Operations**: performs O&M operations on and manages clusters with the project permissions, such as viewing the cluster status.
		  ◦ **Service Operations**: manages services with the service permissions, such as viewing the service list information.
		  ◦ **Machine Operations**: maintains and manages all the machines in Apsara Infrastructure Management Framework, such as viewing the machine status.
		• **Tasks**: A rolling task is generated after you modify the configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects.
		• **Reports**: displays the monitoring data in tables and provides the function of searching for different reports.
		• **Monitoring**: effectively monitors metrics in the process of system operation and sends alert notifications for abnormal conditions. This menu includes the functions of displaying alert status, modifying alert rules, and searching for the alert history.
		• **Tools**: provides the machine tools and the IDC shutdown function.

Area		Description

2	Function buttons in the upper-right corner	• Search box: Supports global search. Enter a keyword in the search box to search for clusters, services, and machines.
		• Move the pointer over the time and then you can view:
		  ◦ **TJDB Sync Time**: the generated time of the data that is displayed on the current page.
		  ◦ **Desired State Calc Time**: the calculation time of the desired-state data that is displayed on the current page.
		After data is generated, the system processes the data at maximum speed. As an asynchronous system, Apsara Infrastructure Management Framework has some latency. The time helps explain why the current data results are generated and determine whether the current system has a problem.
		• 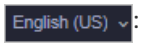 : In the English environment, click this drop-down list to switch to another language.
		• Click the avatar of the logon user and then select **Exit** to log out of Apsara Infrastructure Management Framework.

3	Status section of global resources	Displays the overview of global resources. • **Clusters**: displays the total number of clusters, the percentage of clusters that reach the desired state, and the number of abnormal clusters. • **Instances**: displays the total number of instances, the percentage of instances that reach the desired state, and the number of abnormal instances. • **Machines**: displays the total number of machines, the percentage of machines with the **Normal** state, and the number of abnormal machines. Move the pointer over the section and then click Show Detail to go to the Cluster Operations page, Service Operations page, or Machine Operations page.
4	Task status section	Displays the information of tasks submitted in the last week. Click the number at the right of a task status to go to the My Tasks page and then view tasks of the corresponding status. The top 5 latest tasks are displayed at the bottom of this section and you can click **Details** to view the task details.
5	Quick actions	Displays links of common quick actions, which allows you to perform operations quickly.
6	Expand/collapse button	If you are not required to use the left-side navigation pane when performing O&M operations, click this button to collapse the left-side navigation pane and increase the space of the content area.

5.5.2.4. Project operations

The Project Operations module allows you to search for and view details of a project.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Project Operations**.

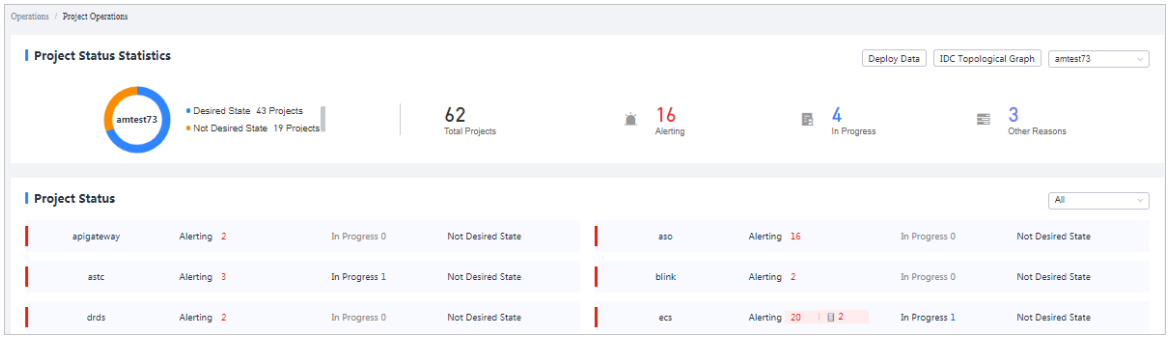

3. On this page, you can:
   - Search for a project

     Click the drop-down list in the upper-right corner of the **Project Status** section. Enter a project name in the search box, and then select the name to search for the project. You can view the numbers of alerts and running tasks for the project and whether the project reaches the desired state.
   - View the details of a project
     - Find the project whose details you are about to view. Click the number at the right of **Alerting**. In the displayed Alert Information dialog box, view the specific monitoring metrics, monitoring types, and alert sources. Click the value in the Alert Source column to view the service details.
     - Find the project whose details you are about to view. Click the number at the right of **In Progress**. In the displayed Tasks dialog box, view the details of Upgrade Service and Machine Change.

5.5.2.5. Cluster operations

This topic describes the actions about cluster operations.

5.5.2.5.1. View the cluster list

The cluster list allows you to view all of the clusters and the corresponding information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. To view the cluster list, you can:
   - On the **Homepage**, move the pointer over the **Clusters** section and then click Show Detail in the upper-right corner.

- In the left-side navigation pane, choose **Operations > Cluster Operations**.

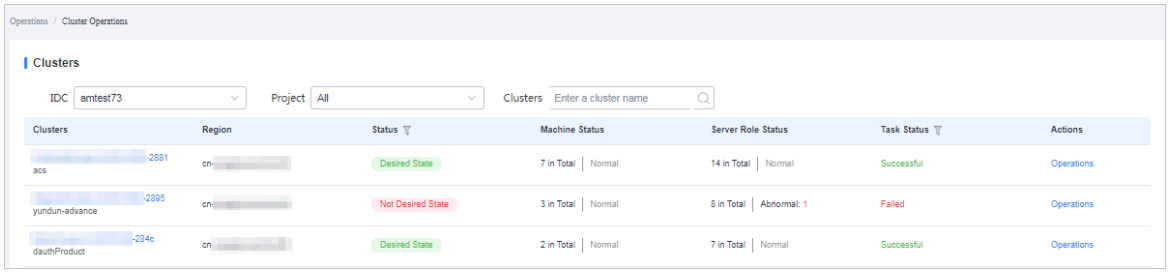

On this page, you can view the following information.

Item	Description
Clusters	The cluster name. Click the cluster name to view the cluster details.
Region	The name of the region where the cluster is located.
Status	Indicates whether the cluster reaches the desired state. Use to filter the clusters. ○ Desired State: All the clusters of a project reach the desired state. ○ Not Desired State: After a project reaches the desired state for the first time, a server role does not reach the desired state because of undefined reasons.
Machine Status	The number of machines and the corresponding status in the cluster. Click the status to go to the Machines tab of the Cluster Details page.

Item	Description
Server Role Status	The number of server roles and the corresponding status in the cluster. Click the status to go to the Services tab of the Cluster Details page. Click **Abnormal** in the Server Role Status column to view all the abnormal server roles in the cluster in the displayed dialog box. Click **View Details** in the upper-right corner of the dialog box to go to the Services tab of the Cluster Details page. 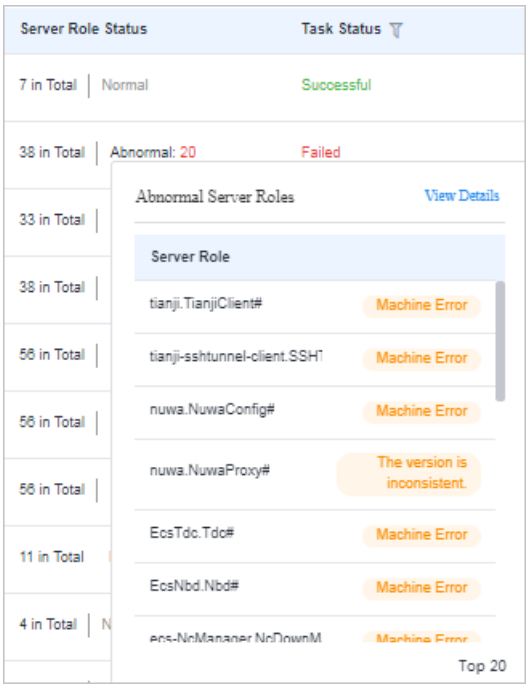
Task Status	The status of the task submitted to the cluster. Use to filter the clusters. Click the status to view the task details.

5.5.2.5.2. View the cluster details

You can view the cluster statistics by viewing the cluster details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. (Optional)Select a project from the Project drop-down list or enter the cluster name in the Clusters field to search for the corresponding cluster.
4. Find the cluster whose configurations you are about to view. Click the cluster name or **Operations** in the **Actions** column at the right of the cluster to go to the **Cluster Details** page.

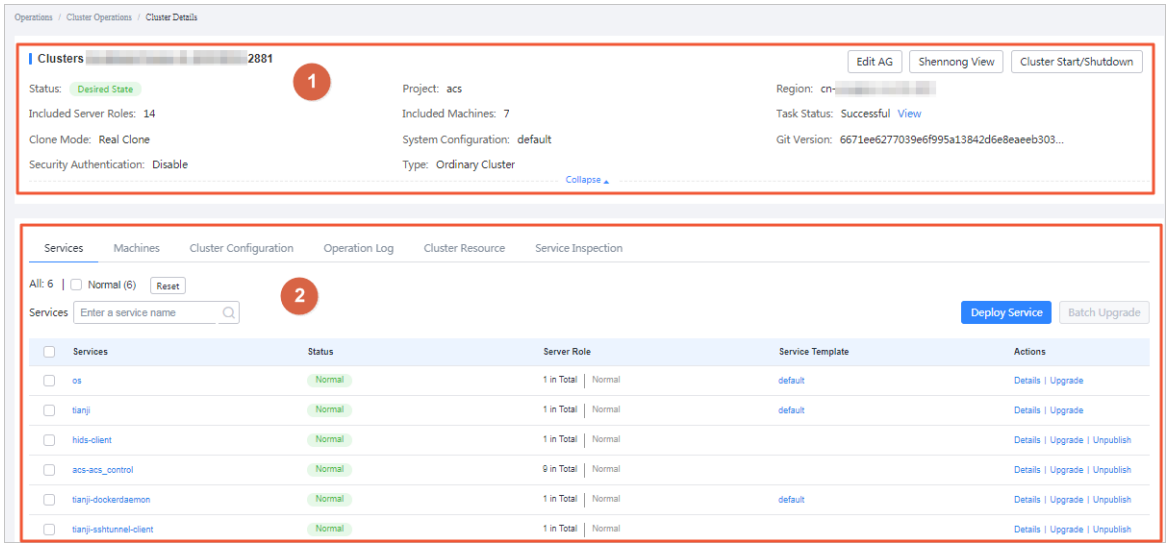

Area	Item	Description
1	**Status**	○ **Desired State**: All the clusters of this project reach the desired state. ○ **Not Desired State**: After the project reaches the desired state for the first time, a server role does not reach the desired state because of undefined reasons.
	Project	The project to which the cluster belongs.
	Region	The region to which the cluster belongs.
	Included Server Roles	The number of server roles included in the cluster.
	Included Machines	The number of machines included in the cluster.
	Task Status	The status of the current task. Click **View** to view the task details. ○ **Successful**: indicates the task is successful. ○ **Preparing**: indicates data is being synchronized and the task is not started yet. ○ **In Progress**: indicates the cluster has a changing task. ○ **Paused**: indicates the task is paused ○ **Failed**: indicates the task failed. ○ **Terminated**: indicates the task is manually terminated.
	Clone Mode	○ **Mock Clone**: The system is not cloned when a machine is added to the cluster. ○ **Real Clone**: The system is cloned when a machine is added to the cluster.

Area	Item	Description
	System Configuration	The name of the system service template used by the cluster.
	Git Version	The change version to which the cluster belongs.
	Security Authentication	The access control among processes. Generally, the non-production environment uses the default configurations and does not perform the verification. In other cases, customize the configurations based on actual requirements to enable or disable the verification.
	Type	◦ **Ordinary Cluster**: an operations unit facing to machine groups, where multiple services can be deployed. ◦ **Virtual Cluster**: an operations unit facing to services, which can centrally manage software versions of machines of multiple physical clusters. ◦ **RDS**: a type of cluster that renders special cgroup configurations according to a certain rule. ◦ **NETFRAME**: a type of cluster that renders special configurations for the special scenario of Server Load Balancer (SLB). ◦ **T4**: a type of cluster that renders special configurations for the mixed deployment of e-commerce. Currently, Alibaba Cloud Apsara Stack only has ordinary clusters.
2	**Services**	View the statuses of all the services in this cluster. You can also upgrade or unpublish a service. ◦ **Normal**: The service works properly. ◦ **Not Deployed**: No machine is deployed on the service. ◦ **Changing**: Some server roles in the service are changing. ◦ **Operating**: No server role is changing, but the machine where server roles are installed is performing the Operation and Maintenance (O&M) operations. ◦ **Abnormal**: No server role is changing or the machine where server roles are installed is not performing the O&M operations, but the server role status is not **GOOD** or the version that the service runs on the machine and the version configured in the configurations are different.
	Machines	View the running statuses and monitoring statuses of all the machines in this cluster. You can also view the details of server roles to which the machine belongs.
	Cluster Configuration	The configuration file used in the cluster.

Area	Item	Description
	Operation Log	View the version differences.
	Cluster Resource	Filter the resource whose details you are about to view according to certain conditions.
	Service Inspection	View the inspection information of each service in the cluster.

5.5.2.5.3. View operation logs

By viewing operation logs, you can obtain the differences between Git versions.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. To view the operation logs of a cluster, you can:
   - Enter a cluster name in the search box in the upper-right corner of the page. Click **Operations** at the right of the cluster to go to the Cluster Details page. Click the **Operation Log** tab.
   - In the left-side navigation pane, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click **Operations** in the **Actions** column at the right of a cluster to go to the Cluster Details page. Click the **Operation Log** tab.

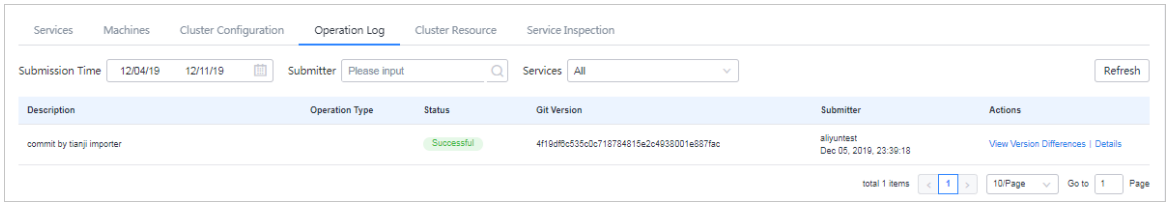

3. On the **Operation Log** tab, view the version differences.
   i. Click **View Version Differences** in the **Actions** column at the right of a log.
   ii. On the **Version Differences** page, select a basic version from the **Versus** drop-down list. Then, the contents of the different file are automatically displayed.
   iii. Select each different file from the **Different File** drop-down list to view the detailed differences.

5.5.2.6. Service operations

5.5.2.6.1. View the service list

The service list allows you to view all of the services and the corresponding information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. To view the service list, you can:
   - On the **Homepage**, move the pointer over the **Instances** section and then click Show Detail in

the upper-right corner.

- In the left-side navigation pane, choose **Operations > Service Operations**.

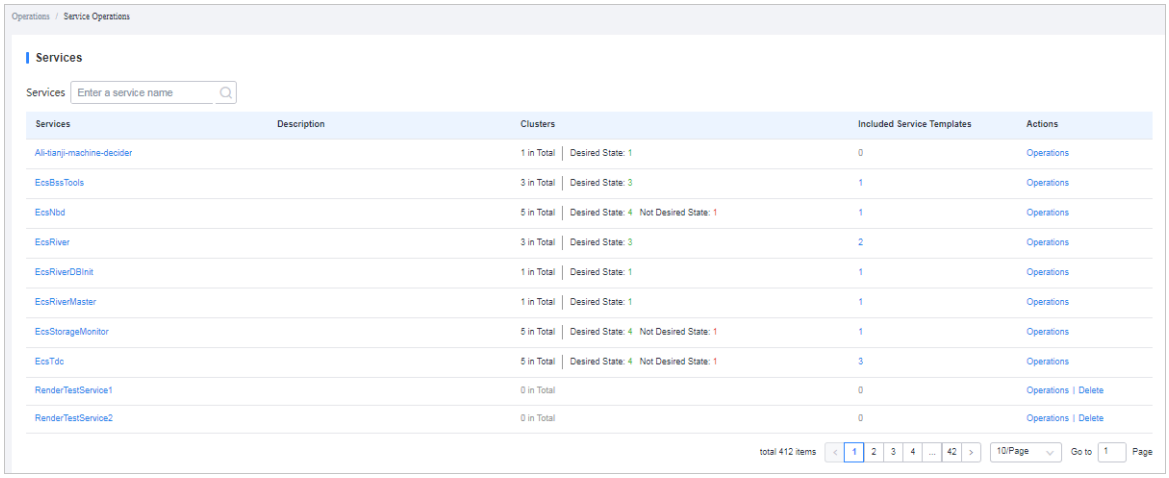

On this page, you can view the following information.

Item	Description
Services	The service name. Click the service name to view the service details.
Clusters	The number of clusters where the service is located and the corresponding cluster status.
Included Service Templates	The number of service templates this service includes.
Actions	Click **Operations** to go to the Service Details page.

3. (Optional)Enter a service name in the search box and then the service that meets the condition is displayed in the list.

5.5.2.6.2. View the server role details

You can view the server role statistics by viewing the server role details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Service Operations**.
3. (Optional)Enter a service name in the search box and then the service that meets the condition is displayed in the list.
4. Click the service name or click **Operations** in the **Actions** column.

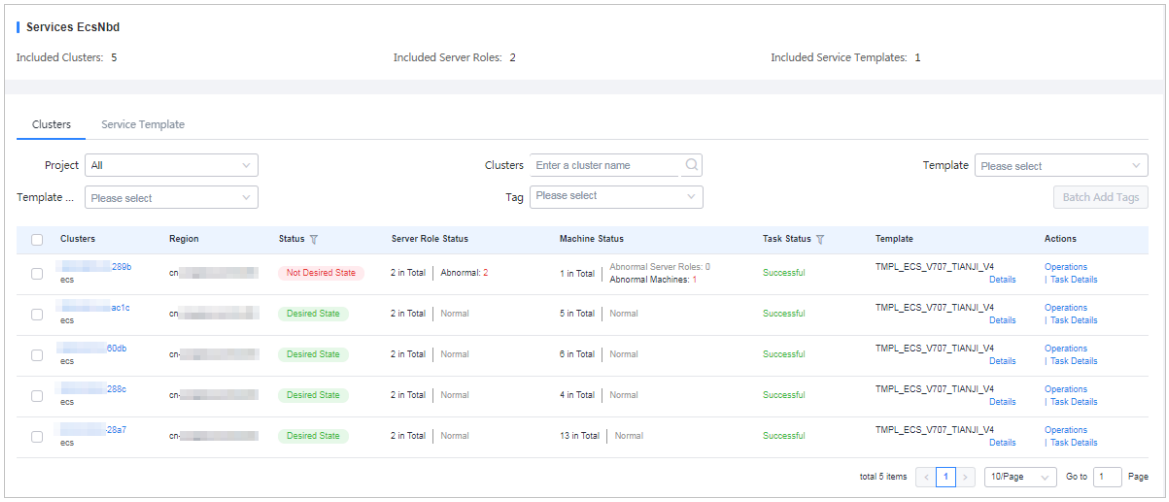

5. On the **Clusters** tab, click the status in the Server Role Status column to view the server roles included in a cluster.

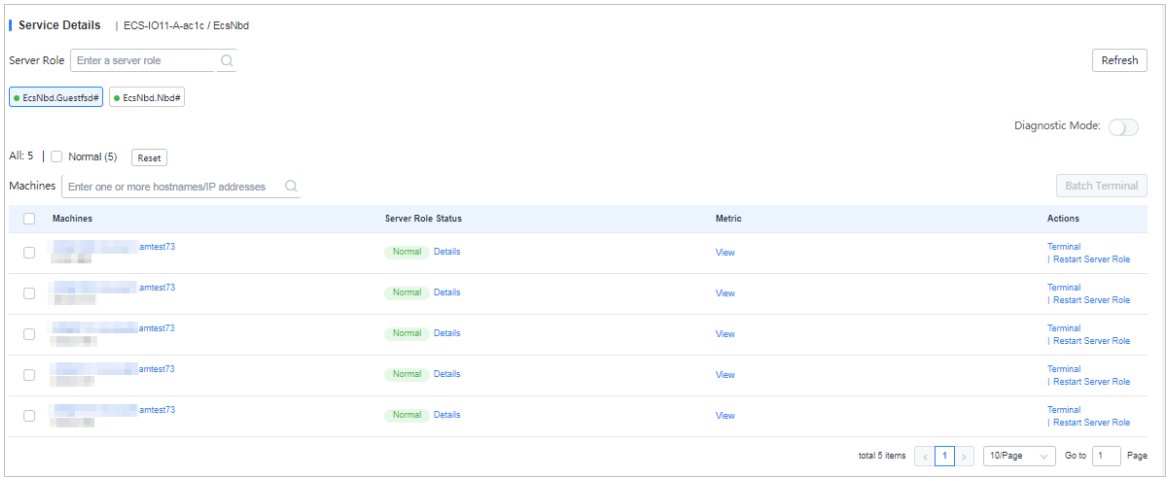

6. Enter a keyword in the search box to search for a server role. Then, the details of the corresponding server role are displayed in the list.

Item	Description
Machines	The machine to which the server role belongs. Click the machine name to view the machine details.
Server Role Status	The status of the server role. Click **Details** to view the basic information, application version information, application process information, and resources of the server role.
Metric	Click **View** to view the statuses of server role metrics and machine metrics.
Actions	○ Click **Terminal** to log on to the machine and perform operations. ○ Click **Restart Server Role** to restart the server role.

5.5.2.7. Machine operations

You can view the machine statistics by viewing the machine list.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. To view the machine list, you can:
   - On the **Homepage**, move the pointer over the **Machines** section and then click Show Detail in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Machine Operations**.

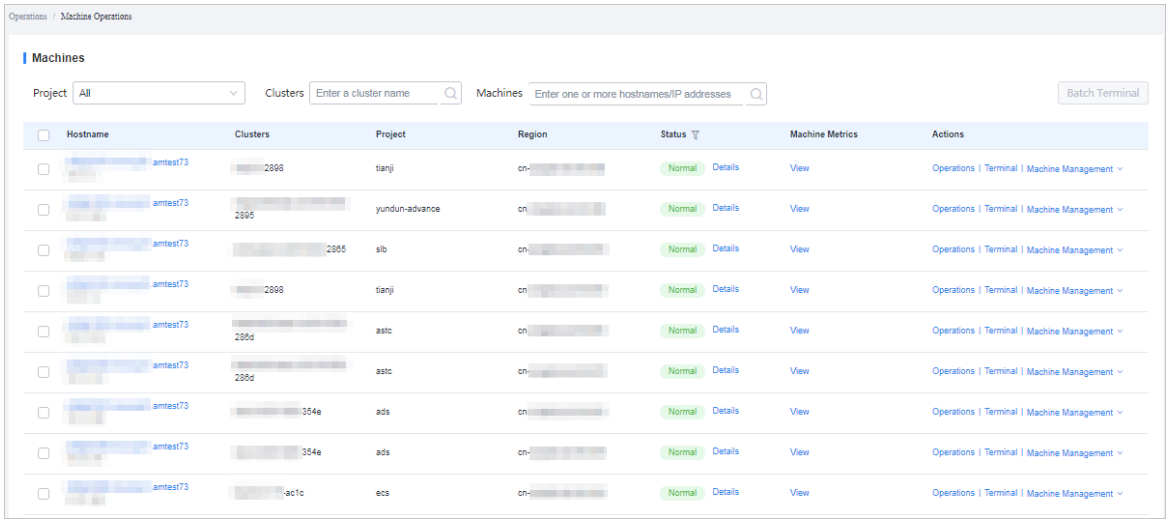

3. (Optional)Select a project or enter the cluster name or machine name to search for the corresponding machine.

Item	Description
Hostname	Click the hostname to go to the Machine Details page.
Status	The current status of the machine. Use to filter the machines. Click **Details** and then the **Status Details of Machine** dialog box appears.

Item	Description
Machine Metrics	Click **View** and then the **Metrics** dialog box appears. 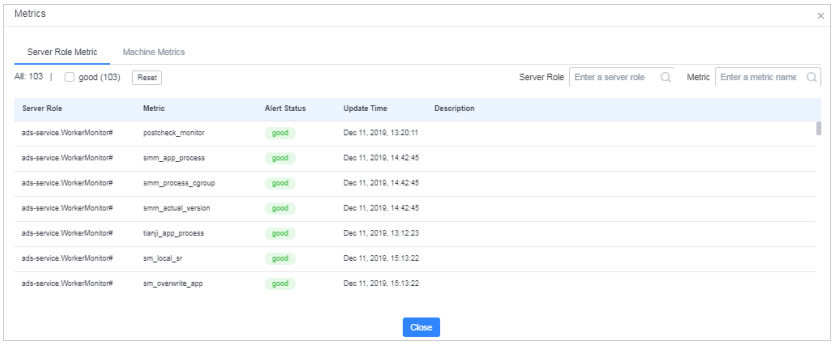 The **Server Role Metric** tab and **Machine Metrics** tab display the corresponding metrics, and you can view the specific alert status and updated time. Enter a keyword in the search box in the upper-right corner to search for a server role or metric. You can also select the status to filter metrics.
Actions	○ Click **Operations** to go to the Machine Details page. ○ Click **Terminal** to log on to the machine and perform operations. You can select multiple machines and then click **Batch Terminal** in the upper-right corner to log on to multiple machines at a time. ○ Click **Machine Management** to perform an out-of-band restart operation on the machine.

5.5.2.8. Monitoring center

You can view the alert status, alert rules, and alert history in the monitoring center.

5.5.2.8.1. View the monitoring instance status

You can view the status of a monitoring instance after it is deployed.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Service Operations**.
3. (Optional)Enter a service name in the search box to search for the corresponding service.
4. Click **Operations** in the **Actions** column at the right of the service.
5. On the **Clusters** tab, configure the conditions and then search for the cluster. Click **Operations** in the **Actions** column.
6. On the **Cluster Details** page, select the server role you are about to view and then click **View** in the **Metric** column. Then, view the statuses of server role metrics and machine metrics.

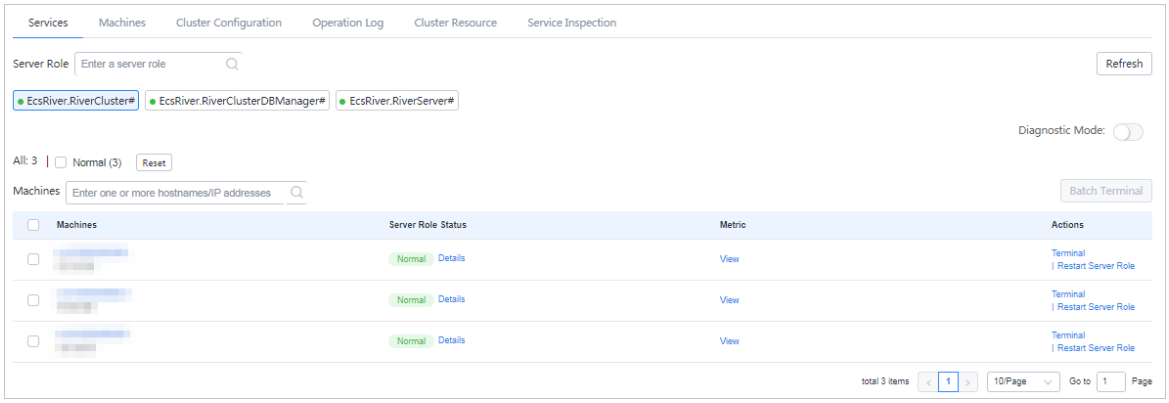

5.5.2.8.2. View the alert status

The Alert Status page allows you to view the alerts generated in different services and the corresponding alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Monitoring**. On the Monitoring page, click **Go** to open the target page.
3. In the top navigation bar, choose **Monitoring > Alert Status**.

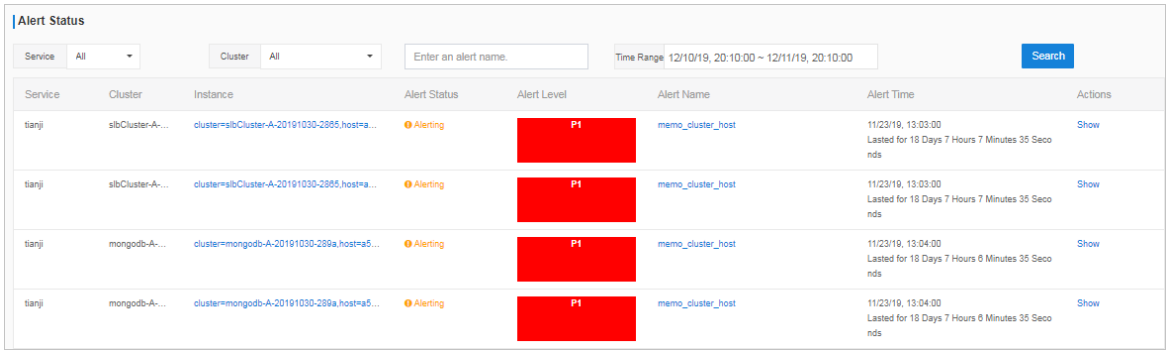

4. (Optional)You can configure the service name, cluster name, alert name, or the time range when the alert is triggered to search for alerts.
5. On the **Alert Status** page, view the alert details. For more information about the alert status descriptions, see the following table.

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Instance	The name of the service instance being monitored. Click the instance to view the alert history of this instance.
Alert Status	Alerts have two statuses: **Restored** and **Alerting**.

Item	Description
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. ○ P1 ○ P2 ○ P3 ○ P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered and how long the alert has lasted.
Actions	Click **Show** to show the data before and after the alert time.

5.5.2.8.3. View alert rules

The Alert Rules page allows you to view the configured alert rules.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Monitoring**. On the Monitoring page, click **Go** to open the target page.
3. In the top navigation bar, choose **Monitoring > Alert Rules**.

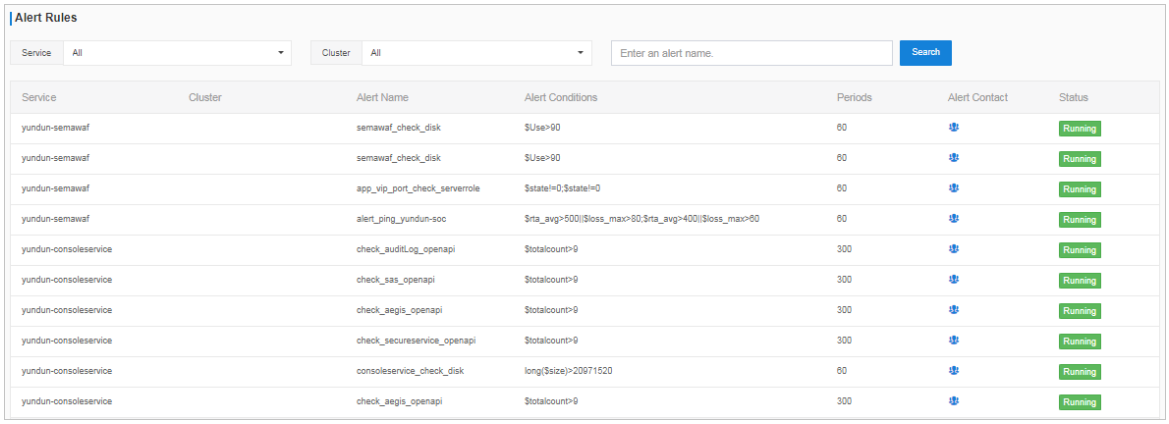

4. (Optional)You can configure the service name, cluster name, or alert name to search for alert rules.
5. On the **Alert Rules** page, view the detailed alert rules. For more information about the alert rule descriptions, see the following table.

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.

Item	Description
Alert Name	The name of the generated alert.
Alert Conditions	The conditions met when the alert is triggered.
Periods	The frequency (in seconds) with which an alert rule is run.
Alert Contact	The groups and members that are notified when an alert is triggered.
Status	The current status of the alert rule. ○ **Running**: Click to stop the alert rule. ○ **Stopped**: Click to run the alert rule.

5.5.2.8.4. View the alert history

The Alert History page allows you to view all the history alerts generated in different services and the corresponding alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose Monitoring. On the **Monitoring** page, click **Go** to open the target page.
3. In the top navigation bar, choose **Monitoring > Alert History**.

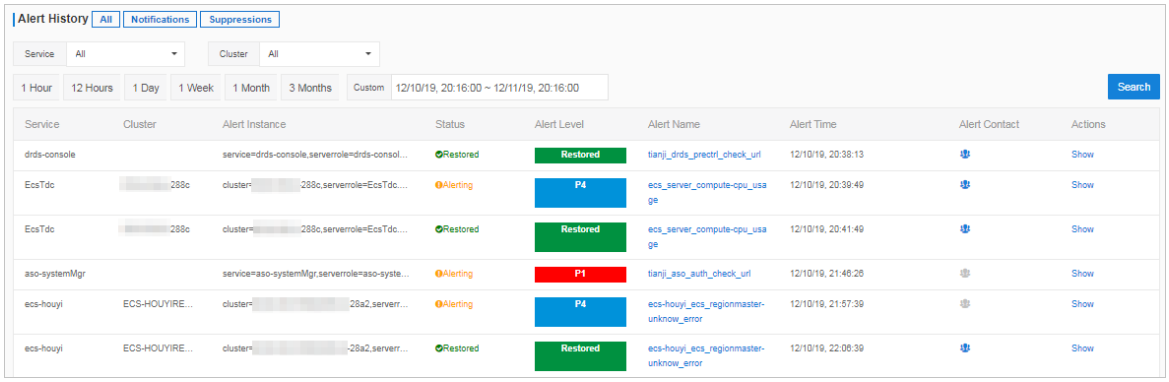

4. (Optional)You can configure the service name, cluster name, time range, or period to search for alerts.
5. On the **Alert History** page, view the history alerts. For more information about the history alert descriptions, see the following table.

Item	Description
Service	The name of the service to which the alert belongs.
Cluster	The name of the cluster where the service is located.

Item	Description
Alert Instance	The name of the resource where the alert is triggered.
Status	Alerts have two statuses: **Restored** and **Alerting**.
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. ○ P1 ○ P2 ○ P3 ○ P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members that are notified when an alert is triggered.
Actions	Click **Show** to show the data before and after the alert time.

5.5.2.9. View tasks

The task list allows you to view the submitted tasks and the corresponding status.

Procedure

1. Log on to Apsara Infrastructure Management Framework
2. To view the task list, you can:
   - In the left-side navigation pane, choose **Tasks > My Tasks**.
   - In the left-side navigation pane, choose **Tasks > Related Tasks**.
3. You can use  to filter tasks in the **Status** column.
4. Find the task whose details you are about to view and then click the task name or click **Details** in the **Actions** column.
5. On the **Task Details** page, view the status and progress of each cluster and server role.

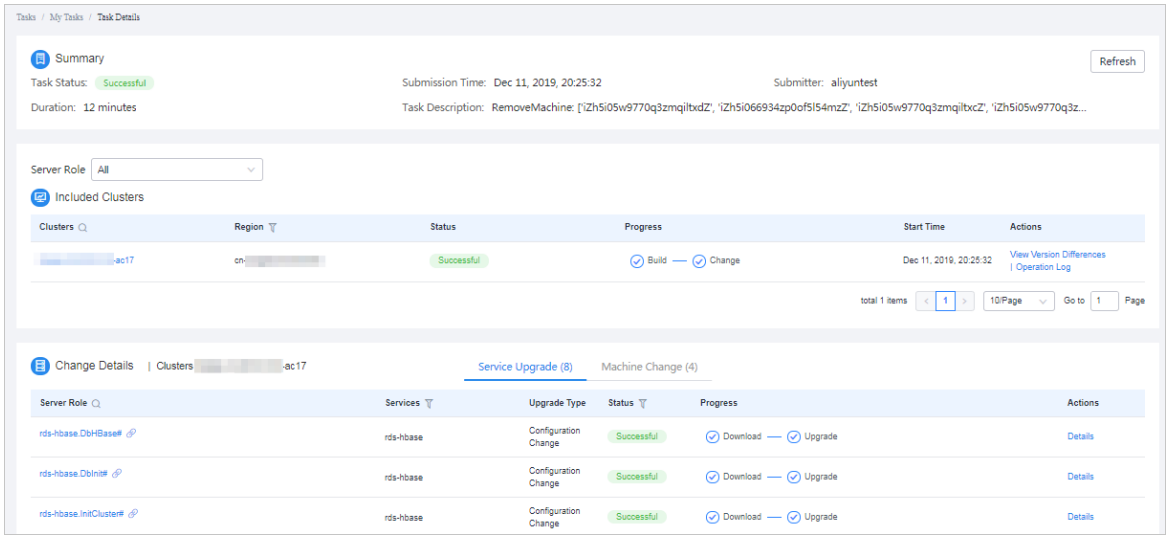

5.5.2.10. Reports

5.5.2.10.1. View reports

The Reports module allows you to view the statistical data.

Context

You can view the following reports on Apsara Infrastructure Management Framework.

- System reports: default and common reports in the system.
- All reports: includes the system reports and custom reports.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Reports**. On the Reports page, click **Go** to open the target page.

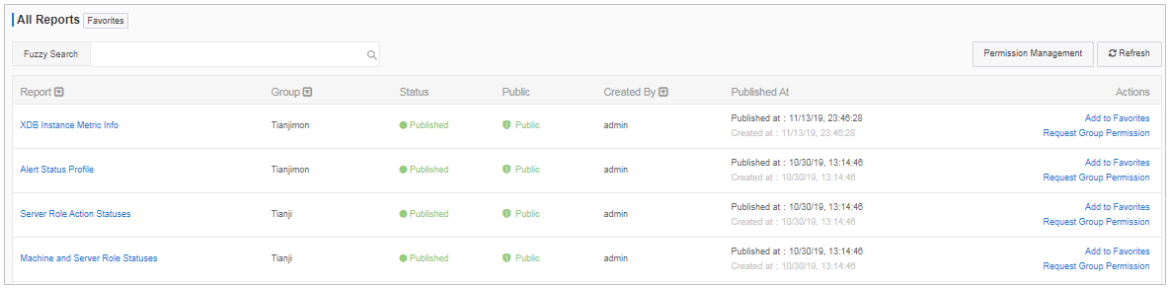

For more information about the report descriptions, see the following table.

Item	Description
Report	The report name. Move the pointer over the down-arrow button next to Report to search for reports by report name.

Item	Description
Group	The group to which the report belongs. Move the pointer over the down-arrow button next to Group to filter reports by group name.
Status	Indicates whether the report is published. ○ Published ○ Not published
Public	Indicates whether the report is public. ○ Public: All of the logon users can view the report. ○ Not public: Only the current logon user can view the report.
Created By	The person who creates the report.
Published At	The time when the report is published and created.
Actions	Click **Add to Favorites** to add this report to your favorites. Then, you can view the report by choosing **Reports > Favorites** in the top navigation bar.

3. (Optional)Enter the name of the report that you are about to view in the search box.
4. Click the report name to go to the corresponding report details page. For more information about the reports, see Appendix.

5.5.2.10.2. Add a report to favorites

You can add common reports to favorites. Then, find them quickly on the Favorites page.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Reports**. On the Reports page, click **Go** to open the target page.
3. (Optional)Enter the name of the report that you are about to add to favorites in the search box.
4. At the right of the report, click **Add to Favorites** in the **Actions** column.
5. In the displayed **Add to Favorites** dialog box, enter tags for the report.
6. Click **Add to Favorites**.

5.5.2.11. Tools

5.5.2.11.1. Machine tools

The Machine Tools module guides operations personnel to perform Operation & Maintenance (O&M) operations in common scenarios.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Tools > Operation Tools > Machine Tools**. On the Machine Tools page, click **Go** to open the target page.
3. Select the operation scene according to actual situations.

Operation scene	Description	Action
Scene 1: NC Scale-out (with existing machines)	Scales out an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 2: Host Scale-out (with existing machines)	Scales out the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 3: NC Scale-in	Scales in an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 4: Host Scale-in	Scales in the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 5: VM Migration	Migrates virtual machines (VMs) from a host to another host.	Select a source host and a destination host. Select the VMs to be migrated in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.

Operation scene	Description	Action
Scene 6: Host Switching	Switches from a standby host to a primary host.	Select a source host and a destination host. Click **Submit** and then click **Confirm** in the displayed dialog box.

5.5.2.11.2. IDC shutdown

If you are about to maintain the data center or shut down all of the machines in the data center, you must shut down the data center.

Prerequisites

Warning This is a high-risk operation, so proceed with caution.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Tools > IDC Shutdown**, and then click **Go** to open the target page.
3. On the **Clusters Shutdown** page, click **Start Shutdown** to shut down all of the machines in the data center with one click.

5.5.2.12. Appendix

5.5.2.12.1. Project component info report

This report displays the name and status for each type of project components, including services, server roles, and machines.

Item	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.

Item	Description
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

5.5.2.12.2. IP list

This report displays the IP addresses of physical machines and Docker applications.

IP List of Physical Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.

IP List of Docker Applications

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.
Docker IP	The Docker IP address.

5.5.2.12.3. Machine info report

This report displays the statuses of machines and server roles on the machines.

Machine Status

Displays all the machines currently managed by Apsara Infrastructure Management Framework and their corresponding statuses. In the **Global Filter** section at the top of the page, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists, and then click **Filter** on the right to filter the data.

Item	Description
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The machine status.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status.
Status Description	The description about the machine status.

Expected Server Role List

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The name of the expected server role on the machine.

Abnormal Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

Server Role Version and Status on Machine

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.

Item	Description
Server Role	The server role name.
Server Role Status	The status of the server role.
Target Version	The expected version of the server role on the machine.
Current Version	The current version of the server role on the machine.
Status Description	The description about the status.
Error Message	The exception message of the server role.

Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

5.5.2.12.4. Rolling info report

This report displays the running and completed rolling tasks and the task-related status.

Choose a rolling action

This list only displays the running rolling tasks. If no rolling task is running, no data is available in the list.

Item	Description
Cluster	The cluster name.
Git Version	The version of change that triggers the rolling task.
Description	The description about the change entered by a user when the user submits the change.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.

Item	Description
Submitted By	The ID of the user who submits the change.
Rolling Task Status	The current status of the rolling task.
Submitted At	The time when the change is submitted.

Server Role in Job

Select a rolling task in the **Choose a rolling action** section to display the rolling status of server roles related to the selected task. If no rolling task is selected, the server role statuses of all historical rolling tasks are displayed.

Item	Description
Server Role	The server role name.
Server Role Status	The rolling status of the server role.
Error Message	The exception message of the rolling task.
Git Version	The version of change to which the rolling task belongs.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Approve Rate	The proportion of machines that have the rolling task approved by the decider.
Failure Rate	The proportion of machines that have the rolling task failed.
Success Rate	The proportion of machines that have the rolling task succeeded.

Server Role Rolling Build Information

The source version and target version of each application under the server role in the rolling process.

Item	Description
App	The name of the application that requires rolling in the server role.
Server Role	The server role to which the application belongs.
From Build	The version before the upgrade.
To Build	The version after the upgrade.

Server Role Statuses on Machines

Select a server role in the **Server Role in Job** section to display the deployment status of this server role on the machine.

Item	Description
Machine Name	The name of the machine on which the server role is deployed.
Expected Version	The target version of the rolling.
Actual Version	The current version.
State	The status of the server role.
Action Name	The Apsara Infrastructure Management Framework action currently performed by the server role.
Action Status	The action status.

5.5.2.12.5. Machine RMA approval pending list

Some Apsara Infrastructure Management Framework actions (such as restart) on machines and server roles can be triggered by users, but this type of actions must be reviewed and approved. This report is used to process the actions that must be reviewed and approved.

Machine

Displays the basic information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
State	The running status of the machine.
Action Name	The action on the machine.
Action Status	The status of the action on the machine.
Actions	The approval button.

Machine Serverrole

Displays the information of server roles on the pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
Serverrole	The server role name.
State	The running status of the server role.
Action Name	The action on the server role.
Action Status	The status of the action on the server role.
Actions	The approval button.

Machine Component

Displays the hard disk information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
Component	The hard disk on the machine.
State	The running status of the hard disk.
Action Name	The action on the hard disk.
Action Status	The status of the action on the hard disk.
Actions	The approval button.

5.5.2.12.6. Registration vars of services

This report displays values of all service registration variables.

Item	Description
Service	The service name.

Item	Description
Service Registration	The service registration variable.
Cluster	The cluster name.
Update Time	The updated time.

5.5.2.12.7. Virtual machine mappings

Use the global filter to display the virtual machines of a specific cluster.

Displays the information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.

Item	Description
Project	The project name.
Cluster	The cluster name.
VM	The hostname of the virtual machine.
Currently Deployed On	The hostname of the physical machine on which the virtual machine is currently deployed.
Target Deployed On	The hostname of the physical machine on which the virtual machine is expected to be deployed.

5.5.2.12.8. Service inspector report

Use the global filter to display the service inspection reports of a specific cluster.

Service Inspector: Data is available only for services with inspection configured.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Description	The contents of the inspection report.
Level	The level of the inspection report.

5.5.2.12.9. Resource application report

In the **Global Filter** section, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists and then click **Filter** on the right to display the corresponding resource application data.

Change Mappings

Item	Description
Project	The project name.
Cluster	The cluster name.
Version	The version where the change occurs.
Resource Process Status	The resource application status in the version.
Msg	The exception message.
Begintime	The start time of the change analysis.
Endtime	The end time of the change analysis.

Changed Resource List

Item	Description
Res	The resource ID.
Type	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.
Parameters	The resource parameters.
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.

Item	Description
APP	The application of the server role.
Name	The resource name.
Type	The resource type.
Status	The resource application status.
Parameters	The resource parameters.
Result	The resource application result.
Res	The resource ID.
Reprocess Status	The status of the interaction with Business Foundation System during the VIP resource application.
Reprocess Msg	The error message of the interaction with Business Foundation System during the VIP resource application.
Reprocess Result	The result of the interaction with Business Foundation System during the VIP resource application.
Refer Version List	The version that uses the resource.
Error Msg	The exception message.

5.5.2.12.10. Statuses of project components

This report displays the status of all server roles in an abnormal status on machines of the project, and the monitoring information (alert information reported by the server role to Apsara Infrastructure Management Framework monitor) of server roles and machines.

Error State Component Table

Only displays the information of server roles that are not in GOOD status and server roles to be upgraded.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.

Item	Description
Need Upgrade	Whether the current version reaches the final status.
Server Role Status	The current status of the server role.
Machine Status	The current status of the machine.

Server Role Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Machine Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Service Inspector Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

5.5.2.12.11. Relationship of service dependency

This report displays the dependencies among server roles. Use the global filter to display the data of a specific cluster in the list.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Dependent Service	The service on which the server role depends.
Dependent Server Role	The server role on which the server role depends.
Dependent Cluster	The cluster to which the dependent server role belongs.
Dependency in Final Status	Whether the dependent server role reaches the final status.

5.5.2.12.12. Check report of network topology

This report checks if network devices and machines have wirecheck alerts.

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Network Instance	The name of the network device.
Level	The alert level.
Description	The description about the alert information.

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Machine Name	The server (machine) name.
Level	The alert level.
Description	The description about the alert information.

5.5.2.12.13. Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Status	The running status of the machine.
Clone Progress	The progress of the current clone process.

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.

Item	Description
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

5.5.2.12.14. Auto healing/install approval pending report

The list structure is the same as the machine RMA approval pending list, whereas this view is used for the approval during the installation. For more information, see Machine RMA approval pending list.

5.5.2.12.15. Machine power on or off statuses of clusters

After a cluster starts or shuts down machines, you can view the related information in this report.

Cluster Running Statuses

If a cluster is starting or shutting down machines, the corresponding data is available in this list. No data indicates that no cluster has machines shut down.

Item	Description
Project	The project name.
Cluster	The cluster name.
Action Name	The startup or shutdown action that is being performed by the cluster.
Action Status	The status of the action.

Server Role Power On or Off Statuses

Displays the power on or off statuses of server roles in the cluster selected in the **Cluster Running Statuses** section.

Select a row in the **Cluster Running Statuses** section to display the information of the corresponding cluster in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.

Item	Description
Action Name	The startup or shutdown action that is being performed by the server role.
Action Status	The status of the action.

Statuses on Machines

Displays the running status of the selected server role on machines.

Select a row in the **Server Role Power On or Off Statuses** section to display the information of the corresponding server role in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
Machine Name	The machine name.
Server Role Status	The running status of the server role.
Server Role Action	The action currently performed by the server role.
Server Role Action Status	The status of the action.
Error Message	The exception message.

Machine Statuses

Displays the running statuses of machines in the selected cluster.

Select a row in the **Statuses on Machines** section to display the information of the corresponding machine in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status of the machine.
Error Message	The exception message.

6.Products

The Products module allows you to click operations and maintenance services of other products on the cloud platform and ISV access configurations to go to the corresponding page.

6.1. Product list

On the Product List page, you can go to the corresponding operations and maintenance page of a product or ISV page by using Single Sign-On (SSO) and redirection.

Prerequisites

To access the ISV page, make sure that the ISV access information is configured on the **ISV Access Configurations** page. For more information about how to configure the ISV access information, see Configure the ISV access information.

Context

After you log on to the Apsara Stack Operations (ASO) console, you can view O&M icons of different products and different ISV icons on the **Product List** page based on your permissions. An operations system administrator can view all the O&M components of the cloud platform.

The read and write permissions for product O&M are separated. Therefore, the system can dynamically assign different permissions based on different roles.

Procedure

1. In the left-side navigation pane, choose **Products > Product List**.
2. On the **Product List** page, you can view the O&M icons of different products and ISV icons based on your permissions.

6.2. ISV access configurations

The **ISV Access Configurations** module allows you to configure, modify, and delete the ISV access information.

6.2.1. Configure the ISV access information

You can configure the ISV access information in the system based on business needs. Then, you can click an icon on the product list page to access the corresponding ISV page.

Procedure

1. In the left-side navigation pane, choose **Products > ISV Access Configuration**.
2. In the upper part of the page, click **Add**.
3. In the **Add** pane, configure the ISV access information.

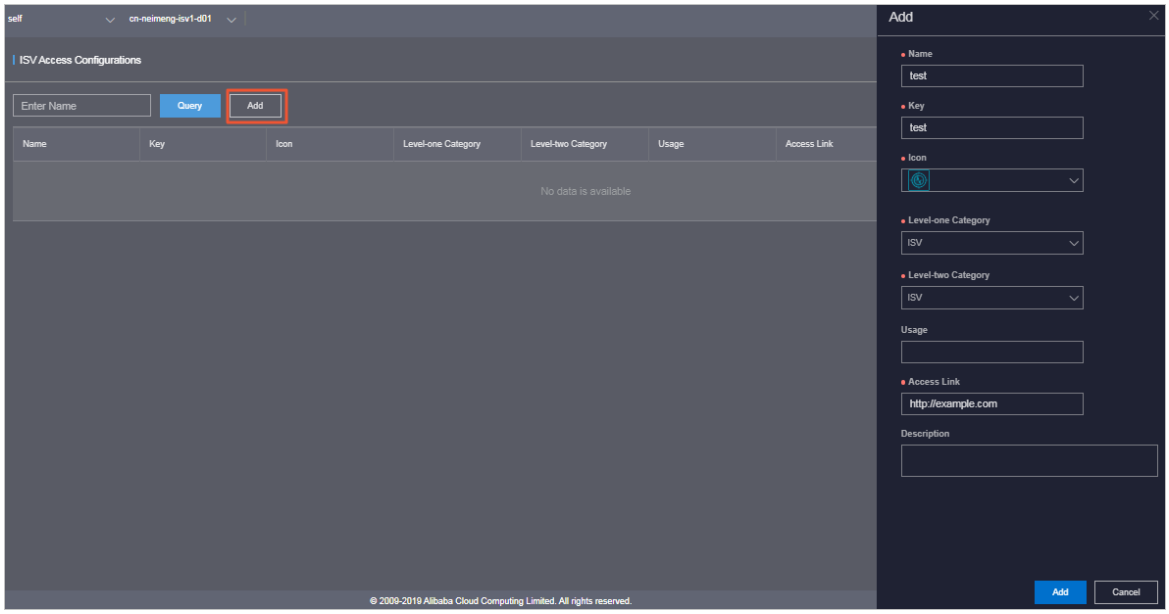

The following table describes the parameters.

Parameter	Description
Name	The name of the ISV to be accessed.
Key	Typically, enter an identifier related to the ISV business as the key.
Icon	Select the icon displayed on the **Product List** page for the ISV to be accessed.
Level-one Category and **Level-two Category**	The category to which the ISV to be accessed belongs on the **Product List** page.
Usage	The function of the ISV to be accessed.
Access Link	The address of the ISV to be accessed.
Description	The description related to the ISV to be accessed.

4. Click **Add**.

Result

You can view the added ISV icon in the Product List page by choosing **Products > Product List**. Click the icon and then you can go to the corresponding page.

6.2.2. Modify the ISV access information

If the ISV information is changed, you can modify the ISV access information.

Procedure

1. In the left-side navigation pane, choose **Products > ASV Access Configuration**.
2. (Optional)In the search box on the page, enter the ISV name, and then click **Query**. Fuzzy search is supported.

3. Find the ISV whose access information is to be modified. Click **Modify** in the **Actions** column.

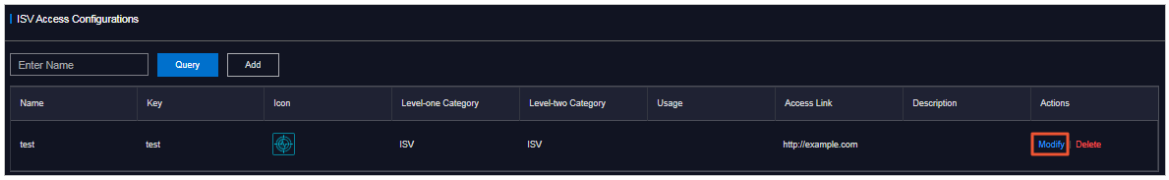

4. In the **Modify** pane, modify the name, key, icon, level-one category, level-two category, usage, access link, or description of the ISV.
5. Click **Modify**.

6.2.3. Delete the ISV access information

You can delete the ISV access information added in the system based on business needs.

Procedure

1. In the left-side navigation pane, choose **Products > ISV Access Configuration**.
2. (Optional)In the search box on the page, enter the ISV name, and then click **Query**. Fuzzy search is supported.
3. Find the ISV whose access information is to be deleted. Click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

Result

The deleted ISV will no longer be displayed in the **Product List**.

7.PaaS operations and maintenance

7.1. PaaS console overview

The PaaS console is designed based on the platform and products. The console is mainly used to view, manage, and upgrade the products deployed in the PaaS console. The PaaS console also provides task management capabilities to support orchestration, O&M, and custom extension.

7.2. Log on to the PaaS console

This topic describes how to log on to the PaaS console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   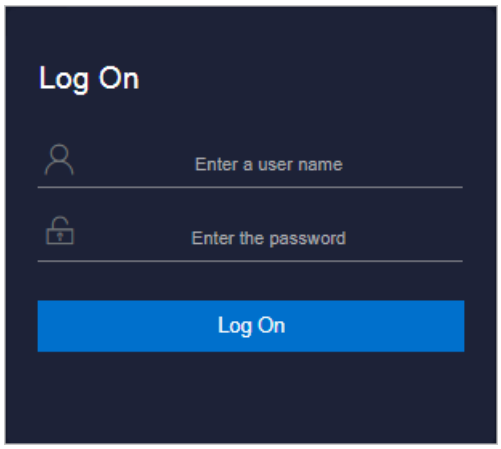

   Note You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

   Note Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**.
6. In the **Apsara Stack O&M > Basic O&M** section, click **PaaS Console**.

7.3. Overview

The overview includes cluster overview and alert events.

7.3.1. Cluster overview

On the Cluster Overview page, you can view information such as the cluster overview, node status, alerts, and cluster events.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Cluster Overview** from the **Overview** drop-down list.
2. Select the target cluster from the drop-down list next to **Cluster Overview**.
3. View the following information:
   - Cluster overview

     The Cluster Overview section displays status statistics of nodes and pods in clusters.

     Click a status value in the **Node** column to go to the Nodes page. You can also choose **Clusters > Nodes** in the left-side navigation pane to go to the Nodes page. Click the **Not Ready** value in the **Pod Status** column to view error messages of abnormal pods.
   - Node status

     The **Node Status** section displays the CPU utilization of each node, allocated size of memory, and usage of shared disks.

     When the CPU utilization of a node or allocated size of memory exceeds 80%, the information of the node is displayed in red.

     > **Note** To view the latest data of nodes, click **Refresh** on the right side of the Node Status section.

     Click **More** on the right side of the Node Status section to go to the **Nodes** page.
   - Alerts

     The **Alerts** section displays the alert events that are triggered by the monitoring system. O&M engineers can view the alerts.

> **Note** To view the latest data of alerts, click **Refresh** on the right side of the Alerts section.

Click **More** on the right side of the Alerts section to go to the Alert Events page. You can also choose **Overview > Alert Event** in the left-side navigation pane to go to the Alert Events page.

- Cluster events

  The **Cluster Events** section displays Kubernetes events by type. By default, alert events in all namespaces are displayed. You can filter events by namespace and type.

  > **Note** To view the latest data of cluster events, click **Refresh** on the right side of the Cluster Events section.

  Click **More** on the right side of the Cluster Events section to go to the Events page. You can also choose **Cluster > Events** in the left-side navigation pane to go to the Events page.

7.3.2. Alert events

The Alert Events page displays all alert events and all aggregated alert events by alert or product name.

7.3.2.1. View aggregated alert events by alert name

You can view aggregated alert events by alert name on the Alert Aggregation tab.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Events** from the **Alerts** drop-down list.
   The **Alert Aggregation** tab is displayed by default.
2. In the upper part of the page, select the target cluster from the drop-down list. By default, all alert events that are aggregated by alert name are displayed.
3. In the alert name view, the Alert Aggregation tab displays all aggregated alert events by alert name. The aggregated alert event list includes the following columns: Alert Name, Details, Total Alerts, Severity, and Actions.

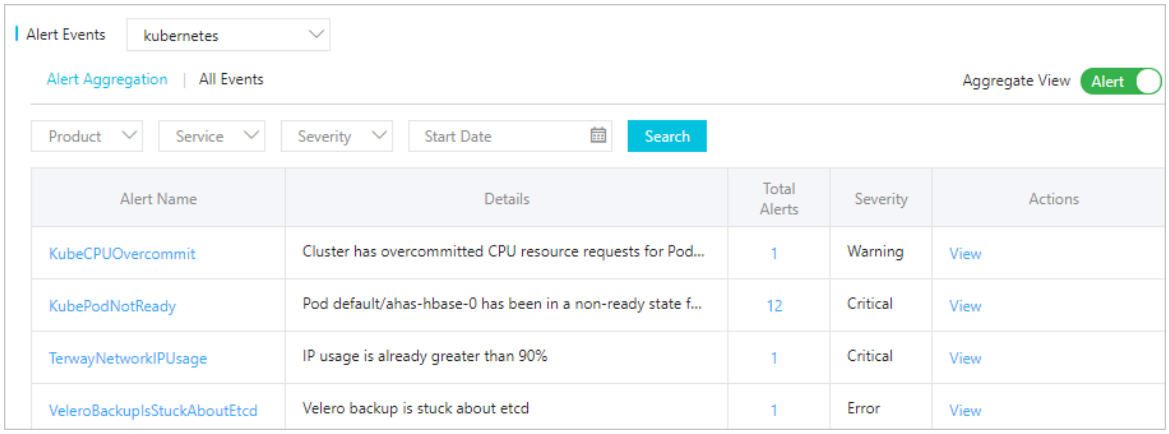

4. (Optional)In the search box at the top of the tab, set Product, Service, Severity, and Start Date,

and then click **Search** to query aggregated alert events that meet the conditions.

5. Find the target aggregated alert events. Click the name in the **Alert Name** column and the number in the **Total Alerts** column, or click **View** in the **Actions** column to view details of individual alert event within the aggregated alert events.

   The alert details include the following columns: Status, Start Time, End Time, Update Time, and Label.

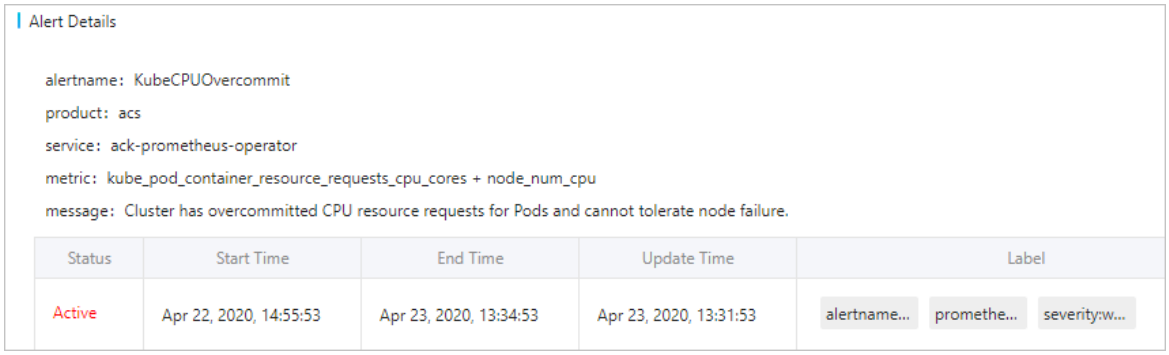

7.3.2.2. View aggregated alert events by product name

You can view aggregated alert events by product name on the Alert Aggregation tab.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Events** from the **Alerts** drop-down list.
   The **Alert Aggregation** tab is displayed by default.
2. In the upper part of the page, select the target cluster from the drop-down list. By default, all alert events that are aggregated by alert name are displayed.
3. Turn off the **Aggregate View** to switch to the product name view.

   In the product name view, the **Alert Aggregation** tab displays all aggregated alert events by product name.

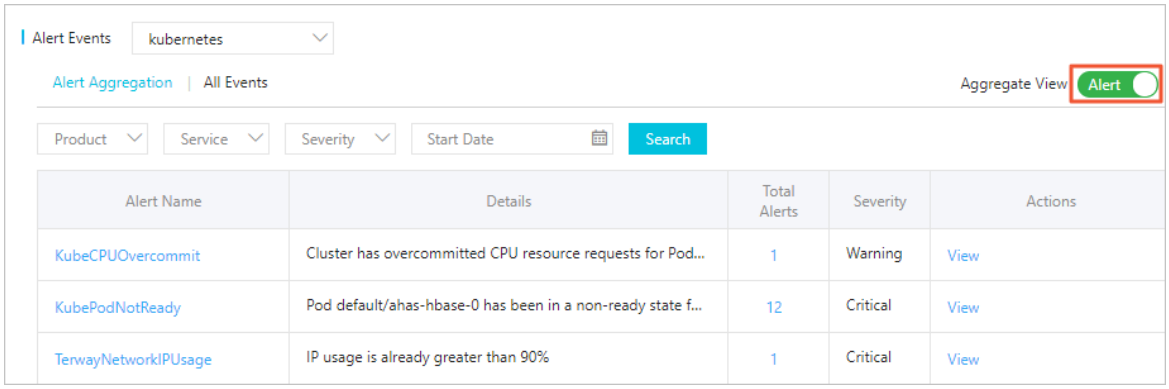

4. (Optional)In the search box at the top of the tab, set Product, Service, Severity, and Start Date, and then click **Search** to query aggregated alert events that meet the conditions.
5. Find the target aggregated alert events. Click the name in the **Alert Name** column and the number in the **Total Alerts** column, or click **View** in the **Actions** column to view details of individual alert events within the aggregated alert events. The alert details include the following columns: Status, Start Time, End Time, Update Time, and Label.

7.3.2.3. View all alert events

On the All Events tab, you can view all alert events generated in the PaaS console.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Events** from the **Alerts** drop-down list.
2. Click the **All Events** tab.
3. All alert events are displayed on the tab. The alert event list includes the following columns: Alert Name, Start Time, End Time, Update Time, Status, Details, Severity, and Label.

7.4. Clusters

7.4.1. View the cluster list

On the Clusters page, you can view the status and kubeconfig connection information of the clusters managed by PaaS.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Clusters** from the **Clusters** drop-down list.
2. On the **Clusters** page, view all clusters managed by PaaS.

3. Find a cluster, and then click **View** in the **Actions** column to view the kubeconfig connection information of the cluster.

7.4.2. Node management

You can add node tags or taints for clusters to manage scheduling policies.

7.4.2.1. Add tags

You can add tags to nodes for subsequent cluster scheduling, configuration, and behavior customization.

Add a tag on the Nodes page

On the Nodes page, you can add tags to one or more nodes.

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.

3. Select one or more nodes to which you want to add a tag. Click **Add Label** in the lower-left corner.
4. Perform the following operations:
   - Add a built-in tag

     In the **Add Label to Node** dialog box, click a tag in the Built-in Labels field. The tag name is automatically filled into the Key field. Set **Value** and then click **OK**.

     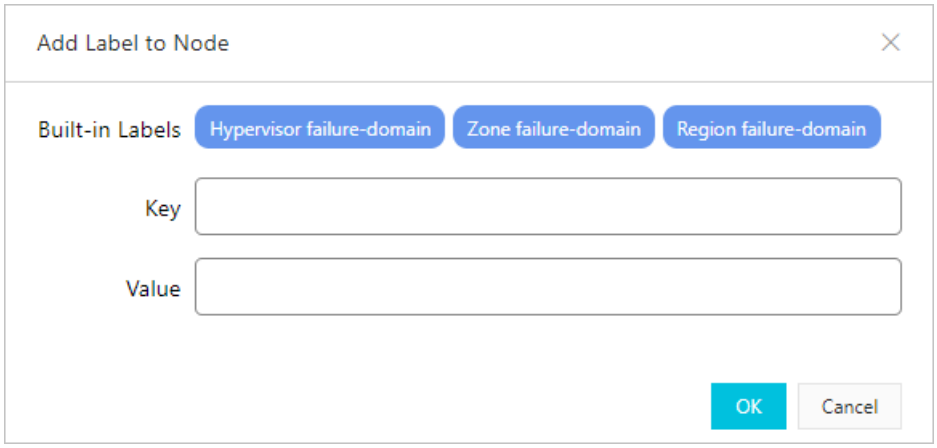

     The following table describes the parameters.

Parameter	Description
Built-in Labels	Existing tags in the system. Valid values: ■ **Hypervisor failure-domain**: During output virtualization, virtual machines are distributed across different physical machines. This tag can be used to distribute pods to different physical machines. ■ **Zone failure-domain**: distributes Kubernetes nodes to different zones. ■ **Region failure-domain**: distributes Kubernetes nodes to different regions.
Key	After you click a tag in the Built-in Labels field, the tag name is automatically filled into the **Key** field. You can also set **Key** to specify a custom tag.
Value	The custom tag value.

   - Add a custom tag

     In the **Add Label to Node** dialog box, set Key and Value, and then click **OK**.

Add a tag on the Node Details page

On the Node Details page, you can add a tag to a node.

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.

3. Select a node to which you want to add a tag. Click **View** in the **Actions** column.
4. On the Node Details page, click the icon next to the tag, and then perform the following operations:
   - Add a built-in tag

     In the **Add Label to Node** dialog box, click a tag in the Built-in Labels field. The tag name is automatically filled into the Key field. Set **Value** and then click **OK**.
   - Add a custom tag

     In the **Add Label to Node** dialog box, set Key and Value, and then click **OK**.

7.4.2.2. Add taints

You can add taints to nodes for subsequent pod scheduling.

Add a taint on the Nodes page

On the Nodes page, you can add a taint to one or more nodes.

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Select one or more nodes to which you want to add a taint. Click **View** in the **Actions** column.
4. Perform the following operations:
   - Add a built-in taint

     In the **Add Taint** dialog box, click a taint in the Built-in Taints field. The taint name is automatically filled into the Key field. Set **Value** and **Effect**, and then click **OK**.

     The following table describes the parameters.

Parameter	Description
Built-in Taints	Existing taints in the system.
Key	After you click a taint in the Built-in Taints field, the taint name is automatically filled into the **Key** field. You can also set **Key** to specify a custom taint.
Value	The custom taint value.
Effect	The effect of the taint. Valid values: ■ **PreferNoSchedule**: Kubernetes avoids scheduling pods that do not tolerate the taint onto the node. ■ **NoSchedule**: Pods that do not tolerate the taint are not scheduled on the node. ■ **NoExecute**: Pods are evicted from the node if they are already running on the node, and are not scheduled onto the node if they are not running on the node.

- Create a custom taint

  In the **Add Taint** dialog box, set **Key**, **Value**, and **Effect**, and then click **OK**.

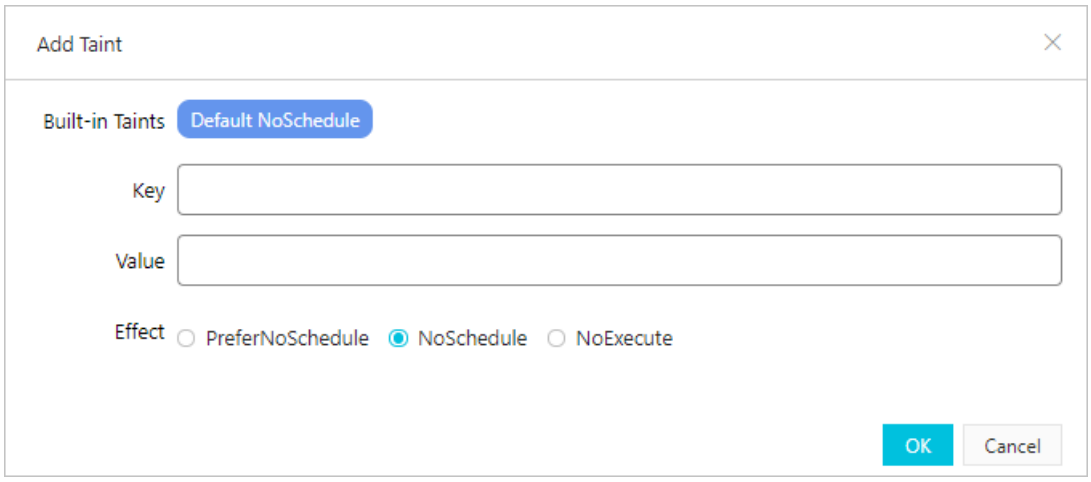

Add a taint on the Node Details page

On the Node Details page, you can add a taint to a node.

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Select one or more nodes to which you want to add a taint. Click **Add Taint** in the lower-left corner.
4. On the Node Details page, click the icon next to Tags and then perform the following operations:
   - Add a built-in taint

     In the **Add Taint** dialog box, click a taint in the Built-in Taints field. The taint name is automatically filled into the Key field. Set **Value** and **Effect**, and then click **OK**.
   - Create a custom taint

     In the **Add Taint** dialog box, set **Key**, **Value**, and **Effect**, and then click **OK**.

7.4.2.3. Query nodes by tag

You can filter nodes by tag to find nodes that have a specified tag.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, enter a tag name or specify a tag in the **key=value** format in the search box and then click the Search icon.

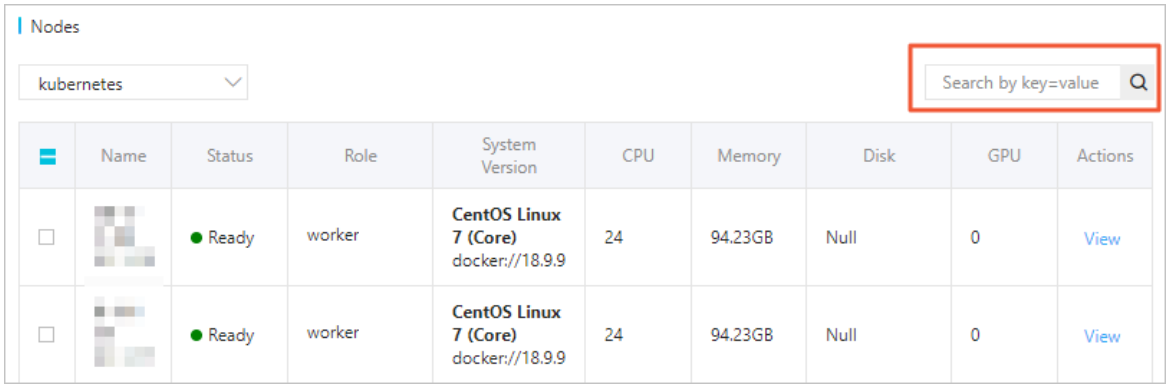

7.4.2.4. Delete a tag or taint

You can delete built-in or custom tags or taints from nodes. Kubernetes-defined tags of nodes cannot be deleted.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Find the target node and the target tag, and then click **View** in the **Actions** column.
4. In the dialog box that appears, move the pointer over the target tag or taint, and then click **Delete**.
5. In the message that appears, click **OK**.

7.4.3. Query events

The Events section displays various Kubernetes events.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Events** from the **Clusters** drop-down list.By default, alert events in all namespaces of Kubernetes clusters are displayed.
2. (Optional)To view the latest data of cluster events, click **Refresh** on the right side of the Cluster Event section.
3. You can filter events by cluster, namespace, and type.

7.5. Product center

7.5.1. Product list

The product list displays the information about all products deployed in the PaaS console, including their names and versions. In the product list, you can perform O&M operations and view product resources or register variables. You can also remove products that are no longer needed.

7.5.1.1. View product details

You can view the details of products deployed in the PaaS console, including their names, versions, and components.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.

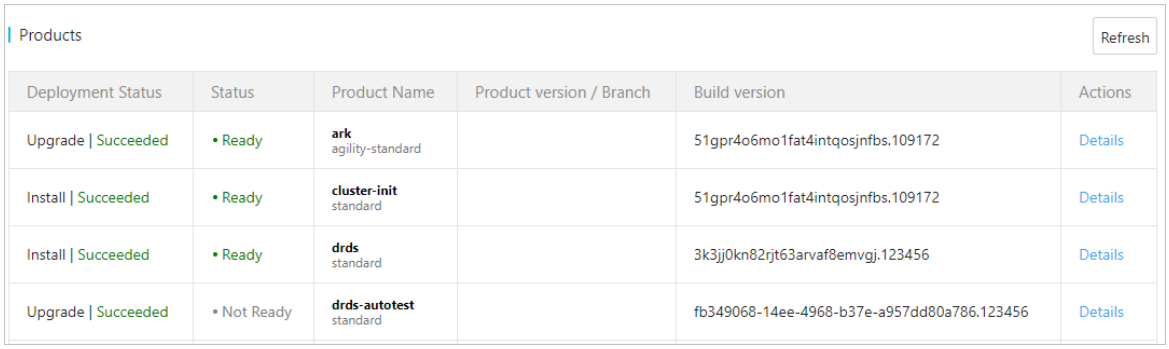

3. On the **Overview** page, view the name, version information, and components of the product.

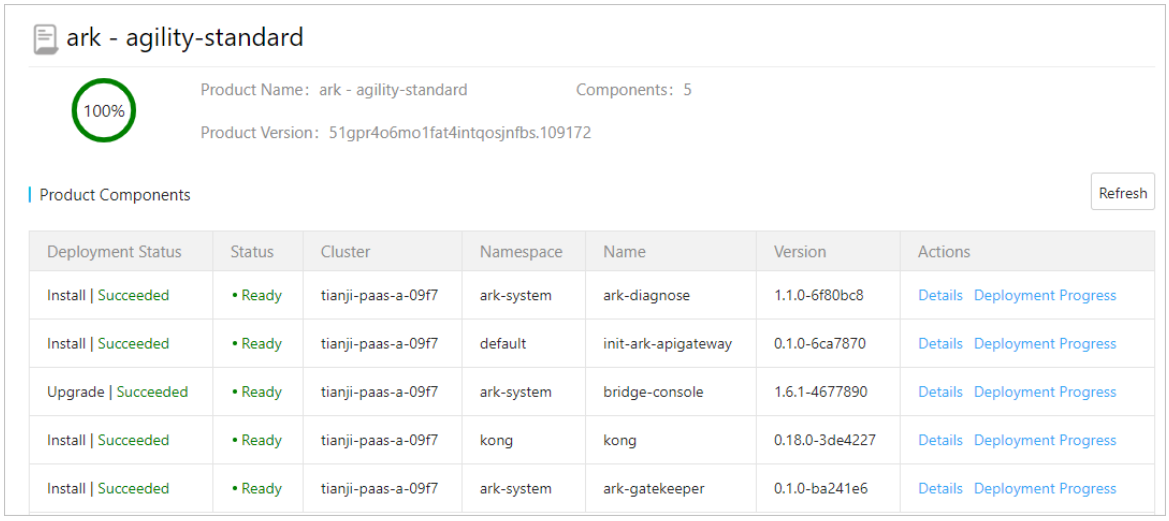

7.5.1.2. View component information

You can view the component details in the Product Components section of the Overview page of a product.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the Product Components section of the **Overview** page, view the deployment information of components, such as the deployment status, component status, cluster, namespace, component name, and component version.

4. Find a component and click **Details** in the **Actions** column to view details of the component.
5. The **Component Details** page contains the following tabs: StatefulSets, Deployments, DaemonSets, Jobs, Services, and Persistence Volume Claims.

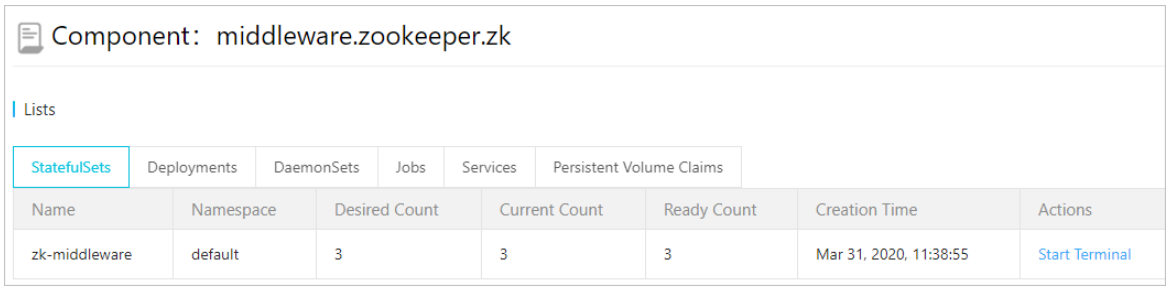

7.5.1.3. View the deployment progress of product components

You can view the deployment progress of product components.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the Product Components section of the **Overview** page, find the target component and click **Deployment Progress** in the **Actions** column.

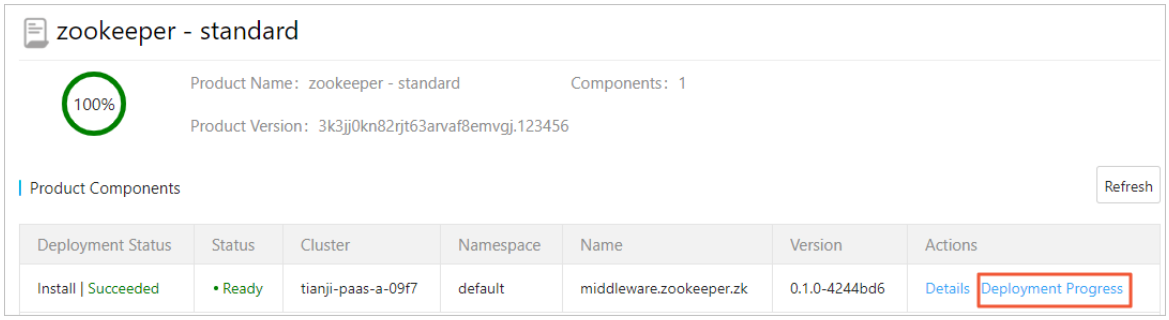

4. On the **Product Component Deployment Progress** page, click the deployment nodes in sequence to view the deployment progress and logs of the current component.

> **Note** You can click **LOGS** in the lower-left corner of the Summary tab to view the deployment logs.

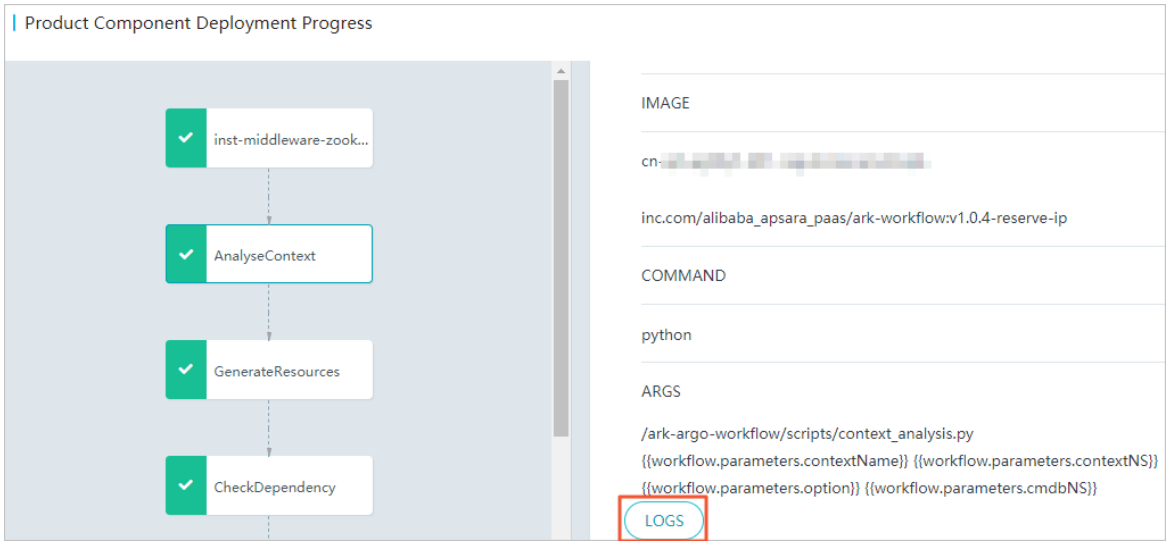

7.5.1.4. Log on to a web terminal

The **StatefulSets** and **Deployments** tabs of the **Component Details** page list available terminals. Browser-based terminals are mainly used for O&M management and troubleshooting.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center**.
3. In the product list, find the target product and click **Details** in the **Actions** column.
4. In the Product Components section of the **Overview** page, find the target component and click **Details** in the **Actions** column.
5. On the **Component Details** page, click the **StatefulSets** or **Deployments** tab.
6. Find the target component. Click **Start Terminal** in the **Actions** column. Available containers that are based on the number of replicas are displayed in the pane.

7. Select the target container and then click **OK** to start the terminal process.

7.5.1.5. View dependency topology of a product

The dependency topology of a product displays the references and dependencies among components of the product. It also displays the running status of each dependent component. This allows you to locate and analyze issues.

Procedure

1. In the left-side navigation pane of the PaaS console, select Products from the **Product Center** drop-down list.

2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the left-side navigation pane of the product page, click **Dependency Topology**.
4. Click the names of components in sequence next to **Components**, and then view the references and dependencies among the components.

   > **Note** Click **Show** or **Hide** to show or hide the names of the components.

   Move the pointer over a component in the topology. The details of the component are automatically displayed, including the name, platform, version, status, and check result.

   Move the pointer over a connection line that links the target component and another component. If two components reference each other, information such as the link type, reference service, and reference variables is displayed. If two components are dependent on each other, information such as the link type, dependent component, and dependency version is displayed.
5. (Optional)Click **Panorama** to view the panorama of all product components.

7.5.1.6. Perform O&M operations

The O&M Actions page displays the O&M operations that are available to a product. You can also perform O&M operations on this page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the left-side navigation pane, click **O&M Actions**.
4. Perform O&M operations that are available to the product.

7.5.1.7. View a resource report

The Resource Report page displays the information of all resources that a product has requested from the PaaS console. The resource type can be cni (ip), db, vip, dns, and accesskey.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the left-side navigation pane, click **Resource Report**.

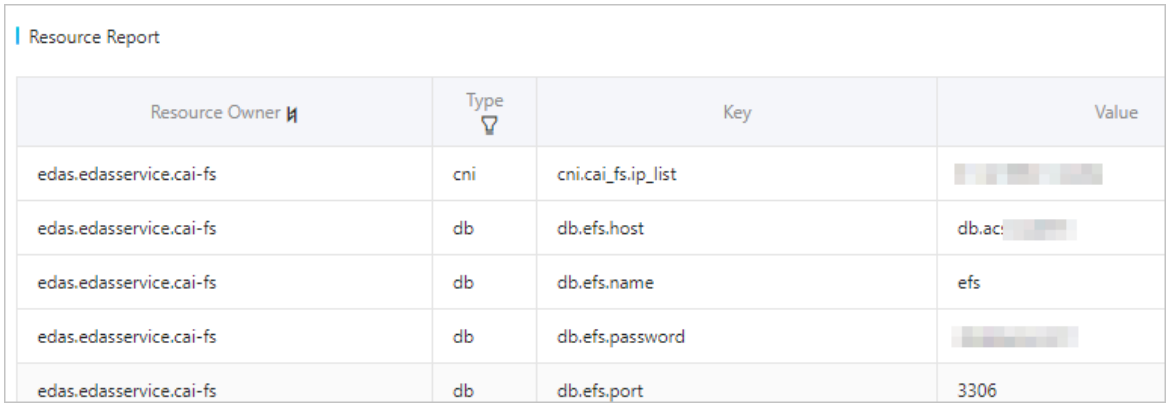
Resource Report

Resource Owner	Type	Key	Value
edas.edasservice.cai-fs	cni	cni.cai_fs.ip_list	
edas.edasservice.cai-fs	db	db.efs.host	db.ac
edas.edasservice.cai-fs	db	db.efs.name	efs
edas.edasservice.cai-fs	db	db.efs.password	
edas.edasservice.cai-fs	db	db.efs.port	3306

4. View the information of resources.

   By default, all resources are displayed. You can click the up and down arrows next to **Resource Owner** to sort resources. You can also click the icon next to **Type** to filter resources.

Field	Description
Resource Owner	The name of the component to which the resource belongs.
Type	The type of the resource.
Key	The attribute name of the resource.
Value	The attribute value of the resource.

7.5.1.8. View service registration variables

The Service Registration Variables page displays the values of all service registration variables. You can view the service registration variables of a product. The service registration variables report for a product lists the variables that the product can deliver to other products or components.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the left-side navigation pane, click **Service Registration Variables**.

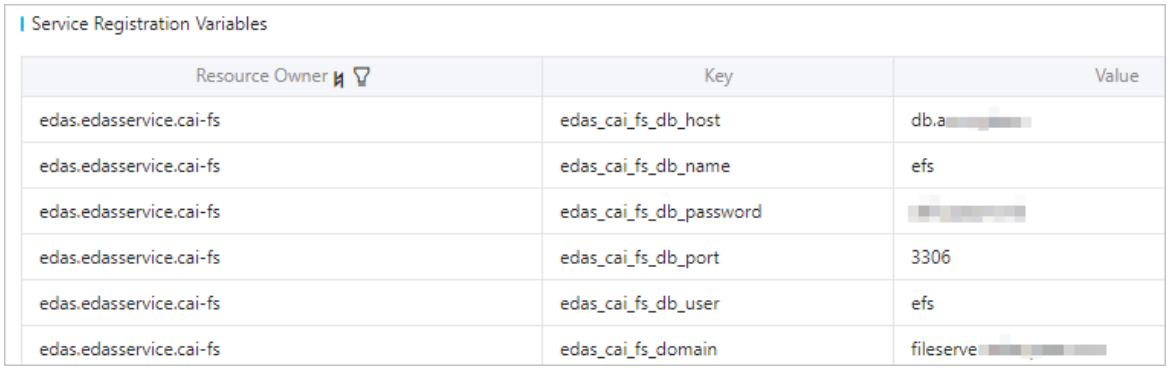
Service Registration Variables

Resource Owner	Key	Value
edas.edasservice.cai-fs	edas_cai_fs_db_host	db.a
edas.edasservice.cai-fs	edas_cai_fs_db_name	efs
edas.edasservice.cai-fs	edas_cai_fs_db_password	
edas.edasservice.cai-fs	edas_cai_fs_db_port	3306
edas.edasservice.cai-fs	edas_cai_fs_db_user	efs
edas.edasservice.cai-fs	edas_cai_fs_domain	fileserve

4. View the information of service registration variables.

By default, all service registration variables are displayed. You can click the up and down arrows next to **Resource Owner** to sort service registration variables. You can also click the icon next to **Resource Owner** to filter service registration variables.

The following table describes the fields for service registration variables.

Field	Description
Resource Owner	The name of the component to which the resource belongs.
Key	The variable name that is registered on CMDB and can be used by this product or other product components.
Value	The variable value that is registered on CMDB.

7.5.2. Deployment and upgrade

This topic describes how to perform batch upgrade and incremental deployment. You can deploy a product by product feature. If the product supports custom configuration, the system automatically goes to the custom configuration page.

Prerequisites

The deployment upgrade package is imported to the PaaS console.

You can follow the import the deployment upgrade package in the following way:

1. Upload the installation disk used for deployment and upgrade to the bootstrap node in the on-site environment.
2. Log on to the bootstrap node over SSH.
3. Run the following command to import deployment packages and generate a deployment package list:

   sh upgrade.sh *{packages -path}.iso*

   Replace *{packages -path}.iso* with the actual storage path of the iso file on the installation disk.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Deploy&Upgrade** from the **Products** drop-down list. The **System Packages** page displays deployment packages that have been imported to the PaaS console.
2. Find the target deployment package on the **System Packages** page.

   > **Note** If multiple deployment upgrade packages exist, you can enter the system ID in the search box to query a deployment package that meets the condition.

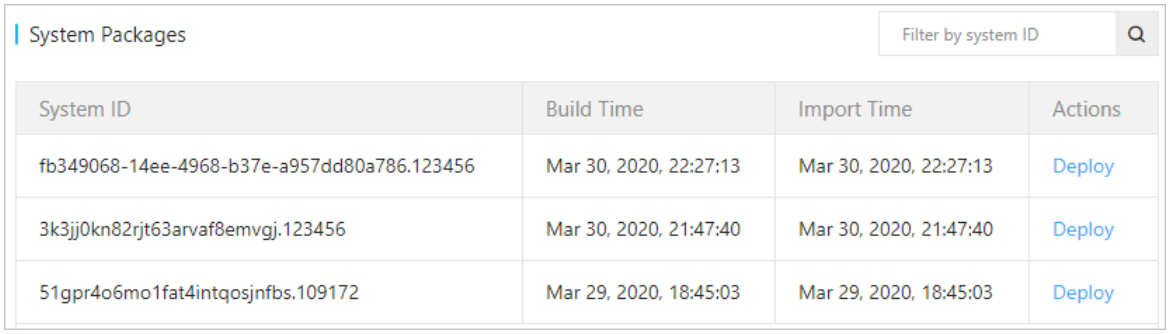

3. Click **Deploy** in the **Actions** column to start the deployment or upgrade process.
4. (Optional)In the **Select Products** step, click the number in the **Components** column to view the components and versions of the current product.

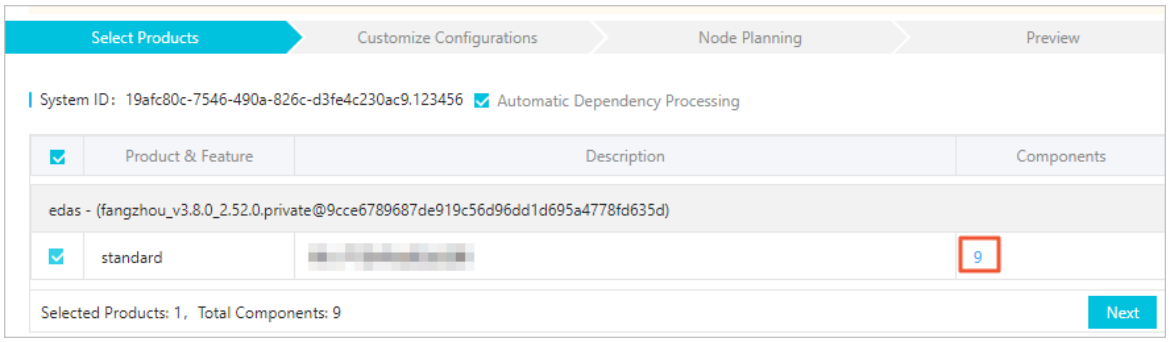

5. Select the required features and click **Next**.

> **Note** The system can automatically parse dependencies among products. When the Automatic Dependency Processing check box is selected, the system automatically checks whether dependencies exist between the deployed products and the products to be deployed and then select dependent products. If you want to manually select the products to be deployed, you can clear the **Automatic Dependency Processing** check box. If you select products that have been deployed, the system upgrades these products. If you select products that have not been deployed, the system performs incremental deployment on these products.

If the custom configuration feature is enabled for a selected product, the Customize Configurations page is displayed. Otherwise, the **Preview** page is displayed.

6. In the **Customize Configurations** step, configure the parameters as prompted and then click **Save**. Click **Next**.
7. In the **Node Planning** step, verify that the node planning is correct and click **Next**.

> **Note** If you want to modify the node planning, click **Reselect**.

8. In the **Preview** step, check the information of the products to be deployed.

   Icons before **Product & Feature** indicate different states of products:

   - : indicates that the product is newly deployed.
   - : indicates that the product has been deployed and does not need to be updated.
   - : indicates that the product has been updated. You can click the icon to check the

differences.

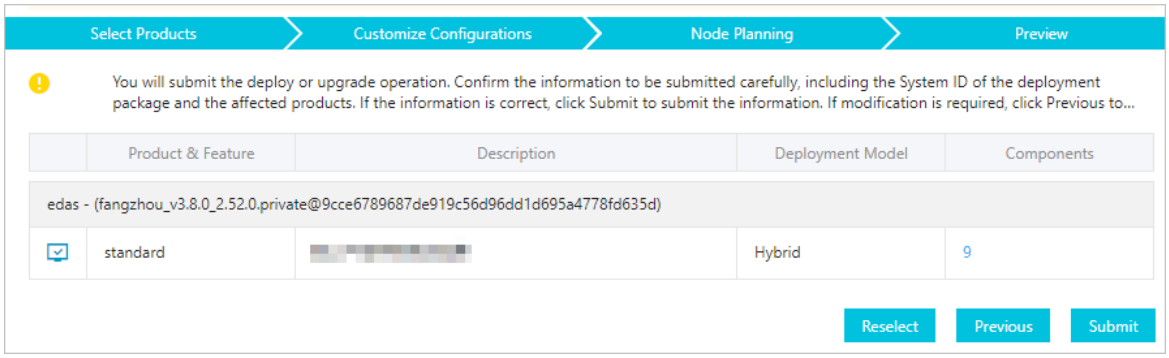

9. Click **Submit** to start the deployment or upgrade process.

   After the deployment process starts, you can view the progress on the Task Instances page by choosing **Task Center > Task Instances**.

7.6. Task center

The Task Center module provides general task management capabilities. You can view and run task templates, and view, suspend, resume, terminate, and delete tasks.

7.6.1. Task templates

The Task Templates page lists all task templates, both imported and preset.

7.6.1.1. View a task template

You can view information of all task templates on the Task Templates page, such as the name, description, parameters, and workflow definition.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Templates** from the **Task Center** drop-down list.
2. Find the target task template. Click **View** in the **Actions** column.

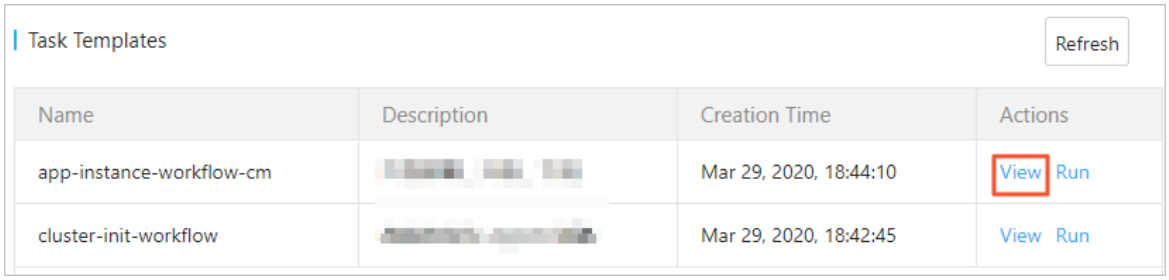

3. In the pane that appears, view the name, description, parameters, and workflow definition of the task template.

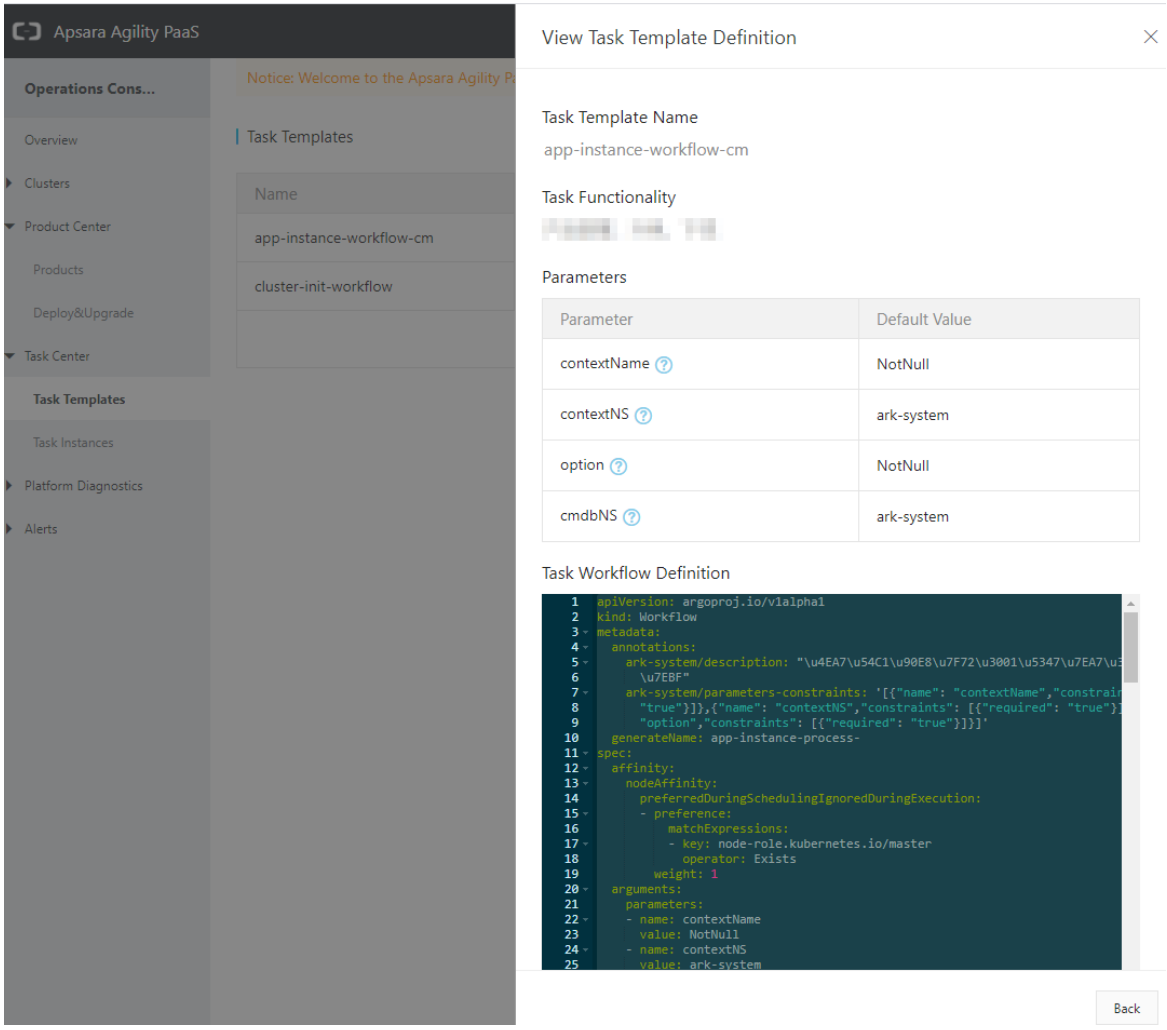

7.6.1.2. Run a task

You can run a task on the Task Templates page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Templates** from the **Task Center** drop-down list.
2. Find the target task template. Click **Run** in the **Actions** column.
3. In the pane that appears, set Task Instance Name and Action Parameters.

> **Note** If the task instance name is not specified, the system automatically generates a task instance name. We recommend that you enter a recognizable name for easy query.

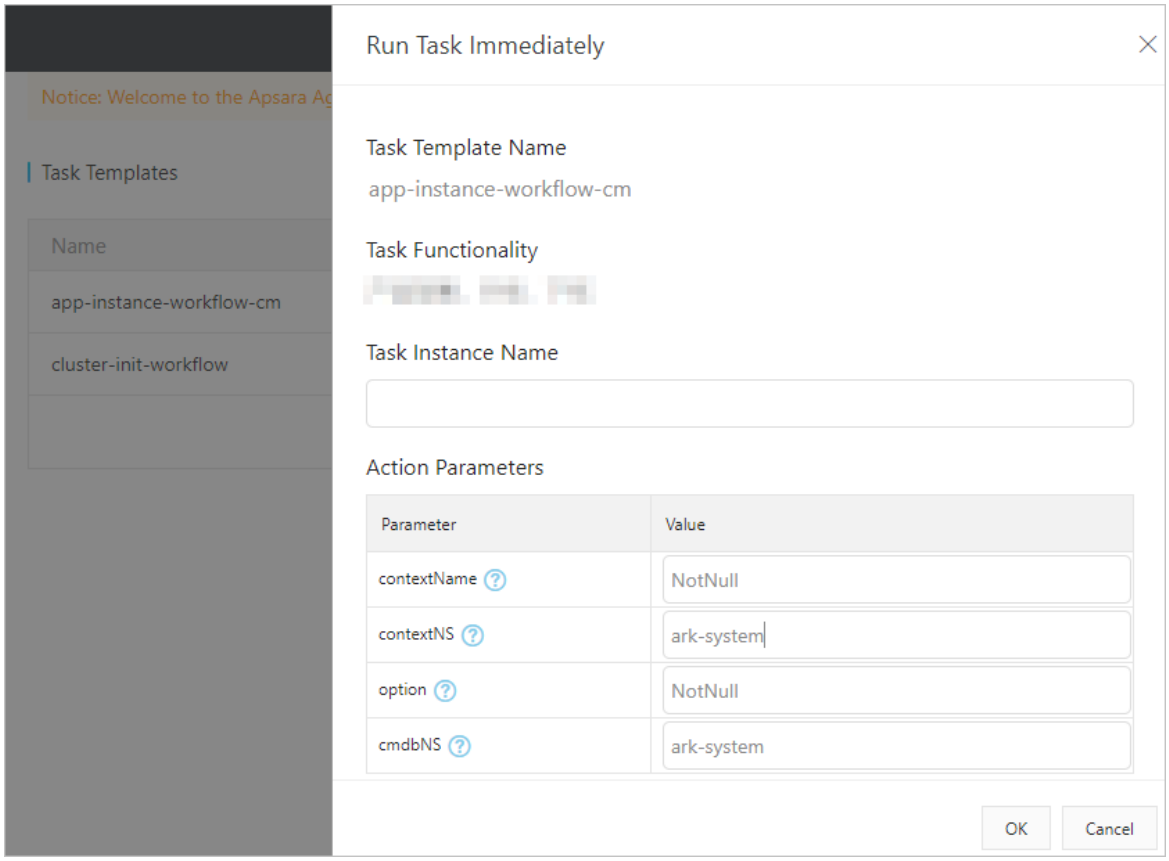

4. Click **OK**.

7.6.2. Task instances

The Task Instances page displays information of all tasks. On this page, you can view, suspend, resume, terminate, retry, and delete tasks.

7.6.2.1. View task details

After you run a task, you can view the progress, logs, and parameters of the task on the Task Instances page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, view the status of all tasks.

   Valid values of the task status:

   - **Succeeded**: indicates that the task has been executed.
   - **Running**: indicates that the task is being executed.
   - **Running (Suspended)**: indicates that the task has been suspended.
   - **Failed**: indicates that the task has failed.
   - **Failed (Terminated)**: indicates that the task has been terminated.
3. Find the target task. Click **View** in the **Actions** column. Then, you are redirected to the **Task**

Instance Details page.

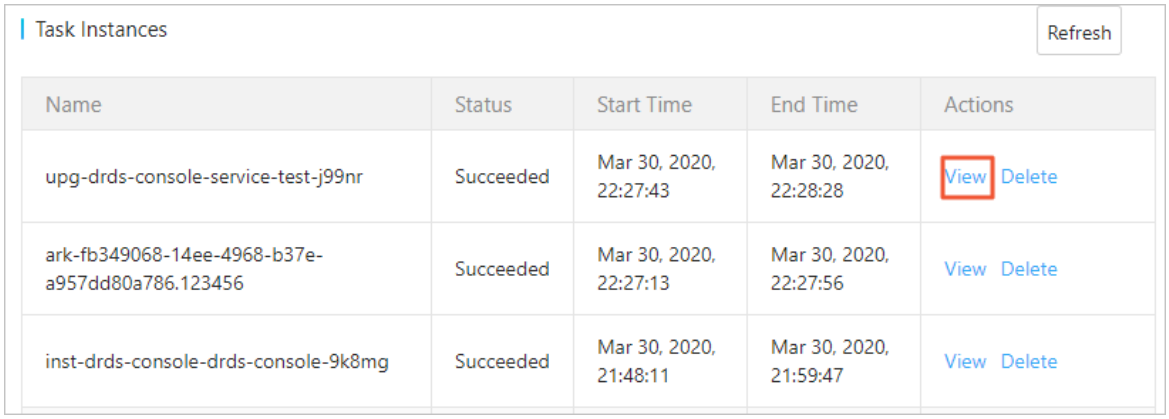

4. On the Task Instance Details page, click the task nodes in sequence to view the information and logs of the current task.

> **Note** You can click **LOGS** in the lower-left corner of the Summary tab to view task logs.

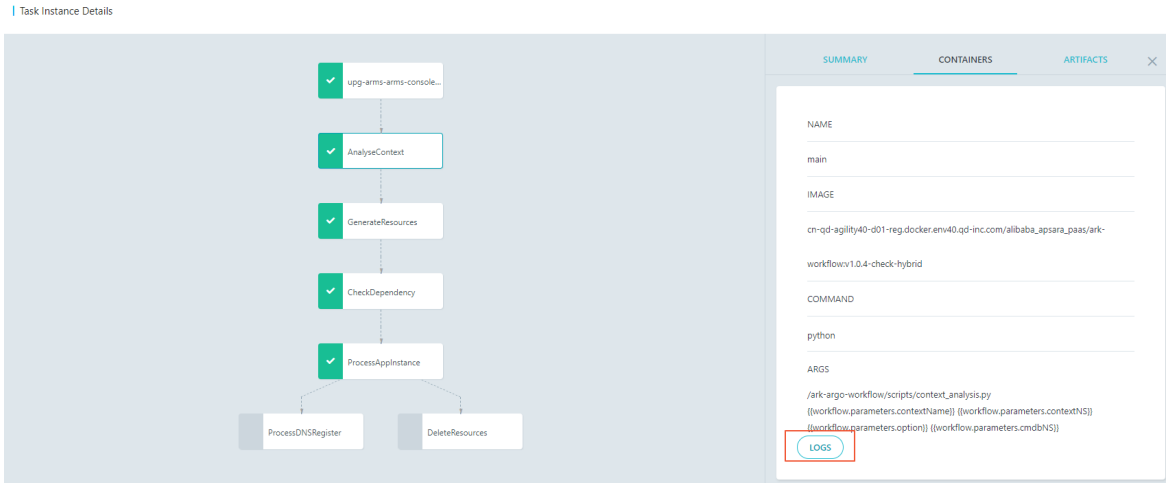

7.6.2.2. Suspend a task

You can suspend a task in the Running state. Then, the task status becomes Running (Suspended).

Prerequisites

The task is in the **Running** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find the task in the **Running** state that you want to suspend. Click **Suspend** in the **Actions** column. After a successful operation, the task status changes from **Running** to **Running (Suspend)** in the **Status** column.

7.6.2.3. Resume a task

After a task is suspended, the task is in the Running (Suspended) state. Then, you can click Resume in the Actions column to resume the task.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find the task in the **Running (Suspended)** state that you want to resume. Click **Resume** in the **Actions** column. After a successful operation, the task status changes from **Running (Suspend)** to **Running** in the **Status** column.

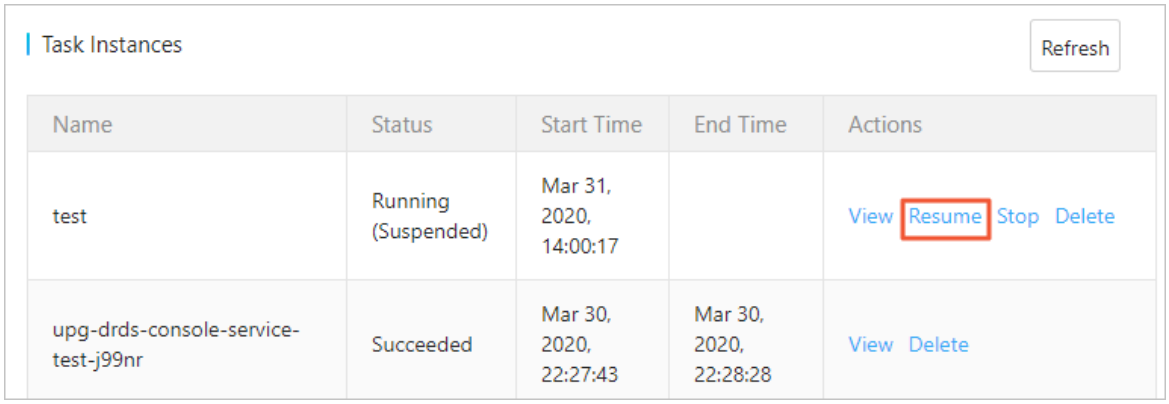

7.6.2.4. Terminate a task

You can terminate a task in the Running (Suspended) or Running state.

Prerequisites

The task is in the **Running (Suspended)** or **Running** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find a task in the **Running (Suspended)** or **Running** state. Click **Stop** in the **Actions** column. For a task in the Running the system immediately terminates the task and the task status becomes **Failed (Terminated)**. For a task whose **Status** is **Running (Suspended)**, the system immediately terminates the task when the task status becomes Running again. Then the task status becomes **Failed (Terminated)**.

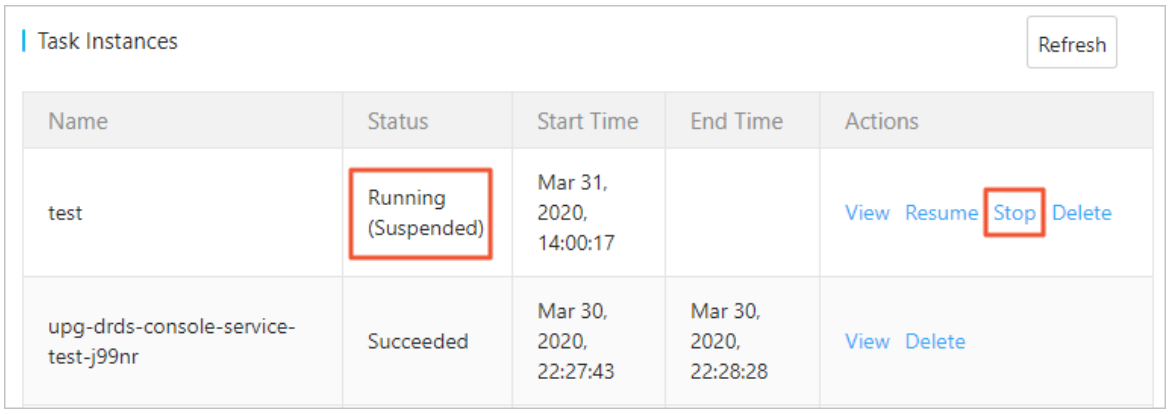

7.6.2.5. Retry a task

You can retry a task in the Failed or Failed (Terminated) state. When a task is retried, the task restarts from the failed or terminated task node.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find a task in the **Failed** or **Failed (Terminated)** state. Click **Retry** in the **Actions** column.

7.6.2.6. Delete a task

You can delete a task in any state. If a task is in the Running state, this operation enables the system to immediately terminate the task and delete the task record. If a task is in a state other than Running, this operation enables the system to immediately delete the task record.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find the target task. Click **Delete** in the **Actions** column.

7.7. Platform diagnostics

The PaaS console provides platform-level diagnostics. This module collects information about the console and products deployed in the console, presents summary diagnostic results, and allows you to download detailed diagnostic results. The module aims to improve user experience of diagnostics.

7.7.1. Diagnostic items

The Diagnostic Items page displays all diagnostic items in the PaaS console. On this page, you can view, execute, and delete diagnostic items.

7.7.1.1. View a diagnostic item

You can view details about the current diagnostic item, such as the name, type, description, start time, deletion protection, and definition.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Items** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic item. Click **View** in the **Actions** column.
3. In the pane that appears, view details of the diagnostic item.

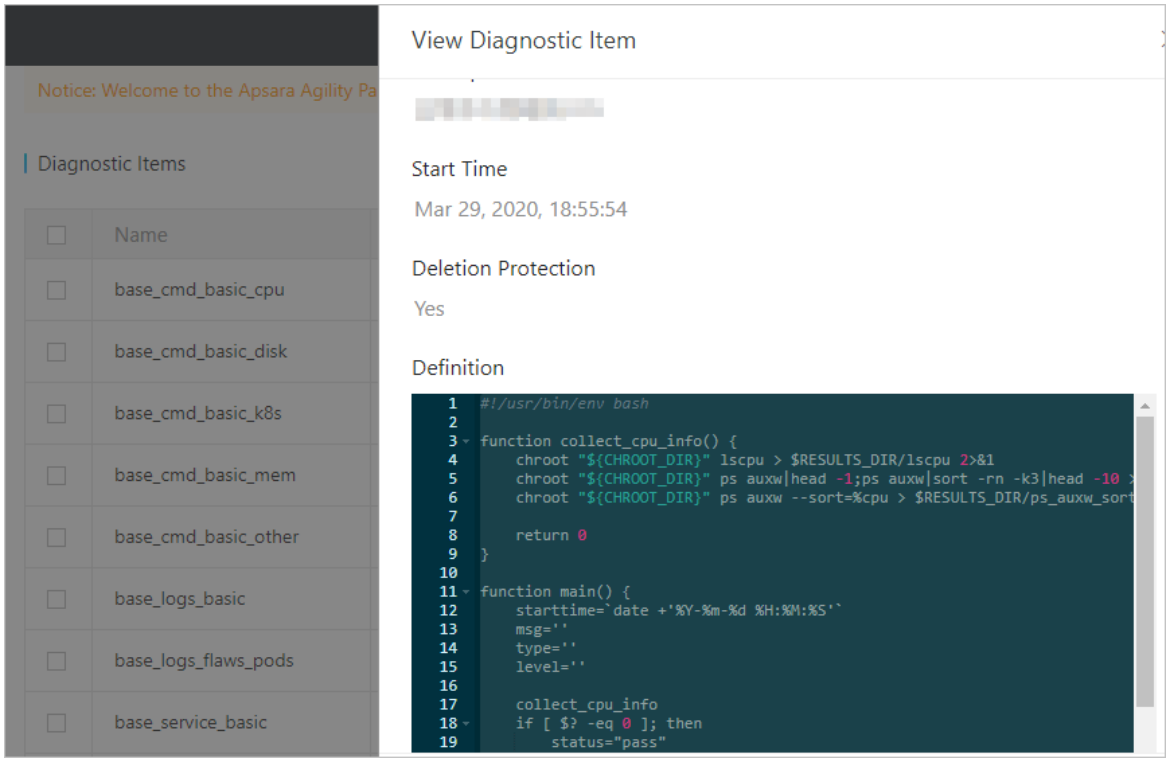

7.7.1.2. Execute diagnostic items

You can execute diagnostic items on the Diagnostic Items page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Items** from the **Platform Diagnostics** drop-down list.
2. Select one or more diagnostic items and click **Submit Diagnosis**.

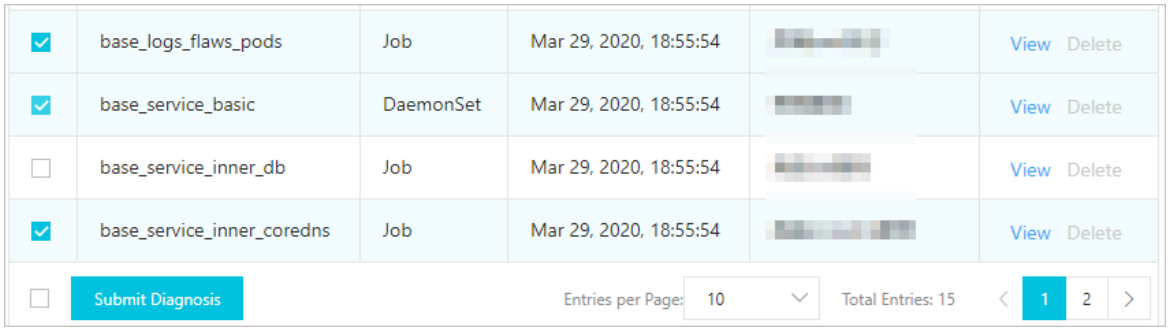

3. In the message that appears, click **OK**.

7.7.1.3. Delete a diagnostic item

You can delete a diagnostic item. You can only delete imported diagnostic items, but not the diagnostic items preset by the system.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Items** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic item. Click **Delete** in the **Actions** column.
3. In the message that appears, click **OK**.

7.7.2. Diagnostic tasks

The Diagnostic Tasks page displays all diagnostic tasks. On this page, you can view diagnostic progress, view diagnostic reports, download diagnostic reports, terminate diagnostic tasks, and delete diagnostic tasks.

7.7.2.1. View diagnostic progress

After you start a diagnostic task, you can view its diagnostic progress on the Diagnostic Tasks page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Click **Diagnostic Progress** in the **Actions** column.
3. On the **Diagnostic Progress** page, click the task nodes in sequence to view the diagnostic progress and logs of the current diagnostic task.

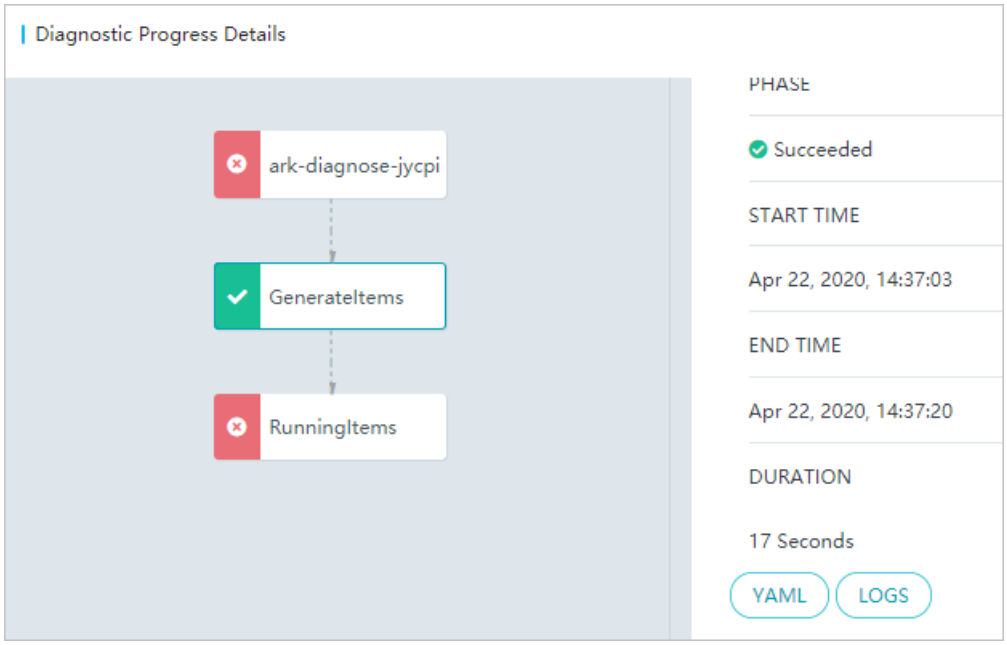

7.7.2.2. View a diagnostic report

After a diagnostic task is complete, you can view its diagnostic report.

Prerequisites

You can view the diagnostic report only for a diagnostic task in the **Succeeded** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Click **View Report** in the **Actions** column.
3. In the pane that appears, view the diagnostic results, such as the name, status, and details.

7.7.2.3. Download a diagnostic report

After a diagnostic task is complete, you can download its diagnostic report to your on-premises machine for offline query and analysis.

Prerequisites

You can download the diagnostic report only for a diagnostic task in the **Succeeded** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Click **Download** in the **Actions** column.

7.7.2.4. Terminate a diagnostic task

You can terminate a diagnostic task in the Running state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Choose **More > Terminate** in the **Actions** column.
3. In the message that appears, click **OK**.

7.7.2.5. Delete a diagnostic task

You can delete a diagnostic task that is no longer needed.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Choose **More > Delete** in the **Actions** column.
3. In the message that appears, click **OK**.

7.8. Alerts

The **Alerts** module implements unified management of alerts in the PaaS console. You can view alert rules, notification channels, and alert events. You can also configure alert rules and notification channels in the **Alerts** module.

7.8.1. Alert rule groups

An alert rule must belong to an alert rule group. You can create alert rule groups and add alert rules to alert rule groups.

7.8.1.1. Create an alert rule group

You can create an alert rule group. When you create an alert rule group, you must add an alert rule to the group.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Create Rule Group**.
4. In the **Create Rule Group** dialog box, configure the parameters.

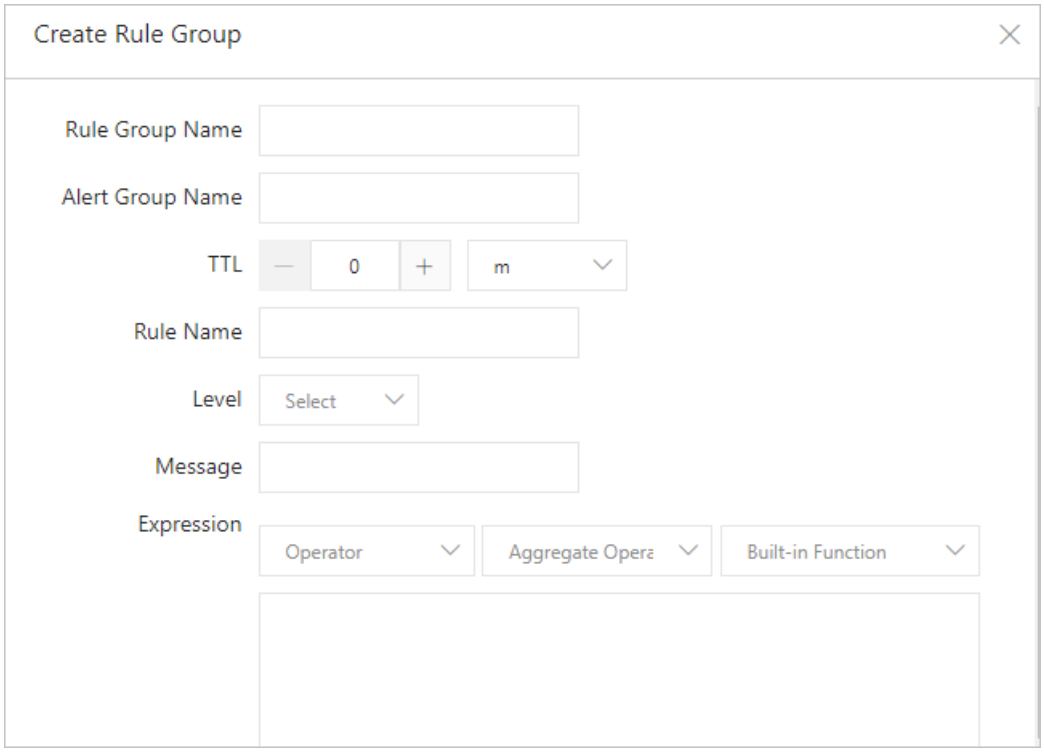

Parameter	Description
Rule Group Name	The globally unique name of the alert rule group.
Alert Group Name	The globally unique name of the alert group. An alert rule group must have an alert group.
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: indicates hours. ○ **m**: indicates minutes. ○ **s**: indicates seconds.
Rule Name	The globally unique name of the alert rule.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

5. Click **Submit**.

7.8.1.2. Create an alert rule

After you create an alert rule group, you can add an alert rule to the group.

Prerequisites

An alert rule group is created. For more information about how to create an alert rule group, see Create an alert rule group.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. Find the target rule group. Click **Modify Rule** in the **Actions** column.
   The **Rules** page is displayed. You can view all alert rules in the alert rule group.

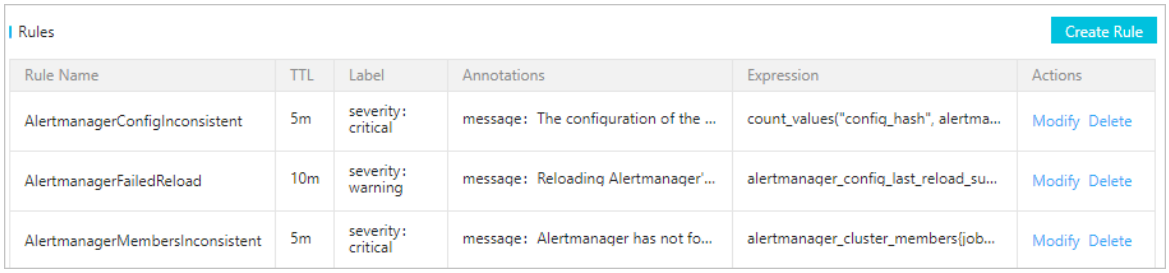

4. In the upper-right corner of the page, click **Create Rule**.
5. In the **Create Rule** dialog box, configure the parameters.

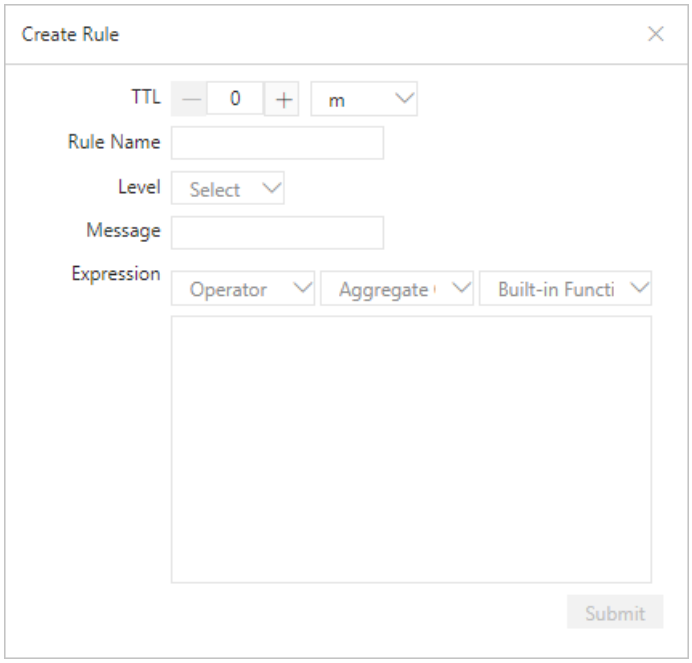

Parameter	Description

Parameter	Description
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: indicates hours. ○ **m**: indicates minutes. ○ **s**: indicates seconds.
Rule Name	The globally unique name of the alert rule.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

6. Click **Submit**.

7.8.1.3. Modify an alert rule

You can modify an alert rule.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. On the **Rule Groups** page, view all alert rule groups defined in the system.
4. Find the rule group for the target rule. Click **Modify Rule** in the **Actions** column.
5. On the **Rules** page, view all alert rules in the rule group.
6. Find the target rule. Click **Modify** in the **Actions** column.
7. Modify the TTL, Level, Message, and Expression parameter settings of the alert rule.

Parameter	Description

Parameter	Description
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: indicates hours. ○ **m**: indicates minutes. ○ **s**: indicates seconds.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

8. Click **Submit**.

7.8.1.4. Delete an alert rule

You can delete an alert rule which is no longer needed from an alert rule group.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Groups**.
3. On the **Rule Groups** page, view all alert rule groups defined in the system.
4. Find the rule group for the target rule. Click **Modify Rule** in the **Actions** column.
5. On the **Rules** page, view all alert rules in this rule group.
6. Find the target rule. Click **Delete** in the **Actions** column.
7. In the message that appears, click **OK**.

7.8.1.5. Delete an alert rule group

You can delete an alert rule group that is no longer needed.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list. The **Rule Groups** page is displayed.
2. In the upper part of the page, select the target cluster from the drop-down list.

3. Find the target rule group. Click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

7.8.2. Notification channels

You can view and modify notification channel settings on the Notification Channels page.

7.8.2.1. View notification channel settings

You can view the current notification channel settings.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the **Global Settings**, **Routing**, and **Receiver** sections, view the relevant information.

7.8.2.2. Modify notification channel settings

You can modify notification channel settings such as global settings, routing, and receivers.

7.8.2.2.1. Modify global settings

You can modify global settings, such as the resolve_timeout, smtp_info, and notifications settings.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Edit**.
4. In the **Global Settings** section, modify the resolve_timeout, smtp_info, and notifications settings.

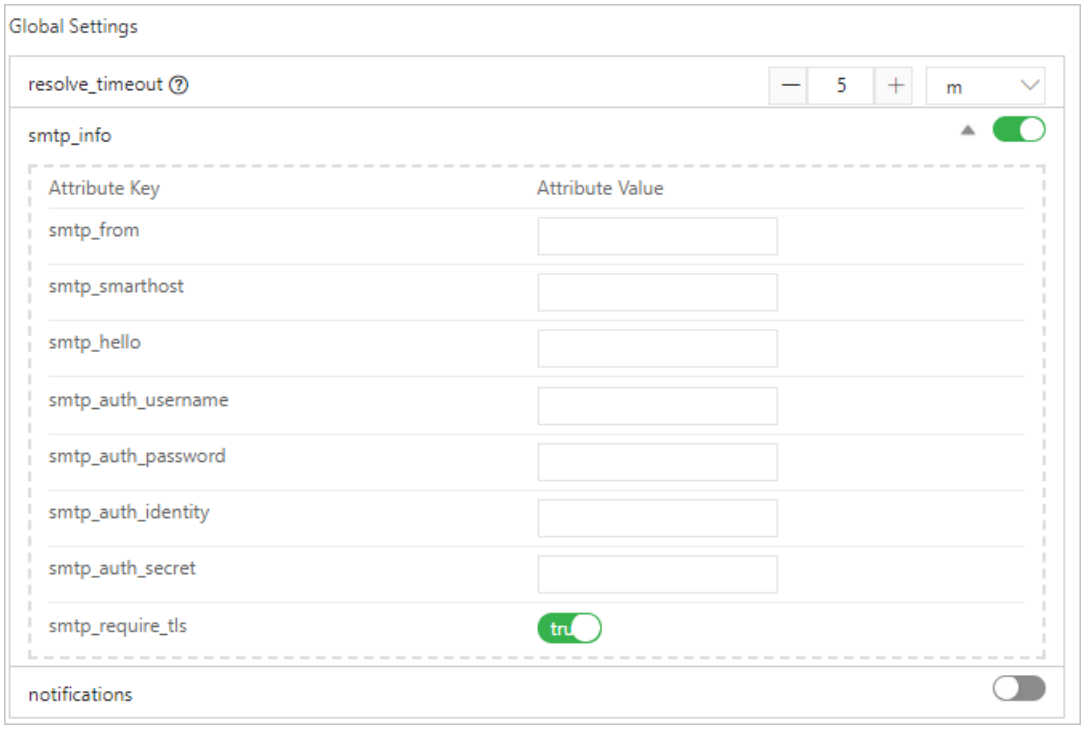

Parameter	Description
resolve_timeout	Specifies the time period before an alert is marked as resolved if the Alertmanager does not receive further notifications of the alert.
smtp_info	Specifies global SMTP information. To modify this item, turn on the switch on the right and then click the Show icon. You can configure the following parameters: ○ **smtp_from**: the source email address used to send alerts. ○ **smtp_smarthost**: the SMTP server address and port number for the source email address used to send alerts. Example: smtp_smarthos t:smtp.example.com:465 ○ **smtp_hello**: the default hostname that identifies the SMTP server. ○ **smtp_auth_username**, **smtp_auth_password**: the username and password for the source email address used to send alerts. ○ **smtp_auth_identity**: specifies the PLAIN SMTP authentication method. ○ **smtp_auth_secret**: specifies the CRAM-MD5 SMTP authentication method. ○ **smtp_require_tls**: the default SMTP TLS configuration. Although the default value is **true**, the parameter is typically set to **false** to avoid starttls errors that occur if the parameter is set to **true**.

Parameter	Description
notifications	The Slack configuration. To modify this item, turn on the switch on the right and then click the Show icon. You can configure the following parameters: ○ **slack_api_url**: the API URL for Slack notifications. ○ **victorops_api_key**: the VictorOps API key. ○ **victorops_api_url**: the VictorOps API URL. ○ **pagerduty_url**: the destination URL for API requests. ○ **opsgenie_api_key**: the Opsgenie API key. ○ **opsgenie_api_url**: the destination URL for Opsgenie API requests. ○ **hipchat_api_url**: the source URL for API requests. ○ **hipchat_auth_token**: the authentication token. ○ **wechat_api_url**: the WeChat API URL. ○ **wechat_api_secret**: the WeChat API key. ○ **wechat_api_corp_id**: the WeChat API corporate ID.

5. In the upper-right corner of the page, click **Save**.

7.8.2.2.2. Modify routing settings

You can modify global routing settings, and create or delete sub-routes.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Edit**.
4. In the **Routing** section, perform the following operations:
   - Modify routing settings

     You can modify the default route or subroutes.

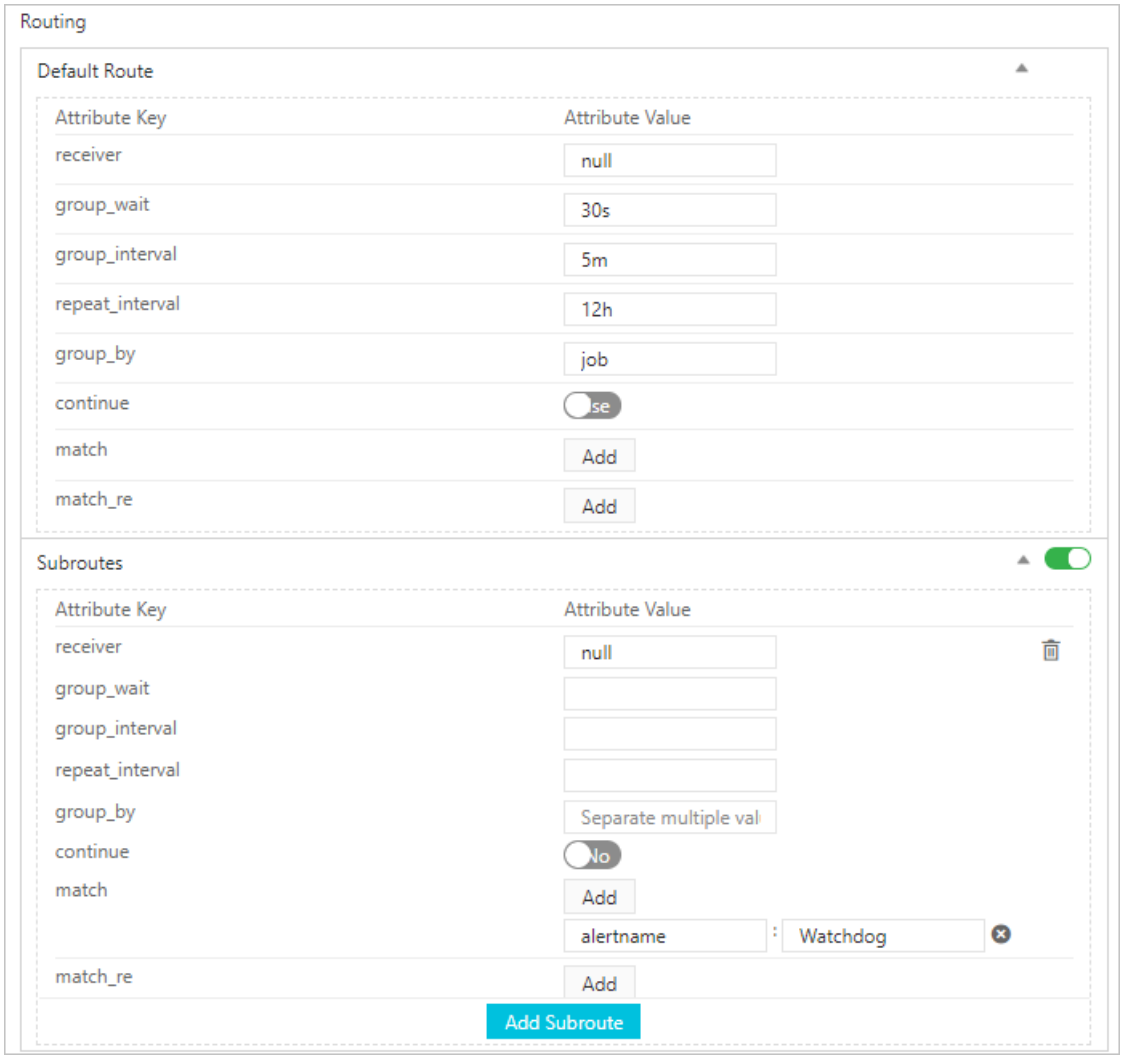

Parameter	Description

Parameter	Description
Default Route	The global route. You can configure the route information based on the actual environment. ■ **receiver**: the name of the alert receiver. ■ **group_wait**: specifies the waiting time to initialize a message when a new alert group is created. This method ensures that the system can have enough time to obtain multiple alerts for the same alert group, and then trigger an alert message. ■ **group_interval**: specifies the waiting time to send a new alert message. ■ **repeat_interval**: specifies the waiting time to resend an alert message. ■ **group_by**: the tag list. It is the regrouping tag list after alert messages are received. For example, all received alert messages that contain the `cluster=A` and `alertname=Latncy High` tags are aggregated into a group. ■ **continue**: specifies whether an alert matches subsequent nodes. ■ **match**: Click **Add** and specify a receiver for matched alerts. ■ **match_re**: Click **Add**. Enter a regular expression and specify a receiver for alerts that match the regular expression.
Subroutes	Configure subroutes in a similar way to the global route, so that you can export an alert type to another location. ■ **receiver**: the name of the alert receiver. ■ **group_wait**: specifies the waiting time to initialize a message when a new alert group is created. This method ensures that the system can have enough time to obtain multiple alerts for the same alert group, and then trigger an alert message. ■ **group_interval**: specifies the waiting time to send a new alert message. ■ **repeat_interval**: specifies the waiting time to resend an alert message. ■ **group_by**: the tag list. It is the regrouping tag list after alert messages are received. For example, all received alert messages that contain the `cluster=A` and `alertname=Latncy High` tags are aggregated into a group. ■ **continue**: specifies whether an alert matches subsequent nodes. ■ **match**: click **Add**. Enter the key and value of a tag and specify a receiver for alerts that match the tag. ■ **match_re**: click **Add**. Enter a regular expression based on the key and value of a tag and specify a receiver for alerts that match the regular expression.

- Create a subroute

  To export an alert type to another location, you can click **Add Subroute** in the lower part of the **Routing** section to configure a new subroute.

- Delete a subroute

  In the **Routing** section, find a subroute that is no longer needed and click the Delete icon to delete the subroute.

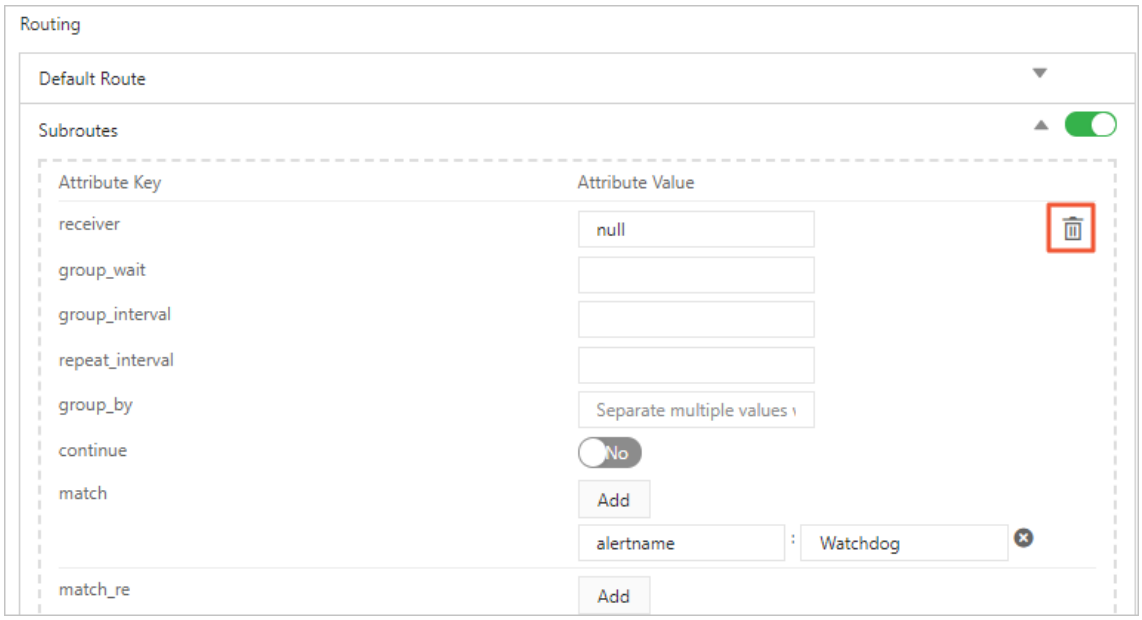

7.8.2.2.3. Modify receiver settings

You can create, modify, or delete alert receiver settings.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Edit**.
4. In the **Receivers** section, perform the following operations:
   - Modify receiver settings

     Modify the name and type of a receiver.

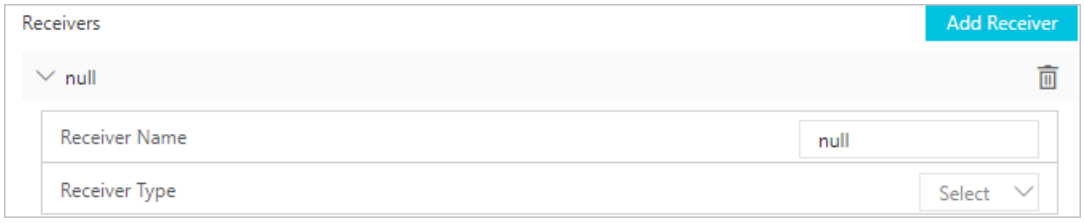

Parameter	Description
Receiver Name	The name of the alert receiver.

Parameter	Description
Receiver Type	Valid values for **Receiver Type**: **webhook** and **email**. If Receiver Type is set to **webhook**, you must configure the following parameters: ■ **url**: the URL of the alert receiver. ■ **send_resolved**: specifies whether to send messages for resolved alerts. Default value: **No**. If Receiver Type is set to **email**, you must configure the following parameters: ■ **send_resolved**: specifies whether to send messages for resolved alerts. Default value: **No**. ■ **to**: the destination email address for alerts. ■ **from**: the source email address used to send alerts. ■ **smarthost**: the server address and port number for the source email address used to send alerts. ■ **hello**: the default hostname that identifies the email server. ■ **auth_username**: the username for the source email address used to send alerts. ■ **auth_password**: the password for the source email address used to send alerts. ■ **auth_secret**: specifies the CRAM-MD5 authentication method. ■ **auth_identity**: specifies the PLAIN authentication method. ■ **require_tls**: the default TLS configuration. Although the default value is **Yes**, the parameter is typically set to **No** to avoid starttls errors that occur if the parameter is set to **Yes**.

- Add a receiver

  In the upper-right corner of the **Receivers** section, click **Add Receiver**. Configure the parameters.

- Delete a receiver

  In the **Receivers** section, find the target receiver and click the Delete icon to delete a receiver that is no longer needed.

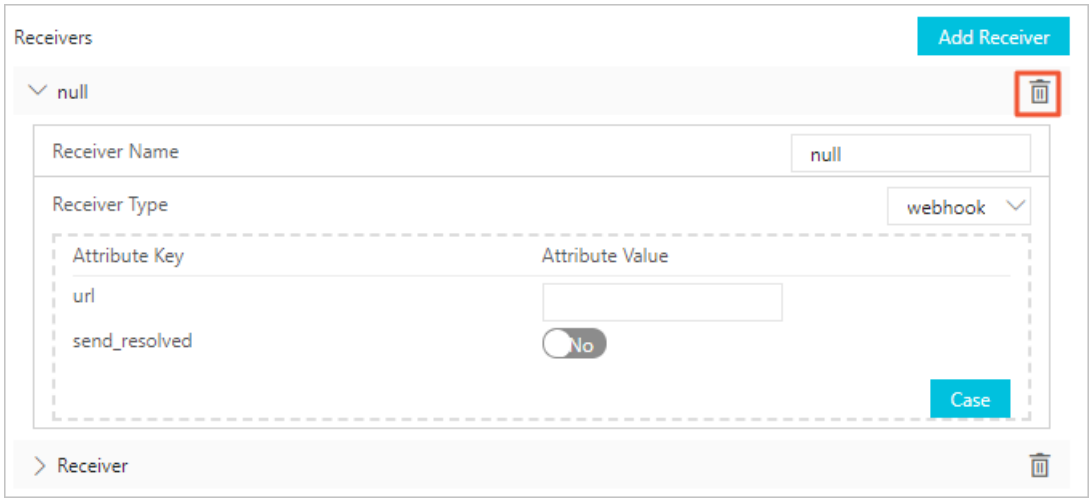
Receivers
Add Receiver
null
Receiver Name
null
Receiver Type
webhook
Attribute Key
Attribute Value
url
send_resolved
No
Case
Receiver

8.Network operations

8.1. Apsara Network Intelligence

8.1.1. What is Apsara Network Intelligence?

Apsara Network Intelligence is a system that can analyze network traffic. It provides data to facilitate resource planning, diagnostic functions, monitoring, system management, and user profiling.

You can use Apsara Network Intelligence to:

- Manage cloud service types.
- Query VPC and SLB instance details with a single click.
- Configure reverse access to cloud services.
- Configure leased lines by using graphical interfaces and set up primary and secondary routes.
- Query the tunnel VIPs of cloud services.
- Create Layer 4 listeners.

8.1.2. Log on to the Apsara Network Intelligence console

This topic describes how to log on to the Apsara Network Intelligence console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

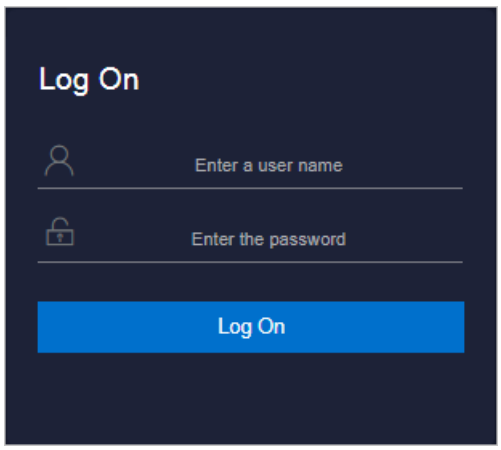

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Infrastructure as a Service (IaaS)** section, click **Apsara Network Intelligence**.

8.1.3. Query information

You can enter an instance ID to query details of the instance.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. Enter the ID of a VPC or an SLB instance to query details.
   - Enter the ID of a VPC to query VPC, VRouter, and VSwitch details.
     - VPC details

       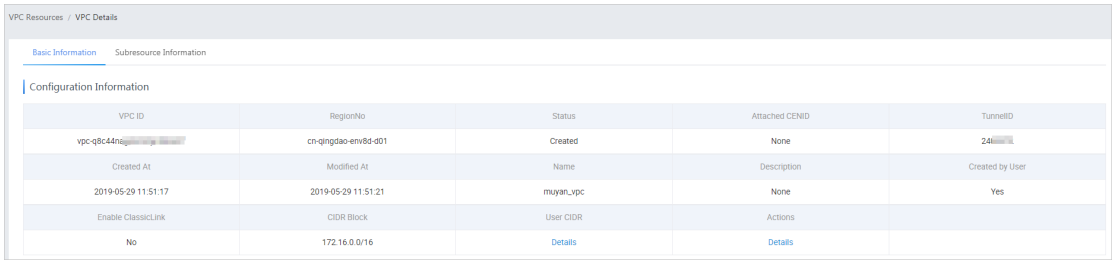

     - Information about VRouters, route tables, router interfaces, and VSwitches .
   - Enter the ID of an SLB instance to query instance details.

- Information about SLB instance configurations, VIPs, specifications, and users

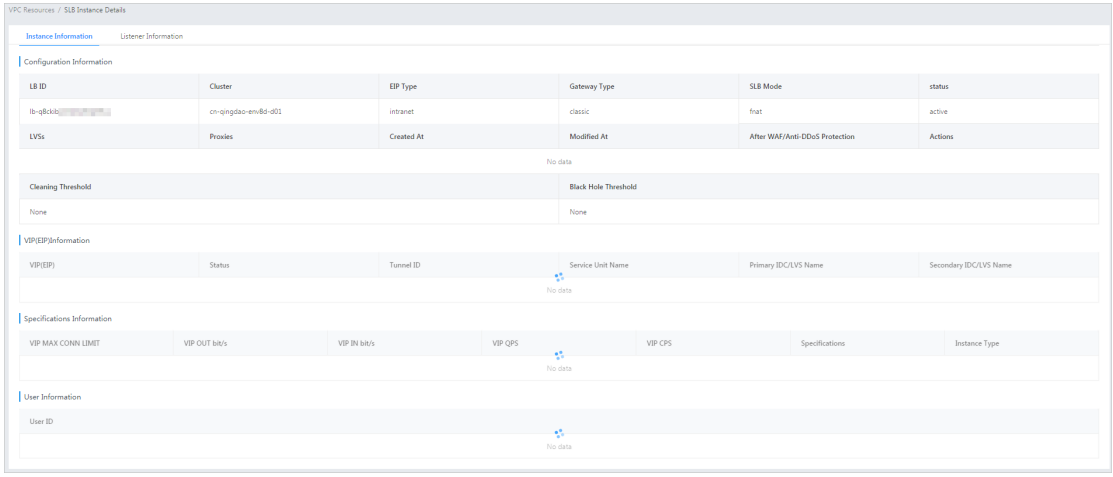

- Listener information

  Click **Show** in the **Back-end Server/Health Check** column to view details on backend servers.

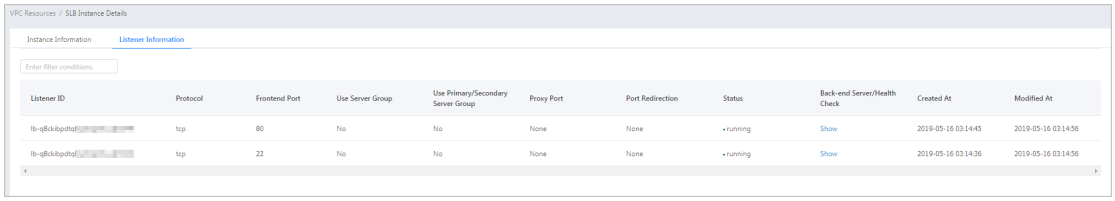

8.1.4. Manage cloud service instances

You can create a cloud service in a region or query the instance information of a region.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. From the **Products** menu, choose **Virtual Private Cloud > VPC Instance Type Management**.
3. Select the region from the **Select Region** drop-down list for which you want to create a cloud service instance. All cloud service instances in the specified region are displayed.
4. Click **Add** to add a cloud service type.

8.1.5. Tunnel VIP

8.1.5.1. Create a Layer-4 listener VIP

You can create Layer-4 listener VIPs to forward traffic for cloud services in your VPC.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. In the top navigation bar, click **Products** and choose **Server Load Balancer > VIP Management**.
3. Click **Create VIP**.
4. In the **Create VPC Instance** step, set the VPC instance parameters.

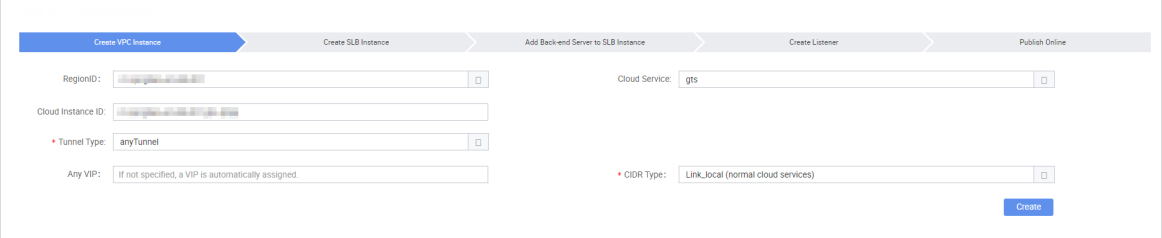

The following tunnel types are available:

- **singleTunnel**: specifies a single tunnel VIP that allows the Elastic Compute Service (ECS) instances in a single VPC to access external cloud services.
- **anyTunnel**: specifies a tunnel VIP that allows the ECS instances in all VPCs to access a specified cloud service.

5. Click **Create**.
6. In the **Create SLB Instance** step, select a primary data center or use the default data center.

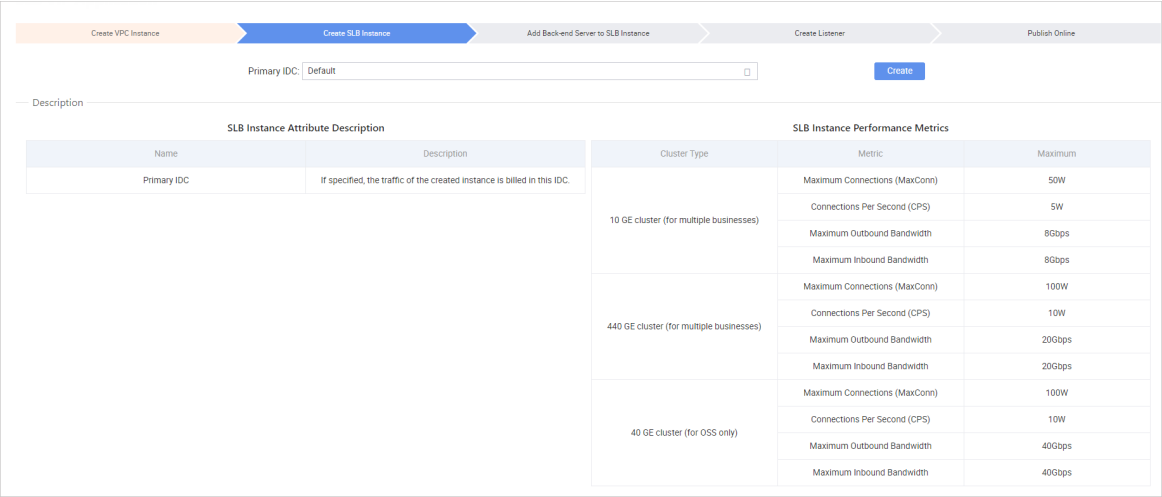

7. Click **Create**.
8. In the **Add Back-end Server to SLB Instance** step, specify the following information:
   - **VPC ID**: Enter the ID of the VPC to which target ECS instances belong. This parameter must be set if the target ECS instances are deployed in a VPC.
   - **Back-end Servers**: Specify the backend servers that you want to add. You can specify the server IP address and weight of only one backend server in each line. You can separate an IP address and the weight value with either a space or a comma (,). If no weight value is specified, the default value 100 is used.

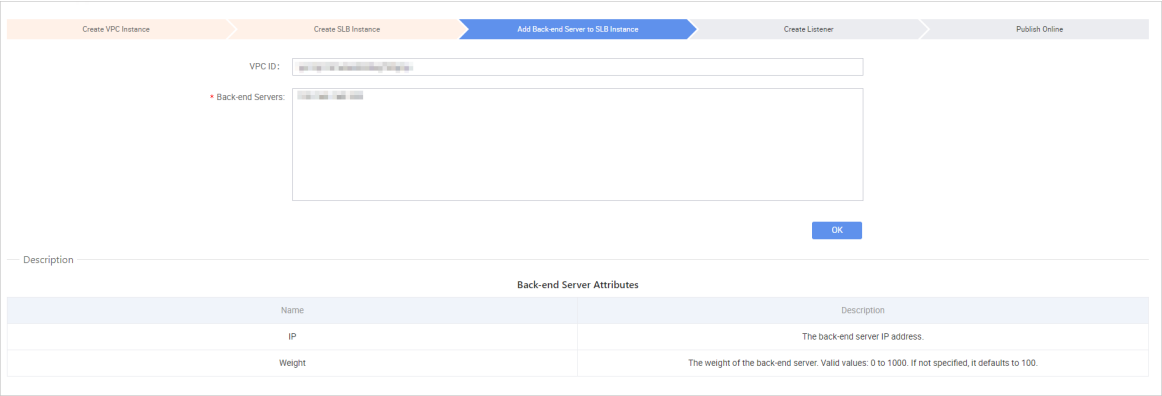

9. Click **OK**.
10. In the **Create Listener** step, click **Add** to configure a User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) listener.

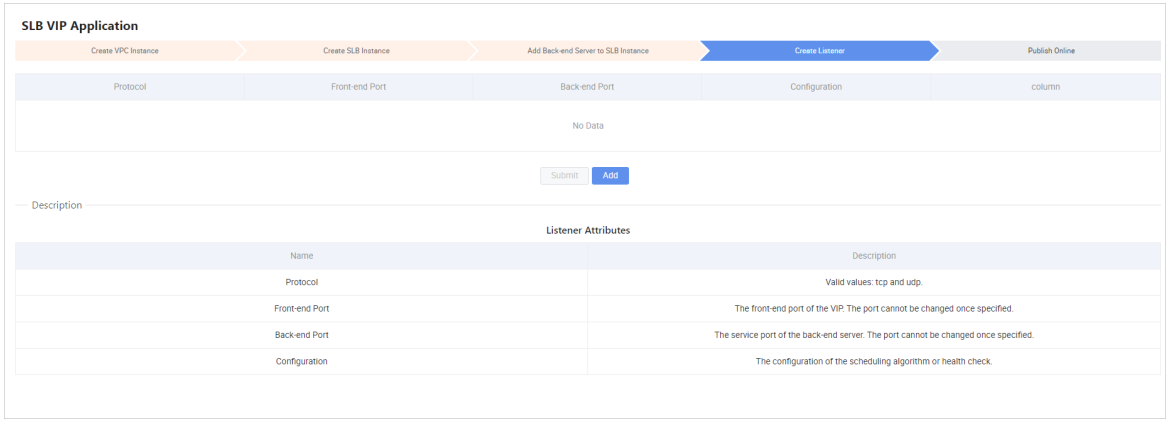

11. Click **Submit**.
12. In the **Publish Online** step, click **Yes** and then click **OK**.

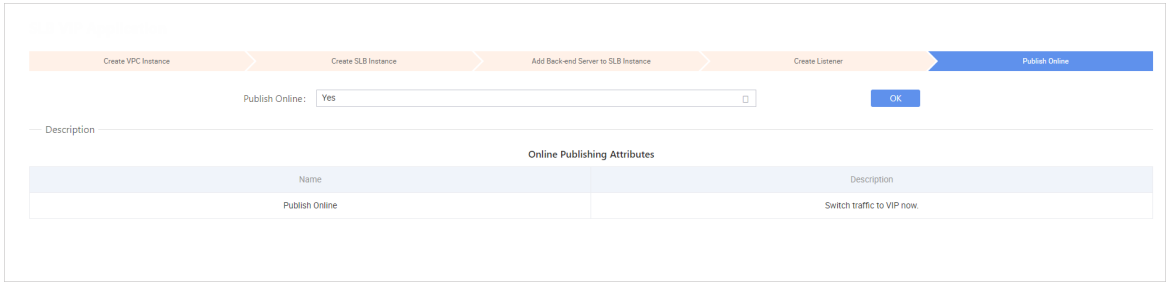

Result

The cloud services for which you have created the VIP can forward traffic through the created Layer-4 listener.

8.1.5.2. Query the tunnel VIP of a cloud service

You can query information such as creation time, connectivity, and VIP for cloud services that have Server Load Balancer (SLB) VIPs.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. From the **Products** menu, choose **Server Load Balancer > VIP Management**.
3. On the **Tunnel VIP Management** page, select Region ID, Cloud Service, and Status. Click **Search**.

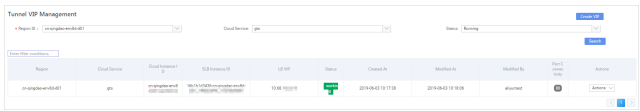

8.1.6. Create a Direct Any Tunnel VIP

You can create Direct Any Tunnel VIPs for cloud services in your VPC to allow traffic forwarding through XGW.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. From the **Products** menu, choose **Server Load Balancer > Direct Any Tunnel VIP Management**.
3. On the **Direct Any Tunnel VIP Management** page, click **Create Direct Any Tunnel VIP**.
4. On the **Create Direct Any Tunnel VIP** page, configure the parameters for the Direct Any Tunnel VIP.

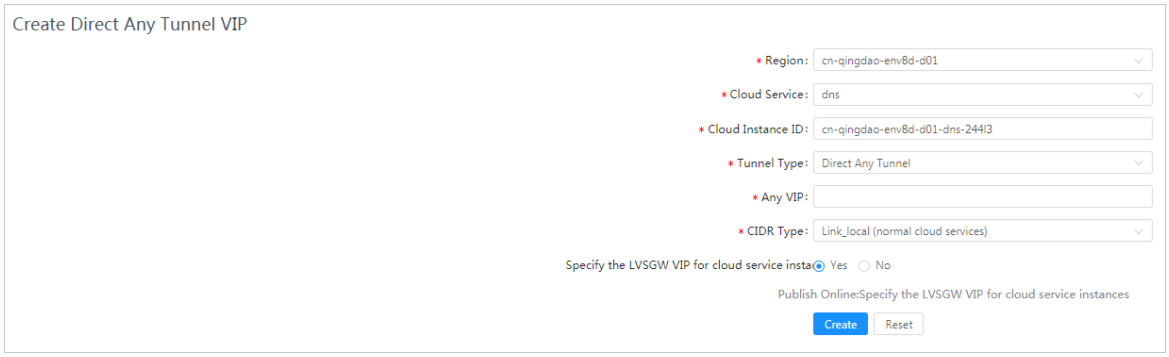

5. Click **Create**. Cloud service instances that have Direct Any Tunnel VIPs can forward traffic through XGW.

8.1.7. Leased line connection

8.1.7.1. Overview

You can connect a VPC to an IDC through a leased line.

Before connecting to a VPC through a leased line, you must confirm the initial CSW configurations meet the following conditions:

- You have uploaded the licenses required for VLAN functions onto the CSWs.
- You have set the management IP address on the loopback 100 interface of each CSW.
- You have configured the CSW uplink interfaces to ensure interoperability with the Layer 3 interfaces used by VPC APIs.
- You have deleted the default configuration of bridge-domain.
- You have enabled NETCONF and STelnet for CSWs. The configuration details are included in the CSW initial configuration template.
- You have configured the service type of CSW interfaces to tunnel.

You must also obtain the following account information:

- BID: specifies the ID of the account group. The BID for Mainland China users is 26842, and the BID for international users is 26888.
- UID: specifies the ID of the account to which the destination VPC belongs.

8.1.7.2. Manage access points

Access points are Alibaba Cloud data centers located in different regions. One or more access points are deployed in each region. This topic describes how to query and modify information about access points of a region.

Query an access point

Perform the following operations to query an access point:

1. Log on to the Apsara Network Intelligence console.
2. In the top navigation bar, click **Products** and choose **Express Connect > Daily Operation**.
3. In the left-side navigation pane, choose **Daily Operation > Access Points**.
4. Select the region and enter the ID of the access point that you want to query.
5. Click **Search**.

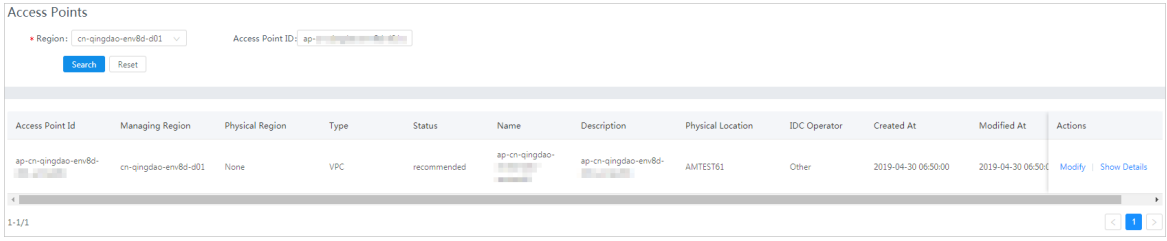

Modify access point information

Perform the following operations to modify the information about an access point:

1. Find the target access point and click **Modify** in the **Actions** column.
2. In the dialog box that appears, modify the information as needed.
3. Click **Modify**.

   Note the following points when you modify access point information:

   - **Access Point Location**: Enter the physical location of the access point. You can set a custom value.
   - **Access Point IDC Operator**: Enter the name of the data center operator.

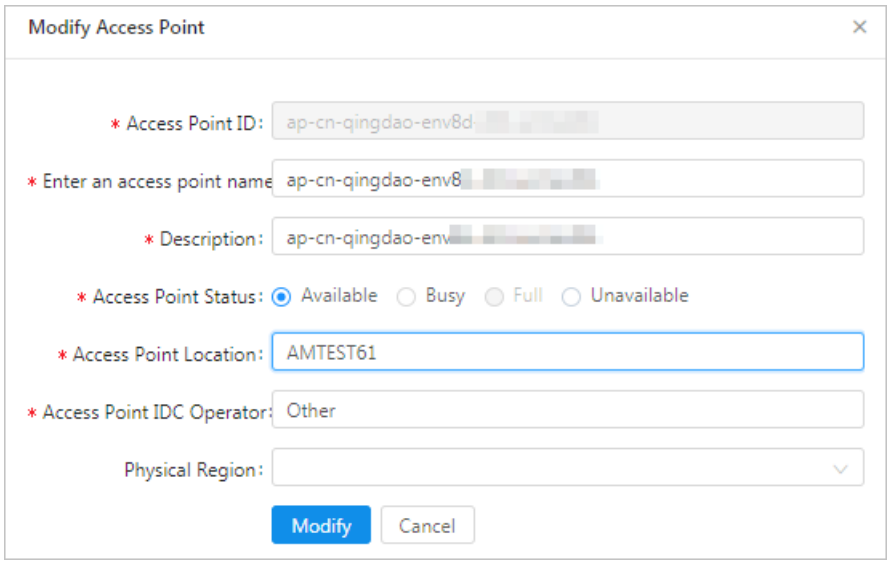

8.1.7.3. Manage access devices

This topic describes how to query and modify information about access devices of a region.

Query an access device

Perform the following operations to query an access device:

1. Log on to the Apsara Network Intelligence console.
2. In the top navigation bar, click **Products** and choose **Express Connect > Daily Operation**.
3. In the left-side navigation pane, choose **Daily Operation > Access Devices**.
4. Select the region and enter the ID of the access device that you want to query.

   > **Note** If Device ID is not set, all devices in the specified region are queried.

5. Click **Search**.

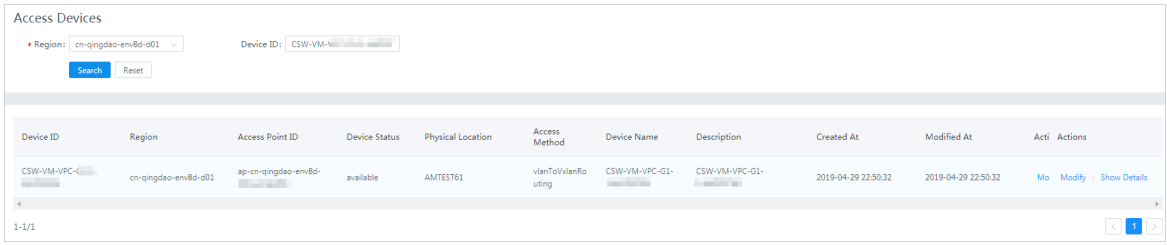

6. Find the target access device and click **Show Details** in the **Actions** column to view details of the access device.

Modify access device information

Perform the following operations to modify the information about an access device:

1. Find the target access device and click **Modify** in the **Actions** column.
2. In the dialog box that appears, modify the device information.

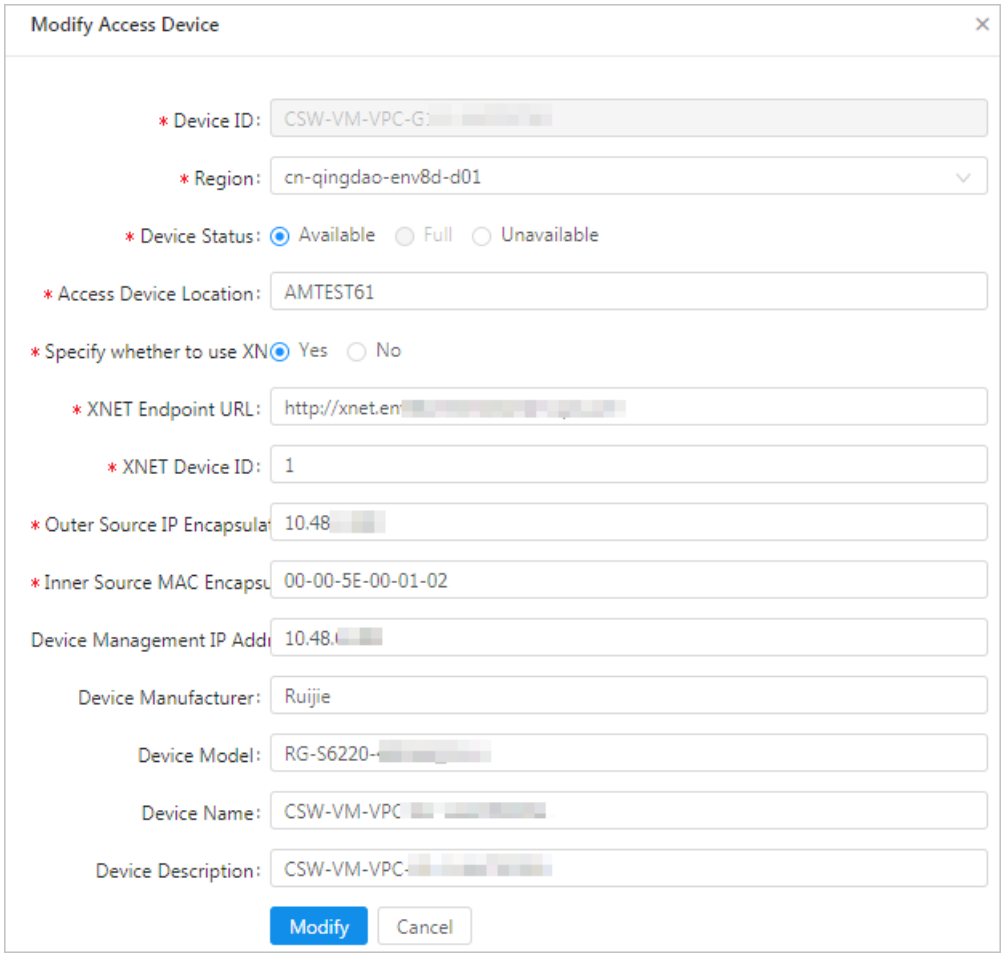

3. Click **Modify**.

8.1.7.4. Establish a leased line connection

A leased line can be obtained from a telecom operator to establish a physical connection between your data center and an Alibaba Cloud access point. This topic describes how to establish a leased line connection and query leased line information of a region.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. In the top navigation bar, click **Products** and choose **Express Connect > Network Environment Management**.
3. In the left-side navigation pane, choose **Function Modules > Leased Lines**. On the page that appears, click **Create Leased Line**.
4. In the dialog box that appears, configure the leased line and click **Create**. Note the following points when you create a leased line:
   - **Device Name**: Optional. If you set a device name, the device name must be the same as the CSW host name.
   - **Device Port**: Optional. If you set a device port, the device port number must be the same as the CSW port number.
   - **UID**: Enter the ID of the account to which the destination VPC belongs.

- **Access Point ID**: Select the ID of the region where your data center is located.
- **Redundant Leased Lines**: Select a previously obtained leased line as the redundant leased line for the leased line you are creating.

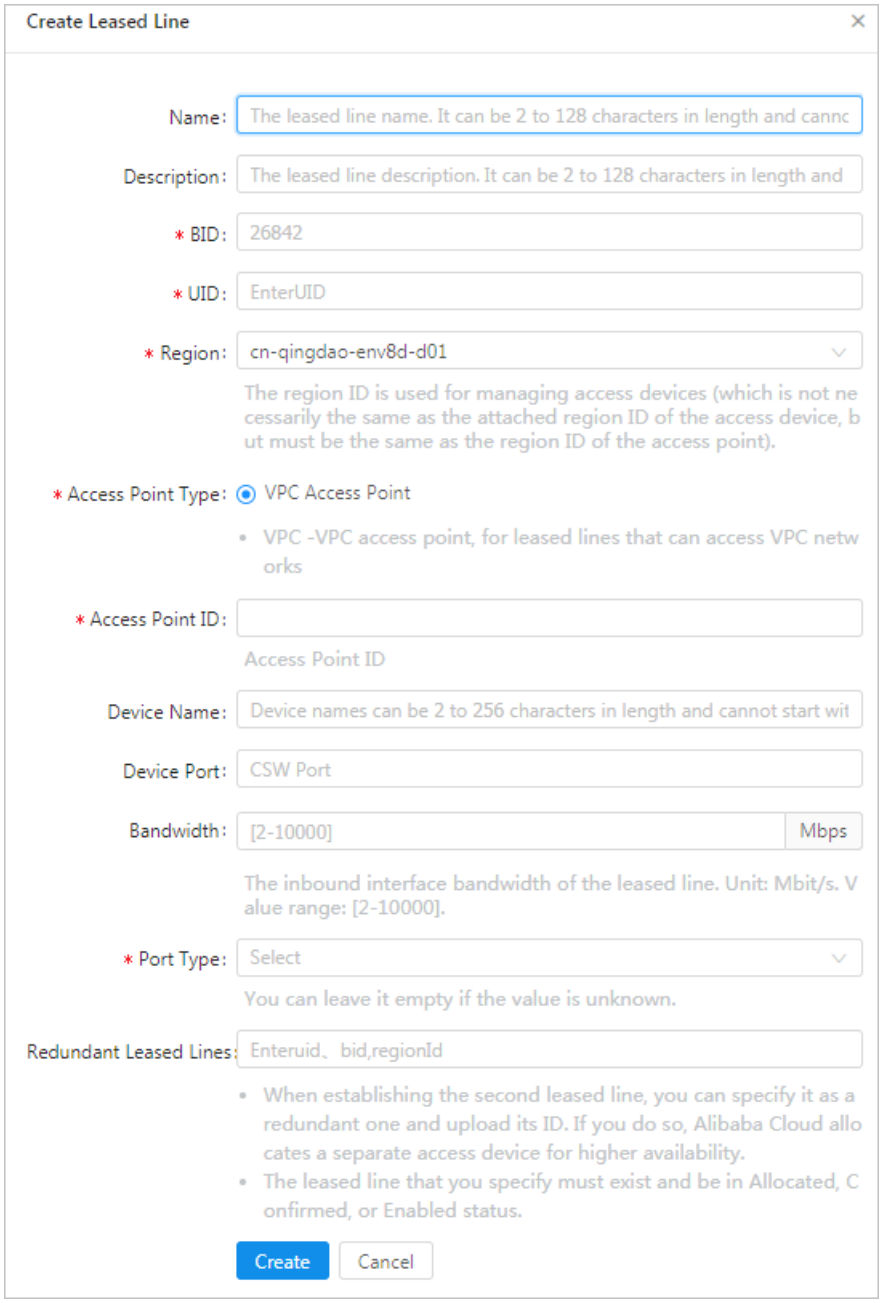

When the leased line state is **Confirmed**, the line is created.

5. On the **Leased Lines** page, find the created leased line and choose **Actions > Enable**.

   If the allocation process for a leased line persists for several minutes after you click Enable, choose **Products > Network Controller > Business Foundation System Flow**. On the page that appears, set Instance ID to the leased line ID, set **Step Status** to **All**, and click Search. Check the flow status in the search results. A red flow indicates that the corresponding task has failed. You can click **Resend** to restart the task and then requery the flow status.

If the second attempt still fails, run the `vpcregiondb -e "select * from xnet_publish_task order by id desc limit 5"` command on the ECS availability group (AG). If an error is returned, you can check service logs in Network Management and Operations to troubleshoot the issue based on the returned error.

8.1.7.5. Create a VBR

A virtual border router (VBR) is a router between customer-premises equipment (CPE) and a VPC, and functions as a data forwarding bridge from a VPC to an on-premises IDC. This topic describes how to create a VBR in a region and query VBR information of the region.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. From the **Products** menu, choose **Express Connect > Network Environment Management**.
3. Choose **Network Environment Management > VBRs**.
4. Click **Create VBR**.

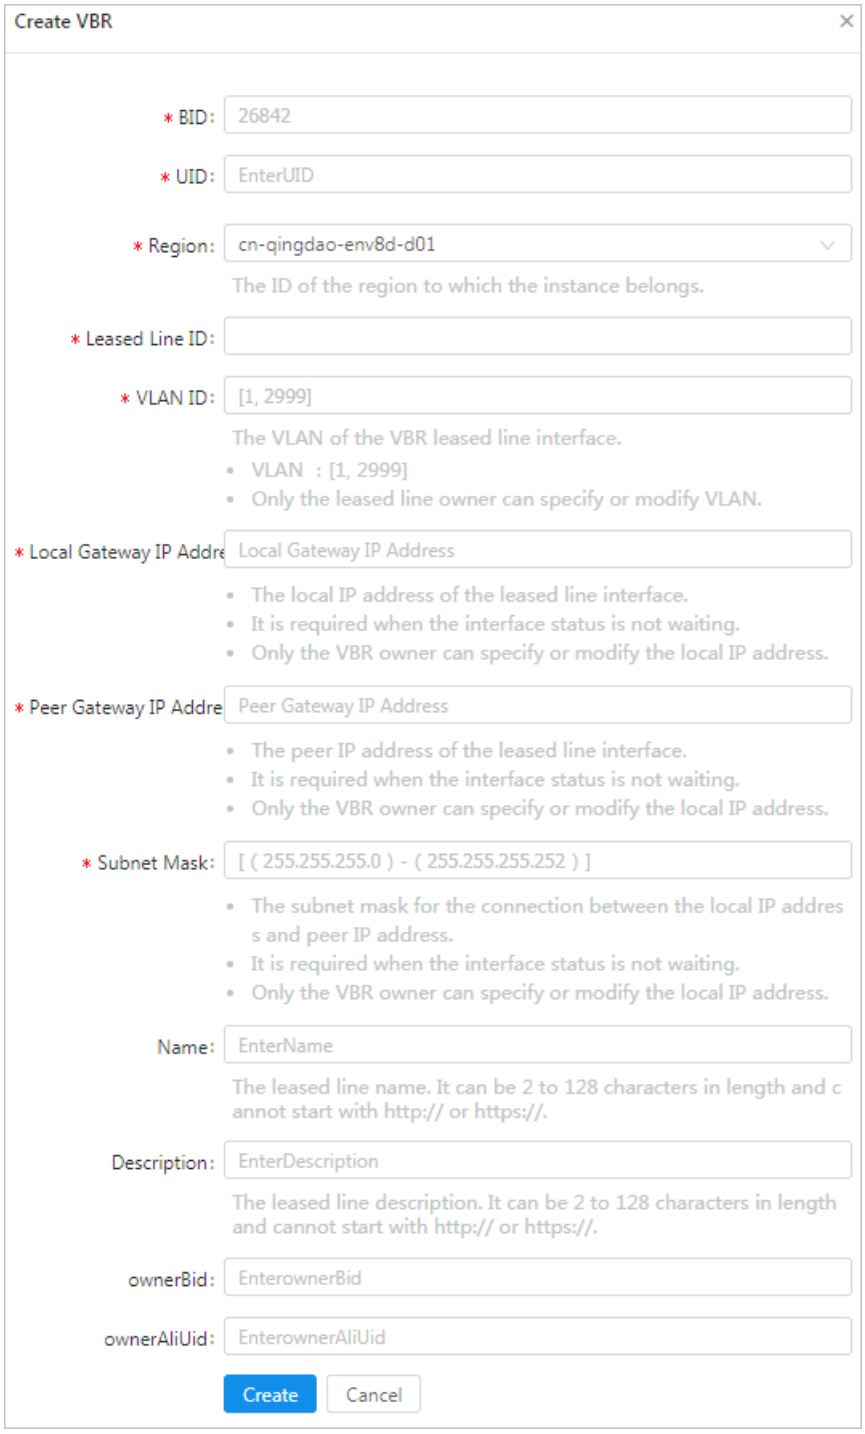

5. Follow the on-screen prompts to configure the VBR parameters.

   The parameters are described as follows:

   - **Leased Line ID**: specifies the ID of the leased line that the VBR connects to.
   - **VLAN ID**: specifies the VLAN ID of the VBR. The value ranges from 0 to 2999.

     When creating router interfaces, you can use VLAN IDs to identify subsidiaries or departments that use the leased line, thus implementing Layer 2 network isolation between them.
   - **Local Gateway IP**: specifies the local IP address of the router interface for the leased line.
   - **Peer Gateway IP**: specifies the peer IP address of the router interface for the leased line.

- **Subnet Mask**: specifies the subnet mask of the leased line between the local IP address and peer IP address.

  Only two IP addresses are required. Therefore, you can enter a longer subnet mask.

6. Click **Create**.
   When the VBR state is **Active**, the VBR is created.

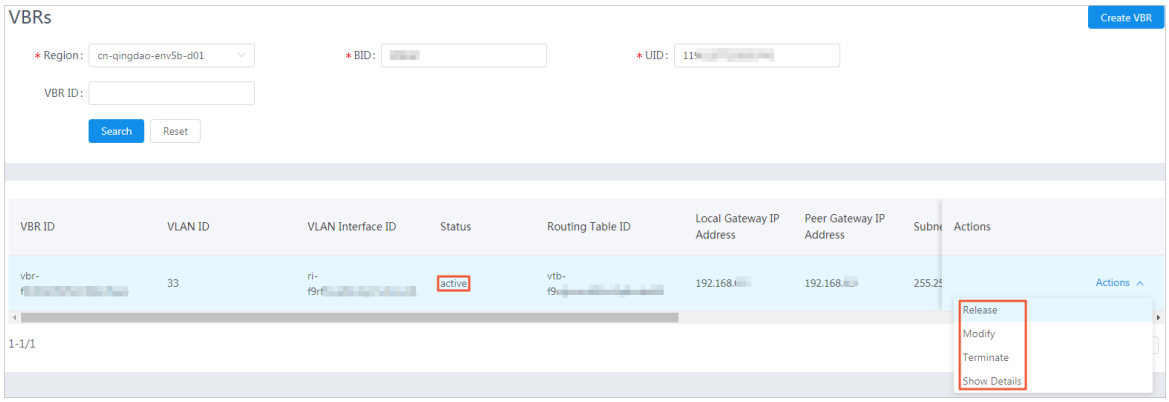

   You can click **Release**, **Modify**, **Terminate**, or **Show Details** in the **Actions** column to manage a VBR.

8.1.7.6. Create router interfaces

After you create a VBR, you must create a pair of router interfaces to connect the VBR and VPC. The connection initiator must be the VBR.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. From the **Products** menu, choose **Express Connect > Network Environment Management**.
3. Choose **Network Environment Management > Router Interfaces**.
4. Click **Create Router Interface**.
5. Configure router interface parameters and click **Submit**.

   Set **Create Router Interface** to **Double**. Configure the local router interface based on the created VBR information, and configure the peer router interface based on the destination VPC information.

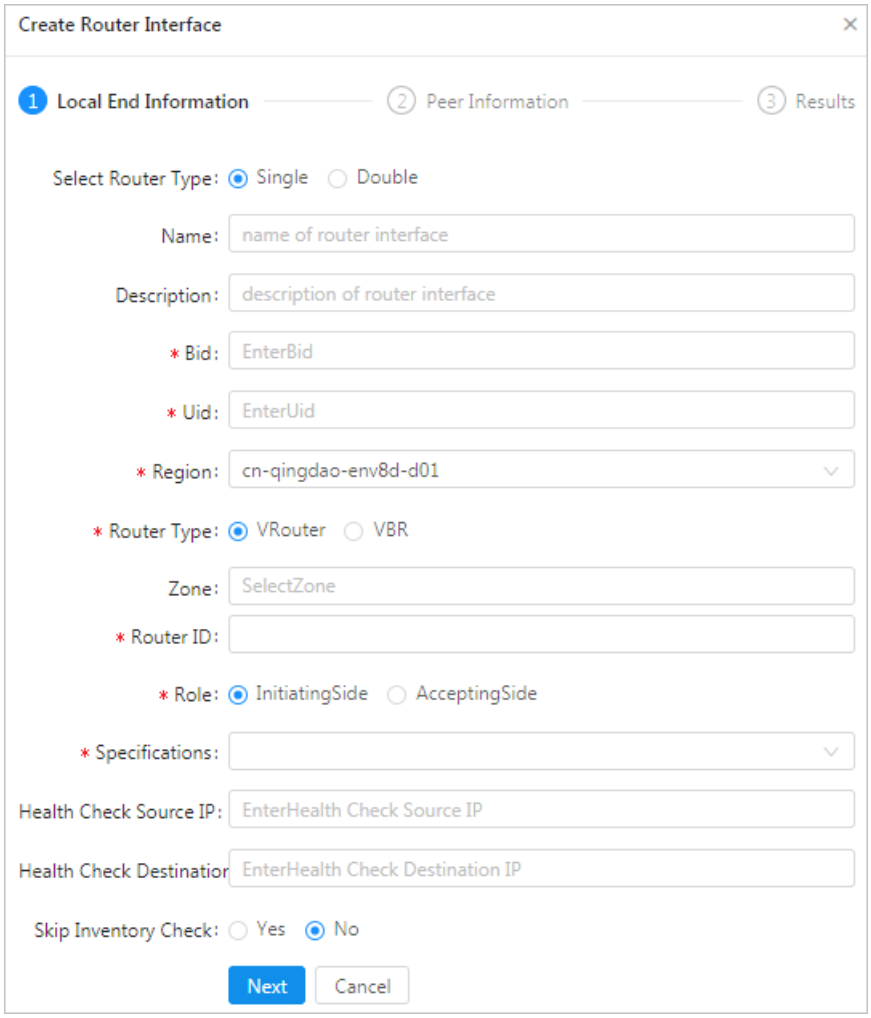

When the router interface state is **Active**, the interface is created.

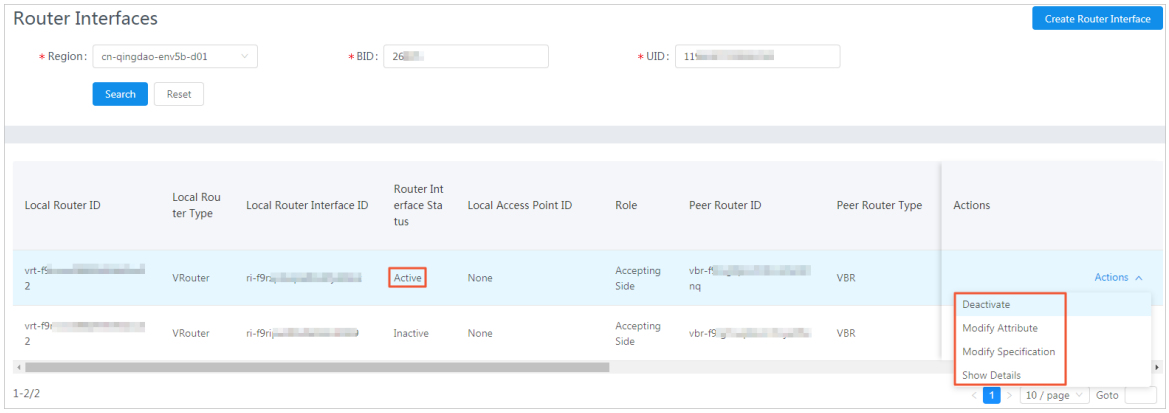

8.1.7.7. Create a routing table

A routing table is a list of route entries on a VRouter. This topic describes how to create routing tables in a region and query the routing table information of a region.

Procedure

1. Perform the following steps to add routes on a VBR destined for a VPC and an IDC:
   i. Log on to the Apsara Network Intelligence console.

i. Log on to the Apsara Network Intelligence console.

ii. From the **Products** menu, choose **Express Connect > Network Environment Management**.

iii. Choose **Function Modules > Routing Tables**.

iv. Set search conditions such as Region, BID, UID, Router Type, Routing Table ID, and Router ID, and click **Search** to query routing tables.

v. Click **Add Route Entry** in the **Actions** column corresponding to a routing table.

vi. Specify a route entry destined for the CIDR block of a destination VPC, and click **Create**.

The parameters are described as follows:

- **Destination CIDR Block**: the destination CIDR block.
- **Next Hop Type**: the next hop type.
- **Next Hop Instance ID**: the ID of the next hop instance for the specified next hop type.

Add a route destined for a destination VPC

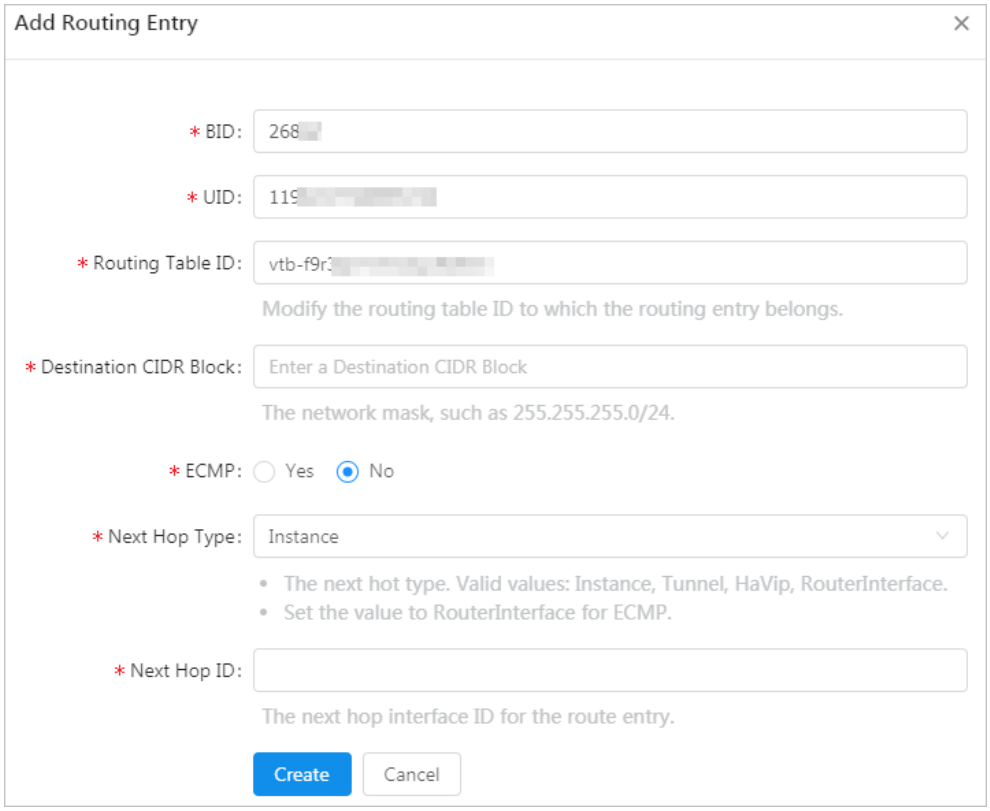

vii. Repeat the preceding steps to add a route destined for a target IDC.

> **Note** You can navigate to the VBRs page and locate the **VLAN Interface ID** area to obtain next hop router interface information.

2. Add a route destined for the router interface of a VBR in the VPC.
3. On the gateway of the on-premises IDC, configure a route destined for the VPC.

8.1.8. Manage Business Foundation System flows in a VPC

You can view the execution state of tasks in a VPC and restart the tasks as needed.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. From the **Products** menu, choose **Network Controller > Business Foundation System Flow**.
3. Query the flow state of the task you want to view.

   Enter a leased line ID in **Instance ID** and set **Step Status** to **All** to check the flow status. A flow in red indicates that the corresponding step has failed. Click **Resend** to restart the task, and then requery the flow status.

Flow Management page

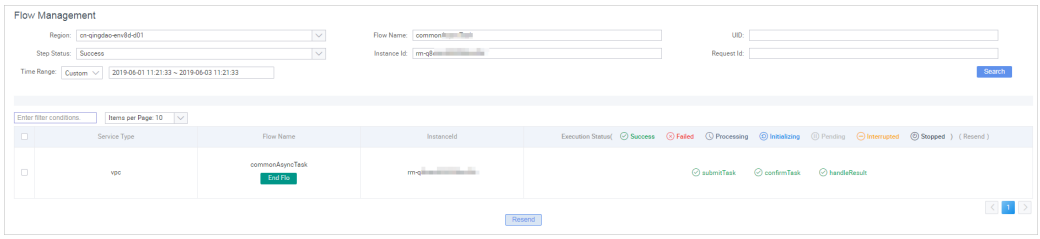

8.1.9. Configure reverse access to cloud services

Cloud services cannot be directly accessed through external networks. You must configure reverse access to allow external networks to access cloud services through ECS instances.

Prerequisites

Log on to the Apsara Stack Cloud Management (ASCM) console. Go to the **Personal Information** page and obtain **AccessKey ID** and **AccessKey Secret**.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. In the top navigation bar, click **Products** and choose **Cloud Service Management > Cloud Service Reverse Access**.
3. On the page that appears, enter **AccessKey ID** and **AccessKey Secret** and click **OK**. The **Cloud Service Reverse Access** page appears.
4. Click **Create Cloud Service Reverse Access**.
5. In the **Allocate Cloud Service ID** step, set Region, Name, and Description.
6. Click **Continue**. The following information is automatically created and displayed in the **Create Address Pool** step: the application ID of the cloud service that allow reverse access and the address pool that is used for reverse access to the cloud service.
7. Click **Continue**. In the **Add Server Address** step, configure the ECS instance to be used for reverse access.

- **VPC ID**: Enter an ID of a VPC, an ECS instance, or a single-tunnel cloud service instance.
- **Server IP**: Enter the IP address of the ECS instance to be used for reverse access.

8. Click **Continue**. In the **Create Mapping IP** step, configure VSwitch ID and Mapping IP of the ECS instance in the destination VPC.
9. Click **Continue**. In the **Complete Authorization** step, specify VPC ID, Server IP, and Instance Port for reverse access.

   The value of Instance Port must be an integer value. You can specify multiple instance ports separated with commas (,). Example: 10,20,30. You can configure up to 10 instance ports.

8.2. Network Management and Operations

8.2.1. Overview

This topic provides an overview of the Apsara Stack Network Operation Platform (NET). NET is a platform where network construction and O&M activities (including planning, design, construction, delivery, maintenance, changing, scheduling, and offlining) are transformed from **offline procedures** into **online automated processes**.

NET allows you to establish connections between physical network devices by using Secure Shell (SSH), Telnet, Simple Network Management Protocol (SNMP), RESTCONF, or gRPC Remote Procedure Calls (gRPC), and facilitates the creation of subsystems in O&M phases.

NET has the following features:

- Network automation

  The network automation feature allows you to establish connections by using protocols such as SSH, SNMP, and Telnet. You can create template scripts in Python to manage and organize upper-layer O&M systems and businesses.

- Change center

  NET provides a change automation engine that operates based on automation templates. You can use the features of the change center to orchestrate the execution sequence and correlations in the templates based on your business scenarios and O&M experience to formulate standardized and reusable change management plans. Such change management plans can be automatically implemented by the system to improve efficiency and reduce the risk of human errors.

Intended users

Network engineers who are responsible for construction, operation, and engineering change management of IT infrastructure.

8.2.2. Log on to the NET console

This topic describes how to log on to the NET console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

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

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

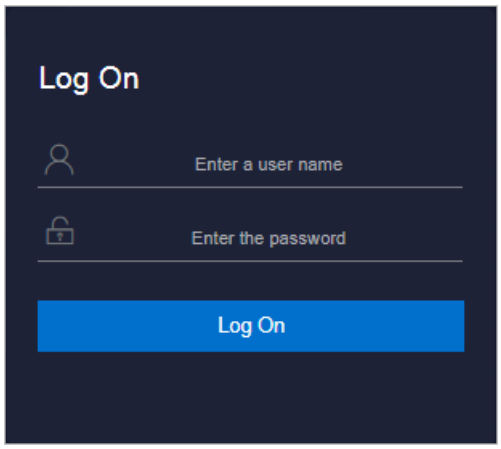

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**, and then click **Network Management and Operations** in the **Infrastructure as a Service (IaaS)** section.

8.2.3. Network Automation

8.2.3.1. Manage devices

This topic describes how to query, add, modify, and delete a device. After the NET system is started, the backend services automatically synchronize the data of physical network devices in the planning information to Apsara Infrastructure Management Framework.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Net Automation > Device Management**. The list of network devices appears on the page.
3. You can perform the following operations on physical switches.
   - Query

     Enter device properties and click **Search** to display the filtered results.
   - Modification

     Find the target device and click the  icon in the **Actions** column. On the **Update Device** page, modify the device settings.

     > **Note** If you do not set a user name and password, your Apsara Stack user name and password will be used for task execution. For more information about accounts, contact the O&M personnel.
   - Addition

     Find the target device and click the  icon in the **Actions** column. On the **Add Device** page, add a new device.

     > **Note** The hostname, IP address, and serial number must be globally unique.
   - Deletion

     Find the target device and click the  icon in the **Actions** column to delete the device.

8.2.3.2. Configure templates

8.2.3.2.1. Overview

This topic provides a brief overview of template configuration on NET. You can create template scripts in Python and configure input parameters.

The following features are supported:

- Using third-party libraries to abstract the capabilities of establishing connections and running command-line interface (CLI) commands to built-in methods (functions).
- Using nested templates to flexibly design device operations and structure business logic.

8.2.3.2.2. Create a device template

This topic describes how to create a device template. Device templates are applied to devices. By default, the system performs device logon and logoff before and after executing the scripts of a device template. The commands contained in a device template are run based on the built-in exec_cli method.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Net Automation > Device Templates**.
3. Add a single device template.
   i. Click **Add Single Device Template**.

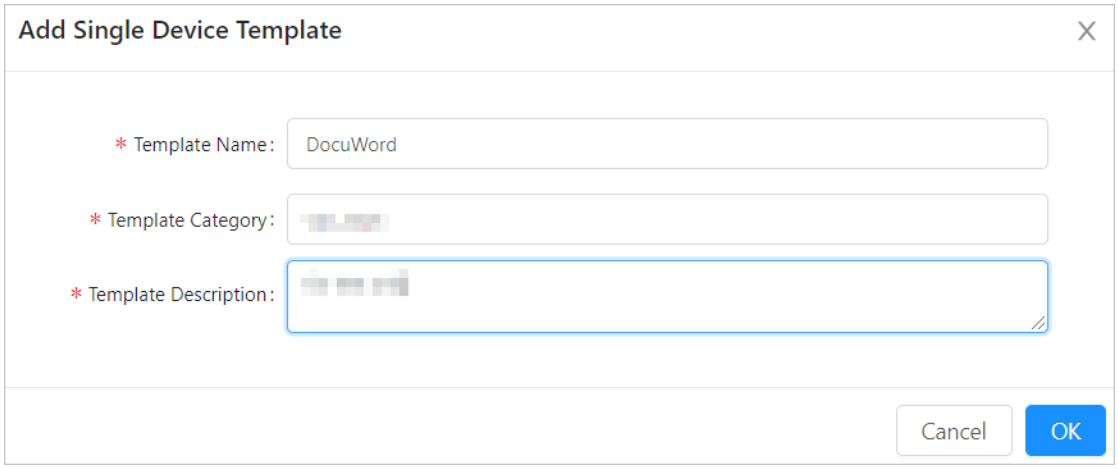

   ii. Configure the device template.
   - **Template Name**: Enter a name for the template.
   - **Template Category**: Specify a category for the device template. Device templates can be classified into different categories for merged queries.

     You can select an existing template category from the drop-down list.

     To create a new template category, enter a name for the category and then click **OK**.
   - **Template Description**: Enter a description for the template.

   iii. Click **OK**.

4. Find the target template and click the icon in the **Actions** column to add a rule to the template.
   i. Click **Add Template**.

ii. Click the **Classified Matching** tab. You can add rules based on the six tuples of different devices.

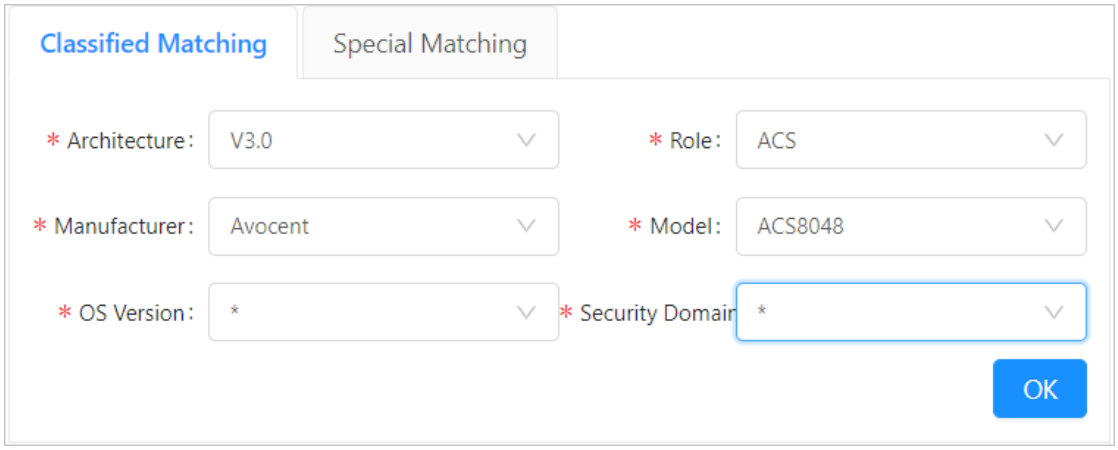

Since configuration commands vary by device model and manufacturer, you can select asterisk (*) from the drop-down list to denote an absence or omission of information.

> **Note** You can click the **Special Matching** tab and enter the hostnames to which this template is applicable. Multiple hostnames must be separated with line breaks.

iii. Click **OK**.

5. On the **Update Rule** page, click the serial number of a script to write template code.

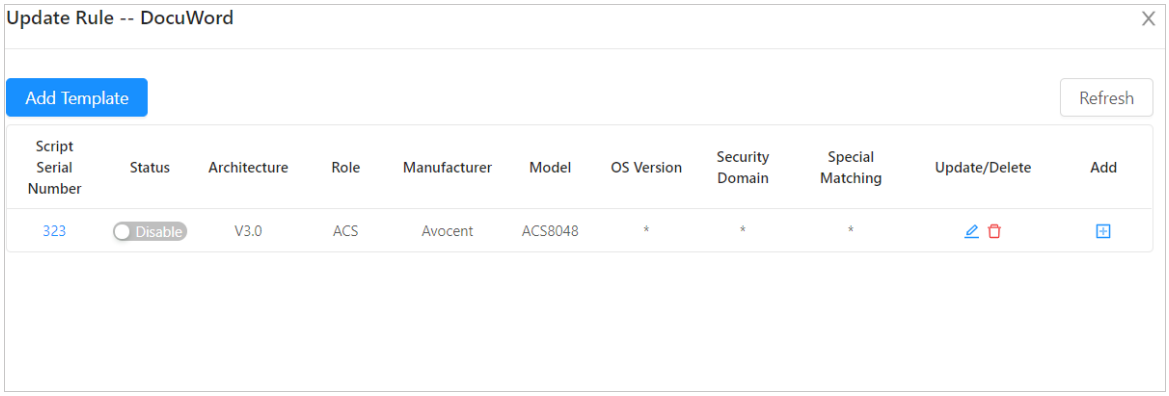

i. On the **Edit Script** tab, write the script.

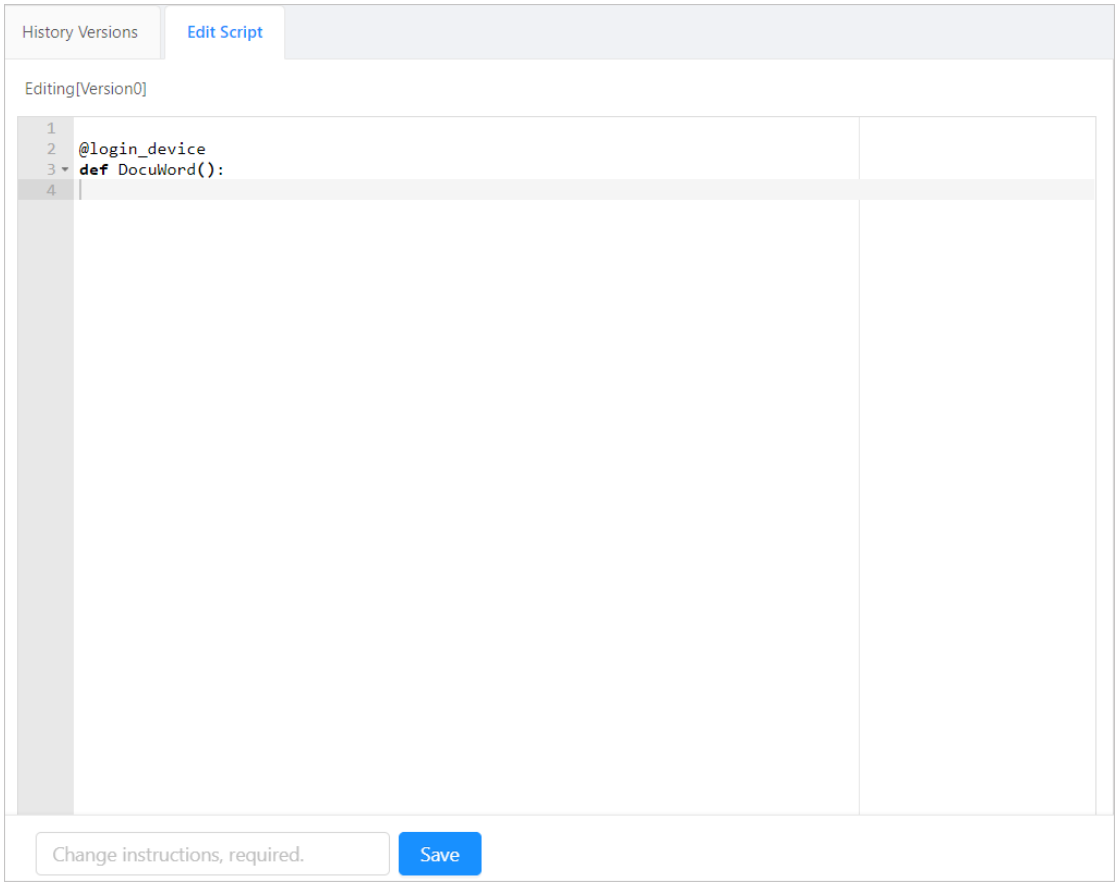

The login decorator (device_login) and function name (template name) of the device template use default values. You only need to import a library, such as re or JSON, as needed.

The following example illustrates how to query the platform version. The exec_cli is a built-in function that can be used directly, which is equivalent to running the `display version` command and returning the version information.

```
import re
@login_device
def get_software_version():
  output = exec_cli("display version")
  version = re.findall(r'[Ss]oftware, Version (\d+.\d+.? \d*)', output)[0]
  return version
```

ii. Enter the change instructions and click **Save**.

iii. On the **History Versions** tab, click **Release**.

> **Note**
>
> Each **Save** operation generates an entry recorded in the **History Versions**. Such scripts are classified as test scripts by default. You can select a test script and click **Release** to change its version from test to release. Only the scripts of the release version can be executed online.

6. (Optional)In the left-side navigation pane, choose **Net Automation > Device Templates**. Find the target template and then click the  icon in the **Actions** column. In the dialog box that appears, modify the name, category, and description of the template.
7. (Optional)In the left-side navigation pane, choose **Net Automation > Device Templates**. Find the target device template and then click the  icon in the **Actions** column to update the rules of the template.
8. (Optional)In the left-side navigation pane, choose **Net Automation > Device Templates**. Find the target device template and then click the  icon in the **Actions** column to delete the device template.

8.2.3.2.3. Create a user template

This topic describes how to create a user template. User templates are not related to devices. You can write user template code in Python, and use the built-in exec_script method to call other templates.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Net Automation > User Templates**.
3. Add a user template.
   i. Click **Add User Template**.

ii. In the **Add User Template** dialog box, configure the user template.

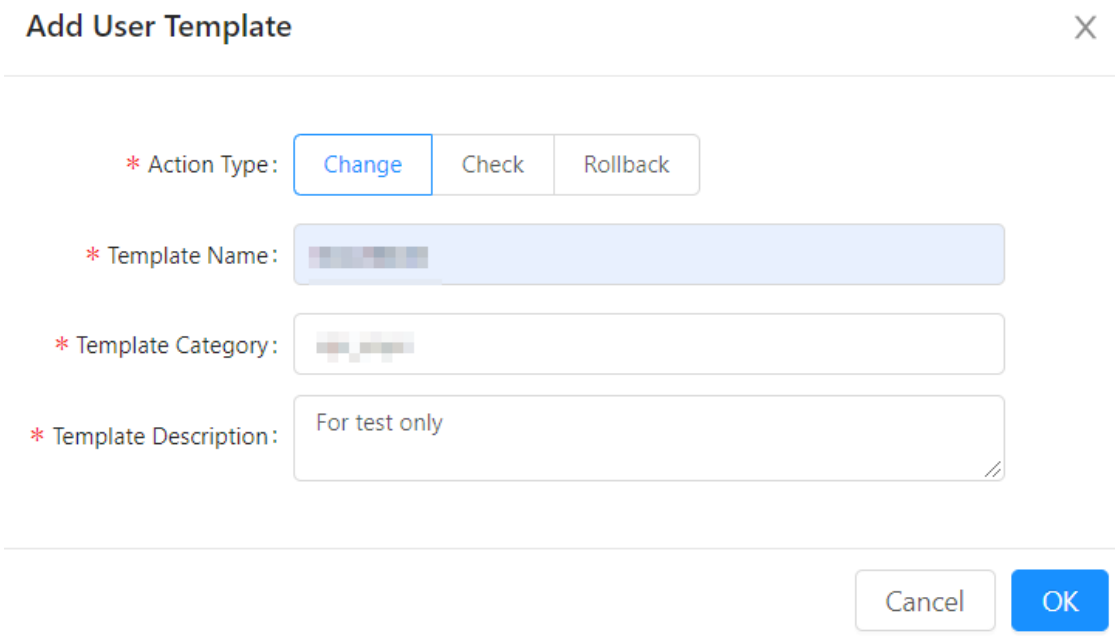

- **Action Type**: Select the action type for the user template. You can select one from **Change**, **Check**, and **Rollback**.
- **Template Name**: Enter a name for the user template.
- **Template Category**: Specify a category for the user template. User templates can be classified into different categories for merged queries.

  You can select an existing template category from the drop-down list.

  You can also enter a category name to create a new template category.
- **Template Description**: Enter a description for the user template.

iii. Click **OK**.

4. On the **User Templates** page, find the target user template and click the corresponding script serial number to edit the script.

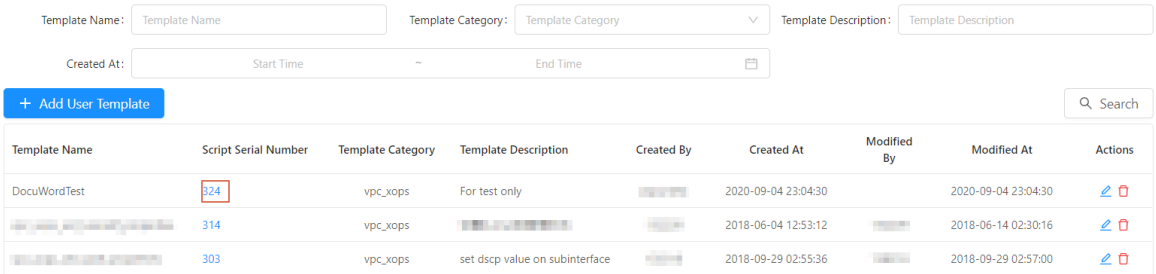

5. On the **Edit Script** tab, add parameters and write the script.

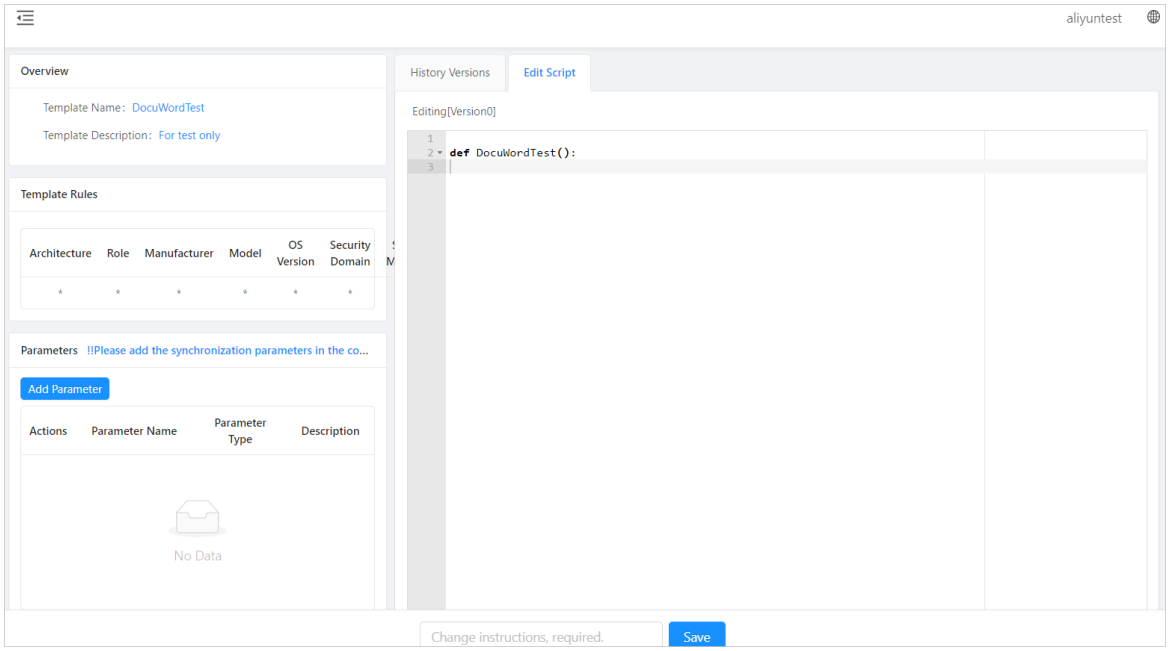

Call get_software_version to check whether the software version of the device meets the delivery requirement:

```
def check_device_version(device, version):
  """
  Check whether the OS version of the device meets the requirement
  : param device: the device name.
  : param version: the version number.
  : return: OK is returned if the version meets the requirement. Otherwise, an abort error is returned.
  """

  # Call get_software_version to retrieve the version number.
  # Call other templates by using the built-in exec_script method.
  # To call a device template, set the first parameter as the device template name and the second param
eter as the device name, which is followed by other parameters of the device template.
  # To call a user template, set the first parameter as the user template name, which is followed by othe
r parameters of the user template.
  dev_ver = exec_script("get_software_version", device)

  # Compare the retrieved version number of the device with the input parameter. If the version numbe
rs match, OK is returned. Otherwise, an abort error is returned.
  # abort is a built-in method, which is similar to raise Exception in Python.
  if version == dev_ver:
    return "OK"
  else:
    abort("version not match, input:%s <-> real:%s" % (version, dev_ver))
```

- The device_login decorator is unavailable because user templates are unrelated to devices.
- You can add parameters to a template function by clicking **Add Parameter** on the left side of the page. Parameters added otherwise will be discarded.
- In the template code, exec_script is a built-in method for nested calls to other templates.
- In the template code, abort is a built-in method that is used to throw an exception when a task fails. A normal return value indicates that the task has succeeded.

6. Enter the change instructions and click **Save**.
7. On the **History Versions** tab, click **Release**.

> **Note**
>
> Each **Save** operation generates an entry recorded in the **History Versions**. Such scripts are classified as test scripts by default. You can select a test script and click **Release** to change its version from test to release. Only the scripts of the release version can be executed online.

8. (Optional)In the left-navigation pane, choose **Net Automation > User Templates**. On the **User Templates** page, find the target user template and click the icon  in the **Actions** column. In the dialog box that appears, modify the action type, and the name, category, and description of the user template.
9. (Optional)In the left-navigation pane, choose **Net Automation > User Templates**. On the **User Templates** page, find the target user template and click the  icon in the **Actions** column to delete the user template.

8.2.3.3. Manage change tasks

This topic describes how to manage change tasks on NET. Change tasks are triggered by platform applications or external API calls. After you specify the information such as the name of the task entry template and necessary parameters, the system automatically executes logic in the background based on the online template script.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Net Automation > Task Management**. The list of change tasks appears on the page.

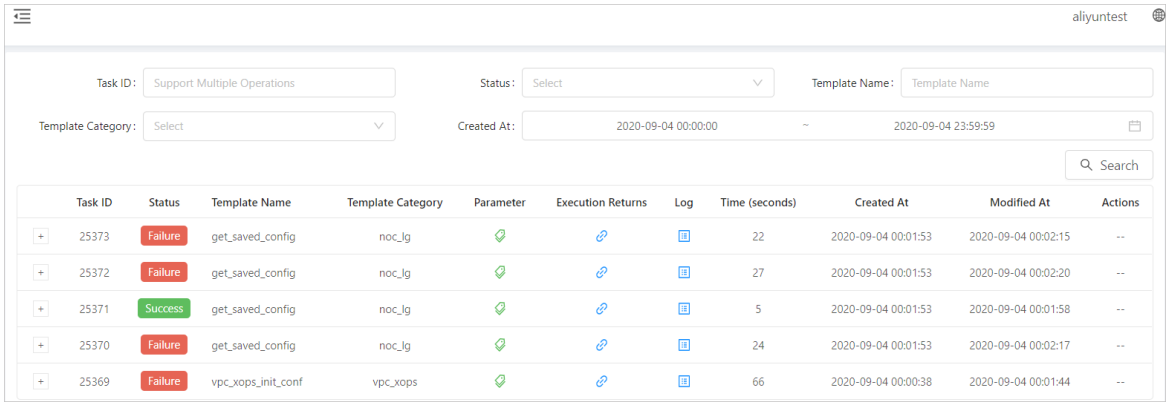

8.2.3.4. Trigger real-time tasks

This topic describes how to trigger real-time tasks by using device templates.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Net Automation > Real-time Tasks**.
3. In the **Real-time Task Execution** section, select a device template and set the parameters for the real-time task.
4. Click **Execute**. After the execution, you can check the execution result in the **Task Execution Results** section.

8.2.3.5. Manage files

This topic describes how to upload the configuration files of devices by using the file management feature of the NET platform.

Procedure

1. Log on to the NET console.
2. In the left-side navigation page, choose **Net Automation > File Management**.
3. Click **Upload**. In the **Upload File** dialog box, specify the file type, hostname, serial number, and description, and then select the file to be uploaded.
4. Click **Submit**.

8.2.4. Network monitoring

8.2.4.1. Dashboards

8.2.4.1.1. Check the status of a device

This topic describes how to check the status of a device by using the status view feature.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Monitor View > Status View**.
3. Select the target device and the event types.
4. Click **Search**.
5. Click the  icon in the **Details** column to check the monitoring details.

8.2.4.1.2. Check the aggregate status

This topic describes how to check the aggregate status. The aggregate status displays the numerical aggregation of status data collected from monitored single devices.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Monitor View > Aggregate Status**.
3. Select the alarm status, aggregate data, and aggregate type, and then enter the data items.
4. Click **Search** to check the aggregate status.

8.2.4.1.3. Check the data view

This topic describes how to view monitoring data by using the data view feature. You can use the data view feature to display the dashboards that visualize the data collected based on monitor settings.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Monitor View > Data View**.
3. Enter one or more keywords in the search bar, and click **Search**.

8.2.4.2. Configuration management

8.2.4.2.1. Add a monitoring item

This topic describes how to configure monitoring by adding monitoring items. For each monitoring item, you can specify a collection type, such as PING or Simple Network Management Protocol (SNMP), define a collection interval, set data items, and add alarm rules.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Configuration > Monitoring Items**.
3. On the **Monitoring Items** page, click **Add Monitoring Item**.
4. On the **Monitoring Item Management** page, configure the monitoring item.The following

configuration parameters are included:

- **Monitoring Item Name**: Specify a name for the monitoring item. The name must be globally unique. We recommend that you use a name that describes the feature of the monitoring item.
- **Description**: Enter a comprehensive description of the feature of the monitoring item. A detailed description helps increase maintenance efficiency.
- **Security Domain**: Security domains are not interconnected. Multiple Server Load Balancer (SLB) instances can be added in each security domain.
- **Collection Type**: Select a collection type from the displayed options. PING and SNMP are the most common collection types.
  - PING: The device is pinged periodically based on the configuration to calculate the packet loss rate and latency. After you select PING as the collection type, you must set values for constant, interval, and packNum.
  - SNMP: Different types of data are collected based on the definitions of object identifiers (OIDs). After you select SNMP as the collection type, you must define the OIDs supported by the device, set a password, and specify the data types to be monitored.
- **Effective**: Select whether to apply the monitoring item. If this parameter is set to NO, the monitoring agent does not collect data as the monitoring item specifies. We recommend that you perform debugging first, and then set the monitoring item to be effective if no anomaly is detected.
- **Execution Interval**: Specify a time interval for periodic data collection. We recommend that you set the interval to one minute for PING or SNMP monitoring to avoid undetectable anomalies across long intervals.
- **Parsing Code**: The system formats and parses the data collected by the agent, and returns the data items in the specified format.
- **Data Item**: Set the data items to be collected for the monitoring item. This is an optional parameter.
- **Alarm Rules**: The system categorizes alarms into five status conditions: normal, warning, critical, error, and waiting.

  You can specify the rules for warning and critical alarms. The system changes the alarm status based on the specified rules.

5. Click **Debugging**.
6. Click **Submit**.

8.2.4.2.2. Add a notification group

This topic describes how to add a notification group to subscribe the intended recipients to receive notifications. You can use notification groups to reduce maintenance costs for subscriptions.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, click **Network Monitor > Configuration > Notification Group**.
3. Click **Add Notification Group**.
4. In the **Add** dialog box, specify the name, description, and contacts of the notification group.
5. Click **OK**.

8.2.4.2.3. Subscription management

8.2.4.2.3.1. Subscribe to single-device notifications

This topic describes how to create a single-device notification subscription. Single-device alarms of different levels are generated based on configured alarm rules for the data collected by monitoring items, and are sent as emails, SMS messages, or DingTalk messages to specified recipients or notification groups.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Configuration > Subscription**.
3. Click the **Subscribe Single Device Alarm** tab.
4. Click **Add Subscription**.
5. In the dialog box that appears, configure the subscription.The following configuration parameters are included:
   - **Subscription Node**: Set a monitoring item.
   - **Alarm Status**: Select the status that triggers the notification.
   - **Continuous Trigger**: You can set this parameter to 2 to prevent excessive false alarms. In this case, a notification is not triggered when the status changes from normal to warning or critical for the first time, and is only triggered if such a change occurs for the second time.
   - **Inhibition Strategy**: Select an inhibition strategy from the drop-down list.
     - No suppression policy is set: A notification is sent whenever the status becomes critical or warning.
     - Iteration slowdown: The notification intervals are prolonged with each successive notification, and a notification interval is the collection interval multiplied by the nth power of the step size.
     - Status change: Notifications are sent when the status becomes critical or warning, and can only be sent up to three consecutive times. When the status returns to normal, the count of notifications is reset back to zero.
   - **Notice Method**: Select a method for sending notifications.
   - **Receiver**: Enter recipient information or select a notification group.
   - **Advanced Configuration**: Configure fine-grained filtering based on the on-premises data center or role to which the device belongs.

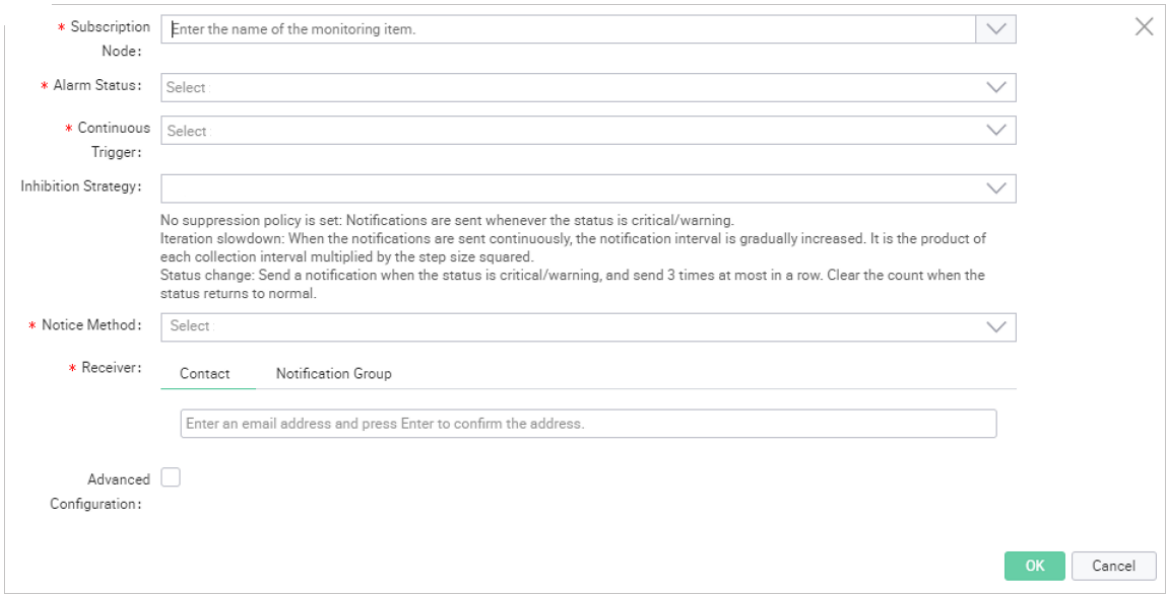

6. Click **OK**.

8.2.4.2.3.2. Subscribe to aggregate notifications

This topic describes how to create an aggregate notification subscription. Aggregate alarms of different levels are generated based on configured alarm rules for the data collected by monitoring items, and are sent as emails, SMS messages, or DingTalk messages to specified recipients or notification groups.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Mnitor > Configuration > Suscription**.
3. Click the **Subscription Aggregate Alarm** tab and then click **Add Subscription**.
4. In the **Add** dialog box, configure the subscription.The following configuration parameters are included:
   - **Subscription Node**: Set a monitoring item.
   - **Aggregate Data**: Select aggregate data.
   - **Alarm Status**: Select the status that triggers the notification.
   - **Inhibition Strategy**: Select an inhibition strategy from the drop-down list.
     - No suppression policy is set: A notification is sent whenever the status becomes critical or warning.
     - Iteration slowdown: The notification intervals are prolonged with each successive notification, and a notification interval is the collection interval multiplied by the nth power of the step size.
     - Status change: Notifications are sent when the status becomes critical or warning, and can only be sent up to three consecutive times. When the status returns to normal, the count of notifications is reset back to zero.
   - **Notice Method**: Select a method for sending notifications.
   - **Receiver**: Enter recipient information or select a notification group.

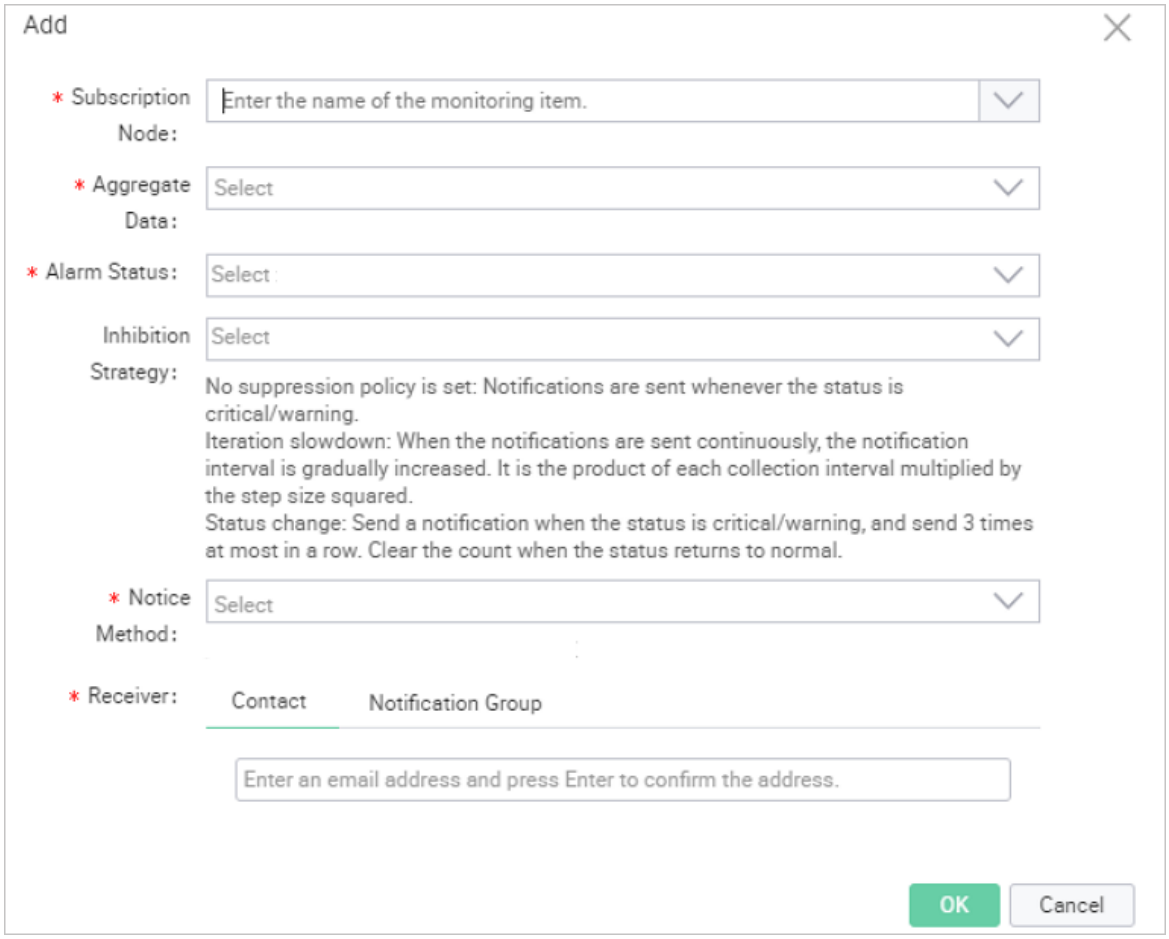

5. Click **OK**.

8.2.4.2.4. Add an aggregate data configuration

This topic describes how to add an aggregate data configuration. Aggregate data is the compilation of data sourced from one or more monitored devices, and is aggregated to calculate the maximum, minimum, average, or sum value of the data objects.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Configuration > Aggregate Data Configuration**.
3. On the **Aggregate Data Configuration** tab, click **Add Aggregate Data Configuration**.
4. In the **Add Aggregate Data Configuration** dialog box, complete the configuration.The following configuration parameters are included:
   - **Name**: Enter a name for the aggregate data. The name must be globally unique.
   - **Monitoring Item**: Select the monitor item where the data to be aggregated is located.
   - **Data Item**: After the monitoring item is specified, select a data item from the drop-down list.
   - **Description**: Enter a description for the aggregate data.
   - **Covering Device**: Select the mode for device coverage.

- Aggregate All Devices: Data collected from all the devices covered by the monitoring item are aggregated.
- Select Some Devices: Only the data collected from one or more specified devices covered by the monitoring item are aggregated.

- **Aggregate Type**: Select one or more aggregate functions. Available options include sum, avg, max, and min.

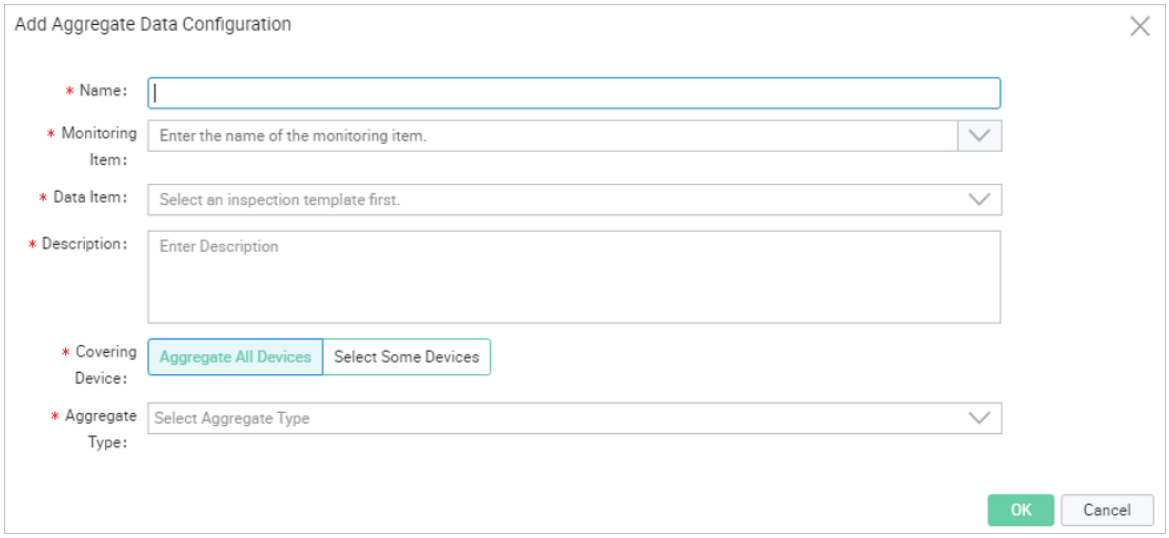

5. Click **OK**.

8.2.4.2.5. Add a port set

This topic describes how to add a port set to collect statistics on inbound and outbound traffic on specified device ports.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Configuration > Port Set**.
3. On the **Port Set Management** tab, click **Add Port Set**.
4. On the **Port Set Management** page, configure the port set.The following configuration parameters are included:
   - **Port Set Name**: Enter a name for the port set. The name must be globally unique.
   - **Description**: Enter a comprehensive description of the feature of the port set. A detailed description helps increase maintenance efficiency.
   - **Add Port**: Select one or more devices and add ports used by the corresponding devices.

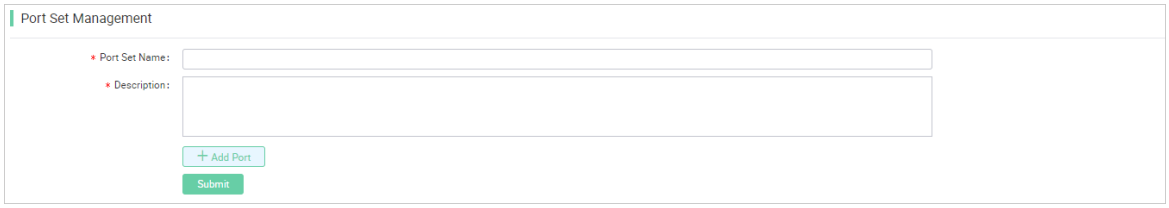

5. Click **Submit**.

8.2.4.2.6. Add a data view

This topic describes how to add and configure a data view as a monitoring dashboard to display aggregated data or port set statistics.

Procedure

1. Log on to the NET console.
2. In the left-side navigation pane, choose **Network Monitor > Configuration > Data View**.
3. On the **Aggregate Data View** tab, click **Add View**.
4. On the **View Management** page, configure the data view.The following configuration parameters are included:
   - **View Name**: Enter a name for the data view. The name must be globally unique.
   - **Description**: Enter a comprehensive description of the feature of the data view. A detailed description helps increase maintenance efficiency.
   - **Add Chart**: Add one or more charts to the data view.
     - **Data Sources**: Select the source of data to be visualized for a chart. Available options include Aggregate Data and Port Set Traffic.
     - **Select Aggregate Data**: Select the aggregate data for the chart to display.
     - **Chart Type**: Choose a chart type for data visualization. Available options include Broken Line Chart and Area Chart.
     - **Chart Width**: Select a ratio from the drop-down list to determine how much width to be taken up by the chart on the web page.
     - **Sort Number**: Specify an order number for the chart. Charts are displayed in the ascending order in the dashboard.

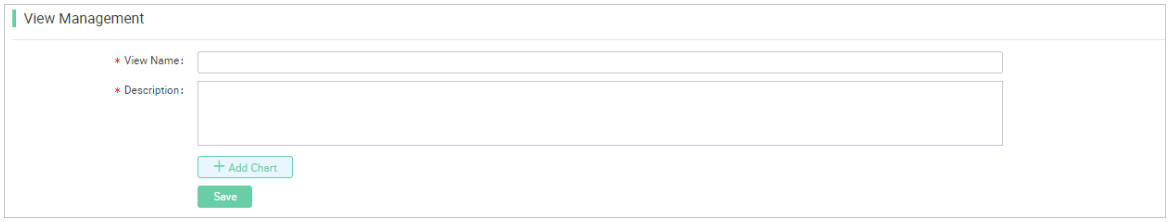

5. Click **Save**.

8.2.5. Sample templates

This topic provides sample templates that are provisioned in production environments. Templates help you automate the implementation of network changes and configuration management with script logic and agile orchestration.

Sample templates

```
import datetime
import requests
import json
import time
```

```
import re

@login_device()
def atom_flow_isolate(restriction):
  """
  Isolate traffic before device OS upgrade
  Return: execution result and default routing peer
  """
  # Exit configuration mode and disable logging.
  exec_cli("return", strict=False)
  exec_cli("undo terminal monitor", strict=False)
  exec_cli("screen-length 0 temporary", strict=False)
  stack_check= exec_cli("display stack", strict=False)
  dis_version=exec_cli("display version", strict=False)
  check_isolate = exec_cli('display curr config bgp | in 2100', strict=False)
  if re.search('filter-policy 2100 export',check_isolate):
    logger.info("The filter-policy 2100 is existed in bgp config,Please check!")
    isolate_res.append([hostname+'(fail)',' filter-policy 2100 exists in BGP configuration.'])
    return default_peers, backup_ospf_stub,interfaces,isolate_res

  # Obtain the Border Gateway Protocol (BGP) number of the device.
  output = exec_cli("display cu configuration bgp | include bgp", strict=False)
  bgp_no = re.search(r"bgp (\d+)", output)
  if bgp_no:
    bgp_no = bgp_no.group(1)
    restriction['rollback_bgp_no']=bgp_no
  exec_cli("sys", strict=False)

  # Obtain the timestamp before isolation.
  time1=time.time()
  get_time()
  logger.info('begin time:%s'%time1)

  #1 Open Shortest Path First (OSPF) traffic isolation (LoadBalance Switch (LSW)).
  if cu_rol == 'LSW':
    output = exec_cli("display cu configuration ospf | include ospf", strict=False)
    ospf_no = re.search(r"ospf (\d+)", output)
    if ospf_no:
      ospf_no = ospf_no.group(1)
      backup_ospf_stub = exec_cli("disp cu conf ospf | in stub", strict=False)
      exec_cli("ospf %s" % ospf_no, strict=False)
```

```
        exec_cli("undo stub-router", strict=False)
        exec_cli("stub-router include-stub summary-lsa external-lsa", strict=False)
        exec_cli("commit", strict=False)
        logger.info('ospf_isolate')
        radar_chek=radar_result(15)
        if radar_chek:
          logger.info('%s ospf isolate successfull'%hostname)
        else:
          return default_peers, backup_ospf_stub,interfaces,isolate_res

  # Destack Access Switch (ASW) and Link Layer Discovery Protocol (LLDP) isolation.
  out1 = exec_cli("display stack", strict=False)
  if cu_rol == 'ASW' and (cu_arc == '5.0L' or '4.2' in cu_arc):
    # if 'ASW' in hostname:

    out2 = exec_cli('display lldp nei brief | ex "DSW|PSW|M"', strict=False)
    out3 = exec_cli('display interface brief | in up', strict=False,timeout = 120)
    int_down = []
    tmp_down = get_info(out2,0,2,3)
    for i in tmp_down:
      int_down.append(i[0])
    if 'Standby' in out1:
      isolate_res.append([hostname+'(fail)',' stack configuration exists for the ASW to be destacked, please c
heck'])
      return default_peers, backup_ospf_stub,interfaces,isolate_res
    else:
      exec_cli('sys', strict=False)
      exec_cli('undo lldp enable', strict=False)
      exec_cli('commit', strict=False)
      radar_chek_lldp=radar_result(30)
      if radar_chek_lldp:
        logger.info('%s lldp isolate successfull'%hostname)
      else:
        return default_peers, backup_ospf_stub,interfaces,isolate_res
      exec_cli('return', strict=False)
      exec_cli("reset interface counters", strict=False)
      exec_cli('sys', strict=False)
      radar_chek_lldp=radar_result(30)
      if radar_chek_lldp:
        logger.info('lldp isolate successfull')
      else:
```

```
    else:
        return default_peers, backup_ospf_stub,interfaces,isolate_res

if ('4.1' in cu_arc or '3.' in cu_arc) and cu_rol == 'ASW':
    logger.info('stack device')
else:
    #2 BGP traffic isolation (all).
    output = exec_cli("display ip interface brief | include Loop", strict=False)
    loop1 = re.search(r"LoopBack1\s+(\d+\.\d+\.\d+\.\d+)/\d+", output)
    loop1 = loop1.group(1)
    loop2 = re.search(r"LoopBack101\s+(\d+\.\d+\.\d+\.\d+)/\d+", output)
    loop2 = loop2.group(1) if loop2 else ""


    #2.1 Configure the access control list (ACL)
    exec_cli("acl number 2100", strict=False)
    exec_cli("rule permit source %s 0.0.0.0" % loop1, strict=False)
    if loop2:
        exec_cli("rule permit source %s 0.0.0.0" % loop2, strict=False)
    exec_cli("commit", strict=False)

    #2.2 Check the ACL
    output = exec_cli('dis acl 2100 | in " %s%s"' % (loop1, ('| '+loop2) if loop2 else ''), strict=False)
    if not output:
        isolate_res.append([hostname + '(fail)',' ACL 2100 configuration failed, please check'])
        return default_peers, backup_ospf_stub,interfaces,isolate_res

#2.2 BGP configuration (PoD Switch (PSW), Distribution Switch (DSW), and LSW).
if cu_rol in ['DSW','PSW','LSW','CSW']:
    exec_cli("bgp %s" % bgp_no, strict=False)
    exec_cli("ipv4-family unicast", strict=False)
    exec_cli("filter-policy 2100 export", strict=False)
    exec_cli("commit", strict=False)
    if cu_rol == 'DSW' and cu_arc ! = '4.0V' and cu_arc ! = '4.1Lv':
        sdn_check = exec_cli('display sdn openflow session', strict=False)
        adv_check = exec_cli('dis cu conf bgp | in advertise lowest-priority ', strict=False)
        restriction['rollback_sdn_check']=sdn_check
        restriction['rollback_adv_check']=adv_check
        if 'REGISTERED' not in sdn_check and 'advertise lowest-priority on-startup' in adv_check:
            exec_cli('undo advertise lowest-priority on-startup ', strict=False)
            exec_cli("commit", strict=False)
```

```
        logger.info('bgp_isolate')

    #2.3 ASW destacking and BGP isolation configuration.
    if cu_rol == 'ASW' and rate_check == True and 'Standby' not in out1 and not slot and (cu_arc == '5.0L' or '4.
2' in cu_arc or '5.1' in cu_arc):
        exec_cli("bgp %s" % bgp_no, strict=False)
        exec_cli("ipv4-family unicast", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("commit", strict=False)
        logger.info('bgp_isolate')

   #2.3 MC BGP isolation.
    if cu_rol == 'MC':
        exec_cli("bgp %s" % bgp_no, strict=False)
        exec_cli("ipv4-family vpn-instance Alimaster", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("ipv4-family vpn-instance Alimaster-public", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("ipv4-family vpn-instance Alipay", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("ipv4-family vpn-instance Alipay-public", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("ipv4-family vpn-instance Aliyun", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("ipv4-family vpn-instance Aliyun-public", strict=False)
        exec_cli("filter-policy 2100 export", strict=False)
        exec_cli("commit", strict=False)
        logger.info('bgp_isolate')
    #3 Delete default route advertisements and close open ports.

    if cu_rol == 'LSW':
        output = exec_cli("disp cu conf bgp | in default-route-advertise", style='verbose', strict=False)
        logger.debug("default route: %s", output)
        default_peers = re.findall(r"peer (. *?) default-route-advertise", output)
        if default_peers:
            exec_cli("bgp %s" % bgp_no, strict=False)
            exec_cli("ipv4-family unicast", strict=False)
            for peer in default_peers:
                exec_cli("undo peer %s default-route-advertise" % peer)
            exec_cli("commit", strict=False)
        # Obtain the port list.
```

```
    output2 = exec_cli("display interface brief | include up | exclude NULL | exclude Loop | exclude M-G | exclu
de MEth", strict=False)
    rvs = re.findall(r"(. *?) \s+up\s+", output2)
    for rv in rvs:
      if '.' in rv:
        continue
      interfaces.append(rv.replace("(10GE)",""))
    # Shut down ports in batches.
    exec_cli("sys", strict=False)
    for interface in interfaces:
      exec_cli("interface %s" % interface)
      exec_cli("shutdown")
    exec_cli("commit", strict=False)
    radar_chek_lsw=radar_result(30)
    if radar_chek_lsw:
      logger.info('%s lsw isolate successfull'%hostname)
    else:
      return default_peers, backup_ospf_stub,interfaces,isolate_res
    # Check the isolation result.
    output3 = exec_cli("display interface brief | include up | exclude NULL0 | exclude Loop | exclude M-G | excl
ude MEth", strict=False)
    rvs = re.findall(r"(. *?) \s+up\s+", output3)
    for rv in rvs:
      if '.' in rv:
        continue
      isolate_res.append([hostname+'(fail)','Physical ports remain open after LSW port isolation, please che
ck')
      return default_peers, backup_ospf_stub,interfaces,isolate_res

# Stack ASW isolation.
  if cu_rol == 'ASW' and 'Standby' in stack_check and slot and action:
    slot1_check=True
    slot2_check=True
    slot1_int=exec_cli("display interface brief | in 1/0", strict=False,timeout=300)
    slot2_int=exec_cli("display interface brief | in 2/0", strict=False,timeout=300)
    if '*down ' in slot1_int and slot == '2' and not ingore_ADM:
      slot1_check=False
      isolate_res.append([hostname+'(fail)',unicode('Add/Drop Multiplexer (ADM) port is detected in slot 1 b
efore slot 2 isolation. Slot 2 is not isolated. Please check:\n%s','utf-8')%slot1_int])
      return default_peers, backup_ospf_stub,interfaces,isolate_res
    if '*down ' in slot2_int and slot == '1' and not ingore_ADM:
```

```
    if 'down' in slot2_int and slot == '1' and not ingore_ADM:
        slot2_check=False
        isolate_res.append([hostname+'(fail)',unicode('ADM port is detected in slot 2 before slot 1 isolation. Sl
ot 1 is not isolated. Please check:\n%s','utf-8')%slot2_int])
        return default_peers, backup_ospf_stub,interfaces,isolate_res
    if 'CE6851' in dis_version and slot == '1':
        exec_cli("sys", strict=False)
        exec_cli("interface M 0/0/0", strict=False)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
        exec_cli("interface range 40GE1/0/3 to 40GE1/0/6", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
        exec_cli("interface range 10GE1/0/1 to 10GE1/0/48", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
    elif 'CE6851' in dis_version and slot == '2':
        exec_cli("sys", strict=False)
        exec_cli("interface M 0/0/0", strict=False)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
        exec_cli("interface range 40GE2/0/3 to 40GE2/0/6", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
        exec_cli("interface range 10GE2/0/1 to 10GE2/0/48", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
    elif 'CE5850' in dis_version and slot == '1':
        exec_cli("sys", strict=False)
        exec_cli("interface  range 10GE 1/0/1 to 10GE 1/0/4", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
        exec_cli("interface range GE 1/0/1 to GE 1/0/48", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
    elif 'CE5850' in dis_version and slot == '2':
        exec_cli("sys", strict=False)
        exec_cli("interface  range 10GE 2/0/1 to 10GE 2/0/4", strict=False,timeout=300)
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
        exec_cli("interface range GE 2/0/1 to GE 2/0/48", strict=False,timeout=300)
```

```
        exec_cli("shutdown", strict=False,timeout=300)
        exec_cli("commit", strict=False,timeout=300)
    # Destack ASW and shut down uplink and downlink ports.
    if cu_rol == 'ASW' and '5.1' in cu_arc and 'CE6865' in dis_version:
      exec_cli("interface range 25GE 1/0/1 to 25GE 1/0/48", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False,timeout=300)
      # exec_cli("interface range 100 1/0/1 to 100 1/0/8", strict=False,timeout=300)
      # exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False,timeout=300)
    check_2100 = exec_cli('display curr config bgp | in 2100', strict=False)
    if cu_rol == 'ASW' and '5.0L' in cu_arc and 'CE6860' in dis_version and '2100' in check_2100:
      exec_cli("interface range 25GE 1/0/1 to 25GE 1/0/48", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False)
      exec_cli("interface range 100 1/0/1 to 100 1/0/12", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False)
    if cu_rol == 'ASW' and '4.2L' in cu_arc and 'CE6851' in dis_version and '2100' in check_2100:
      exec_cli("interface range 10GE1/0/1 to 10GE1/0/48", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False)
      exec_cli("interface range 40GE1/0/1 to 40GE1/0/4", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False)
    if cu_rol == 'ASW' and '4.2M' in cu_arc and 'CE6851' in dis_version and '2100' in check_2100:
      exec_cli("interface range 10GE1/0/1 to 10GE1/0/48", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False)
      exec_cli("interface range 40GE1/0/3 to 40GE1/0/6", strict=False,timeout=300)
      exec_cli("shutdown", strict=False,timeout=300)
      exec_cli("commit", strict=False)
    exec_cli("commit", strict=False)
    return default_peers,backup_ospf_stub,interfaces,isolate_res
```

9.Operations of basic cloud products

9.1. Elastic Compute Service (ECS)

9.1.1. ECS overview

Elastic Compute Service (ECS) is a user-friendly computation service featuring elastic processing capabilities that can be managed more efficiently than physical servers. You can create instances, resize disks, and release any number of ECS instances at any time based on your business needs.

An ECS instance is a virtual computing environment that includes basic components such as the CPU, memory, and storage. Users perform operations on ECS instances. Instances are the core concept of ECS, and are operated from the ECS console. Other resources such as block storage, images, and snapshots can be used only after they are integrated with ECS instances. For more information, see ECS instance.

ECS instance

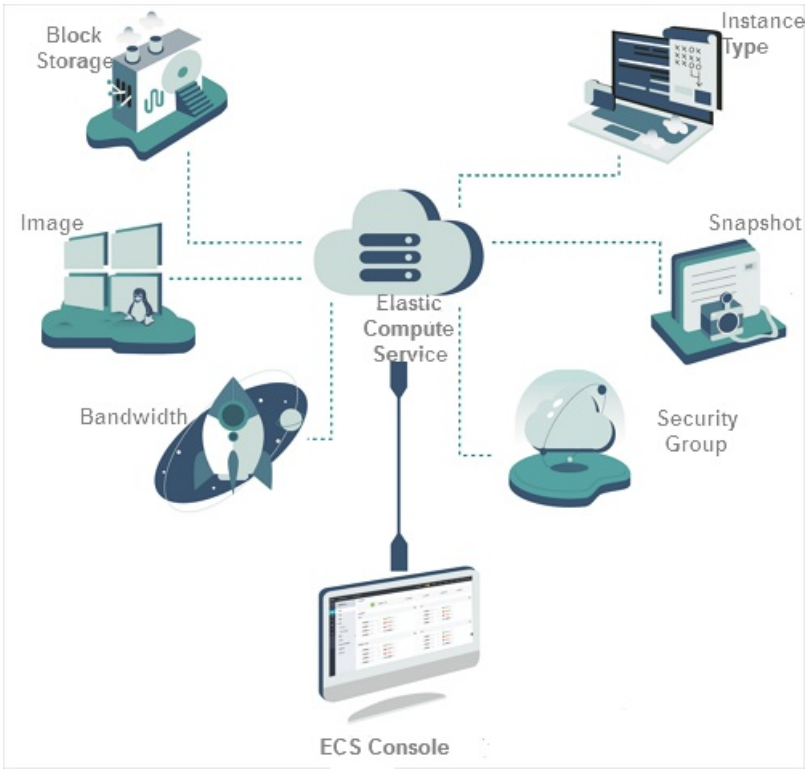

9.1.2. Log on to the Apsara Stack Operations console

This topic describes how to log on to the Apsara Stack Operations (ASO) console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   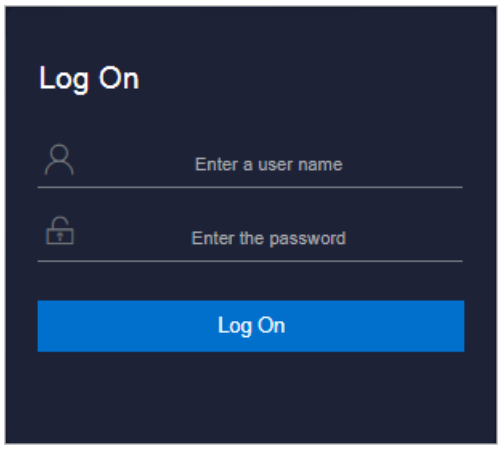

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.

9.1.3. ECS operations and maintenance

9.1.3.1. Overview

The ECS Operations and Maintenance Platform is a platform for support engineers to operate and monitor ECS instances, help users troubleshoot problems with ECS instances, and ensure that ECS instances are properly operated and utilized.

9.1.3.2. VM

9.1.3.2.1. Overview

On the ECS Operations and Maintenance Platform page, the existing ECS VM information and available O&M functions are displayed. You can search for, start, and migrate a VM as needed.

9.1.3.2.2. Search for VMs

You can view the list of existing VMs and their information in the ASO console.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.
6. In the VM list, click a VM ID. You can view the VM information in the **VM Details** pane.

9.1.3.2.3. Start a VM

In the ASO console, you can start a VM in the same way as you start a real server.

Prerequisites

The VM to be started must be in the **Stopped** state.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.
6. In the VM list, select the VM to be started. Click **Start** above the list.
7. In the Start VM dialog box, set Start. You can select **Normal** or **Repair**.

   > **Note** If you want to reset the network settings of the VM, set Start to **Repair**. Otherwise, set Start to **Normal**.

8. Set Operation Reason. Click **OK**.

9.1.3.2.4. Stop a VM

In the ASO console, you can stop a VM in the same way as you stop a real server.

Prerequisites

The VM to be stopped must be in the **Running** state.

Context

This operation will interrupt the programs that are running on the VM. Perform this operation during off-peak hours to minimize the impact on services.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.
6. In the VM list, select the VM to be stopped. Click **Stop** above the list.
7. In the Stop VM dialog box, set Shutdown Policy. You can select **Non-force Shutdown** or **Force Shutdown**.

   > **Note** When Force Shutdown is selected, the VM is stopped regardless of whether its processes have been stopped. We recommend that you do not select Force Shutdown unless Non-force Shutdown does not work.

8. Set Operation Reason. Click **OK**.

9.1.3.2.5. Restart a VM

In the ASO console, you can restart a VM in the same way as you restart a real server.

Prerequisites

The VM to be restarted must be in the **Running** state.

Context

This operation will interrupt the programs that are running on the VM. Perform this operation during off-peak hours to minimize the impact on services.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.

6. In the VM list, select the VM to be restarted. Click **Reboot** above the list.
7. In the Reboot VM dialog box, set Start and Shutdown Policy.
   - You can set Start to **Normal** or **Repair**.
   - You can set Shutdown Policy to **Non-force Shutdown** or **Force Shutdown**.
8. Set Operation Reason. Click **OK**.

9.1.3.2.6. Cold migration

In the ASO console, you can perform cold migration on a VM to implement failover.

Prerequisites

Cold migration requires that the VM be taken offline. Make sure that the VM is in the **Stopped** state before you migrate it.

Context

If a VM or an NC fails, you must fail over the VM by stopping the VM and migrating it to a new NC. Failover can be performed only within the same zone. Cross-zone failover is not allowed.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.
6. In the VM list, select the VM to be migrated. Click **Stop and Migrate** above the list.
7. In the Stop and Migrate VM dialog box, configure the parameters described in the following table.

Parameter	Description
Switchable NC	The destination NC to which to migrate the VM.
Switchover Policy	The switchover policy. Valid values: ○ **Force Migrate** ○ **Active Migrate**
Start	The startup mode. Valid values: ○ **Normal** ○ **Repair**

Parameter	Description
Recover	The recovery mode. Valid values: ○ **Start After Migration** ○ **Stop After Migration** ○ **Status Unchanged After Migration** **Status Unchanged After Migration** takes effect only on VMs that are in the Pending state.

8. Set Operation Reason. Click **OK**.

9.1.3.2.7. Hot migration

In the ASO console, you can perform hot migration on VMs.

Context

- You can use hot migration to migrate a VM in the **Running** state from one NC to another without interrupting normal services. Hot migration can be used for load balancing or other purposes. If a failure occurs, you must perform cold migration instead of hot migration. For more information, see Cold migration.
- Security risks may arise if you perform hot migration. Exercise caution when you perform hot migration.
- Hot migration will not interrupt services that are running on the VM.
- Hot migration can be performed only within the same zone. Cross-zone hot migration is not allowed.

Prerequisites

You can perform hot migration only on VMs in the **Running** state.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.
6. In the VM list, select the VM to be migrated. Choose **More > Online Migrate** above the list.
7. Set Throughput Limit. The value of Throughput Limit ranges from 1 MByte/s to 1,000 MByte/s. The default value is 20 MByte/s.
8. Set Operation Reason. Click **Online Migrate**. The destination NC is automatically selected during migration. You can view the ID of the destination NC in the migration result.

9.1.3.2.8. Reset a disk

You can reset disks to restore them to their initial status.

Prerequisites

- When you reset a disk, applications that are installed on the disk are lost. Before you perform a reset operation, make sure that you have backed up your data.
- To reset a disk, make sure that the VM to which it is attached is in the **Stopped** state.

Context

After a disk is reset, it is restored to its initial status but is not reformatted. The image that is used to create the disk still exists after the disk is reset.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set search conditions and click **View**.
6. In the VM list, select the VM to which the disk to be reset is attached. Choose **More > Reset Disk** above the list.
7. In the Reset Disk dialog box, select the disk to be reset and set Operation Reason. Click **OK**.

9.1.3.3. Disks

9.1.3.3.1. Overview

In an ECS instance, cloud disks can be considered as physical disks. You can mount, detach, and create snapshots for disks.

9.1.3.3.2. Search for disks

In the ASO console, you can view the list of existing disks and their information.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set search conditions and click **View**.

9.1.3.3.3. View snapshots

In the ASO console, you can view the list of snapshots that are created for a disk and their information.

Procedure

1. Log on to Apsara Stack Operations.

2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set search conditions and click **View**.
6. Find the disk whose snapshots you want to view and choose  **> View Snapshot**.

   The information of all snapshots on the disk is displayed.

9.1.3.3.4. Attach a disk

After a disk is created, you must attach the disk to a VM.

Context

Only disks that are in the **Available** state can be attached.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set search conditions and click **View**.
6. Find the disk to be attached and choose  **> Mount**.
7. In the Mount Disk dialog box, set VM ID and Operation Reason. Click **OK**.

9.1.3.3.5. Detach a disk

In the ASO console, only data disks can be detached. System disks and local disks cannot be detached.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set search conditions and click **View**.
6. Find the disk to be detached and choose  **> Detach**.
7. In the Detach Disk dialog box, set Operation Reason. Click **OK**.

9.1.3.3.6. Create a snapshot

In the ASO console, you can manually create snapshots for disks.

Context

Snapshots can be created only for system disks.

Procedure

1. Log on to Apsara Stack Operations。
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set search conditions and click **View**.
6. Find the disk for which you want to create a snapshot and choose  **> Take Snapshot**.
7. In the Disk Snapshot dialog box, set Snapshot Name, Snapshot Description, and Operation Reason. Click **OK**.

9.1.3.4. Snapshots

9.1.3.4.1. Overview

A snapshot stores the data stored on a disk for a certain point in time. Snapshots can be used to back up data or create a custom image.

When using disks, note the following points:

- When writing or saving data to a disk, we recommend that you use the data on one disk as the basic data for another disk.
- Although the disk provides secure data storage, you must still ensure that stored data is complete. However, data can be stored incorrectly due to an application error or malicious usage of vulnerabilities in the application. For these cases, a mechanism is required to ensure that data can be recovered to the desired state.

Alibaba Cloud allows you to create snapshots to retain copies of data on a disk for specific points in time.

9.1.3.4.2. Search for snapshots

In the ASO console, you can view the list of existing snapshots and their information.

Prerequisites

The AliUid of the disk for which the snapshot is taken is obtained. For more information, see Search for disks.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Snapshots** tab.

5. On the Snapshots tab, set search conditions and click **View**. AliUid is a required search condition.

9.1.3.4.3. Delete a snapshot

In the ASO console, you can delete snapshots that are no longer needed.

Prerequisites

The AliUid of the disk for which the snapshot is taken is obtained. For more information, see Search for disks.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Snapshots** tab.
5. On the Snapshots tab, set search conditions and click **View**.AliUid is a required search condition.
6. Find the snapshot to be deleted and choose  **> Delete**.
7. In the Delete Snapshot dialog box, set Operation Reason. Click **OK**.

9.1.3.4.4. Create an image

You can create a custom image from a snapshot. The image includes the operating system and environment variables of the snapshot.

Prerequisites

The AliUid of the disk for which the snapshot is taken is obtained. For more information, see Search for disks.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Snapshots** tab.
5. On the Snapshots tab, set search conditions and click **View**.AliUid is a required search condition.
6. Find the snapshot from which you want to create an image and choose  **> Create Image**.
7. In the Create Image dialog box, set Image Name, Image Version, Image Description, and Operation Reason. Specify whether the system disk for which the snapshot was taken is based on a public image or a custom image. Click **OK**.

9.1.3.5. Images

9.1.3.5.1. Overview

An ECS image is a template that contains software configurations such as the ECS instance operating system and the programs and servers for applications. You must specify an ECS image to create an instance. The operating system and software provided by the image will be installed on the instance that you create.

9.1.3.5.2. Search for images

In the ASO console, you can view the list of existing images and their information.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Images** tab.
5. On the Images tab, set search conditions and click **View**.

> **Note** If you set Image Type to Custom Image, you must also set AliUid.

9.1.3.6. Security groups

9.1.3.6.1. Overview

A security group is a virtual firewall that provides Stateful Packet Inspection (SPI). Security groups provide virtual firewall-like functionality and are used for network access control for one or more ECS instances. They are important means of network security isolation and are used to divide security domains on the cloud.

Security group rules can permit the inbound and outbound traffic of the ECS instances associated with the security group. You can authorize or cancel security group rules at any time. Changes to security group rules are automatically applied to ECS instances that are members of the security group.

When you configure security group rules, ensure that the rules are concise and easy to manage. If you associate an instance with multiple security groups, hundreds of rules may apply to the instance, which may cause connection errors when you access the instance.

9.1.3.6.2. Search for security groups

In the ASO console, you can view the list of current security groups and their information.

Context

After an ECS instance is added to a security group, you can add security group rules to allow or deny public or internal network traffic to and from the ECS instance. You can add or delete security group rules at any time. Changes to security group rules are automatically applied to ECS instances in the security group.

> Note
>
> - If the two security group rules differ only in Authorization Policy, the deny rules takes precedence over allow rules.
> - No rule in a security group can allow outbound traffic from an instance while denying inbound traffic to the instance.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Security Groups** tab.
5. On the Security Groups tab, set search conditions and click **View**.

9.1.3.6.3. Add security group rules

You can add rules to security groups to control access to or from the instances in the security groups.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Security Groups** tab.
5. On the Security Groups tab, set search conditions and click **View**.
6. Find the security group to which you want to add a security group rule and choose  **> Add Rule**.
7. In the Add Rule dialog box, configure the parameters.The following table describes the parameters.

Parameter	Description
Protocol	○ TCP ○ UDP ○ ICMP ○ GRE ○ ALL: All protocols are supported.
Rule Priority (1-100)	A smaller value indicates a higher priority.

Parameter	Description
Network Type	○ Public: the Internet ○ Internal: the internal network
Authorization Policy	○ Accept: grants access. ○ Drop: discards the packet on access. ○ Reject: denies the packet on access.
Port Number Range	Valid values: 1 to 65535. Example: 1/200, 80/80, or -1/-1.
Access Direction	○ Ingress: allows inbound traffic. ○ Egress: allows outbound traffic.
IP Address Range	Enter an IP address or a CIDR block. Only IPv4 addresses are supported. Example: 10.0.0.0, 0.0.0.0/0, or 192.168.0.0/24.
Security Group ID	Enter the ID of the security group which you want to allow or deny access to the current security group.
Operation Reason	Optional. Enter a reason for the operation.

8. Click **OK**.

9.1.3.7. Custom instance types

9.1.3.7.1. Add custom instance types

When existing instance types cannot meet your business requirements, you can add custom instance types in the ASO console and create instances of the custom instance types.

Context

Custom instance types are of the ecs.anyshare instance family. You can set the instance type name, number of vCPUs, and memory size. Parameters such as base bandwidth, network packet forwarding rate, and NIC queues are generated by the system. For more information, see Instance types in *ECS Product Introduction*.

> **Note** All custom instance types are shared instance types.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Custom Instance Type** tab.

5. Click **Add**.
6. In the **Add Instance Type** pane, set the Instance Type, vCPUs, and Mem (GiB) parameters. Then, click **OK**.

Result

The new custom instance type is displayed in the custom instance type list. After you add a custom instance type, you can create ECS instances of the custom instance type. Select ecs.anyshare as the instance family. For more information, see Create an instance in *ECS User Guide* .

9.1.3.7.2. Query custom instance types

In the ASO console, you can view the custom instance types that you have added and their information.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Custom Instance Type** tab.
5. View information about the custom instance types. If the custom instance type list is not refreshed, click **Search**.

9.1.3.7.3. Modify custom instance types

If you want to retain a custom instance type but the specifications of this custom instance type do not meet your requirements, you can modify the number of vCPUs and memory size of the custom instance type.

Prerequisites

A custom instance type is added.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Custom Instance Type** tab.
5. Find the custom instance type that you want to modify and click **Modify** in the **Actions** column.
6. In the **Modify Instance Type** pane, set the vCPUs and Mem (GiB) parameters. Then, click **OK**.

9.1.3.7.4. Delete custom instance types

In the ASO console, you can delete custom instance types that are no longer needed. After you delete a custom instance type, you cannot select it when you create a new instance. However, existing instances of this custom instance type can continue to be used.

Procedure

1. Log on to Apsara Stack Operations.
2. In the top navigation bar, select an environment version and a region.
3. Choose **Products > ECS > ECS Operations and Maintenance Platform**.
4. Click the **Custom Instance Type** tab.
5. Find the custom instance type that you want to delete and click **Delete** in the **Actions** column.
6. In the **Deleted** message, click **OK**.

Result

The custom instance type is removed from the custom instance type list.

9.1.4. EBS

EBS provides the following features: EBS dashboard, block master O&M, block server O&M, snapshot server O&M, block gcworker O&M, device O&M, rebalance, IO HANG fault analysis, slow IO analysis, and inventory configuration.

9.1.4.1. EBS dashboard

The EBS Dashboard module shows the data overview and trend charts of cluster usage of EBS clusters.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > EBS Dashboard**.On the page that appears, cluster overview information and trend charts of cluster of all EBS clusters are displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. View the following information:
   - The **Overview** section shows data overview information of the selected cluster, including the storage space, server information, and health information.

     In the **Health** section, when the value of **Abnormal Cloud Disks**, **Abnormal Masters**, **Abnormal Block GcWorker**, or **Abnormal Block Servers** is greater than 0, it is displayed in red.
   - The **Trend Chart of Cluster Usage** section shows the storage usage curve of the cluster for the last 30 days.

9.1.4.2. Block master operations

The Block Master Operations module shows the block master node information of EBS clusters, including the IP address and role. The module also allows you to switch the role of a node to LEADER as well as query and configure flags.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Block Master Operations**.On the page that appears, the master node list and cluster information of the first cluster in the **Cluster Name** drop-down list are displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:

- View the master node list

  You can view the master node information of the selected cluster, including the IP address, role, log ID, and status.

- Switch to LEADER

  A LEADER role for a master node has the same functions as a FOLLOWER role, including controlling and scheduling resources, as well as controlling deployment and service configurations.

  If a node in the master node list assumes a FOLLOWER role, you must switch its role to LEADER. Click **Switch to LEADER** in the **Actions** column. In the message that appears, click **OK**.

- Query a flag

  In the master node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, set flag_key, and then click **Submit**. The deployment and service configurations of the block master node are displayed.

  Perform the following steps to query the flag_key value:

  a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.
  b. Enter EBS in the **Cluster** search box.
  c. Find the EBS cluster and click the cluster name.
  d. Click the **Configure** tab.
  e. Find the *pangu_blockmaster_flag.json* file in */services/EbsBlockMaster/user/pangu_blockmaster*.

     The flag_key values of all block master nodes are stored in the *pangu_blockmaster_flag.json* file.

- Configure a flag

  If you want to modify the deployment and service configurations of a block master node, you can configure a flag and assign it to the node.

  In the master node list, click **Configure Flag** in the **Actions** column corresponding to a LEADER node. In the dialog box that appears, configure the parameters, and then click **OK**.

  The following table describes the parameters of a flag.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_blockmaster_flag.json* file.
flag_value	This value is customized.
flag_type	Select a flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Check the maser node status

  In the master node list, choose **More > Check Master Status** in the **Actions** column corresponding to a node.

- Query the version information

  In the master node list, choose **More > Query Version Information** in the **Actions** column corresponding to a node.

- Query the cluster overview information

  You can query the disk size, number of segments, total storage size, and storage usage of the cluster.

9.1.4.3. Block server operations

The Block Server Operations module shows the block server node information of EBS clusters, including the IP address, status, and real-time server load. The module also allows you to query and modify flags, configure server node status, as well as add nodes to and delete nodes from blacklists.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Block Server Operations**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the server node list

     You can view server node information of the cluster including the IP addresses, status, number of segments, and real-time load (read IOPS, write IOPS, read bandwidth, write bandwidth, read latency, and write latency).

   - Query a flag

     In the server node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, set flag_key, and then click **Submit**. The deployment and service configurations of the block server node are displayed.

     Perform the following steps to query the flag_key value:

     a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.
     b. Enter EBS in the **Cluster** search box.
     c. Find the EBS cluster and click the cluster name.
     d. Find the *pangu_blockserver_flag.json* file in */services/EbsBlockServer/user/pangu_blockserver*.

        The flag_key values of all block server nodes are stored in the *pangu_blockserver_flag.json* file.

   - Configure a flag

     In the server node list, click **Configure Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, flag_value, and flag_type, and then click **OK**.

The following table describes the parameters of a flag.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_blockserver_flag.json* file.
flag_value	This value is customized.
flag_type	Select a flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure server node status

  In the server node list, choose **More > Set Server Status** in the **Actions** column corresponding to a node. In the dialog box that appears, specify server node status, and then click **OK**.

  The following table describes the server node status.

Status	Description
NORMAL	Indicates that the node is normal.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	The node has been disabled.
UPGRADE	The node has been upgraded.
RECOVERY	The node has been restored.

- Query the version information

  In the server node list, choose **More > Query Version Information** in the **Actions** column corresponding to a node. In the dialog box that appears, view the version information of the block server node.

- Add a block server node to the blacklist

  In the upper-right corner of the **Block Server Blacklist** section, click **Add**. In the dialog box that appears, select the IP address of the block server node that you want to add to the blacklist, and then click **OK**.

  The block server node that is added to the blacklist is disabled and will no longer provide services.

- View the block server blacklist

You can view all block server nodes that are added to the blacklist in the **Block Server Blacklist** section.

- Remove a block server node from the blacklist

  In the **Block Server Blacklist** section, find the block server node that you want to remove from the blacklist, and then click **Delete** in the **Actions** column. In the message that appears, click **OK**.

  The block server node that is removed from the blacklist can continue to provide internal and external services.

9.1.4.4. SnapShotServer

The SnapShotServer module shows the snapshot server node information of EBS clusters, including the IP address, status, and other performance parameters. The module also allows you to query and modify flags and configure snapshot server node status.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > SnapShotServer**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the snapshot server node list

     You can view snapshot server node information of the cluster including node IP address, status, loading rate, and number of uploads, replicas, and delayed loadings.
   - Query a flag

     In the snapshot server node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, set flag_key, and then click **Submit**. The deployment and service configurations of the snapshot server node are displayed.

     Perform the following steps to query the flag_key value:

     a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.

     b. Enter EBS in the **Cluster** search box.

     c. Find the EBS cluster and click the cluster name.

     d. Find the *pangu_snapshotserver_flag.json* file in */services/EbsSnapshotServer/user/pangu_snapshotserver*.

     The flag_key values of all snapshot server nodes are stored in the *pangu_snapshotserver_flag.json* file.
   - Configure a flag

     In the snapshot server node list, click **Configure Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, flag_value, and flag_type, and then click **OK**.

     The following table describes the parameters of a flag.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_snapshotserver_flag.json* file.
flag_value	This value is customized.
flag_type	Select a flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the snapshot server node status

  In the snapshot server node list, choose **More > Set snapshotserver Status** in the **Actions** column corresponding to a node. In the dialog box that appears, select the snapshot server node status, and then click **OK**.

  The following table describes the snapshot server node status.

Status	Description
NORMAL	Indicates that the node is normal.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	Indicates that the node has been disabled

- Query the version information

  In the snapshot server node list, choose **More > Version Information** in the **Actions** column corresponding to a node. In the dialog box that appears, view the version information of the node.

9.1.4.5. Block gcworker operations

The Block Gcworker Operations module allows you to view the IP addresses and statuses of block gcworker nodes in EBS clusters. You can also query and modify flags, configure the gcworker node status, and query version information.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Block GcWorker Operations**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:

- View the gcworker node list

  You can view the IP addresses and statuses of the block gcworker nodes in the selected cluster.

- Query a flag

  In the gcworker node list, click **Query Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, and then click **Submit**. The deployment and service configurations of the block gcworker node are displayed.

  Perform the following steps to query the flag_key value:

  a. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, choose **Operations > Cluster Operations**.
  b. Enter EBS in the **Cluster** search box.
  c. Find the EBS cluster and click the cluster name.
  d. Click the **Configure** tab.
  e. Find the *pangu_blockgcworker_flag.json* file in */services/EbsBlockGCWorker/user/pangu_blockgcworker*.

     The flag_key values of all block server nodes are stored in the *pangu_blockgcworker_flag.json* file.

- Configure a flag

  In the gcworker node list, click **Configure Flag** in the **Actions** column corresponding to a node. In the dialog box that appears, specify flag_key, flag_value, and flag_type, and then click **OK**.

  The following table describes the parameters.

Parameter	Description
flag_key	This value is obtained from the service template of the EBS cluster that is stored in the *pangu_blockgcworker_flag.json* file.
flag_value	This value is customized.
flag_type	Select a flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the gcworker node status

  In the gcworker node list, choose **More > Configure gcworker Status** in the **Actions** column corresponding to a node. In the dialog box that appears, specify the gcworket node status and click **OK**.

  The following table describes the gcworker status.

Status	Description
NORMAL	Indicates that the node is running normally.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	Indicates that the node has been disabled.

- Query the version information

  In the gcworker node list, choose **More > Query Version Information** in the **Actions** column corresponding to a node. In the dialog box that appears, view the version information of the block gcworker node.

9.1.4.6. Device operations

The Device Operations module allows you to view disk information in EBS clusters such as the disk ID, status, capacity, and type. You can also perform flush operations, modify disk configurations, query segment information, and open, close, delete, and restore devices.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Device Operations**.On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
2. Select a cluster from the **Cluster Name** drop-down list.
3. Perform the following operations:
   - View the device list

     You can view the total number of devices, the total logical space of devices, and information about each device in the cluster, including the device ID, status, logical capacity, number of segments, mode, and flags.
   - Global check segments

     In the upper-right corner of the **Device List** section, click **Global Check Segment**. You can view all the segments in the selected cluster and their indexes and statuses.
   - Check the status of disks

     In the upper-right corner of the **Device List** section, click **Check Cloud Disk Status**. You can view the number of invalid disks in the selected cluster.
   - Query device information

     In the device list, click **Query Device Information** in the **Actions** column corresponding to a device. In the dialog box that appears, view configuration information of the disk such as the disk ID, disk status, and disk capacity.
   - Delete a device

     In the device list, click **Delete** in the **Actions** column corresponding to a device.

After you delete the disk, its status becomes **DELETING**, and the disk is unavailable. You are not allowed to perform operations such as enabling the device and modifying the configuration.

- Restore a device

  In the device list, find the device whose status is **DELETING** and click **Restore** in the **Actions** column. In the dialog box that appears, click **OK**.

  After you restore the disk, it becomes available. You can perform operations such as enabling the disk and modifying the configuration.

- Enable a device

  In the device list, choose **More > Turn On** in the **Actions** column corresponding to a device. In the dialog box that appears, configure the parameters and click **Submit**.

  > **Note** You can perform read and write operations on a disk only after the disk is enabled.

  The following table describes the parameters for enabling a device.

Parameter	Description
client_ip	Optional. Specifies the client where the disk is enabled. The client IP address is the IP address of the block server. If the client IP address is not specified, the IP address of the local server is used by default.
token	Specifies a string as a token to be used to disable the device.
mode	Specifies the disk mode. Valid values: ■ **ro**: read-only ■ **rw**: read/write Default value: **rw**

- Disable a device

  > **Notice** After a disk is disabled, data can no longer be read from or written to the disk. Proceed with caution.

  In the device list, choose **More > Turn Off** in the **Actions** column. In the dialog box that appears, configure the parameters and click **Submit**.

  The following table describes the parameters for disabling a device.

Parameter	Description
client_ip	Specifies the client IP address of the disk to be disabled. If the client IP address is not specified, the IP address of the local server is used by default.

Parameter	Description
token	Specifies the token for disabling the device, which is configured when the device is enabled. You can query the token by running the **dev -query** command.
open_ver	Specifies the current openversion of the device if the client IP address is not specified. If you specify a client IP address, you do not need to specify openversion. You can query openversion by running the **dev -query** command.

- Flush

  In the device list, choose **More > Flush** in the **Actions** column. In the dialog box that appears, configure the parameters and click **Submit** to clear the current disk or the segment transaction logs on the disk.

  The following table describes the parameters.

Parameter	Description
segment	Select the segment to be flushed. If you do not select any segments, all segments are flushed.
ifnsw	Valid values: ■ **0**: specifies that the index file is flushed during the flush. ■ **1**: specifies that the index file is not flushed during the flush.
dfnsw	Valid values: ■ **0**: specifies that data files are flushed during the flush. ■ **1**: specifies that data files are not flushed during the flush.

- Global flush

  You can perform the flush operation to clear disks or the transaction logs of segments.

  In the upper-right corner of the **Device List** section, click **Global Flush**. In the dialog box that appears, select ifnsw and dfnsw, and then click **OK** to clear all the disks or the transaction logs of segments in the selected cluster.

- Query configuration status

In the device list, choose **More > Query Configuration Status** in the **Actions** column corresponding to a device. In the dialog box that appears, specify config_ver and click **OK**. You can determine whether the disk can be configured based on the query result.

You can obtain the config_ver value from the device information.

- Modify device configurations

  You can modify the configurations of a disk, such as specifying whether to enable data compression, compression algorithms, and storage modes.

  In the device list, choose **More > Modify Device Configurations** in the **Actions** column corresponding to a device. In the dialog box that appears, modify the parameters and click **OK**.

  The following table describes the parameters.

Parameter	Description
compress	Select whether to enable data compression. Valid values: ■ **enable** ■ **disable**
algorithm	Select a data compression algorithm. Valid values: ■ **0**: indicates that no data compression algorithms are used. ■ **1**: indicates that the snappy data compression algorithm is used. ■ **2**: indicates that the lz4 data compression algorithm is used.
ec	Select whether to enable the ec storage mode. Default value: disable. Valid values: ■ **enable** ■ **disable**
data_chunks	Specifies the number of data chunks. Default value: 8.
parity_chunks	Specifies the number of parity chunks. Default value: 3.
packet_bits	Specifies the size of single data block in ec mode. Default value: 15.
copy	Specifies the number of data replicas. Default value: 3.
storage_mode	Specifies the storage mode of the disk.
cache	Select whether to enable the cache mode. Default value: 0. Valid values: ■ **0**: disabled ■ **1**: enabled

Parameter	Description
storage_app_name	Specifies the data storage name.
simsuppress	Select whether to enable the delay simulation feature. Default value: disable. Valid values: ■ **enable** ■ **disable**
baselatency	Specifies the basic latency. Default value: 300.
consumespeed	Specifies the processing speed. Default value: 256 bit/μs.
lat80th	Specifies the quantile jitter control of the latency as 80%.
lat90th	Specifies the quantile jitter control of the latency as 90%.
lat99th	Specifies the quantile jitter control of the latency as 99%.

- Query segment information

  In the device list, choose **More > Segment Information** in the **Actions** column corresponding to a device. In the dialog box that appears, view the segment information such as index and status.

- Query segments

  In the device list, choose **More > Check Segment** in the **Actions** column corresponding to a device. In the dialog box that appears, select the segment to be checked and click **Submit**. You can view the information of the selected segment such as index and status.

9.1.4.7. Enable or disable Rebalance

When segments are unevenly distributed among a block server, you can enable the Rebalance feature to redistribute the segments. After the redistribution, you can disable Rebalance.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > EBS > Rebalance**.
2. Click **Enable Rebalance** or **Disable Rebalance**.

   After you click **Enable Rebalance**, the status of Rebalance changes to **running**.

   After you click **Disable Rebalance**, the status of Rebalance changes to **stopped**.

9.1.4.8. IO HANG fault analysis

The IO HANG Fault Analysis module allows you to view the affected virtual machine (VM) list, VM cluster statistics, and device cluster statistics.

Procedure

1. Log on to the ASO console.

2. In the left-side navigation pane, choose **Storage Operation Center > EBS > IO HANG**. By default, the system displays the affected VM list, VM cluster statistics, and device cluster statistics in the last 24 hours.
3. Select the time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or customize the time range) that you are about to view and then click **Search**. View the following information:
   - **Affected VM List**

     The **Affected VM List** section displays the IO HANG start time and recovery time of all the VMs, and the cluster name and user ID to which these VMs belong.

     To view the information of a cluster, user, or VM, enter the cluster name, user ID, or VM name in the search box to perform a fuzzy search.

     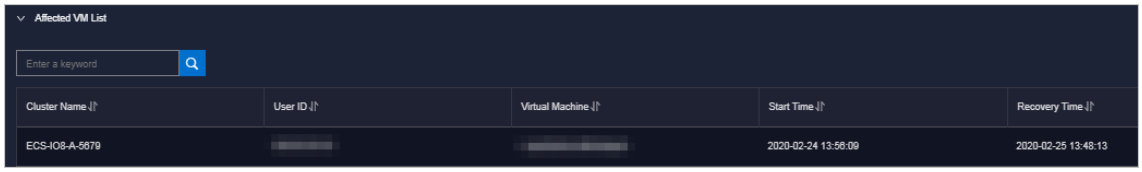

   - **VM Cluster Statistics**

     The **VM Cluster Statistics** section displays the number of affected VMs in a cluster.

     To view the VM statistics of a cluster, enter the cluster name in the search box to perform a fuzzy search.

     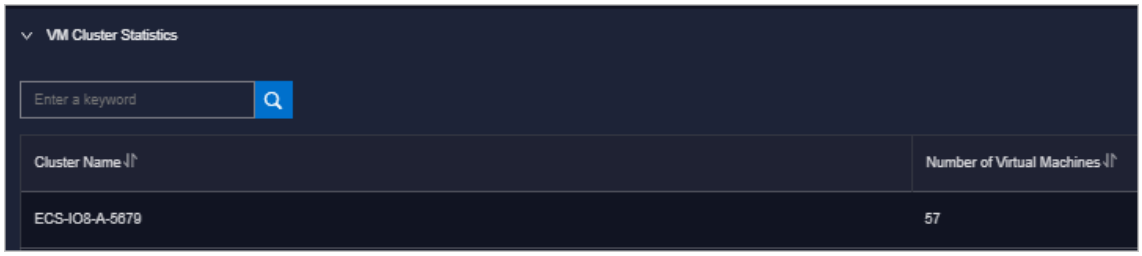

   - **Device Cluster Statistics**

     The **Device Cluster Statistics** section displays the number of affected devices in a cluster.

     To view the device statistics of a cluster, enter the cluster name in the search box to perform a fuzzy search.

     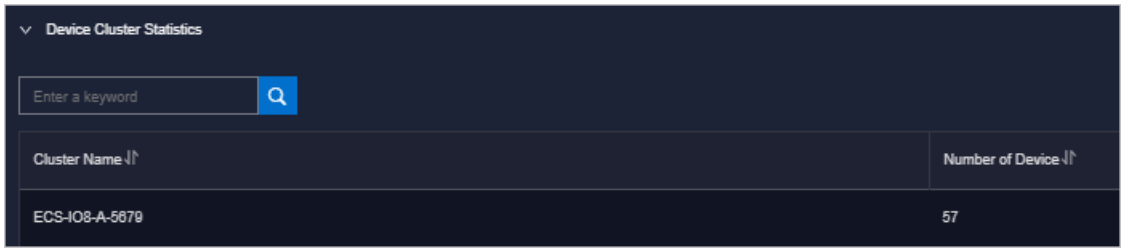

9.1.4.9. Slow IO analysis

The Slow IO Analysis module allows you to view the Slow IO list, top 10 NCs, cluster statistics, top 5 cluster statistics, and reasons.

Procedure

1. Log on to the ASO console.

2. In the left-side navigation pane, choose **Storage Operation Center > EBS > Slow IO**. By default, the system displays the Slow IO list, top 10 NCs, cluster statistics, top 5 cluster statistics, and reasons in the last 24 hours.
3. Select the time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or customize the time range) that you are about to view and then click **Search**. View the following information:
   - **Slow IO List**

     The **Slow IO List** section displays the following Slow IO-related data: cluster name, NC IP address, virtual machine, device ID, storage type, start time, recovery time, number of Slow IO, and reason.

     To view the information of a cluster, NC, or block device, you can enter the cluster name, NC IP address, or device ID in the search box to perform a fuzzy search.

     You can also sort by Cluster Name, NC IP, Virtual Machine, Device ID, Storage Type, Start Time, Recovery Time, Number of Slow IO, and Reason as needed.

     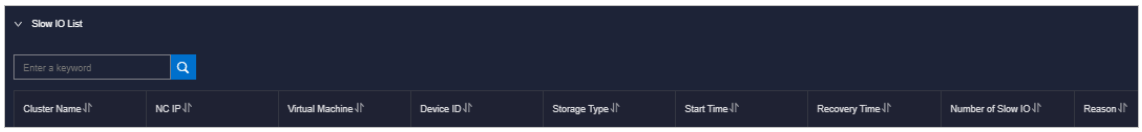

   - **Top Ten NC**

     The system displays the information of top 10 NCs by using a graph and a list.

     Where,

     - The **Graphic Analysis** section displays the proportion for the number of Slow IO in each cluster of the top 10 NCs by using a pie chart.
     - The **Top Ten NC** section displays the NC IP address, cluster name, Slow IO, percentage, and major reason of the top 10 NCs with the most Slow IO by using a list.

       To view the information of a cluster or NC, you can enter the NC IP address or cluster name in the search box to perform a fuzzy search.

       You can also sort by NC IP, Cluster Name, Slow IO, and Major Reason as needed.

       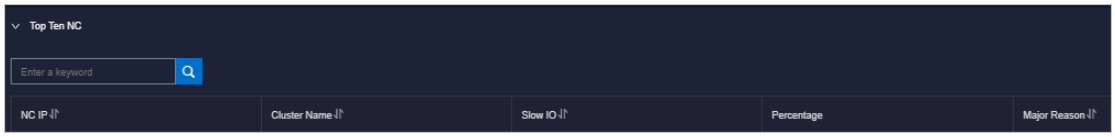

   - **Cluster Statistics**

     The **Cluster Statistics** section displays the cluster name, number of devices, number of Slow IO, percentage, and major reason of a cluster with Slow IO.

     To view the information of a cluster, you can enter the cluster name in the search box to perform a fuzzy search.

     You can also sort by Cluster Name, Number of Device, Number of Slow IO, and Major Reason as needed.

   - **Top Five Cluster Statistics**

     The system displays the statistics of top 5 clusters by using a graph and a list.

     Where,

- The **Top Five Cluster Statistics** section displays the cluster name, number of devices, number of Slow IO, percentage, and major problem of the top 5 clusters with the most Slow IO by using a list.

  To view the information of a cluster, you can enter the cluster name in the search box to perform a fuzzy search.

  You can also sort by Top Five Cluster, Number of Device, Number of Slow IO, and Major Problem as needed.

- The **Graphic Analysis** section displays the proportion for the number of Slow IO in each of the top 5 clusters by using a pie chart.

- **Reason**

  The system displays the reason statistics by using a graph and a list.

  Where,

- The **Reason** section displays the number of Slow IO from the dimension of reasons.

  To view the information of a reason, you can enter the reason information in the search box to perform a fuzzy search.

  You can also sort by Reason and Number of Slow IO as needed.

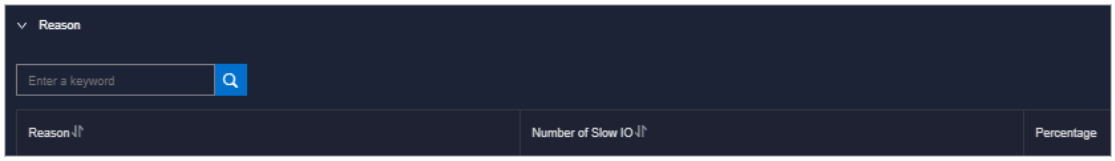

- The **Graphic Analysis** section displays the proportion of reasons by using a pie chart.

9.1.4.10. Inventory settings

The **Inventory Settings** module allows you to view the sales status of a cluster, configure the oversold ratio of a cluster, and configure whether a cluster is on sale.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **Storage Operation Center > EBS > Inventory Settings**. By default, the system displays the data, namely the cluster name, oversold ratio, and sales status, of all the clusters in the current environment.

3. Complete the following configurations:
   - Select a cluster. Enter a number in the **Adjust Setting Oversell Ratio(%)** field, and then click **Confirm** to configure the oversold ratio of the cluster.
   - Select a cluster. Turn on or off the **Adjustment of sales status** switch to configure whether the cluster is on sale.

9.1.5. VM hot migration

9.1.5.1. Overview

Hot migration is the process of migrating a running VM from one host to another. During migration, the VM runs normally and its services are not aware that any migration task is occurring. However, these services can detect a very short interruption between 100 and 1,000 ms.

Scenarios

During system operations and maintenance, hot migration is typically used for the following scenarios:

- **Active O&M**: The host is faulty and must be repaired, but the fault does not affect the operation of the system. You can use hot migration to migrate the VM to another host and repair the faulty host in offline mode.
- **Server load balancing**: When a host is experiencing a high load, you can migrate some of its VMs to other idle hosts to reduce resource consumption on the source host.
- Other scenarios where a VM must be migrated without affecting its business operations.

9.1.5.2. Limits on hot migration

Before performing hot migration, you must understand the limits.

The hot migration feature of Apsara Stack is subject to the following limits:

- Only the go2hyapi command can be used to implement hot migration in the KVM virtualization environment. ECS Operations and Maintenance Platform does not support hot migration.
- Only standard ECS instances support hot migration. ECS provides a list of migratable images. Alibaba Cloud does not take any responsibility for errors that occur when migrating a VM that is not included in the list of migratable images.
- If a VM is used as an RS to provide SLB or as a client to access SLB, the previous session will be closed after hot migration. New sessions created after migration are not affected.
- Migration can only be performed between hosts of the same type. Furthermore, each host must be

running the same versions of software.

- Hot migration is not supported in DPDK avs scenarios.
- VMs using local storage solutions do not support hot migration. This is because after a VM is migrated to another host, it can no longer access the previous local storage space.
- VMs that use GPU, FPGA, or other (passthrough or SR-IOV) devices do not support hot migration.

> **Note** VMs created in Apsara Stack versions earlier than V3.3 do not support hot migration. Hot migration becomes available after you restart the VMs.

9.1.5.3. Complete hot migration on AG

In Apsara Stack Operations, you can start and cancel hot migration operations as needed through the command line interface.

Trigger hot migration

After hot migration is triggered, you can run the `go2which` command or use ECS Operations and Maintenance Platform to check that the VM enters the migrating state. When hot migration is completed, the VM restores the running state.

The `go2which` command output is as follows:

```
go2hyapi live_migrate_vm == Functions usage: == |- live_migrate_vm <vm_name> [nc_id] [rate] [no_check_image] [no_check_load] [downtime]== Usage: == houyi_api.sh <function_name> [--help|-h] [name=value]
```

Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
nc_id	Designates the destination NC to migrate the VM to.	If the NC does not support the specifications of the VM, the migration will fail.	N/A
rate	The amount of host bandwidth to be allocated for migration tasks.	The migration will use the bandwidth resources of the hosts.	• 10 GB network: 80 MB • 1 GB network: 40 MB
downtime	The maximum allowable downtime caused by migration. The default value is 300 ms.	The service downtime caused by migration is affected.	200 ms to 2,000 ms

Parameter	Function	Impact	Value
no_check_image	Forcibly migrates the images that are not supported.	Performing this operation may violate the SLA.	false
no_check_load	Forcibly migrates images even when the load threshold requirements are not met.	Downtime cannot be controlled when this parameter is set to false.	false

Cancel hot migration

Run the following command to cancel a hot migration task:

```
go2hyapi cancel_live_migrate_vm == Usage: == houyi_api.sh <function_name> [--help|-h] [name=value] ==
Functions usage: == |- cancel_live_migrate_vm <region_id> <vm_name>
```

Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
region_id	The ID of the region where the target VM is located.	N/A	N/A

9.1.5.4. Modify the position of the NC where the VM is located

When an exception occurs during hot migration and the migration cannot be rolled back through ECS Operations and Maintenance Platform, you can modify the VM state to trigger rollback.

Trigger rollback

If an exception occurs during hot migration, run the following command to trigger rollback:

```
go2hyapi call_api manually_change_migration_status == Functions usage: == |- call_api manually_change_
migration_status <vm_name> <region_id> <where>
```

Parameter description

Parameter	Function	Impact	Value

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
region_id	The ID of the region where the target VM is located.	N/A	N/A
where	The ID of the NC where the VM is located.	N/A	N/A

9.1.5.5. FAQ

This topic lists common problems that you may encounter during hot migration and how to resolve them.

- **Which parameters are required to call the Server Controller API to perform a hot migration?**
  - Vm_name: VM name
  - nc_id
- **What preparations should I make before performing a hot migration operation?**
  - Confirm that the VM is in the running state.
  - Confirm the destination of the VM migration.
- **Can hot migration be canceled? How can I cancel hot migration?**

  Yes. If the API request is successful and the migration has not completed, run the `go2hyapi cancel_live_migrate_vm vm_name=[vm_name] region_id=[region_id]` command to cancel the hot migration. If the VM has completed its migration to the destination NC, it is too late to cancel the hot migration.

  You can get the value of region_id by running the `go2which [vm_name]` command to view region_info.
- **The VM is still in the migrating state after the hot migration has completed, and the cancel_live_migrate_vm command is not working. What should I do?**

  You can run the `virsh query-migrate [domid]` command on the source NC of the VM to check whether the VM is still being migrated. If the VM is still being migrated, a piece of JSON information will be returned. If the VM has finished migration, run the following command on the AG to modify the state of the VM:

  ```
  go2hyapi manually_change_migration_status vm_name=[vm_name] where=[nc_id for the VM] region_id=[region_id]
  ```

  domid is the name of the VM instance. You can run the `virsh list|grep vm_name` command to view it.
- **How can I confirm whether the VM is migrated successfully?**

  On the destination NC of the VM, run the `sudo virsh list|grep [vm_name]` command. If the VM instance exists and is not in the running state, the migration is successful.

- **When an exception occurs during hot migration, which logs should I refer to?**
  - View the Libvirt bottom layer migration log on the NC.

    Run the `/var/log/libvirt/libvirt.log` command to view information about the migration process, such as vport offline, detach, delete, and relay route.
  - Run the following command to view the API management log of Server Controller on the AG:

    ```
    /var/log/houyi/pync/houyipync.log
    ```
  - View the Qemu log.
  - Run the following command to view the regionmaster log on the VM:

    ```
    regionmaster/logs/regionmaster/error.log
    ```
- **A VM fails to start after hot migration. Is the VM still in the pending state?**

  If **error vport update nc conf by vpc master fails dest_nc_id:xxx** is returned, it indicates that a VPC fault has occurred and the underlying task is interrupted.
- **During hot migration, the API returns the following error message: distributed lock fail. What are the possible causes of this issue?**

  The API has been called too many times within a short period of time. Wait several minutes and then try again.
- **What are some common scenarios where migration fails? How can I resolve these issues?**

  ### Hot migration issues

Scenario	Cause	Solution
The load is too high and the VM migration does not pass the pressure inspection.	Long service interruption.	You can run no_check_load=true to skip this inspection.
The VM fails to pass image inspection.	It is not an Alibaba Cloud-specified image.	You can run no_check_image=true to skip this inspection. Be aware of the risks involved.

9.1.6. Hot migration of disks

9.1.6.1. Overview

Hot migration seeks to facilitate operations and maintenance of online clusters and improve service operation. Hot migration provides online migration capabilities for virtual disks. This function can also quickly copy data to new locations, enhancing the flexibility of services.

9.1.6.2. Limits

Before performing hot migration on a disk, you need to understand the limits.

Limits

- Only disks of the river type support hot migration.
- The source and destination clusters for hot migration must belong to the same OSS domain.
- Disk sharing is not supported.
- Hot migration is not supported on disks whose capacity is greater than 2 TB.
- Format and capacity changes are not supported.
- Hot migration is only supported within the same zone.
- Due to how hot migration is implemented internally, the names of the source and destination clusters must be less than 15 bytes in length.

> Note
>
> - The data of the original source disk will remain on the disk after hot migration has completed. You can use the pu tool to delete the remaining data. Job recycling is unavailable.
> - During migration, an I/O latency of less than 1 second is considered normal.
> - Migration cannot be rolled back.
> - Migration will consume network bandwidth, so you must take measures to limit concurrent traffic during migration.

Migration operation

For more information about the APIs related to disk hot migration, see "**Disk hot migration**" in *ECS Developer Guide* .

9.1.6.3. O&M after hot migration

The original source disk data remains on the source disk after hot migration and data backup operations are completed. To release disk space, delete the data from the source disk. After the data is deleted from the source disk, the space will be released at a later time.

Procedure

1. On the compute cluster AG, run the `go2houyiregiondbrnd -e 'select task_id from device_migrate_log where status="complete"'` command to obtain *task: allTaskIds*.
2. On the compute cluster AG, run the `go2riverdbrnd -e 'select task_id,src_pangu_path,dst_pangu_path from migration_log where task_id in ($allTaskIds) and status=2 and src_recycled=0 and DATE(gmt_finish) < DATE_ADD(CURDATE(), INTERVAL -1 DAY)'` command.
3. Perform the following operations for each set of <task_id,src_pangu_path,dst_pangu_path>:
   i. Run the `/apsara/deploy/bsutil rlm --dir=$dst_pangu_path|grep 'not-loaded'|wc -l` command on the host that runs the bstools role in the storage cluster. If the command output is not 0, proceed to the next step.
   ii. Run the `/apsara/deploy/bsutil delete-image --dir=$src_pangu_path` command on the host that runs the bstools role in the storage cluster.
   iii. Run the `/apsara/river/river_admin migrate recycle $task_id` command on the host that runs the river role in the storage cluster.

9.1.7. Upgrade solution

9.1.7.1. Overview

For both hot and cold migration of GPU and FPGA clusters, you must understand the limitations that apply to cluster upgrades.

9.1.7.2. Limits on GPU clusters

Before upgrading a GPU cluster, you must understand the limits.

The upgrade of GPU clusters in Apsara Stack are subject to the following limits:

- GPU clusters are only supported in Apsara Stack 3.3 or later versions.
- To upgrade a GPU cluster, you must restart the NC server.
- VMs that use GPU, FPGA, or other passthrough or SR-IOV devices do not support hot migration.
- The GN5I, GN5E, and GN4 type GPU clusters do not have the specifications of local disk instances and only support offline cold migration.
- When you perform a forced cold migration on GN5 and GA1 type GPU clusters that have specifications of local disk instances, the local disk will be reformatted, resulting in data loss. These disks must be backed up before they can be migrated.

9.1.7.3. Limits on FPGA clusters

Before upgrading an FPGA cluster, you must understand the limits.

The upgrade of FPGA clusters in Apsara Stack are subject to the following limits:

- FPGA clusters are only supported in Apsara Stack 3.5 or later versions.
- VMs in an FPGA cluster must be shut down before the cluster can be upgraded.
- The FPGA service relies on Redis to a great extent. If the Redis service is interrupted during the hot upgrade of Apsara Stack, the FPGA service will be interrupted. The FPGA service will recover after the Redis service is restored. However, if a Redis instance fails to be created, you must restart the FPGA service after the Redis service is restored.

9.1.8. Handle routine alarms

9.1.8.1. Overview

This topic describes the definition of each key metric and how to handle alerts.

The metrics monitored in ECS can be categorized into three types:

- Basic metrics: These metrics are used to monitor the CPU, memory, and correlated service processes of hosts.
- Connectivity metrics: These metrics are used to monitor the connectivity between different components and the connectivity between different networks.
- Service metrics: These metrics are used for service monitoring, such as the state of various types of API requests.

Description of metric types

Metric type	Function	Solution
Basic metric/service availability metric	Monitors the basic performance of the host and the availability of the services on the host. This kind of metrics includes CPU, memory, and handle count.	When CPU utilization is too high: identify which process consumes a large amount of CPU resources. If it is a key process, evaluate whether it can be restarted.
		When the memory usage is too high (for key services): dump the memory data, request the back-end R&D team to analyze the data, and restart the application.
Connectivity metric	Checks the connectivity between each module and its related modules.	• First, check the health status of the corresponding modules. For example, check whether the host works normally and whether services, ports, and domain names are normal. • If two modules that are connected to each other are healthy, check the network connectivity between them.
Service metric	Monitors aspects of key request calls such as the latency, total number, failures of API requests, and database SQL exceptions.	• In case of an API request failure, you must view the corresponding logs to identify the cause of the failure. • In case of a database SQL exception, check whether the exception was caused by a database exception (system breakdown or high connection count) or a problem with the application. If it is an application problem, forward the error information to the back-end R&D team for troubleshooting.

9.1.8.2. API proxy

This topic describes the metrics of API proxy.

Metric description

Metric	Alert item	Description
check_apiproxy_dns	Database HA switchover occurs or not	Checks whether Server Controller database switchover occurs. If so, nginx will be reloaded automatically.
check_apiproxy_conn_new	check_apiproxy_conn_new	Checks the connectivity to the Server Controller database.
		Checks the connectivity to the API Server: • Checks whether the API Server is down. • Checks the network connectivity.

Metric	Alert item	Description
check_apiproxy_proc_new	check_apiproxy_proc_new	Checks the memory usage and CPU utilization for nginx and memcache processes.

9.1.8.3. API Server

The topic describes the metrics of the API Server.

Metric description

Metric	Alert Item	Solution
check_API Server_proc_new	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage
check_API Server_conn_new	Checks the connectivity between the API Server and Server Controller database.	Checks whether the corresponding component is down. If the corresponding component is down, fix the issue by taking necessary O&M measures. If the database is down, contact DBA to fix the issue. Checks whether the VIP is connected to the corresponding component. If not, contact the network engineer to fix it.
	Checks the connectivity between the API Server and TAIR.	
	Checks the connectivity between the API Server and RegionMaster.	
	Checks the connectivity between the API Server and the RMS.	
check_API Server_perf	Monitors metrics for API requests, such as the latency, total number of API requests, and number of failed API requests.	It is primarily used to identify faults.
check_API Server_errorlog	Checks database exceptions and instance creation failures.	• If an exception occurs to the database, contact DBA to check whether the database is normal. • If the creation of an instance fails, locate the cause of the failure.

9.1.8.4. RegionMaster

This topic describes the metrics of RegionMaster.

Metric description

Metric	Alert item	Description
check_regionmaster_proc	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage.
check_regionmaster_work	rms_connectivity	Checks the connectivity to RMS.
	regiondb_connectivity	Checks the connectivity to the houyiregiondb database.
	houyi_connectivity	Checks the connectivity to the Server Controller database.
	tair_connectivity	Checks the connectivity to TAIR.
check_zookeeper_work	status	Checks the operating state of the Zookeeper process on the Server Controller.
check_regionmaster_errorlog	errorlog_for_db	Checks whether the SQL statements are properly executed.
	check_regionmaster_errorlog	
check_workflow_master	Checks the operating state of the master in the workflow process.	-
check_workflow_worker	Checks the operating state of the worker in the workflow process.	-

9.1.8.5. RMS

This topic describes the metrics of RMS.

Metric description

Metric	Alert item	Description
check_rms_proc	Checks the process status, CPU utilization, and memory usage of RMS.	-
check_rabbitmq_proc	Checks the process status, CPU utilization, and memory usage of the rabbitmq cluster.	-
check_rabbitmq_status	Checks the number of queues, exchanges, and bindings in the rabbitmq cluster.	Follow the maintenance guide for the rabbitmq cluster.

Metric	Alert item	Description
check_rabbitmq_queues	Checks whether messages are accumulated.	If messages are accumulated, it will also check for the cause.
	Check whether there are consumers.	If there are no consumers, check whether Regionmaster and APIserver are operating normally. If they are operating normally, check whether there is a problem with the rabbitmq cluster.

9.1.8.6. PYNC

This topic describes the metrics that are monitored for PYNC.

Metric description

Metric	Alert item	Description
check_vm_start_failed	Checks the causes of a VM startup fault.	You do not need to handle it immediately. It is typically caused by custom images.
check_pync	Checks the CPU utilization and memory usage of PYNC.	-
	PYNC has too many open file handles.	-
	PYNC process count.	PYNC must have four processes.
	It has been long since pyncVmMonitor.LOG was last updated at ${pync_monitor_log_last_updated}.	Checks for reasons why a log has not updated for a long period of time, such as: • Whether a PYNC process has encountered a problem. • Whether the NC is running a key process called Uninterruptible Sleep.

9.1.8.7. Zookeeper

This topic describes the metrics of Zookeeper.

Metric description

Metric	Alert item	Description
		The process does not exist.

cneck_zookeeper_proc Metric	proc Alert item	Description
		The memory usage or CPU utilization is too high.

9.1.8.8. AG

This topic describes the metrics of AGs.

Metric description

Metric	Alert item	Description
disk_usage	apsara_90	*/apsara* disk usage.
	homeadmin_90	Usage of */home/admin*.
check_system_ag	mem_85	Memory usage.
	cpu_98	CPU utilization.
	df_98	Disk usage of the root directory.
check_ag_disk_usage	check_ag_disk_usage	Disk usage.
check_nc_down_new	check_recover_failed	Checks the causes of a VM migration fault. Possible causes include: • No resources are available in the cluster. • A VM does not belong to any cluster.
	check_repeat_recovered	Continuous VM migration.
	check_continuous_nc_down	Checks continuous NC downtime.
	check_nc_down_with_vm	The state of the NC in the database is nc_down, but there are still VMs operating normally on the NC. Checks the NC for hardware faults: • If a hardware fault occurs, you must perform operations and maintenance to resolve the fault. • If no hardware fault is detected, restore the NC and change its state to locked.
check_ag_fhtd_new	Checks whether the FHT downtime migration tool, mostly used by local disks, is operating normally.	If the tool does not exist, download the FHT downtime migration tool.

9.1.8.9. Server groups

This topic describes the metrics that are monitored for server groups.

Metric description

Metric	Alert item	Description
check_pync	pync_mem	Monitors the memory usage of PYNC.
	pync_cpu	Monitors the CPU utilization of PYNC.
	pync_nofile	Monitors the number of PYNC handles.
	pync_nproc	Monitors the number of PYNC processes.
	pync_monitor_log_not_updated	Monitors the status of PYNC scheduled tasks.

9.1.9. Inspection

9.1.9.1. Overview

ECS inspection includes cluster basic health inspection and cluster resources inspection.

9.1.9.2. Cluster basic health inspection

9.1.9.2.1. Overview

Cluster basic health inspection includes monitoring inspection, inspection of basic software package versions, and basic public resources inspection.

9.1.9.2.2. Monitoring inspection

This topic describes basic monitoring inspections and connectivity monitoring inspections.

9.1.9.2.3. Inspection of basic software package versions

This topic describes the version inspections of Server Controller components, Apsara system, virtualization packages, and basic service packages.

9.1.9.2.4. Basic public resources inspection

This topic describes ISO inspections and basic image inspections.

ISO inspection

ECS Operations and Maintenance System provides two basic ISO files for each region:

- linux-virt-release-xxxx.iso
- windows-virt-release-xxxx.iso

You can run the following command to search the database for relevant information:

```
$ houyiregiondb
mysql>select name,os_type,version,path,oss_info from iso_resource where os_type! ='\G
```

Parameters in the command are as follows:

- *name*: the name of the ISO file, such as xxxx.iso.
- *os_type*: the operating system (OS) type of an image.
- *path*: the path on the Apsara Distributed File System cloud disk where the ISO file is stored. You can run the `/apsara/deploy/pu meta $path` command to check whether the ISO exists in the files of Apsara Distributed File System.
- *oss_info*: the path on the local OSS disk where the ISO file is stored. To search for this path, you must provide relevant information to OSS support engineers for inspection.

Basic image inspection

- Run the following command to check the state of a basic image in the database:

  ```
  houyiregiondb
  mysql>select image_no,status,visibility,platform,
  region_no from image;
  ```

- Check whether the basic image is usable. You can call the create_instance API to use relevant images to create a VM and manually check whether the VM can operate normally.

9.1.9.3. Cluster resource inspection

9.1.9.3.1. Overview

Cluster resource inspection includes cluster inventory inspection and VM inspection.

9.1.9.3.2. Cluster inventory inspection

This topic describes the inspections of cluster inventory resources. Cluster inventory resources are specified by the number of VMs that can be created by using the remaining resources in the cluster. You can use the database to obtain the cluster inventory resources.

Suppose you need to inspect the inventory resources of a cluster based on 16-core 64 GB VMs. Run the following command to obtain the inventory resources of the cluster:

```
$ houyiregiondb
mysql> select sum( least ( floor(available_cpu/16),floor(available_memory/64/1024))) from nc_resource,nc w
here nc.cluster_id=$id and nc.biz_status='free' and nc.id=nc_resource.id;
```

If the current cluster contains a relatively large VM, ensure that the cluster has enough free resources to handle the VM, as well as an available host with sufficient resources for backup. This host will be the migration destination of the large VM in case the current host goes down. Otherwise, the large VM cannot be migrated when its host goes down, and you will have to either use hot migration to transfer resources or release redundant VMs in the cluster.

NC state inspection

NC state inspection mainly checks whether the state of a host is normal in the database and Apsara Infrastructure Management Framework.

- A host can be in one of the following states in Apsara Infrastructure Management Framework:
  - Good: indicates that the host is in a normal working state.
  - Error: indicates that the host has an active monitoring alert.
  - Probation: indicates that the host is in the probationary period and may fail.
  - OS _error: indicates that the host has failed and is being cloned.
  - Hw_error: indicates that the hardware of a host has failed and is being repaired.
  - OS _probation: indicates the host is recovering from a fault or hardware failure and is in a probationary period. If the host recovers within the probationary period, the state will change to probation. If the host fails to recover within the probationary period (an error is reported), the state will change to OS _error.

  > **Note** The Good state is considered to be the stable state, and all other states are considered to be unstable states.

- Cluster definitions for Apsara Infrastructure Management Framework:
  - Default cluster: the cluster where NCs are placed when they go offline.
  - Non-default cluster: the cluster for online NCs.

  An NC that is operating normally is placed in a non-default cluster, and is in the Good state.

The mappings of host states between the ECS database and Apsara Infrastructure Management Framework are described in Mappings of host states between the ECS database and Apsara Infrastructure Management Framework.

Mappings of host states between the ECS database and Apsara Infrastructure Management Framework

Host states in ECS database	Cluster	Host state	Scenario
mlock	Non-default cluster	Unstable	A host that goes online is immediately and proactively locked.
locked	Non-default cluster	Unstable	An NC needs to be unlocked.
free	Non-default cluster	Stable	A host operates normally.

Host states in ECS database	Cluster	Host state	Scenario
nc_down	Non-default cluster	Unstable	A host operates normally or is in downtime.
offline	Default cluster	Unstable	A host goes offline from business attributes.

9.1.9.3.3. VM inspection

This topic describes pending VM inspections, VM state inspections, and VM resource inspections.

Pending VM inspection

This type of inspection focuses on VMs that have been in the pending state for a long period of time. When a VM has been in the pending state for a long period of time, it is considered a redundant resource. Contact the user to handle it.

VM state inspection

This type of inspection focuses on the VM state consistency. For example, a VM is displayed as stopped in the database, but is displayed as running in NC. During the inspection, the VM states recorded in the database and on the host are checked. If the VM states are inconsistent, corresponding operations are performed.

- Run the following command to obtain the VM state in a database:

  ```
  houyiregiondb -Ne "select status from vm where name='$name'"
  ```

- Run the following command to obtain the VM state on a host:

  ```
  sudo virsh list | grep $name
  ```

VM resource inspection

After the configuration of a VM is changed, the system checks whether the configuration of the VM recorded in the database is consistent with that used on the host.

- Run the following command to obtain the VM configuration in a database:

  ```
  houyiregiondb -Ne "select vcpu, memory from vm where name='$name'"
  ```

- Run the following command to obtain the VM configuration on a host:

  ```
  sudo virsh list | grep $name
  ```

  Obtain information about CPU and memory by viewing the corresponding fields.

9.2. Container Service for Kubernetes

9.2.1. Components and features

9.2.1.1. Console

The Container Service for Kubernetes console provides a graphical user interface (GUI) that serves as an entry for all operations on Container Service for Kubernetes. The console adopts the deployment mode that applies to standard Java applications on Alibaba Cloud. Each console instance contains a Tengine server and a Jetty container.

Log on to the console container

1. Log on to the Apsara Stack Operations console. In the left-side navigation pane, choose **Products > Product List** to go to the Product List page. In the Apsara Stack O&M pane, click **Apsara Infrastructure Management Framework**.
2. The **Infrastructure Operation Platform** console appears. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, search for a target Container Service for Kubernetes cluster.

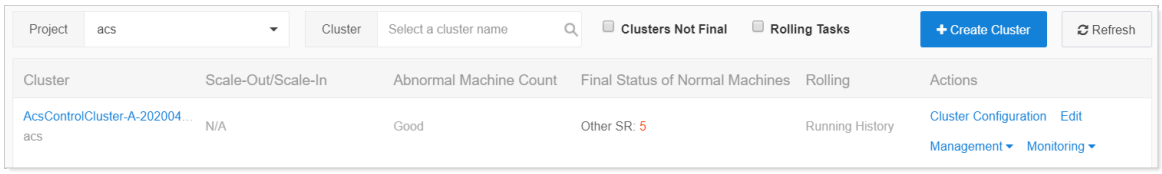

5. Click **Cluster Configuration** in the **Actions** column for the cluster. On the Cluster Configuration page, find the CosConsoleAliyunCom server role in the **Server Role** section and check the hosts for the server role.
6. In the middle of the left-side navigation pane, enter the hostname in the Machine search box. Move the pointer over the More icon next to the hostname and select Terminal from the shortcut menu. This allows you to log on to the host by using a terminal session. On the command line, enter `docker ps` to obtain the ID of the cos-console-aliyun-com container.
7. On the command line, enter `sudo docker exec -it container_id bin/bash` to access the container.
8. Go to the specified directory to find Tengine and Jetty.

O&M commands

- Restart Tengine: `/etc/rc.d/init.d/tengine restart`
- Restart Jetty: `/etc/init.d/jetty restart`

Directories

- Root directory of web applications: `/alidata/www/`
- WAR directory of applications: */alidata/www/wwwroot/cos-console-aliyun-com*

Application log files

- The root directory that stores log files: */alidata/www/logs*
- The path to Jetty: */alidata/www/logs/*jetty

- The path to application log files: */alidata/www/logs/java/cos-console-aliyun-com/applog*

9.2.1.2. Troopers

This topic describes the features and usage of Troopers.

The Troopers daemon is used to create clusters and hosts. You can also use Troopers to manage the clusters and hosts in Container Service for Kubernetes.

Troopers is programmed in Go. Each container runs only the Troopers daemon and does not use any other daemons.

To use Troopers, perform the following steps:

1. Log on to the Apsara Stack Operations console. In the left-side navigation pane, choose **Products > Product List** to go to the Product List page. In the Apsara Stack O&M pane, click **Apsara Infrastructure Management Framework**.
2. The **Infrastructure Operation Platform** console appears. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, search for the target Container Service for Kubernetes cluster. Click **Cluster Configuration** in the Actions column for the cluster. On the Cluster Configuration page, find the Troopers server role in the Server Role section and check the hosts for the server role.
5. In the middle of the left-side navigation pane, enter the hostname in the Machine search box. Move the pointer over the More icon next to the hostname and choose Terminal from the shortcut menu. This allows you to log on to the host by using a terminal session. On the command line, enter `docker ps` to obtain the ID of the Troopers container.
6. On the command line, enter `sudo docker exec -it container_id bin/bash` to access the container.

The following list describes the structures of specific directories of the container:

- */usr/aliyun/acs/troopers*: the root directory of the application.
  - troopers: the main program of Troopers.
  - troopers.json: the configuration file of Troopers.
  - troopers.ym: the configurations of certificate encryption.
  - start.sh: the entry script used to start Troopers. If the Troopers daemon already exists, do not run the *start.sh* script.
- */opt/aliyun/install/check_health.sh*: the script that is used to run health checks.
- */usr/aliyun/acs/certs/control*: the directory that stores a certificate. Troopers uses the certificate to access the Region Controller (RC). You can use OpenSSL to verify the certificate.

Troopers log files are exported to the stdout stream. No log files are stored in the container. To view log records, run the `docker logs` command outside the container.

9.2.1.3. Mirana

This topic describes the deployment modes and features of Mirana.

Log on to the console container

1. Log on to the Apsara Stack Operations console. In the left-side navigation pane, choose **Products > Product List** to go to the Product List page. In the Apsara Stack O&M pane, click **Apsara Infrastructure Management Framework**.
2. The **Infrastructure Operation Platform** console appears. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, search for the target Container Service for Kubernetes cluster. Click **Cluster Configuration** in the Actions column for the cluster. On the Cluster Configuration page, find the Mirana server role in the Server Role section and check the hosts for the server role.

Log query

In the middle of the left-side navigation pane, enter the hostname in the Machine search box. Move the pointer over the More icon next to the hostname and choose Terminal from the shortcut menu. This allows you to log on to the host by using a terminal session. On the command line, enter `docker ps` to obtain the ID of the Mirana container. Enter `docker logs contianer_id` to view the log data.

The Mirana container is stateless. You can try to restart the container if the service is unavailable. On the command line, enter `docker restart ${container_id}` to restart the container.

Deployment mode

- A Mirana container is deployed in each cluster. The deployment mode of the Mirana container is similar to that of the Commander container.
- Mirana containers are deployed on control hosts and use HTTPS to provide services. Mirana requires the Kubernetes API certificate that is provided by Troopers.

Features

- Provides the Kompose tool to convert the Compose file into a YMAL deployment file.
- Uses the Helm client to manage orchestration templates.
- Supports the blue-green deployment of APIs.
- Serves as the proxy for API operations of Kubernetes clusters.

9.2.2. System restart

9.2.2.1. Restart a control node

A container control node runs a Docker container where the services such as CosConsoleAliyunCom, Troopers, and etcd are deployed. To restart a control node, perform the following steps:

Procedure

1. Log on to the Apsara Stack Operations console. In the left-side navigation pane, choose **Products > Product List** to go to the Product List page. In the Apsara Stack O&M pane, click **Apsara Infrastructure Management Framework**.
2. The **Infrastructure Operation Platform** console appears. In the left-side navigation pane, click **Reports**.

3. On the **Reports** page, click **Go**.
4. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, search for the target Container Service for Kubernetes cluster. Click **Cluster Configuration** in the Actions column for the cluster. On the Cluster Configuration page, find the target server role in the Server Role section and find the host where the control node is deployed.
5. On the command line, enter `docker ps|grep [app]` to obtain the container ID.

   `[app]` specifies the name of the application that is deployed in the container. You can obtain the container ID based on the application name.
6. On the command line, enter `docker restart container_id` to restart the container.

9.3. Auto Scaling (ESS)

9.3.1. Log on to the Apsara Stack Operations console

This topic describes how to log on to the Apsara Stack Operations (ASO) console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

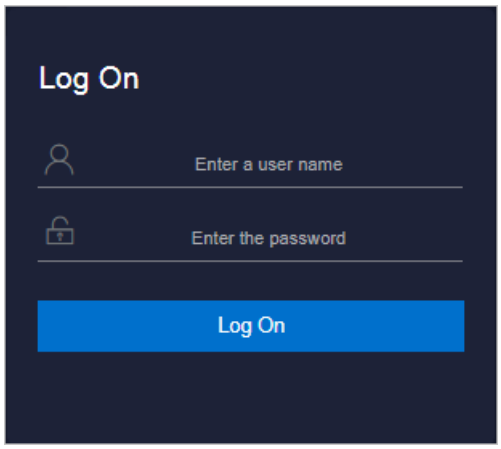

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.

9.3.2. Product resources and services

9.3.2.1. Application deployment

All the applications in the ESS Business Foundation System are stateless. You must restart the applications by running the docker restart command.

- **ess-init**

  It first initializes the database service, and then pushes all API configuration files of ESS to the pop configuration center to initialize OpenAPI Gateway.

- **Trigger (dependent on ess-init)**
  - Trigger executes tasks such as checking health status, checking the maximum and minimum instance numbers, and deleting scaling groups.
  - Triggers scheduled tasks and monitoring tasks.

- **Coordinator**

  Coordinator is the open API layer that provides public-facing services. It maintains persistent requests and issues tasks.

- **Worker**
  - Worker executes all scaling-related tasks, such as creating ECS instances, adding instances to SLB backend server groups and RDS whitelists, and synchronizing CloudMonitor group information.
  - It retries failed tasks and provides the rollback mechanism.

- **service_test**

  It is used for regression tests on the overall application running status. It contains over 60 regression test cases to test the integrity of functions.

9.3.2.2. Troubleshooting

This topic describes how to troubleshoot issues related to product resources and services.

Prerequisites

When issues related to Business Foundation System occur, you can submit tickets on the AliCloud Business Center Platform and check the status of related services in the Apsara Infrastructure Management Framework console.

Procedure

1. Submit a ticket.
2. Check the status of services that depend on Business Foundation System in the Apsara Infrastructure Management Framework console.

   If a service cannot be executed, it affects the running of ESS Business Foundation System. For more information, see Unavailable services and their impacts.

Unavailable services and their impacts

Service	Key impact
middleWare.dubbo	Deployment is affected, and the service is unavailable.
middleWare.tair	Deployment is affected, and the service is unavailable.
middleWare.metaq (message midddleware)	Deployment is affected.
middleWare.zookeeper	Deployment is affected, and the service is unavailable.
middleWare.jmenvDiamondVips	Deployment is affected, and the Diamond configuration item cannot be obtained.
ram.ramService (RAM)	The RAM user is unavailable.
webapp.pop (API Gateway)	The OpenAPI service is unavailable.
ecs.yaochi (ECS Business Foundation System)	All ECS creation requests become invalid.
slb.yaochi (SLB Business Foundation System)	All SLB association requests become invalid.
rds.yaochi (RDS Business Foundation System)	All ApsaraDB for RDS association requests become invalid.
tianjimon (Monitoring System)	Some services are unavailable.

9.3.3. Inspection

9.3.3.1. Overview

ESS inspection monitors the basic health conditions of clusters.

The inspected basic health conditions include the following aspects:

- Monitoring inspection
- Basic software package version inspection

9.3.3.2. Monitoring inspection

The monitoring inspection includes the basic monitoring and connectivity monitoring inspection.

9.3.3.3. Basic software package version inspection

The basic software package version inspection includes the version inspection for trigger, coordinator, worker, and base services.

9.4. Resource Orchestration Service (ROS)

9.4.1. ROS component O&M

9.4.1.1. API Server

The API Server is used to receive ROS requests, send requests to RabbitMQ clusters, and send the responses returned by the Engine Server to callers. The API Server is used to connect the frontend and backend services.

- Components

  The Engine Server and API Server share three servers, all of which are attached to a special Server Load Balancer (SLB) instance.

- O&M methods
  - The storage path of the API Server information is /home/admin/ros-server/bin/.
  - Basic operations of the API Server: `#/usr/local/ros-python/bin/python/home/admin/ros-service/bin/ros-api{stop|status|--daemon}`
    - `stop` : stops the API Server.
    - `status` : queries the status of the API Server.
    - `--daemon` : starts the API Server in daemon mode.

- Health criteria
  - Intrinsic availability: The CPU usage and system memory are within the normal range. The API Server is running normally.
  - Associated component availability: ROS is available.

9.4.1.2. Engine Server

The Engine Server is used to process stack requests. It shares the three servers with the API Server.

- O&M methods
  - The storage path of the API Server information is /home/admin/ros-server/bin/.

- Basic operation of the Engine Server: `/usr/local/ros-python/bin/python /home/admin/ros-service/bin/ros-engine {stop|status|--daemon}`
  - `stop`: stops the Engine Server.
  - `status`: queries the status of the Engine Server.
  - `--daemon`: starts the Engine Server in daemon mode.

- Health criteria
  - Intrinsic availability: The CPU usage and system memory are within the normal range. The Engine Server is running normally.
  - Associated component availability: ROS is available.

9.4.1.3. RabbitMQ clusters

RabbitMQ clusters are used to receive requests from the API Sever and responses from the Engine Server.

- Components

  RabbitMQ clusters are composed of nodes.

  RabbitMQ clusters are used for messaging. Nodes in the clusters use disks for non-persistent storage. Messages are written into the queues that correspond to the nodes. Nodes in a cluster can communicate with each other. Typically, to ensure data accuracy, the minimum number of working nodes is set to [Total number of nodes/2] rounded up. If data of nodes are inconsistent, the secondary nodes synchronize queue messages from the primary nodes.

- O&M methods

  The storage path of the RabbitMQ information is /opt/rabbitmq-server/.

  Common RabbitMQ commands are as follows:

  - You can run the following command to query the cluster status: `sudo /usr/local/sbin/rabbitmq-server/sbin/rabbitmqctl cluster_status`

```
[root@                    /]
#/usr/local/sbin/rabbitmqctl cluster_status
Cluster status of node ros_rabbit@docker011165194088 ...
[{nodes,[{disc,[ros_rabbit@docker011165194088]},
        {ram,[ros_rabbit@docker011165194091]}]},
 {running_nodes,[ros_rabbit@docker011165194091,ros_rabbit@docker011165194088]},
 {cluster_name,<<"ros_rabbit@docker011165194088">>},
 {partitions,[]}]
```

    - `Nodes`: indicates the nodes in the cluster.
    - `Disc`: indicates that the cluster uses disks for storage.
    - `Mem`: indicates that the cluster uses memory for non-persistent storage.
    - `Running_nodes`: indicates the information of the running nodes in the cluster.
    - `Partition`: indicates the partitions of the cluster. If the value field is brackets [], the cluster has no partitions. If this parameter is not empty, the cluster nodes are divided into several partitions.

- You can run the following command to query the virtual hosts in a cluster: `sudo /usr/local/sbin/rabbitmqctl list_vhosts`

  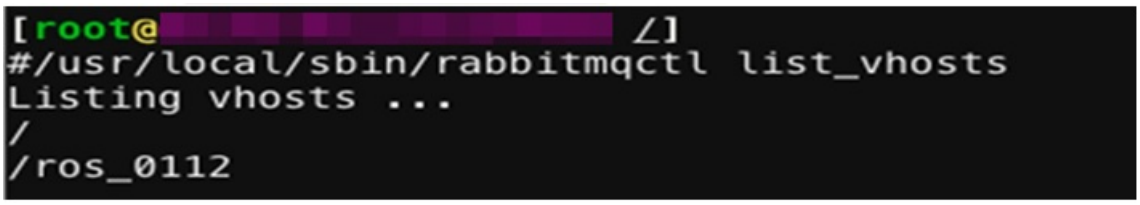

  Typically, there are two virtual hosts. One is displayed as a forward slash (/), and the other is named based on the region where it resides.

- Health criteria
  - Intrinsic availability: The CPU usage and system memory are within the normal range. RabbitMQ is running normally, which indicates that clusters have no partitions, queues are properly processed, and messages are properly consumed.
  - Associated component availability: ROS is available.

9.4.1.4. Notify Server

The Notify Server is the proxy server for ECS instances that reside in a VPC. It sends the execution status and information of operations on ECS instances to ROS.

- Components

  The Notify Server consists of three servers, all of which are attached to a special SLB instance.

- O&M methods

  For example, the virtual IP address of the SLB instance is 10.152.XX.XX. You can run curl `http://10.152.XX.XX:80/health-check` to check whether the Notify Server is running.

- Health criteria
  - Intrinsic availability: The CPU usage and system memory are within the normal range.
  - Associated component availability: ROS is available.

9.5. Object Storage Service (OSS)

9.5.1. Log on to the Apsara Stack Operations console

This topic describes how to log on to the Apsara Stack Operations console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

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

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

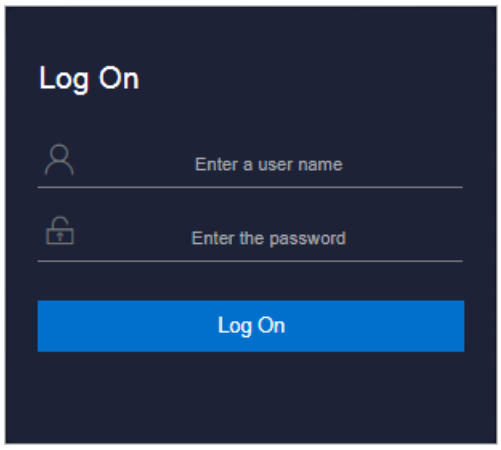

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.

9.5.2. OSS operations and maintenance

9.5.2.1. User data

9.5.2.1.1. Basic bucket information

You can query basic bucket information such as the cluster deployment location, configuration information, current capacity, and object count of a bucket. You can also view this information in a table.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > User Data**.
3. On the **Bucket Basic Information** tab, select the bucket you want to view.
4. Click **View**, as shown in the following figure.

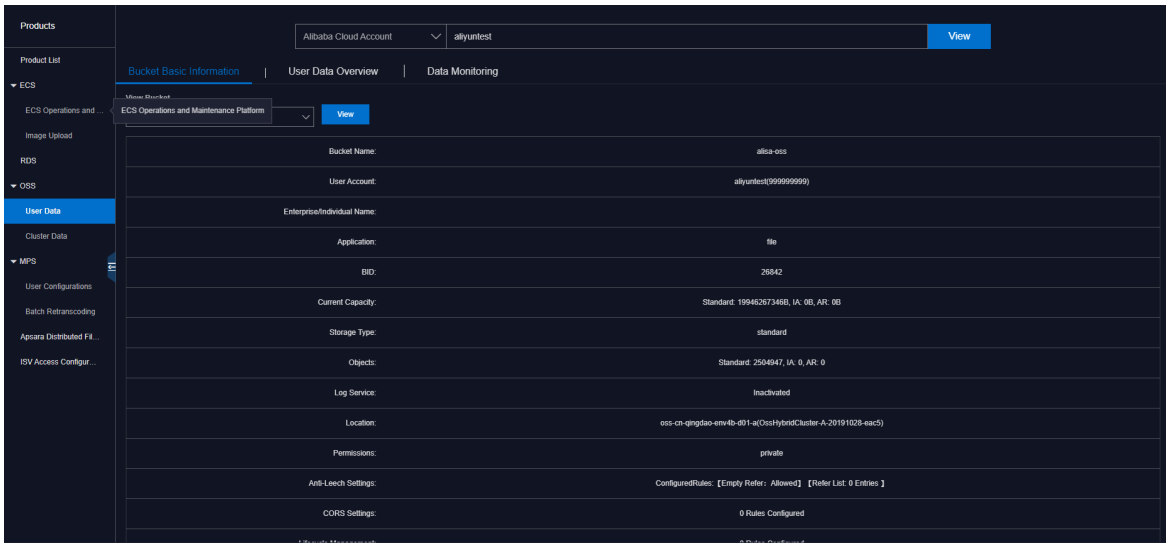

9.5.2.1.2. User data overview

You can query data statistics and trends, including resource usage and basic attributes of resources by **UID**, **Alibaba Cloud Account**, **Bucket Name**, or **Bucket MD5**.

Context

The User Data Overview tab is displayed only when you search by UID or Alibaba Cloud account. On the User Data Overview tab, you can specify a date to view total usage of various resources in all buckets owned by the user account.

You can collect resource statistics by total storage capacity, total inbound or outbound traffic through the public network, internal network, or CDN, or total charged requests.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > User Data**.
3. On the **User Data Overview** tab, you can view resource usage such as total storage capacity, total inbound and outbound traffic, and total charged requests by **Alibaba Cloud Account** or **UID**.
4. Set **Date**. Click OK. Click **View**, as shown in the following figure.

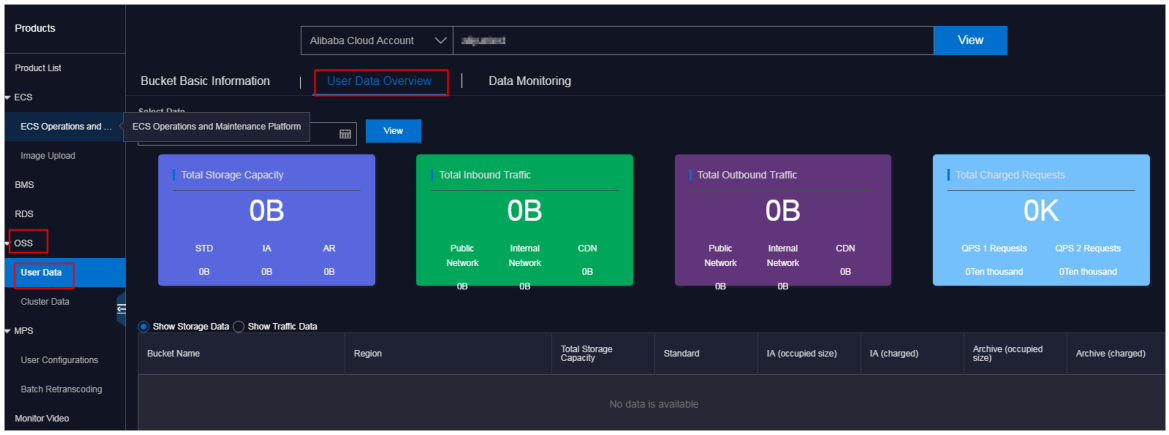

9.5.2.1.3. Data monitoring

This topic describes how to monitor OSS data in the Apsara Stack Operations console.

Context

You can query resource running statuses and usage such as the storage capacity, traffic, SLA, HTTP status, latency, QPS, and image processing capacity by **UID**, **Alibaba Cloud Account**, **Bucket Name**, or **Bucket MD5**. You can also query the resource usage and trends based on a specified time range.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > User Data**.
3. On the **Data Monitoring** tab, set **Bucket Name**, **Specify Time Range**, and **Monitoring Items**.

> **Note** Metric descriptions:
>
> - SLA: indicates the service level availability metric for OSS. Formula: SLA = Non-5xx request count per 10s or hour/Total valid request count × 100%.
> - HTTP Status: collects statistics for the percentages of the numbers of 5xx, 403, 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
> - Latency: collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.
> - Storage Capacity: collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.
> - Image Processing Capacity: collects statistics for the number of processed images.
>
> **Note** By default, this metric is not displayed. You can select this metric from the **Monitoring Items** drop-down list.
>
> - Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
> - QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.

4. Click **View**. The following example describes typical operations on the data monitoring trend chart:
   - If you query data monitoring information by user, you can click the bucket name in the trend chart to show or hide the curve.

     Data monitoring 1

     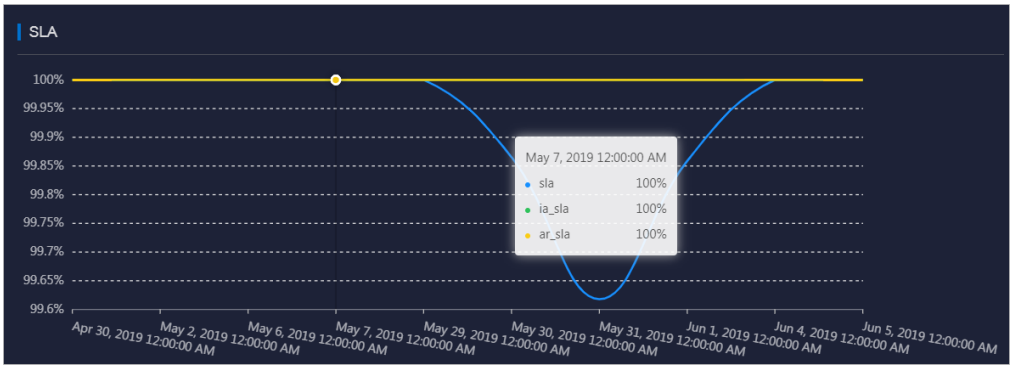

   - Move the pointer over the trend chart to display data at a specific point in time.

     Data monitoring 2

     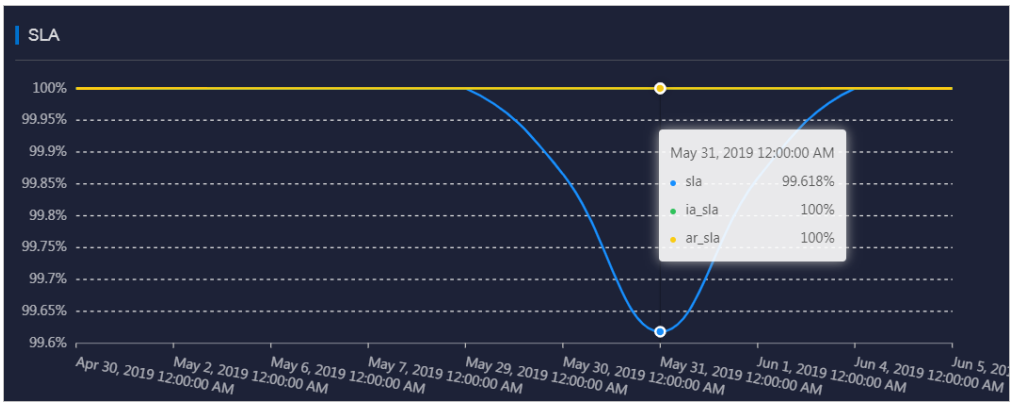

9.5.2.2. Cluster data

9.5.2.2.1. Inventory monitoring

Metrics of inventory monitoring include the total capacity, available capacity, used capacity, backup ratio, and inventory usage.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > Cluster data**.
3. On the **Inventory Monitoring** tab, you can view statistics by Apsara Distributed File System, metric data, or KV data usage.

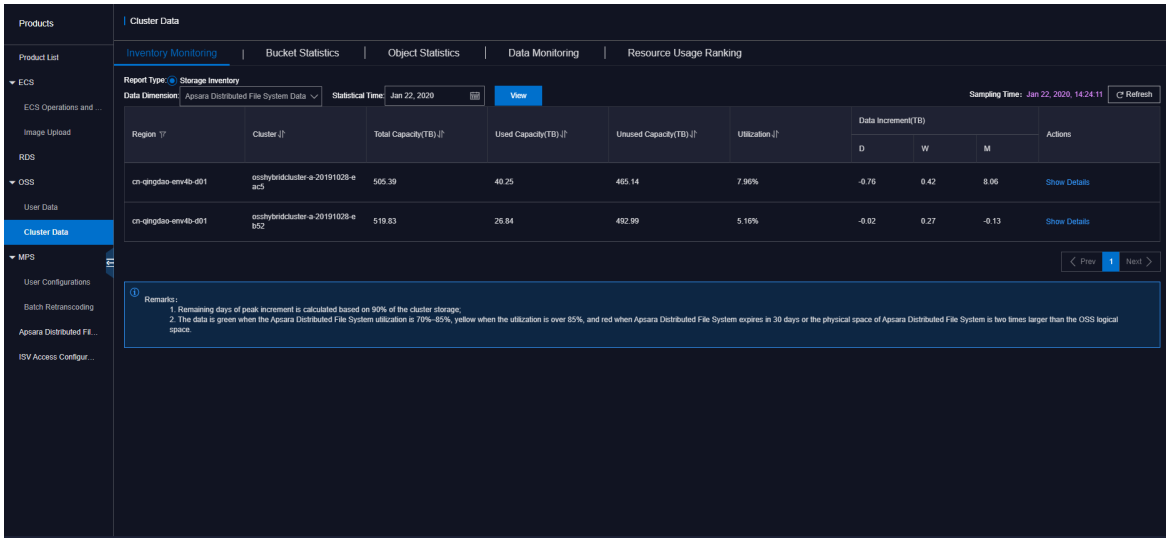

Aside from basic cluster information such as the cluster name and region, you can also view metrics based on the following dimensions:

- Apsara Distributed File System Data: includes the actual total capacity for storage (including the total capacity for multiple data backups), used capacity, remaining capacity (available), usage, and backup ratio.
- Metric Data: includes the bucket storage used by users who use ECS instances and other instances.
- KV Data: includes the logic KV data, KV data in the recycle bin, and data increment (by day, week, or month).

9.5.2.2.2. Bucket statistics

This topic describes how to collect statistics for the number of buckets by cluster.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > Cluster Data**.
3. On the **Bucket Statistics** tab, select **Report**, **Current Overall Statistics**, or **Growth Trend** to view bucket statistics.

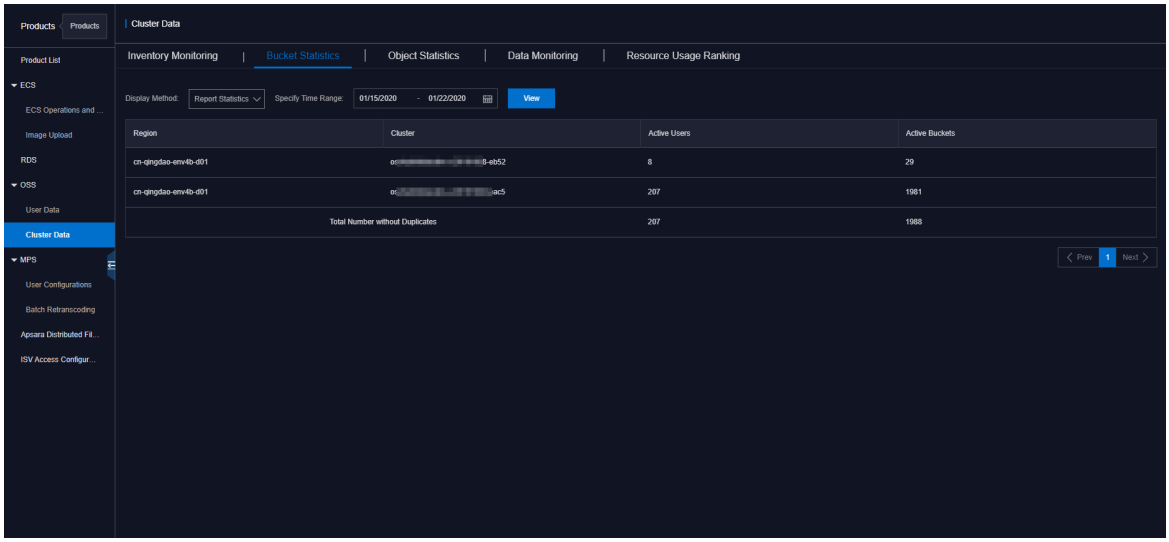

- If you select **Report**, specify the time range.
- You can select **Current Overall Statistics** to query statistics of last hour.
- If you select **Growth Trend**, you can specify a time range of *seven days*, *30 days*, *three months*, *six months*, or *one year*.

4. Click **View**.

9.5.2.2.3. Object statistics

This topic describes how to view the statistics for the number and trend of objects by cluster.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > Cluster Data**.
3. On the **Object Statistics** tab, select **Current Overall Statistics** or **Growth Trend** to view object statistics.

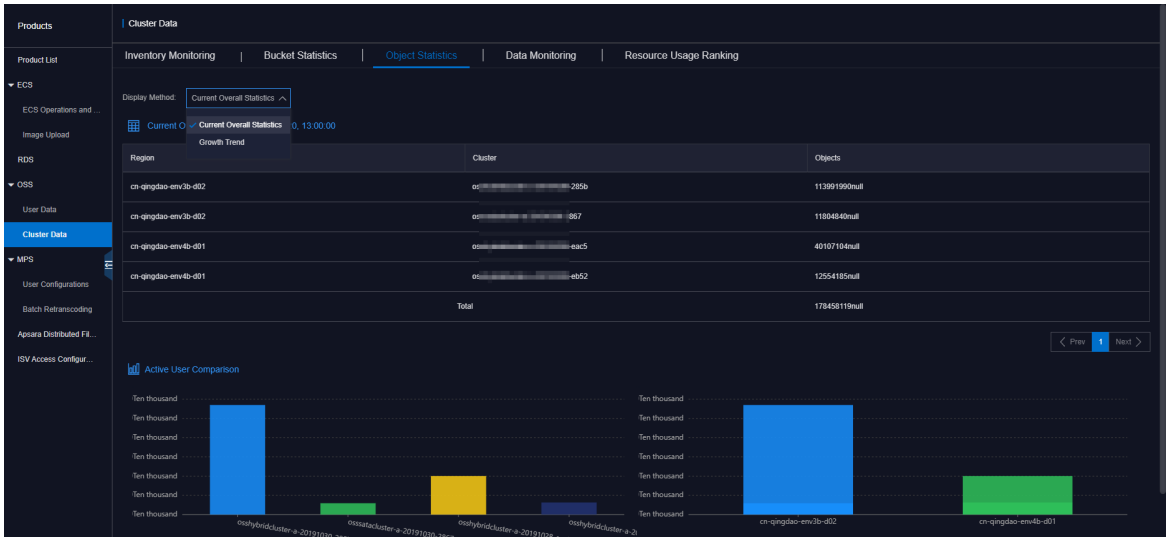

- You can select **Current Overall Statistics** to query statistics of last hour.
- If you select **Growth Trend**, you can specify a time range of *seven days*, *30 days*, *three months*,

six months, or *one year*.

4. Click **View**.

9.5.2.2.4. Data monitoring

This topic describes how to collect statistics for each metric by cluster.

Context

Cluster data metrics are similar to user data metrics except that the object of cluster data metrics is the data collected by cluster.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > Cluster Data**.
3. On the **Data Monitoring** tab, set **Monitoring Items** and **Specify Time Range**. Click **View**.

> **Note** Metric descriptions:
>
> - SLA: indicates the service level availability metric for OSS. Formula: SLA = Non-5xx request count per 10s or hour/Total valid request count × 100%.
> - Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
> - QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.
> - Latency: collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.
> - HTTP Status: collects statistics for the percentages of the numbers of 5xx, 403, 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
> - Storage Capacity: collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.

4. Move the pointer over the trend chart to display data at a specific point in time.

Data monitoring 1

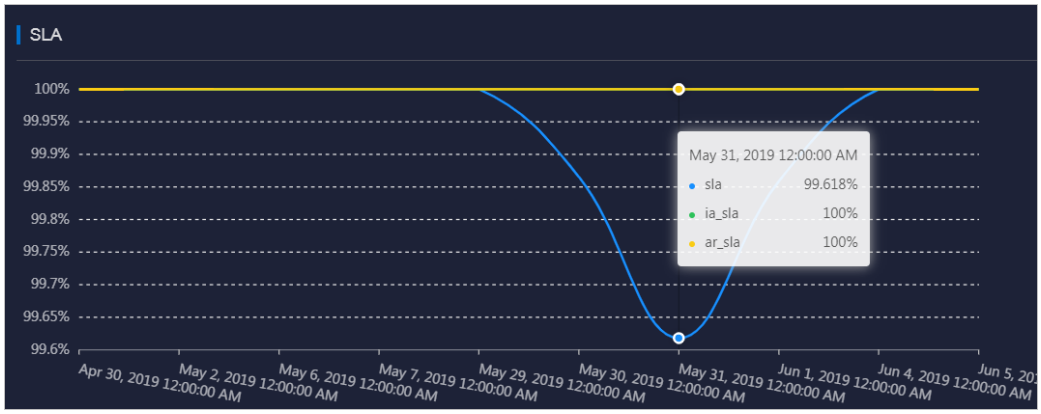

Metric descriptions:

- SLA: indicates the service level availability metric for OSS. Formula: SLA = Non-5xx request count per 10s or hour/Total valid request count × 100%.
- HTTP Status: collects statistics for the percentages of the numbers of 5xx, 403, 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
- Latency: collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.
- Storage Capacity: collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.
- Image Processing Capacity: collects statistics for the number of processed images.

  > **Note** By default, this metric is not displayed. You can select this metric from the **Monitoring Items** drop-down list.

- Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
- QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.

The following example describes typical operations on the data monitoring trend chart:

- If you query data monitoring information by user, you can click the bucket name in the trend chart to show or hide the curve.

  Data monitoring 2

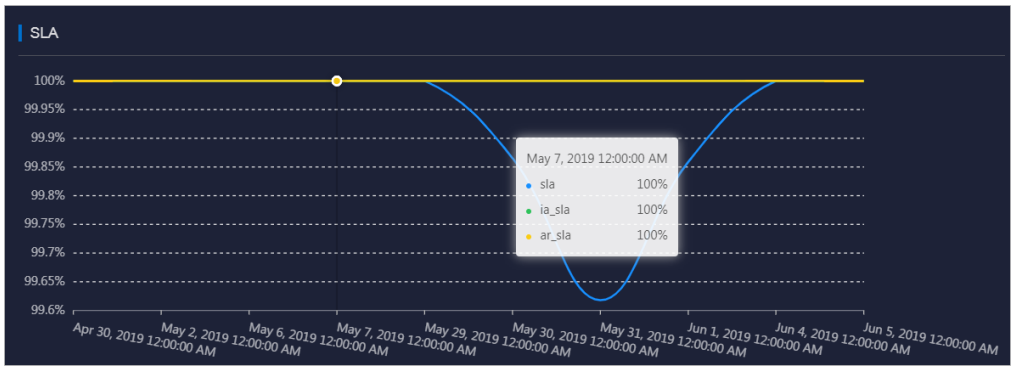

- Move the pointer over the trend chart to display data at a specific point in time.

  Data monitoring 2

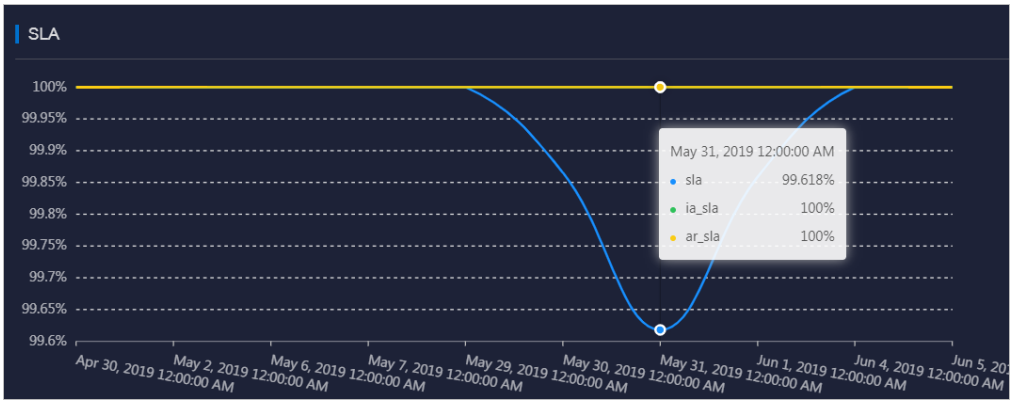

9.5.2.2.5. Resource usage rankings

This topic describes how to collect usage of resources by cluster. This way, administrators can monitor users that consume more resources.

Context

Data resources can be ranked based on the following metrics:

- Total Requests
- Request Errors
- Public Inbound Traffic and Public Outbound Traffic
- Internal Inbound Traffic and Internal Outbound Traffic
- CDN Uplink Traffic and CDN Downlink Traffic
- Storage Capacity, Storage Increment, and Storage Decrement

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > OSS > Cluster Data**.
3. On the **Resource Usage Ranking** tab, select **Report** or **Trend** from the Display Mode drop-down list. Select a number from the **Top** drop-down list. Set **Specify Time Range** and **Monitoring Items** to view resource usage.

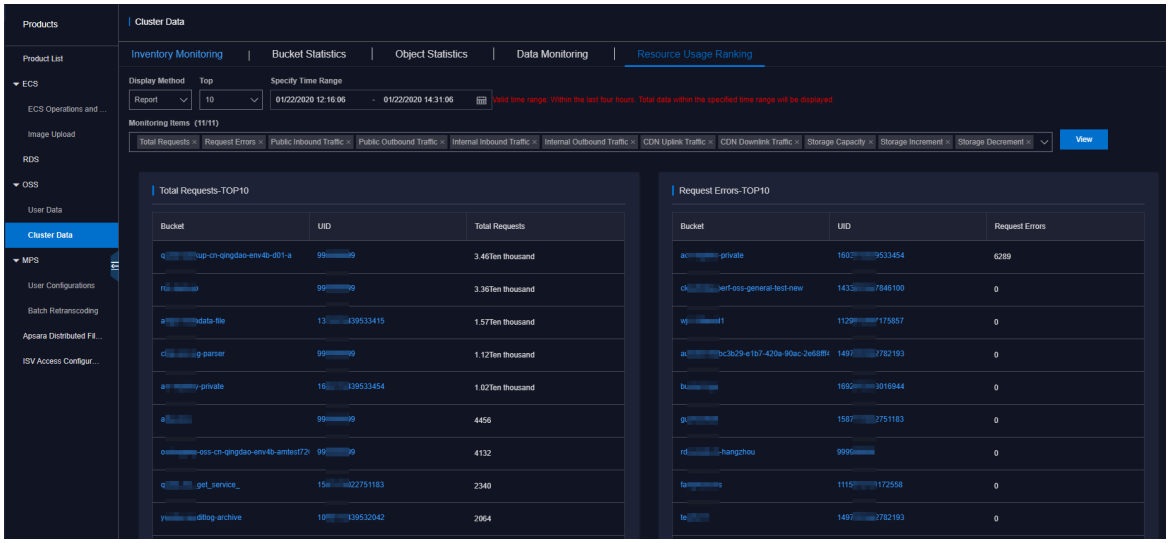

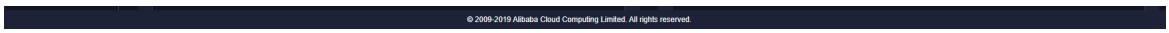

- In **report** mode, you can view the top **10**, **30**, or **50** buckets by resource usage.
- In **trend** mode, you can view the top **10** buckets by resource usage.

4. Click **View**.

9.5.3. Tools and commands

9.5.3.1. Typical commands supported by tsar

You can use tsar to perform operations and maintenance on OSS. This topic describes typical commands supported by tsar.

tsar allows you to run the following commands:

- View help details of tsar

  Command: **tsar -help**

- View the NGINX operation data of each minute from the past two days

  Command: **tsar -n 2 -i 1 -nginx**

  In this command, *-n 2* indicates the data generated in the past two days. *-i 1* indicates one result record generated each minute.

- View the tsar load status and operation data of each minute from the past two days

  Command: **tsar --load -n 2 -i 1**

9.5.3.2. Configure tsar for statistic collection

You can configure tsar to collect data generated when NGINX runs.

Run the following command to configure tsar for statistic collection:

```
cat /etc/tsar/tsar.conf |grep nginx
```

The following figure shows that the status of mod_nginx is *on*.

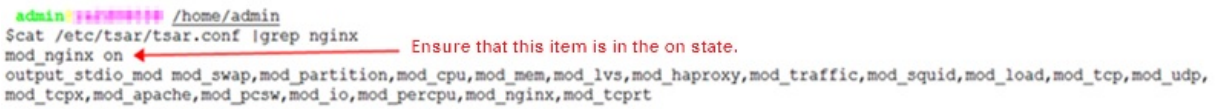

9.6. Tablestore

9.6.1. Tablestore Operations and Maintenance System

9.6.1.1. Overview

This document describes the features, domain name, and modules of Tablestore Operations and Maintenance System.

Tablestore Operations and Maintenance System helps find problems during operations and maintenance (O&M) and notifies you of the current running status of the services. Appropriate use of Tablestore Operations and Maintenance System can significantly improve O&M efficiency.

The domain name of Tablestore Operations and Maintenance System is in the following format: chiji.ots.${global:intranet-domain}.

Tablestore Operations and Maintenance System consists of the following modules: user data, cluster management, inspection center, monitoring center, system management, and platform audit. These modules provide comprehensive O&M functions to meet different requirements.

9.6.1.2. User data

9.6.1.2.1. Instance management

You can query an instance list by specifying the region, and cluster name. You can also query an instance by specifying the filter conditions and view details of the instance and tables.

Instance management provides the following features:

- Query an instance list by specifying the region, and cluster name.

  You can specify a region, and a cluster to view the instances, and the basic information of each instance in the specified cluster.

  - Query the list of instances in a cluster.
  - View basic information of instances in the instance list.
  - View details of an instance by clicking the instance name.
  - Update and delete an instance in the instance list.

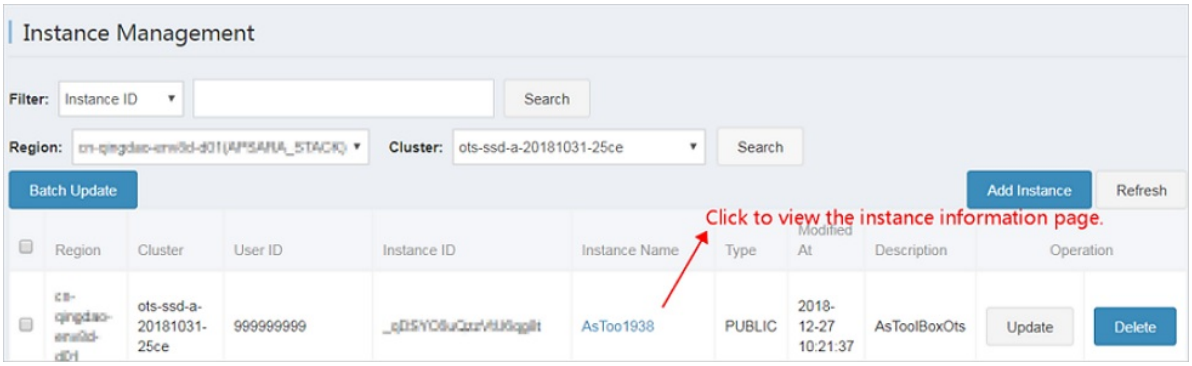

- Search for instances based on specified conditions

  You can search for an instance based on the instance name, instance ID, user ID or Apsara Stack tenant account in all clusters of all regions.

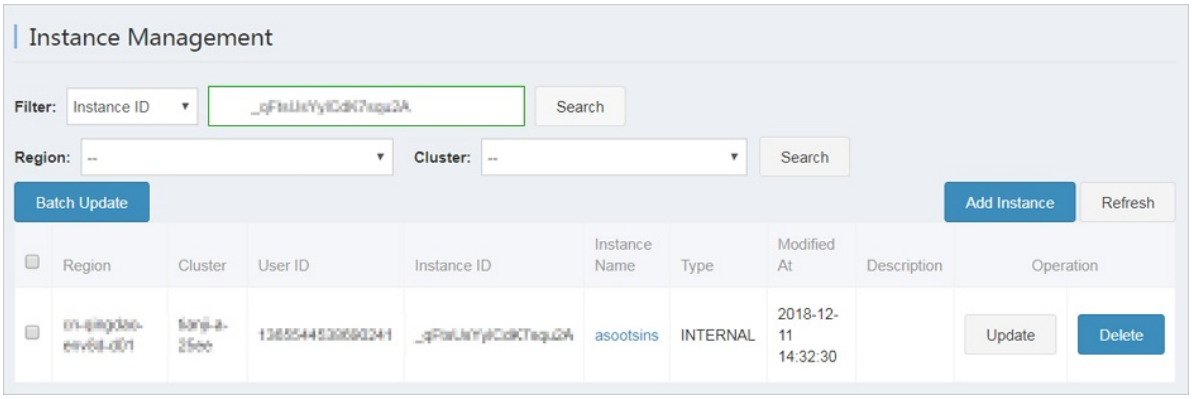

- View instance details
  - Instance overview

    Click the instance name. On the **Details** tab, you can view instance details such as the link for instance monitoring, the IP address of the instance for the Internet and internal network, and the statistics information of tables in the instance.

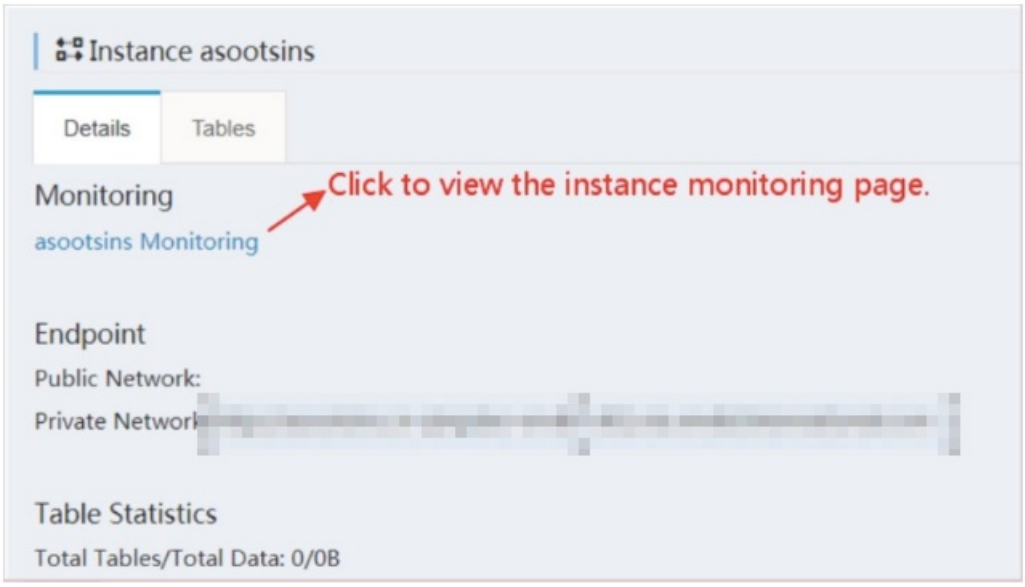

  - Tables information

    Click the instance name. On the **Tables** tab, you can view the maxVersion, ttl, readCU, writeCU, and timestamp of tables.

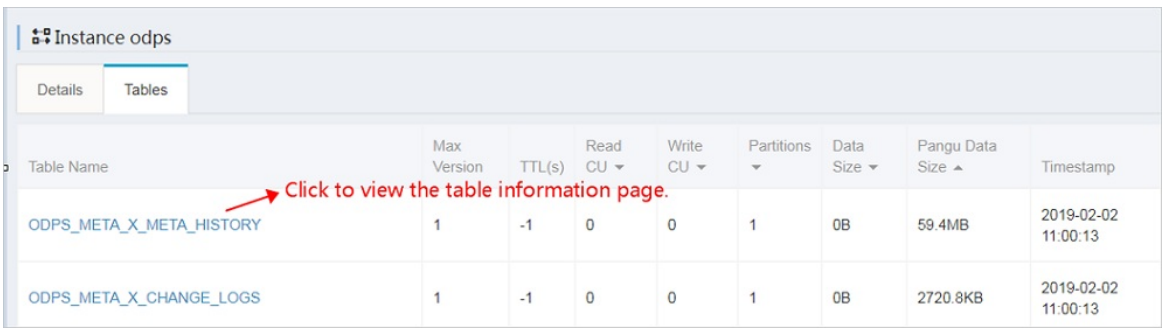

- View table details

- Details

  Click the table name. On the **Details** tab, you can view the overview information of the table, such as the number of partitions and the table size.

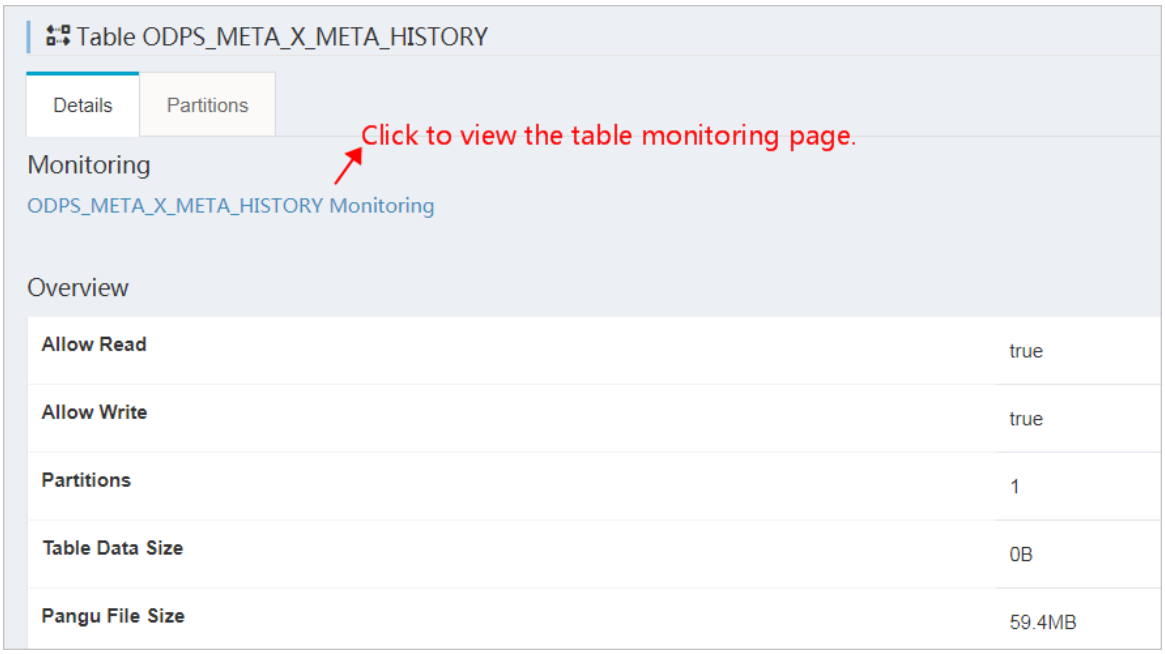

- Partitions

  Click the table name. On the **Partitions** tab, you can obtain the basic information of an partition, such as the partition ID and Worker information. You can also search for partitions based on the Worker name that is listed in the table or the partition ID.

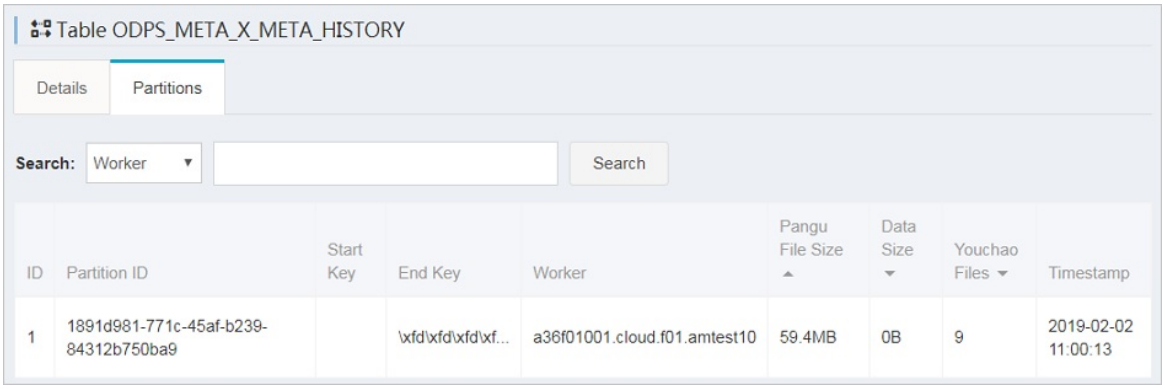

9.6.1.3. Cluster management

9.6.1.3.1. Cluster information

You can obtain the list of clusters, view cluster usage and top requests based on cluster information.

You can perform the following operations based on the cluster information:

- Clusters

  You can query a list of clusters in all regions or in a specified region. Perform the following operations:

- OCM cluster synchronization: An OCM service is deployed in each region of Tablestore. The OCM service contains all cluster information of a region. This function synchronizes OCM clusters with their corresponding regions in Tablestore Operations and Maintenance System to obtain all clusters in the regions.
- Cluster deletion: You can use this function to remove a cluster from Tablestore Operations and Maintenance System after you confirm that the cluster is taken offline.

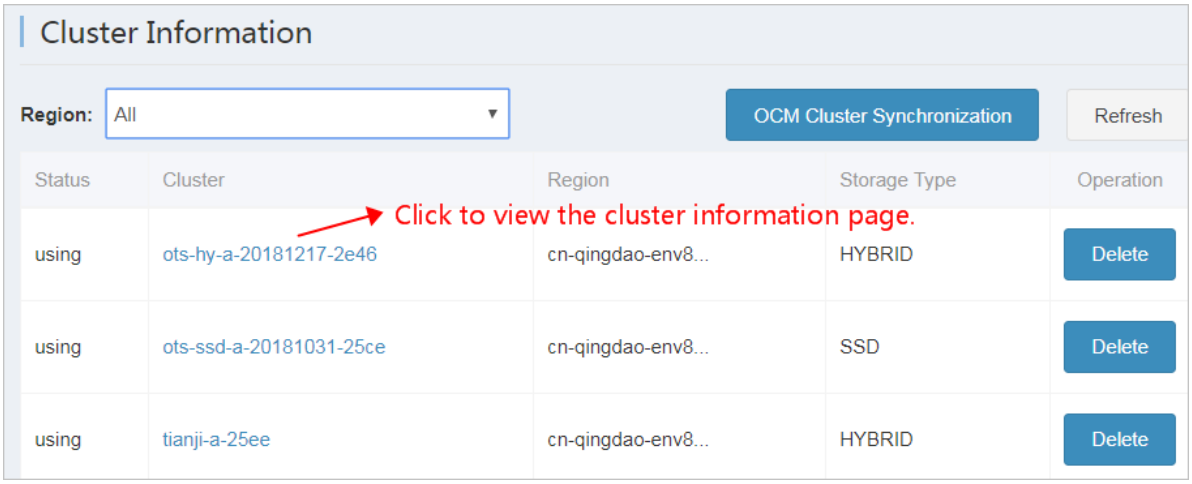

- Cluster details

  Click a cluster name in the Cluster column to go to the cluster details page. You can view the detail information of the cluster, including the overview, top request, and cluster usage.

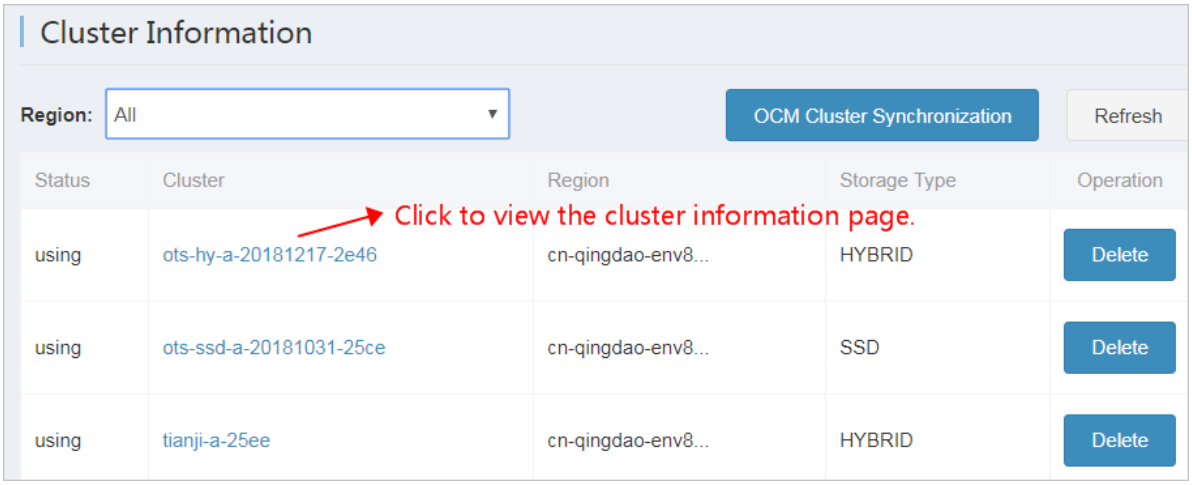

- Overview: provides the basic information of a cluster.

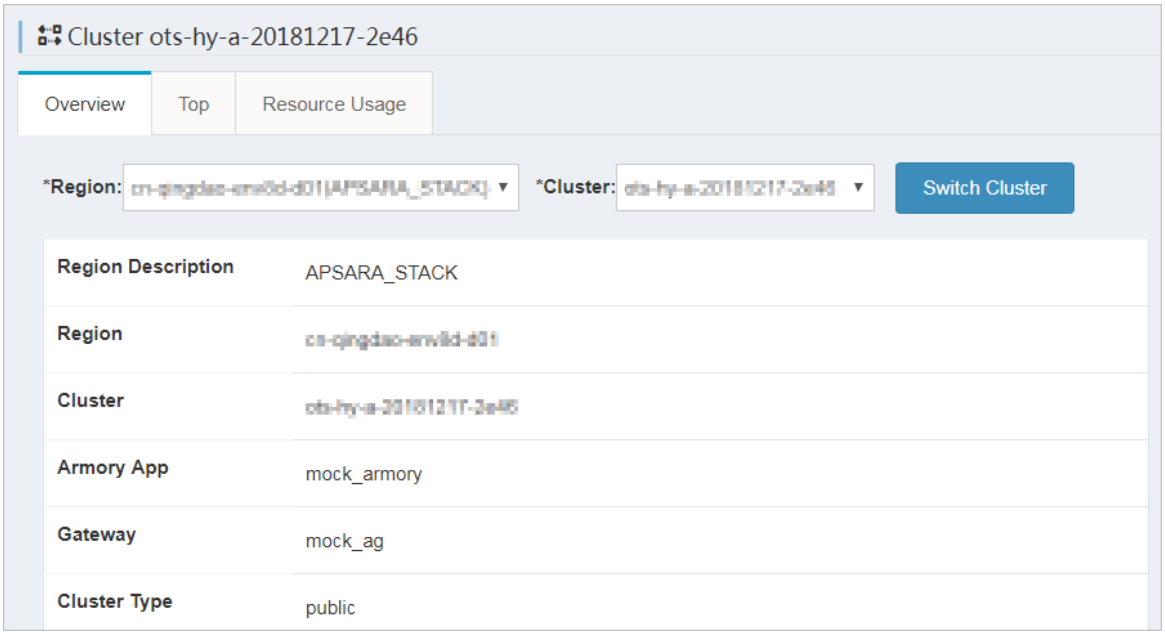

- Top: provides top request information of partitions and tables.

  Click an instance name in the InstanceName column to go to the instance details page, where you can view detailed information of the instance. Click a table name in the TabelName column to go to the table details page, where you can view detail information of the table. Click a partition ID in the PartitionID column to go to the partition details page, where you can view detail information of the partition. Click **More** and you can view detail information of the top request.

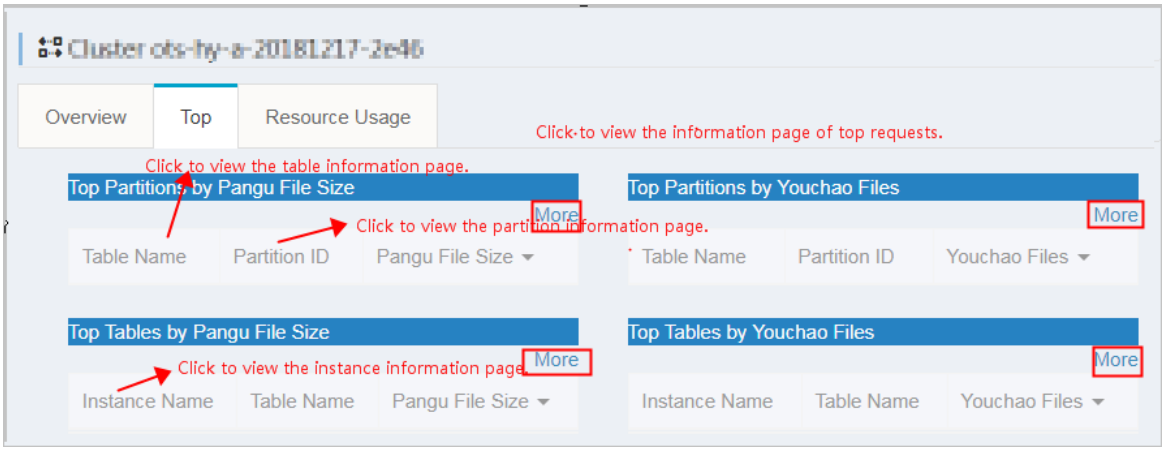

- Resource Usage: provides cluster usage details. The usage statistics collection task is automatically triggered in the background at specific intervals. In special cases, you can click **Collect Data** to manually trigger the usage statistics collection task. After the usage statistics collection task is completed, refresh the page to display the latest usage statistics.

  > **Note** The usage check either succeeds or fails. In addition, you must pay special attention to the cause of a usage check failure. (The usage check failure is caused by the failure to obtain storage space, as shown in the following figure.)

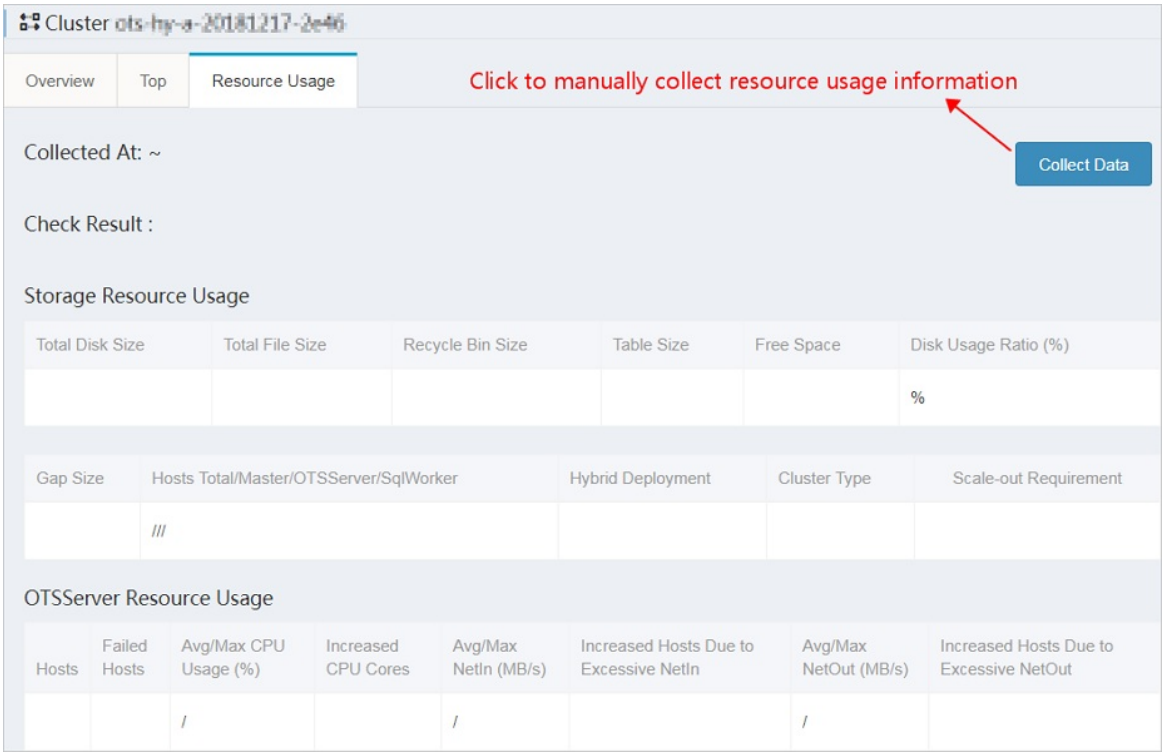

9.6.1.4. Inspection center

9.6.1.4.1. Abnormal resource usage

You can click Abnormal Resource Usage in the left-side navigation pane to find all cluster abnormalities and their causes.

You can click Abnormal Resource Usage in the left-side navigation pane to inspect cluster abnormalities in all regions. Abnormalities are displayed in red, which allows you to find abnormal clusters.

The usage statistics collection task is automatically triggered in the background at specific intervals. In special cases such as a failure in background task execution, you can click **Collect Data** to manually trigger usage statistics collection. The collection action is performed asynchronously. After the usage statistics collection task is completed, refresh the page to display the latest usage statistics.

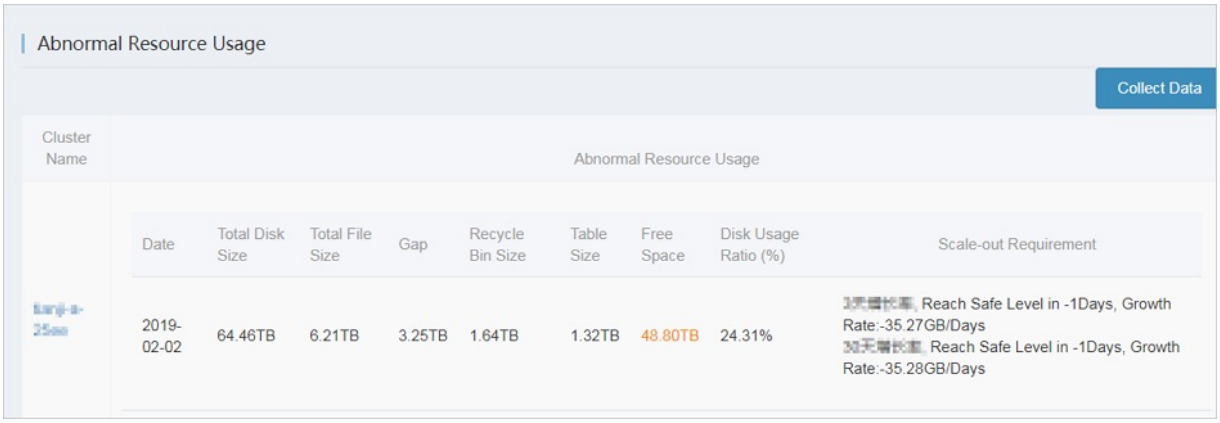

9.6.1.5. Monitoring center

9.6.1.5.1. Cluster monitoring

You can determine the service status of a cluster based on a series of metrics such as cluster-level monitoring information.

You can query the cluster service metrics within a specified time range, and determine whether a cluster service is healthy based on the metrics in the following dimensions.

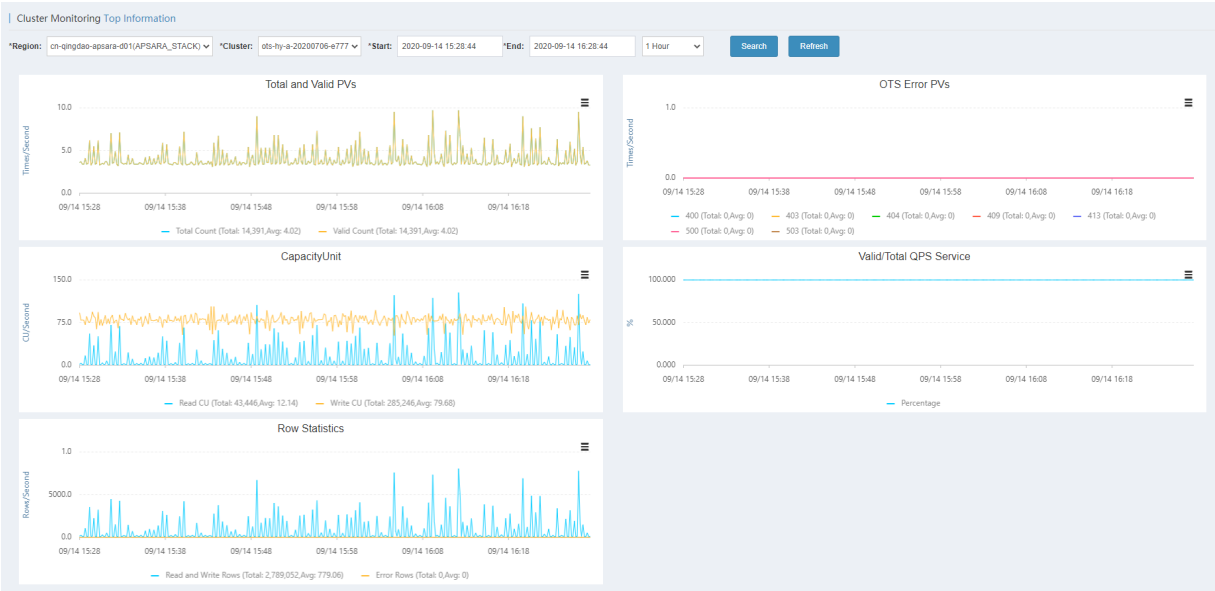

9.6.1.5.2. Application monitoring

You can check the instance-level and table-level metrics to determine whether a service that belongs to a user is abnormal.

You can check the following metrics to determine whether a service for a specified user is in the healthy state.

> **Note** The Instance field is required. The Table and Operation fields are optional.

9.6.1.5.3. Top requests

You can view the top request distribution of clusters by monitoring level or dimension.

The following monitoring levels are supported for top requests: Instance, Instance-Operation, Instance-Table, and Instance-Table-Operation. You can view the top request details of a cluster based on 13 different metrics such as the total number of requests and the total number of rows.

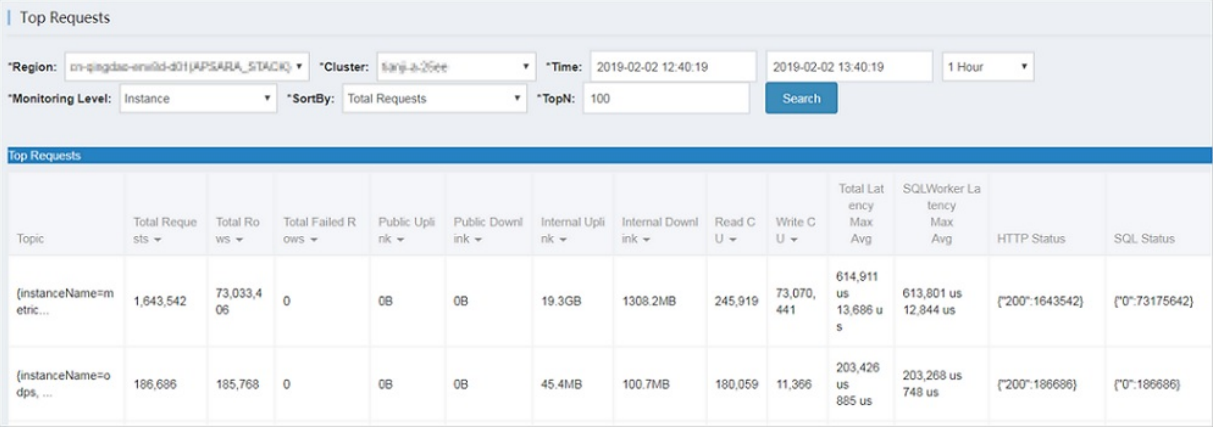

9.6.1.5.4. Request log search

You can search for request logs by using request IDs to assist problem investigation.

You can query all logs associated with a region, cluster, and request ID.

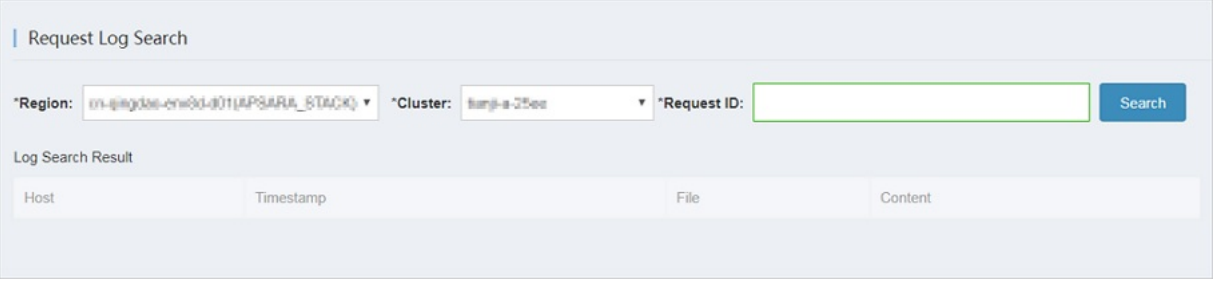

9.6.1.6. System management

9.6.1.6.1. Task management

Background tasks are managed by Tablestore Operations and Maintenance System.

After Tablestore Operations and Maintenance System is deployed in the Apsara Stack environment, the background tasks that collect usage statistics are automatically integrated.

You can perform the following operations on background tasks:

- View task details such as the specific parameters and running time of each task.

  Click **Details** corresponding to a task to view the task details. The following figure shows a monitoring rule displayed on the task details page. The task collects usage statistics at 02:00:00 every day.

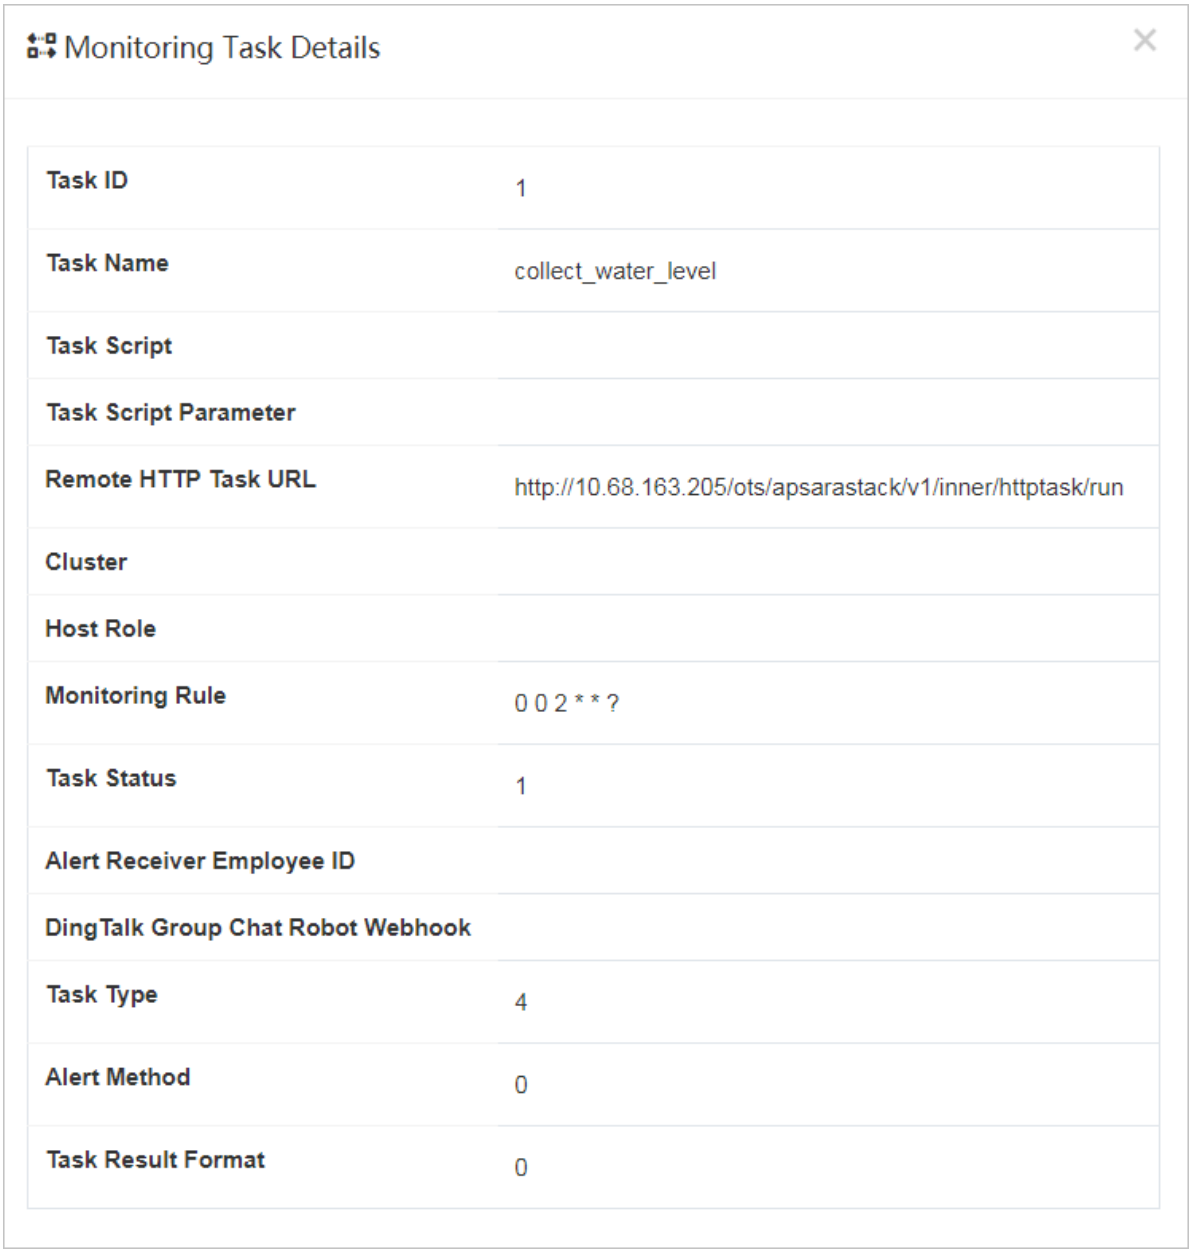

- Enable or disable a task.

  Note Disabled tasks no longer run automatically.

- Run a task immediately.

9.6.1.6.2. View tasks

You can view the execution status of background tasks and find the causes of task exceptions.

The following figure shows the execution status of background tasks in Tablestore Operations and Maintenance System. You can view the succeeded or failed tasks.

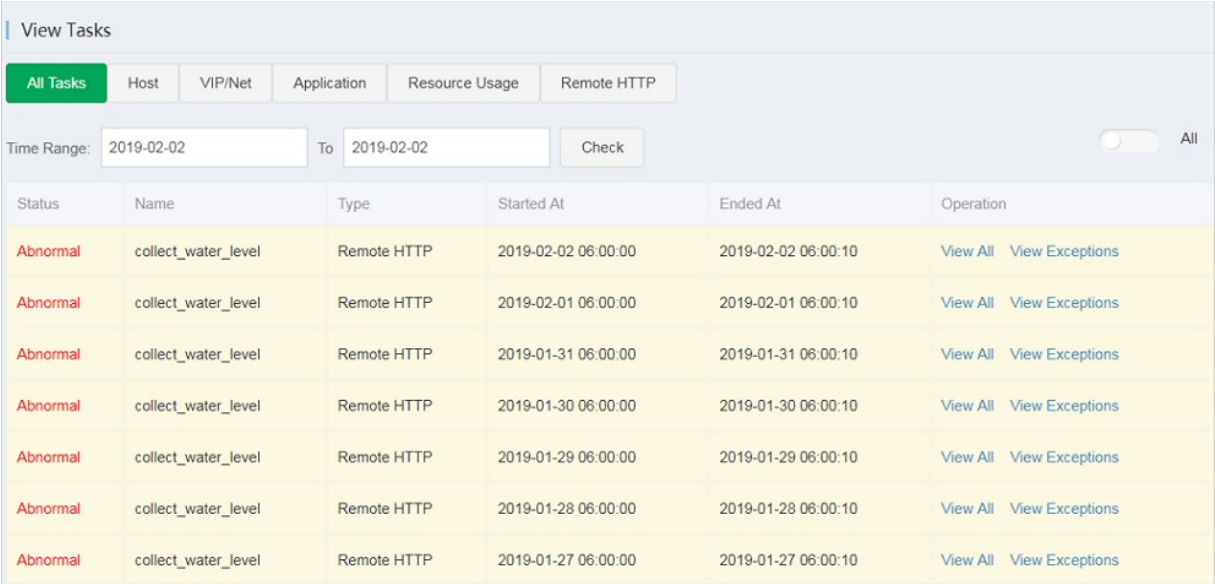

Click **View All** or **View Abnormal** in the Operation column corresponding to the abnormal task to view the specific cause of a task failure, as shown in the following figure.

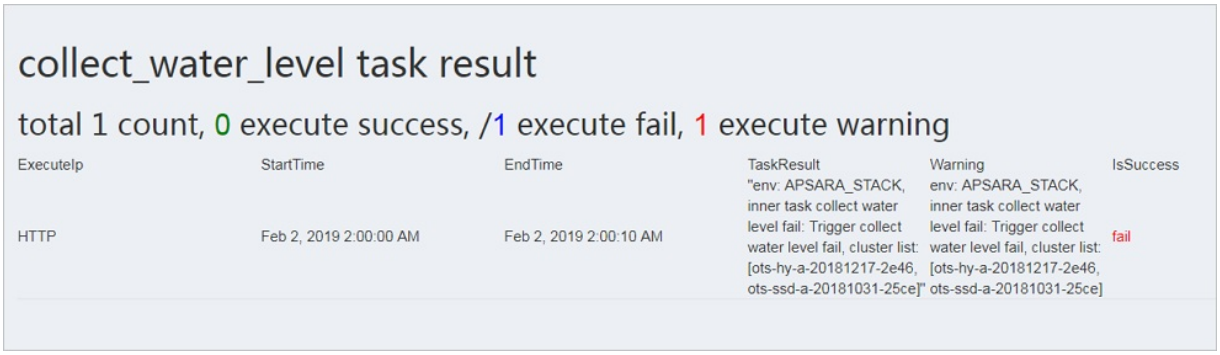

9.6.1.7. Platform audit

9.6.1.7.1. Operation logs

You can view the management and control operation logs of Storage Operations and Maintenance System.

The **Operation Log** page provides the operation logs of Tablestore Operations and Maintenance System. You can query audit records generated within a specified time range and filter the records to obtain information about the platform status.

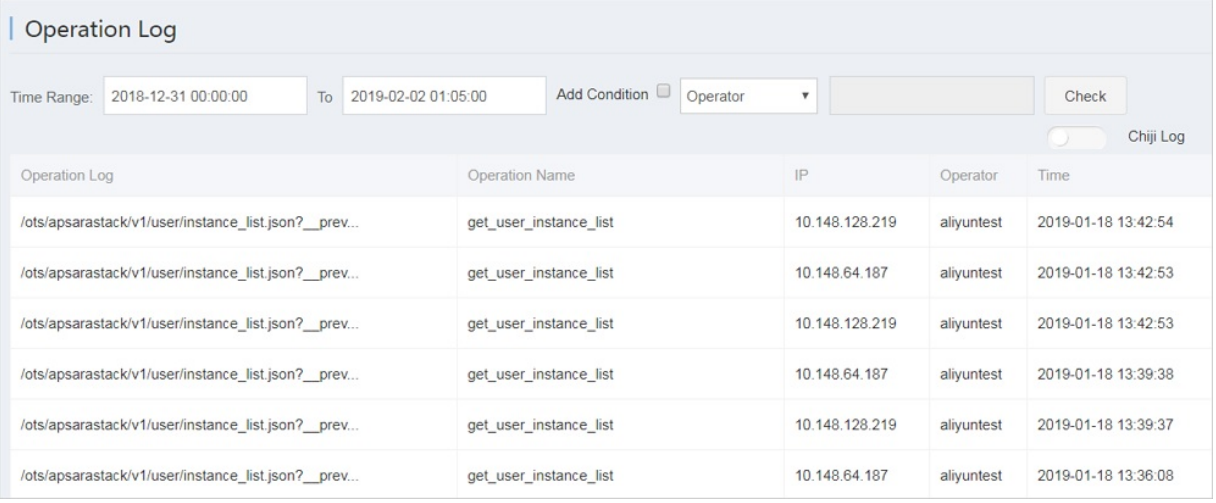

9.6.2. Cluster environments

This topic describes the environment and service information of Tablestore.

Two environments are provided for Tablestore: the internal environment for cloud services such as MaxCompute, Log Service, or StreamSQL, and the external environment deployed for users.

Some cloud services use both environments. For example, metadata of StreamSQL is stored in the internal environment, but its user data is stored in the external environment.

Tablestore services include TableStoreOCM, TableStoreInner/TableStore, TableStorePortal, chiji, and TableStoreSqlInner/TableStoreSql.

- TableStoreOCM: the tool used to manage information about clusters, users, and instances
- TableStoreInner/TableStore: the Tablestore data service node
- TableStorePortal: the background of the Tablestore O&M platform
- chiji: the Tablestore O&M platform used for fault location
- TableStoreSqlInner/TableStoreSql: the Tablestore background tool

9.6.3. System roles

This topic describes the functions of system roles.

- TableStoreOCM
  - OCMInit: the OCM initialization tool used to create tables and bind POP APIs
  - OCM: the service node of OCM
  - ServiceTest: the service test image of OCM
- TableStoreInner/TableStore
  - InitCluster: the process of adding cluster information to OCM, including the domain name, cluster type, and pre-configured Tablestore account information
  - LogSearchAgent: the log collection node of Tablestore
  - MeteringServer: the metering node that is available only in Tablestore
  - MonitorAgent: the data collection node of the Tablestore monitoring system
  - MonitorAgg: the data aggregation node of the Tablestore monitoring system

  - OTSAlertChecker: the alerting module of Tablestore
  - OTSFrontServer: the frontend server of Tablestore, which can be NGINX, OTS Server, or Replication Server
  - OTSServer: the fronted service of Tablestore
  - OTSTEngine: the NGINX service for Tablestore frontend servers
  - PortalAgServer: the background service of Tablestore Operations and Maintenance System
  - ServiceTest: the test service that runs scheduled smoke tests
  - SQLOnlineReplicationServer: the disaster recovery service of Tablestore
  - SQLOnlineWorker: the application that was used to generate alerts but no longer provides services
  - TableStoreAdmin: all O&M tools of Tablestore, including the splitting and merging tools
- TableStorePortal
  - PortalApiServer: the background service of Tablestore Operations and Maintenance System
- TableStoreSqlInner/TableStoreSql
  - Tools: the background tools of Tablestore such as sqlonline_console
  - UpgradeSql: the background hot upgrade tool of Tablestore

9.6.4. Pre-partition a table

9.6.4.1. Pre-partitioning

This topic describes the rules and methods of pre-partitioning.

When you create a table, Tablestore automatically creates a partition for the table. This partition can be configured to automatically split based on the data size or data access load when your business develops. A table that has only one partition may be unable to provide sufficient service capabilities during a stress test or data import. In this scenario, you must pre-partition the table.

Pre-partitioning rules

You can estimate the required number of partitions based on the standard size of 10 GB per partition. However, other factors such as the number of hosts and concurrent write operations by developers must be considered. We recommend that the total number of partitions do not exceed 256. If data can be written into the table evenly, you can partition the table equally based on the number of partitions required.

> **Note** When data is written into the table, the system automatically splits the table to ensure sufficient partitions are available when the data increases.

Pre-partitioning methods

You can use split_merge.py to pre-partition a data table. You can obtain split_merge.py from */apsara/TableStoreAdmin/split* on the host of TableStoreAdmin in TableStoreInner.

You can use any of the following methods to partition a data table:

> **Note** You can also use the following methods to partition a table that already has data.

- Specify a split point

  ```
  python2.7 split_merge.py split_table -p point1 point2 ... table name
  ```

- Specify the number of partitions and the partition key format
  - The partition key is of the int type.

    ```
    python2.7 split_merge.py split_table -n: number of partitions --key_digit: table name
    ```

  - The partition key starts with an MD5 hash in lowercase. The MD5 hash can contain digits and lowercase letters from a to f.

    ```
    python2.7 split_merge.py split_table -n: number of partitions --key_hex_lower: table name
    ```

  - The partition key starts with an MD5 hash in uppercase. The MD5 hash can contain digits and uppercase letters from A to F.

    ```
    python2.7 split_merge.py split_table -n: number of partitions --key_hex_upper: table name
    ```

  - The partition key is Base64-encoded, and can contain the plus sign (+), forward slash (/), digits and letters.

    ```
    python2.7 split_merge.py split_table -n: number of partitions --key_base64: table name
    ```

  - -- only_plan: generates split points but does not split the table. -- force: directly splits the table without manual confirmation.

    ```
    python2.7 split_merge.py split_table -n: number of partitions --key_digit --only_plan: table name
    ```

- Split a partition based on the existing data

  ```
  python2.7 split_merge.py split_partition -n PART_COUNT (number of partitions) partition_id
  ```

9.6.4.2. View partitions

You can view the partitions of a data table in Tablestore Operations and Maintenance System.

On the Tablestore Operations and Maintenance System, find a table in the specified instance. Click the table name to view details of the table. On the **Partitions** tab, you can view the information of all partitions in the table. The information contains the partition ID, range, worker, Apsara Distributed File System file size, and data size. The partition size displayed may not be the current partition size because the data is updated only after the system merges files. The Apsara Distributed File System file size is the compressed data size. The actual storage space is three times the file size because the data is stored in three copies.

9.7. ApsaraDB for RDS

9.7.1. Architecture

9.7.1.1. System architecture

9.7.1.1.1. Backup system

ApsaraDB for RDS can back up databases at any time and restore them to any point in time based on the backup policy, which makes the data more traceable.

Automatic backup

ApsaraDB RDS for MySQL supports both physical and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can create temporary backup files when necessary. Temporary backup files are retained for seven days.

Log management

ApsaraDB RDS for MySQL automatically generates binlogs and allows you to download them for local incremental backup.

Instance cloning

A cloned instance is a new instance with the same content as the primary instance, including data and settings. This feature allows you to restore data of the primary instance or create multiple instances that are the same as the primary instance.

9.7.1.1.2. Data migration system

ApsaraDB for RDS provides Data Transmission Service (DTS) to help you migrate databases.

Replicate databases between instances

ApsaraDB for RDS allows you to migrate databases from one instance to another.

Migrate data to or from RDS instances

ApsaraDB for RDS provides professional tools and migration wizards to help you migrate data to or from RDS instances.

Download backup files

ApsaraDB for RDS retains backup files for seven days. During this period, you can log on to the RDS console to download the files.

9.7.1.1.3. Monitoring system

RDS provides multi-dimensional monitoring services across the physical, network, and application layers to ensure business availability.

Performance monitoring

RDS provides nearly 20 metrics for system performance monitoring, such as disk capacity, IOPS, connections, CPU utilization, network traffic, TPS, QPS, and cache hit rate. You can obtain the running status information for any instances within the past year.

SQL auditing

The system records the SQL statements and related information sent to RDS instances, such as the connection IP address, database name, access account, execution time, and number of records returned. You can use SQL auditing to check instance security and locate problems.

Threshold alerts

RDS provides alert SMS notifications if status or performance exceptions occur in the instance.

These exceptions can be involved in instance locking, disk capacity, IOPS, connections, and CPU. You can configure alert thresholds and up to 50 alert recipients (of which five are effective at a time). When an instance exceeds the threshold, an SMS notification is sent to the alert recipients.

Web operation logs

The system logs all modification operations in the RDS console for administrators to check. These logs are retained for a maximum of 30 days.

9.7.1.1.4. Control system

If a host or instance does not respond, the RDS high-availability (HA) component checks for exceptions and fails over services within 30 seconds to guarantee that applications run normally.

9.7.1.1.5. Task scheduling system

You can use the RDS console or API operations to create and delete instances, or switch instances between the internal network and Internet. All instance operations are scheduled, traced, and displayed as tasks.

9.7.2. Log on to the Apsara Stack Operations console

Prerequisites

- 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.

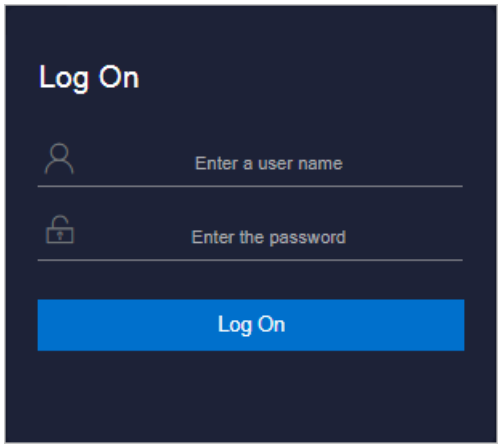

> **Note** You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.

> **Note** Obtain the username and password used to log on to ASO from the deployment personnel or the administrator.

   - 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.

9.7.3. Manage instances

You can view instance details, logs, and user information.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > RDS**.
3. On the **Instance Management** tab of the **RDS** page, you can perform the following operations:
   - View instances

     View instances that belong to the account on the **Instance Management** tab, as shown in Instances.

     Instances

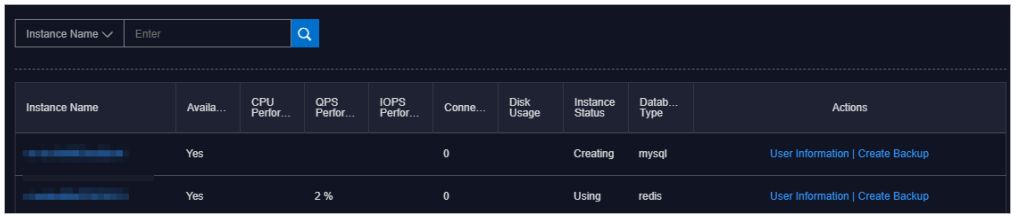

- View instance details

  Click the ID of an instance to view details, as shown in Instance details. You can switch your service between primary and secondary instances and query historical operations on this page.

  > **Note** We recommend that you do not perform forced switchover, because it may result in data loss if data is not synchronized between the primary and secondary instances.

  Instance details

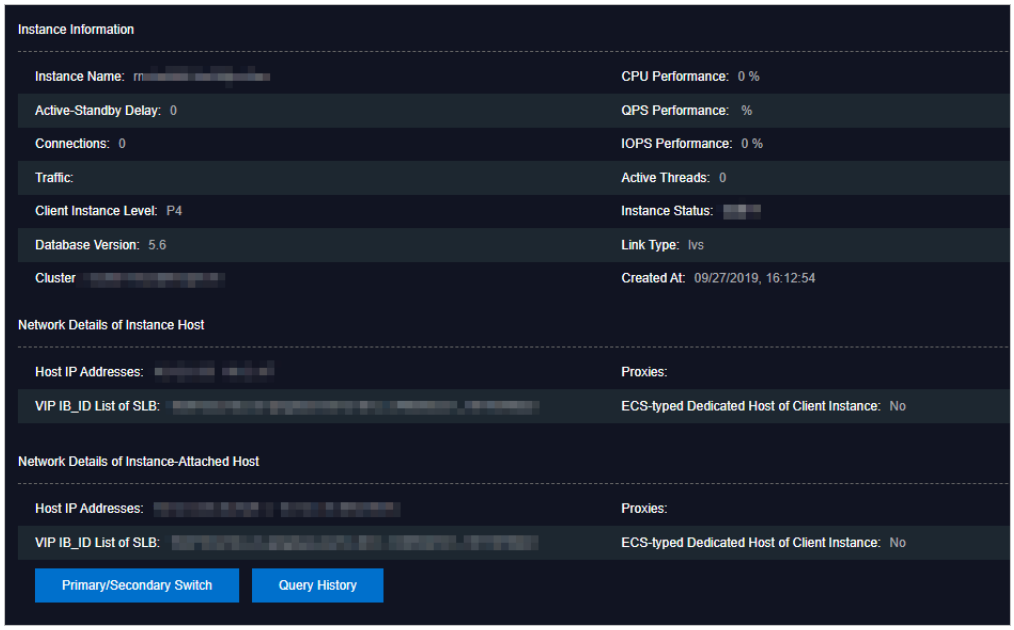

- **View user information**

  Click **User Information** in the **Actions** column corresponding to an instance, as shown in User information.

  User information

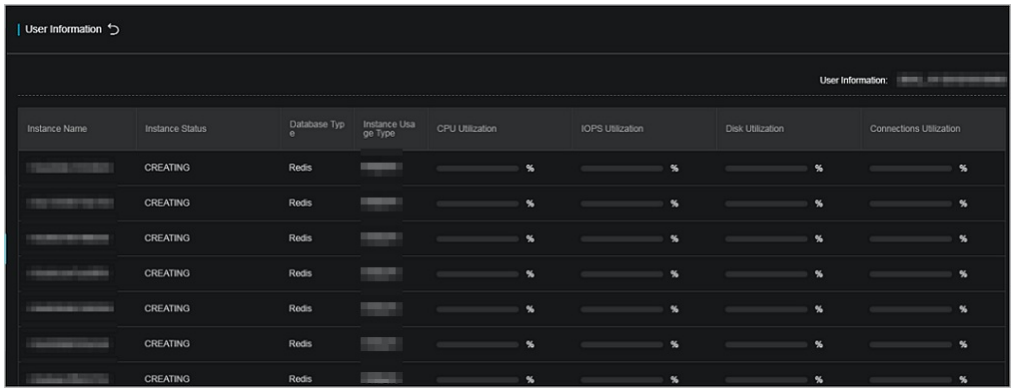

- **Create backups**

  For ApsaraDB RDS for MySQL instances, click **Create Backup** in the **Actions** column to view the backup information, as shown in Backup information. You can also click **Create Single Database Backup** on the Backup Information page to back up a single database.

  Backup information

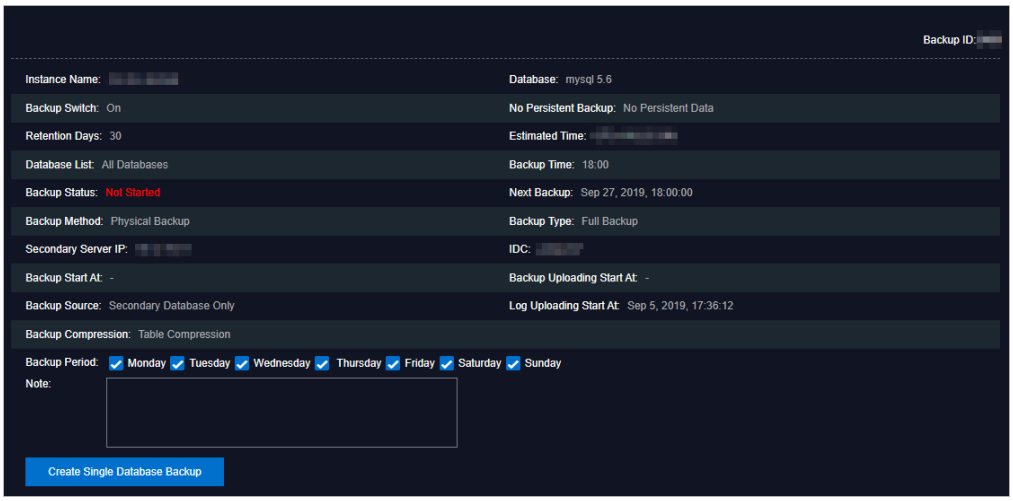

9.7.4. Manage hosts

You can view and manage hosts.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > RDS**.
3. On the **Host Management** tab of the **RDS** page, you can view the information of all hosts.

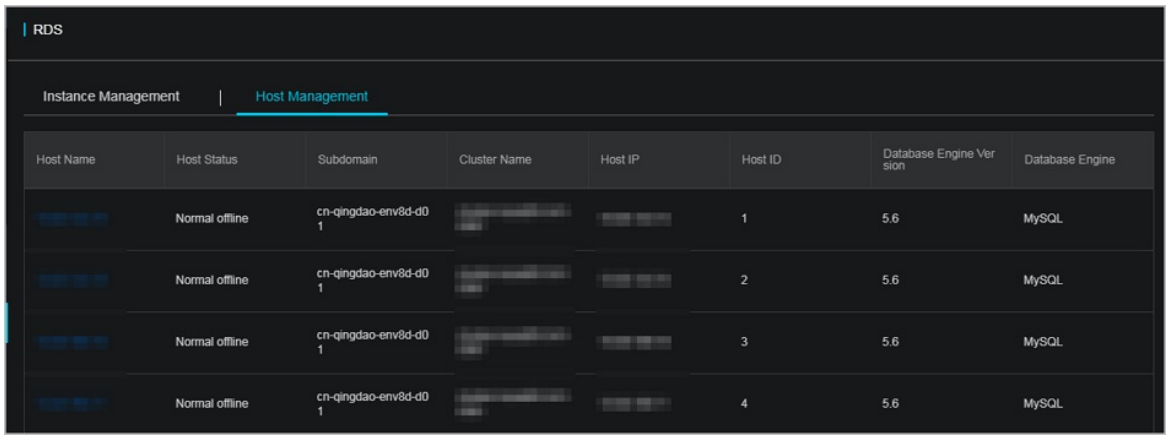

4. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.

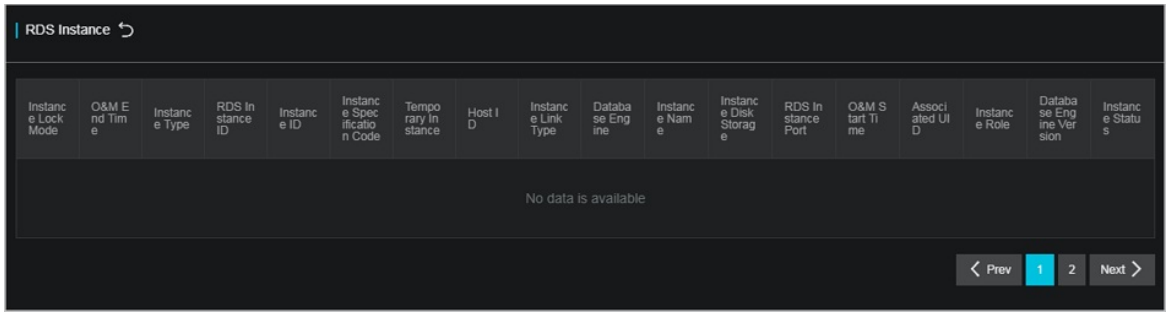

9.7.5. Security maintenance

9.7.5.1. Network security maintenance

Network security maintenance consists of device and network security maintenance.

Device security

Check network devices and enable their security management protocols and configurations of devices.

Check for timely updates to secure versions of network device software.

For more information about the security maintenance method, see the device documentation.

Network security

Based on your network considerations, select the intrusion detection system (IDS) or intrusion prevention system (IPS) to detect abnormal Internet and Intranet traffic and protect against attacks.

9.7.5.2. Account password maintenance

Account passwords include RDS system passwords and device passwords.

To ensure account security, you must periodically change the system and device passwords, and use passwords with high complexity.

9.8. AnalyticDB for MySQL

9.8.1. What is AnalyticDB for MySQL?

AnalyticDB for MySQL is a real-time online analytical processing (RT-OLAP) service that is developed by Alibaba Cloud to analyze large amounts of data at high concurrency. AnalyticDB for MySQL can analyze hundreds of billions of data records across multiple dimensions within milliseconds and provide you with data-driven insights into your business.

AnalyticDB for MySQL can compute large amounts of data quickly and enable you to explore and find data value. You can also embed AnalyticDB for MySQL into a business system to provide users with analysis services.

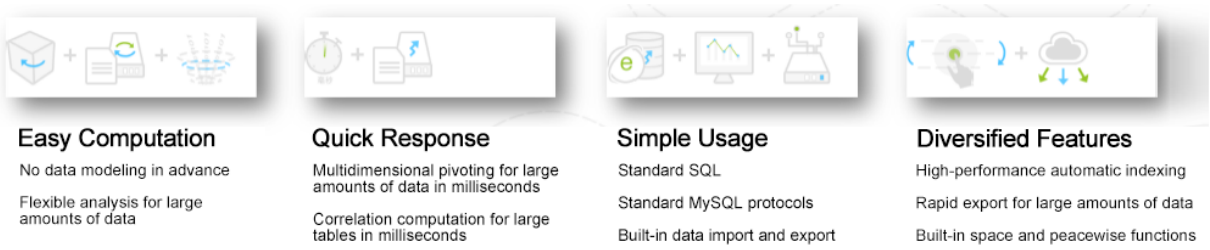

AnalyticDB for MySQL can perform low latency, high concurrency, and real-time online processing and retrieval for large amounts of data. AnalyticDB for MySQL is widely used by enterprises for real-time multidimensional analysis, by businesses for customer group selection and analysis, and by government agencies for flexible big data retrieval and statistics. AnalyticDB for MySQL is used in Internet business systems that have hundreds of thousands to tens of millions of users, such as Data Cube, Taobao Index, Kuaidi Dache, Alimama DMP, and Taobao Groceries.

Key storage technology

AnalyticDB for MySQL provides two storage modes:

- **High-performance storage**: delivers good query and concurrent processing performance but requires high storage costs with SSDs. This storage mode is suitable for scenarios where large amounts of data is flexibly analyzed or queried with high concurrency.
- **Large-capacity storage**: features low storage costs but provides lower query and concurrent processing performance than high-performance storage. This storage mode is suitable for scenarios where details are queried from large amounts of data or low concurrency and high latency analysis are required.

AnalyticDB for MySQL uses column store to store data of tables. Each table can contain more than 1,000 columns. Column store has the following features:

- **Advantage**: Only a few I/O resources are required when data analysis or statistics is performed or when a small number of columns are queried within a wide table.
- **Disadvantage**: Highly distributed data requires excessive I/O resources when many columns are queried.
- **Unique feature**: Data presorting of aggregate columns can help mitigate the downsides.

Scenarios

- **Application type**

  This type of business provides simple queries which only return small amounts of data and do not require multiple tables to be joined.

- **BI type**

  This type of business provides a real-time data warehouse where real-time tables are joined with multiple dimension tables to sort and divide data into many groups.

- **Ad-hoc type**

  This type of business provides complex analysis by joining multiple real-time tables to sort and group data or return more than 500 data records.

9.8.2. Architecture

9.8.2.1. System architecture

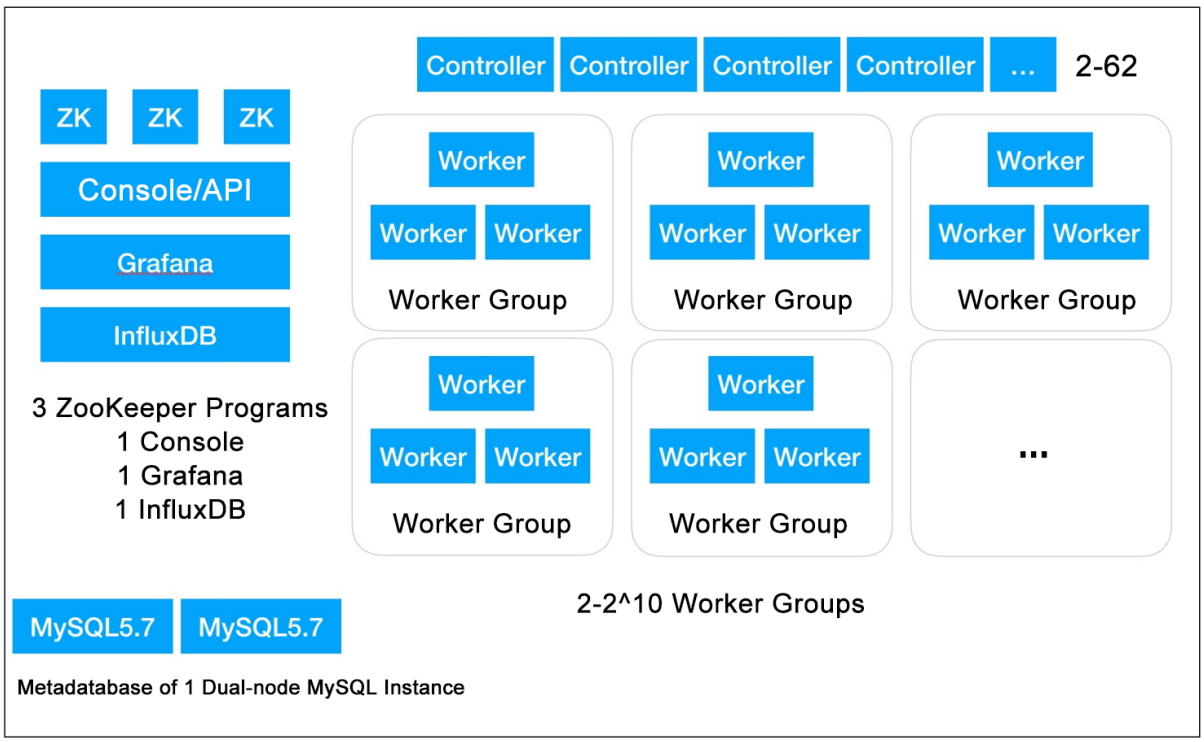

- AnalyticDB for MySQL is compatible with the MySQL protocol and supports JDBC, ODBC, and RESTful APIs. It is also compatible with third-party user data analysis applications, Apsara Stack Quick BI and DataV, and commercial BI tools such as Tableau and QlikView.
- AnalyticDB for MySQL can exchange data with MaxCompute, ApsaraDB for RDS, and OSS in Apsara Stack.
- Controllers parse, plan, and optimize SQL statements. Server Load Balancer (SLB) can be deployed at the frontend for load balancing.

  Workers compute and store data. A worker group consists of three workers. Each cluster can consist of more than two worker groups.
- AnalyticDB for MySQL functional modules include controllers, worker groups, ZooKeeper, InfluxDB, Grafana, console, and MySQL.
- AnalyticDB for MySQL clusters store metadata for control and scheduling.
- ZooKeeper manages the configurations of AnalyticDB for MySQL modules and elects primary and secondary nodes of the modules.
- DMS provides access to the data management console.

9.8.2.2. Components and features

Online resource scheduling module

The online resource scheduling module is an online service module of Job Scheduler. This module provides cluster application, management, and scheduling for online services.

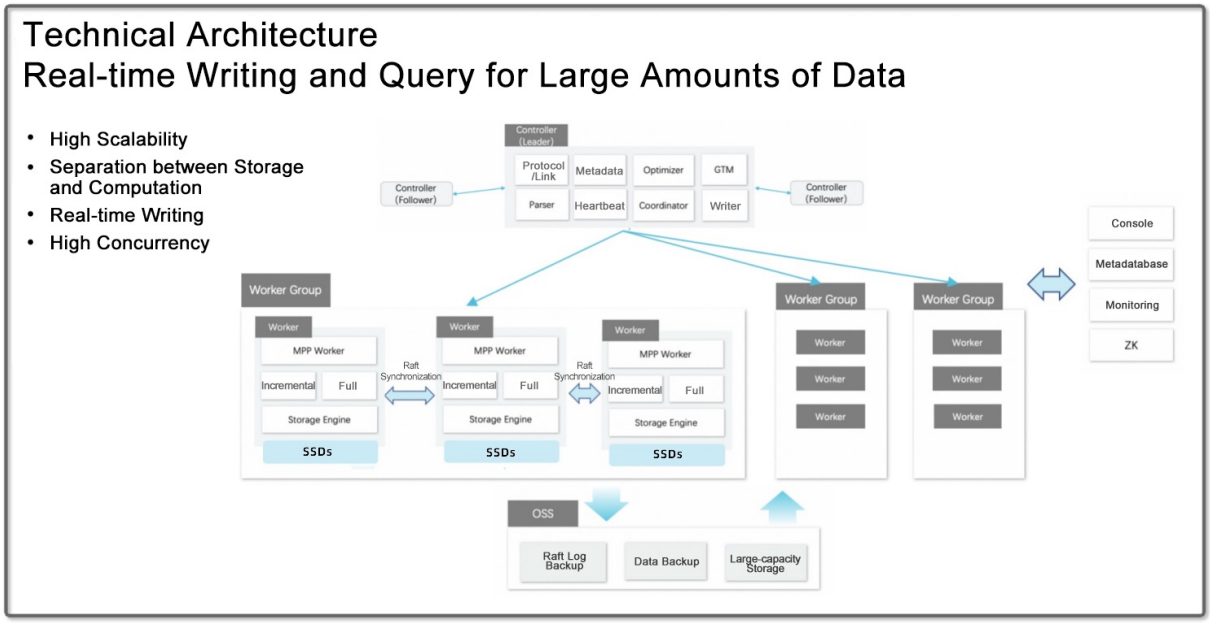

Kernel: SQL Execution

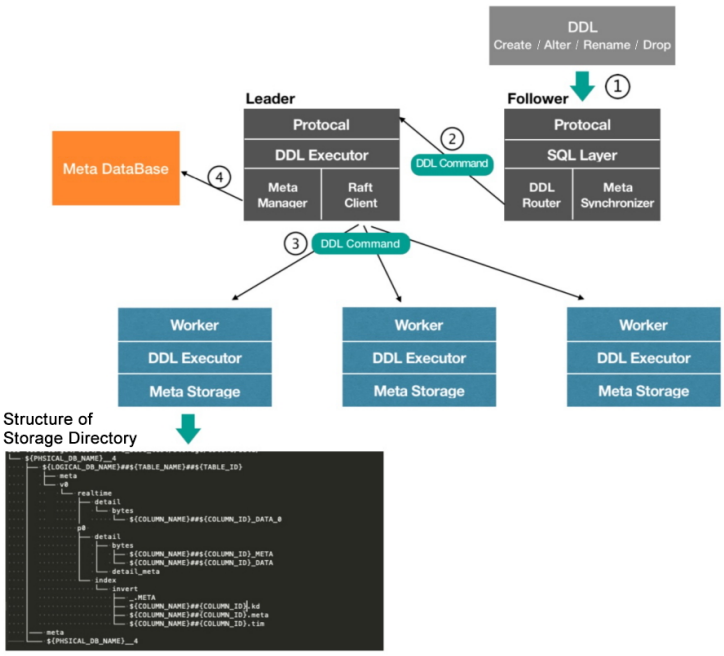

- **Cluster application**: applies for computing and storage resources from the resource pool based on cluster specifications.
- **Cluster management**: starts, stops, restarts, or deletes workers.
- **Cluster scheduling**: handles worker failover and backs up data and logs.

AnalyticDB for MySQL functional modules

AnalyticDB for MySQL functional modules enable you to query, write, modify, and batch import data.

Controller

- Authenticates database users.
- Supports JDBC and ODBC protocols. The controller proxy service supports RESTful APIs.
- Reads table schemas from ApsaraDB for RDS when you create or query tables.
- Parses SQL statements, delivers statements for other functional modules to execute, aggregates query results, and returns them to clients.
- Delivers the CREATE and DROP operations of databases and tables for the resource manager (RM) to execute.
- Reads the ZooKeeper configuration upon startup and selects the primary discovery server in MPP mode.
- Parses, plans, optimizes, and executes SELECT statements, reads indexes and data from a local file system, and returns results.
- Extracts worker-imported data or worker-created indexes from the worker cache or file system.

Worker

- Works in a mode with one primary/secondary coordinator and multiple workers.
- Processes data imported in real time. A worker processes one or more table partitions.
- Executes INSERT statements from controllers, logs INSERT operations into the file system for primary/secondary replication and disaster recovery, and pushes inserted data to the primary compute node so that data is visible in real time.
- Reads the ZooKeeper configuration upon startup and selects the primary coordinator.
- Uses MaxCompute MapReduce to import offline data and creates partitions, row groups, metadata, and indexes for imported data.
- Uses AnalyticDB for MySQL MapReduce to periodically merge the baselines of imported data and creates partitions, row groups, metadata, and indexes for imported data.

RM

- Schedules resources through Gallardo APIs to:
  - Assign or cancel the controller and worker services for new and deleted databases.
  - Start or stop services.
  - Isolate resources such as CPU and memory.
- Schedules data to allocate partition metadata for real-time tables of workers and saves the data to ApsaraDB for RDS.
- Checks system health status.
- Upgrades or rolls back the system online.
- Scales the system in or out.
- Reads the ZooKeeper configuration upon startup and selects the primary RM.

9.8.2.3. Node group specifications

Node groups are the basic unit to distribute storage and computing resources in AnalyticDB for MySQL. Node groups provide the following resources:

- CPU: the available CPU cores
- Memory: the available memory size

- Disk space: the available disk space

9.8.3. AnalyticDB for MySQL console

The AnalyticDB for MySQL console allows you to create or delete database clusters, change specifications of clusters, and manage database accounts.

9.8.3.1. Cluster management

9.8.3.1.1. Log on to the console

Enter a username and password to log on to the console, as shown in the following figure.

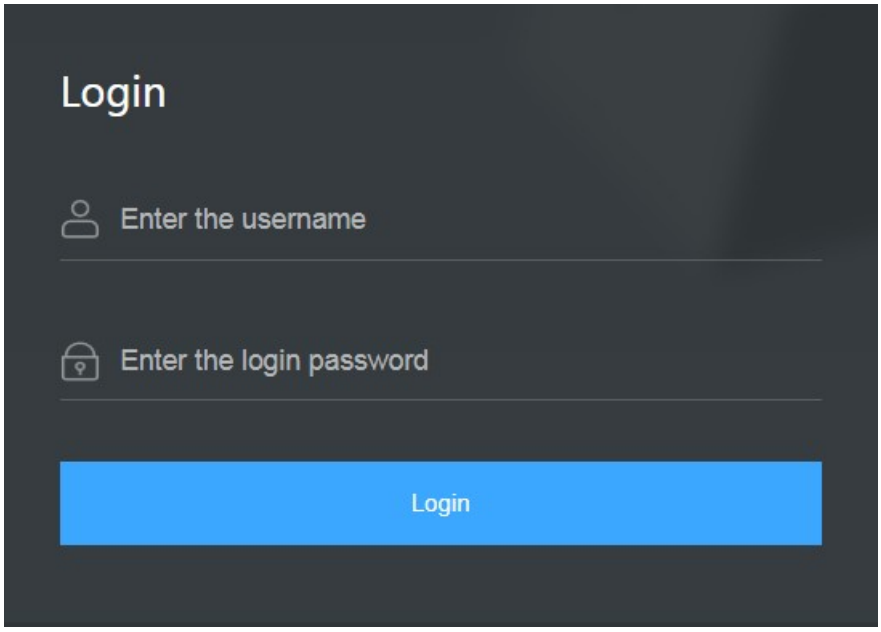

9.8.3.1.2. Manage a cluster

In the top navigation bar, choose **Products > Database Service > AnalyticDB for MySQL** to log on to the AnalyticDB for MySQL console. You can view the list of clusters and their statuses.

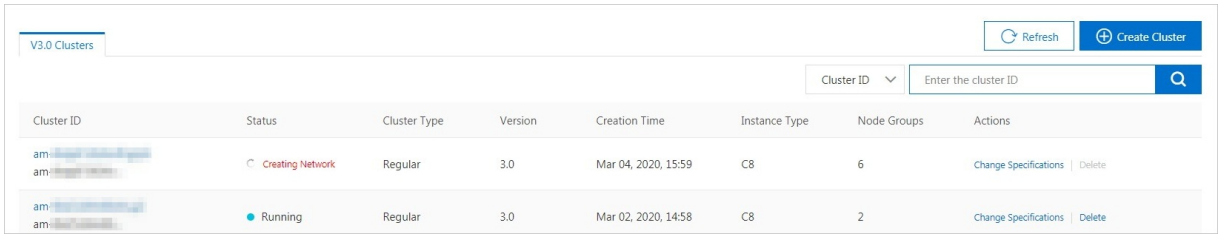

The console provides the following information:

- Clusters: lists all clusters and their statuses.
- Create a Cluster: allows you to create a database cluster.
- Actions: allows you to change specifications of a cluster or delete a cluster.

9.8.3.1.3. Create a database cluster

This topic describes how to create an AnalyticDB for MySQL cluster.

Procedure

1. Log on to the AnalyticDB for MySQL console.
2. Click **Create a Cluster** in the upper-right corner of the page. On the page that appears, configure the following parameters.

Parameter	Description
Region	The region where the cluster resides. You cannot change the region after the cluster is created. We recommend that you select a region that is closest to the geographic area of your business to improve access speed and stability.
Zone	An independent physical area within a region. All zones in a region provide the same services.
Organization	The organization to which the cluster belongs.
Resource Set	The resource set of the cluster.
Version	Only version **3.0** is supported.
Edition	The edition of the cluster. Only **Basic** is supported.
Network Type	AnalyticDB for MySQL supports two types of networks. ○ *VPC*: A VPC helps you build an isolated network environment in Apsara Stack. In a VPC, you can customize the route table, CIDR blocks, and gateway. We recommend that you select VPC for higher security. ○ *Classic Network*: Cloud services on a classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service.
Specifications	The ECU specifications.
Node Groups	The number of node groups. By default, each node group consists of three replicas.
Storage	The storage space of a node group.

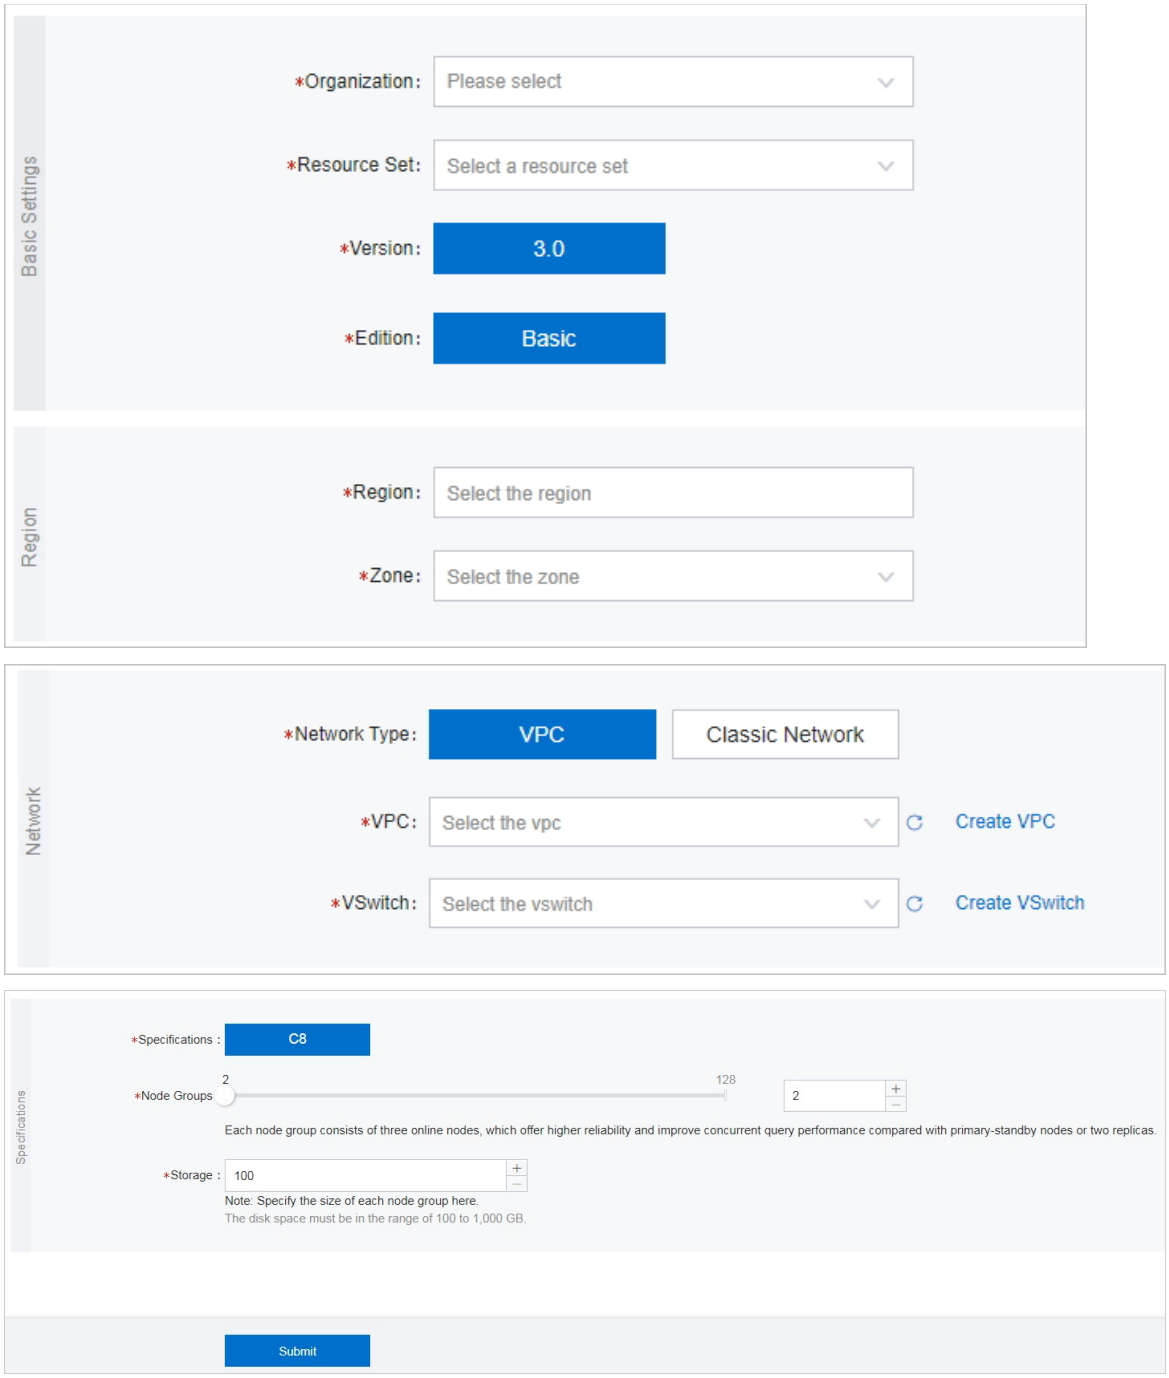

3. After you configure the preceding parameters, click **Submit**.

9.8.3.1.4. View monitoring information

You can view monitoring information of a cluster in real time from the AnalyticDB for MySQL console.

Procedure

1. Log on to the AnalyticDB for MySQL console.
2. In the upper-left corner of the page, select the region where the cluster resides.
3. On the Clusters page, click **Cluster ID** corresponding to the cluster for which you want to view the monitoring information.

4. In the left-side navigation pane, click **Monitoring Information** to view monitoring information of the cluster.

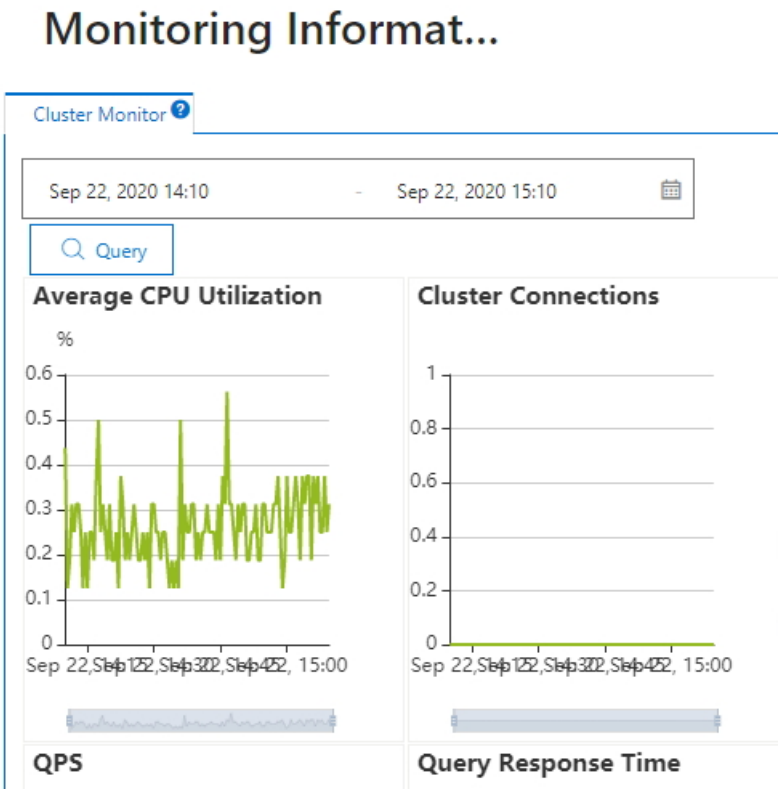

Monitoring information includes the CPU utilization, cluster connections, QPS, and query response time.

9.8.3.2. Account management

9.8.3.2.1. Create a database account

This topic describes types of database accounts and how to create accounts in AnalyticDB for MySQL.

Database account types

AnalyticDB for MySQL provides two types of database accounts: privileged accounts and standard accounts.

Database account types

Database account type	Description

Database account type	Description
Privileged account	• You can create and manage privileged accounts in the AnalyticDB for MySQL console only. • You can create only one privileged account for a cluster. Privileged accounts have permissions to manage all standard accounts and databases within a cluster. • You can use the privileged account to disconnect standard accounts from AnalyticDB for MySQL. You can use the privileged account to manage fine-grained permissions to suit your business needs. For example, you can grant each standard account permissions to query specific tables. • A privilege account in AnalyticDB for MySQL is equivalent to a root account in MySQL.
Standard account	• You can use only SQL statements to create and manage standard accounts. • You can create up to 256 standard accounts for a cluster. • You must manually grant specific database permissions to standard accounts. • You cannot use a standard account to disconnect other accounts from the cluster.

Create a privileged account

1. Log on to the AnalyticDB for MySQL console.
2. In the upper-left corner of the page, select the region where the cluster resides.
3. On the Clusters page, click Cluster ID corresponding to the cluster for which you want to create a privileged account. In the left-side navigation pane, click **Accounts**. On the **Accounts** page, click **Create Account**.
4. In the **Create Account** pane, configure the following parameters.

Parameter	Description
Account	The name of the privileged account. The name must be 2 to 16 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter and end with a lowercase letter or digit.
Account Type	The type of the account. The value is Privileged Account. This value cannot be changed.

Parameter	Description
Password	The password of the privileged account. The password must be 8 to 32 characters in length and contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include ! @ # $ % ^ & * ( ) _ + - =
Confirm Password	Re-enter the password of the privileged account.
Description	Optional. The description of the database account.

5. Click **OK**.

Create a standard account

For information about how to create a standard account and grant permissions, see `CREATE USER` in *User Guide*.

9.8.3.2.2. Create a database account and grant permissions

You can use SQL statements to create database accounts and grant permissions.

- For information about how to create a RAM user, see `CREATE USER` in *User Guide*.
- For information about how to grant permissions to a RAM user, see `GRANT` in *User Guide*.
- For information about how to revoke permissions from a RAM user, see `REVOKE` in *User Guide*.
- For information about how to change an account name, see `RENAME USER` in *User Guide*.
- For information about how to delete a user, see `DROP USER` in *User Guide*.

9.8.4. Security maintenance

9.8.4.1. Network security maintenance

Network security maintenance helps you ensure device and network security.

Device security

Check network devices and enable security management protocols and configurations of devices.

Check frequently for up-to-date versions of network device software and update to more secure versions in a timely manner.

For more information about the security maintenance method, see the product documentation of each device.

Network security

You can select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check public and internal traffic, and defend the network against abnormal behaviors and attacks.

9.8.4.2. Account password maintenance

Account passwords include AnalyticDB for MySQL system and device passwords.

To ensure account security, you must use complex passwords for your systems and devices and change these passwords on a regular basis.

9.8.5. Troubleshooting

9.8.5.1. Fault emergency mechanism

A fault emergency mechanism must be established to minimize the impact on businesses.

Response mechanism

A fault emergency response mechanism must be established, so that you can contact the relevant maintenance personnel and rectify the faults as soon as possible.

Stock-up mechanism

A stock-up mechanism must be established for fragile hardware devices, so that faulty ones that cannot be repaired can be replaced in a short period of time.

9.8.5.2. Stock-up mechanism

A stock-up mechanism must be established for fragile hardware devices, so that hardware faults can be rectified quickly.

9.8.5.3. Troubleshooting methods

Faults must be promptly located and rectified based on their symptoms. This topic describes the troubleshooting procedure.

When you detect a system fault during routine maintenance, you can use ApsaraDB for RDS Operations and Maintenance System to check the fault details, analyze its causes, and rectify the fault based on fault logs.

If the fault cannot be rectified, collect fault information (such as system information and fault symptoms), contact Alibaba Cloud technical support engineers, and troubleshoot the fault under the guidance of the engineers.

After the fault is rectified, verify the solution, review the troubleshooting method, and make improvements.

9.8.5.4. Common failure troubleshooting

9.8.5.4.1. Insufficient disk space

If disk space is insufficient, some operations may fail to be performed.

Cause: The monitoring logs are too large in size.

Solution: Monitoring logs are usually stored in the */usr/local/rds/log* directory. You can delete earlier logs to free up space.

9.8.5.4.2. Insufficient swap space

Cause

Too many clusters exist or programs run abnormally. As a result, too much memory is consumed.

Solution

- If too many clusters run on the host, we recommend that you migrate some of the clusters to another host.
- If this fault is caused by other reasons, perform the following steps:
  i. If the fault exists in the primary database, fail services over to the secondary database.
  ii. Run the following command to release the cache:

  ```
  echo 3 > /proc/sys/vm/drop_caches;
  ```

  iii. Run the following command to clear the swap space:

  ```
  swapoff -a ; date ; swapon –a;
  ```

9.8.5.4.3. Overhigh load

Cause

Data is being backed up, too many programs are running concurrently, or the operating efficiency of SQL is low.

Solution

- If data is being backed up, ignore this fault.
- If too many programs run concurrently, contact R&D engineers to reduce the amount.
- If the operating efficiency of SQL is low, run the explain command to view the usage of indexes. If you cannot use indexes to optimize the operating efficiency of SQL, contact R&D engineers.
- Decrease the limit on CPU and I/O resources.

9.8.5.4.4. Copy latency

Cause

The system is backing up data, no primary key is contained in the tables, an I/O bottleneck exists, or the network is experiencing heavy traffic.

Solution

- If the latency is less than 5,000 milliseconds and continues to decrease, the network may be experiencing slowdown due to heavy traffic. No action is required.
- Enable the concurrent handling function.

- Log on to the cluster and run the following commands:

```
stop slave;
```

```
set global slave_parallel_workers=8;
```

```
start slave;
```

9.8.5.4.5. Process exceptions

Cause

The process does not automatically start or unexpectedly quits. An alert in Cloud Monitor shows the process name.

Solution

View the process logs. If the process does not automatically start, run the following command to start the process:

```
/usr/local/rds/$Program/package/service.sh start
```

9.8.5.4.6. Module exceptions

Cause

A module is not loaded.

Solution

Run one of the following commands to load the module:

```
insmod $module file.ko
```

```
modprode $module file.ko
```

9.8.5.5. Hardware troubleshooting

9.8.5.5.1. Disk failure

RAID 5 is configured on database nodes. Faulty disks can be directly replaced without affecting your business.

9.9. AnalyticDB for PostgreSQL

9.9.1. Overview

Purpose

This guide summarizes possible problems that you may encounter during O&M operations and provides solutions for you.

If you encounter system problems not covered in this guide, you can submit a ticket to Alibaba Cloud for technical support.

Requirements

You must possess IT skills including computer network knowledge, computer operation knowledge, problem analysis, and troubleshooting.

Additionally, you must pass the pre-job training of the Alibaba Cloud system to learn necessary Alibaba Cloud system knowledge, including but not limited to system principles, networking, features, and the use of maintenance tools.

Note that during maintenance operations, you must comply with operating procedures to ensure personal and system security. User data must be kept strictly confidential and must not be copied or disseminated without the written consent of the users.

Precautions

To ensure a stable system and avoid unexpected events, you must follow the following guidelines.

- Hierarchical permission management

  Permissions on networks, devices, systems, and data are granted based on the services and roles of the O&M personnel to prevent system faults caused by unauthorized operations.

- System security

  Before performing any system operations, you must be aware of their impacts.

  You must record all problems encountered during operations for problem analysis and troubleshooting.

- Personal and data security
  - You must take safety measures in accordance with the device manuals when operating electrical equipment.
  - You must use secure devices to access the business network.
  - Unauthorized data replication and dissemination are prohibited.

Support

You can contact Alibaba Cloud technical support for help.

9.9.2. Architecture

This topic describes the architecture of AnalyticDB for PostgreSQL.

Physical architecture of a cluster

The following figure shows the physical architecture of an AnalyticDB for PostgreSQL cluster.

Physical cluster architecture

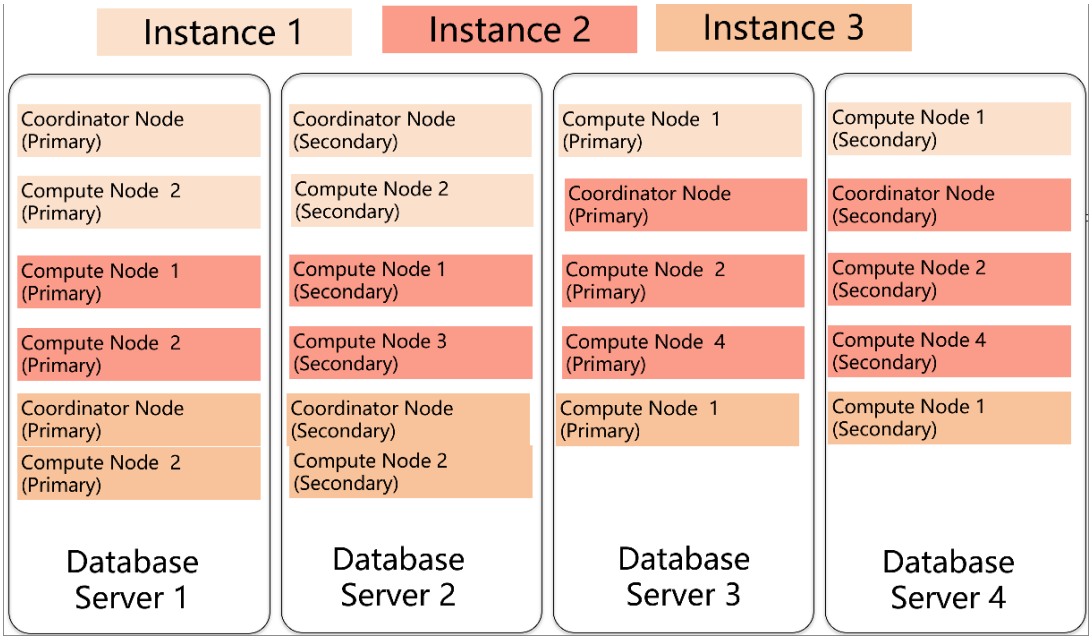

You can create multiple instances in a physical cluster of AnalyticDB for PostgreSQL by using the management and control system. Each instance consists of a coordinator node and multiple compute nodes.

- The coordinator node is used for access from applications. It receives connection requests and SQL query requests from clients and dispatches computing tasks to compute nodes. The cluster deploys a secondary node of the coordinator node on an independent physical server and replicates data from the primary node to the secondary node for failover. The secondary node does not accept external connections.
- Compute nodes are independent instances in AnalyticDB for PostgreSQL. Data is evenly distributed across compute nodes by hash value or RANDOM function, and is analyzed and computed in parallel. Each compute node uses a primary/secondary architecture for automatic failover.

Logical architecture of an instance

You can create multiple instances within an AnalyticDB for PostgreSQL cluster. The following figure shows the logical architecture of an AnalyticDB for PostgreSQL instance.

Logical architecture of an instance

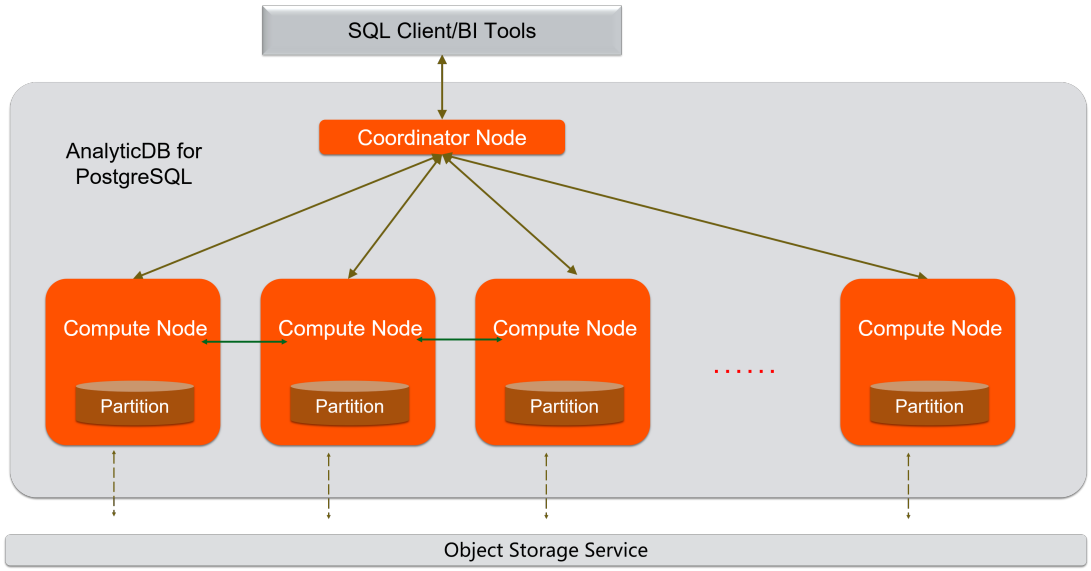

Data is distributed across compute nodes by hash value or RANDOM function of a specified distribution column. Each compute node uses a primary/secondary architecture to ensure dual-copy storage. High-performance network communication is supported across nodes. When the coordinator node receives a request from an application, the coordinator node parses and optimizes SQL statements to generate a distributed execution plan. After the coordinator node sends the execution plan to the compute nodes, the compute nodes perform massively parallel processing of the plan.

9.9.3. Routine maintenance

9.9.3.1. Check for data skew on a regular basis

You must check for data skew on a regular basis during maintenance to prevent the instance from being read-only due to excessive data in some compute nodes.

You can use the following methods to locate data skew. The procedure is as follows.

1. For a single table or database, you can view the space occupied within each compute node to determine whether data has been skewed.
   i. Execute the following statement to determine whether the data in a database has been skewed:

   ```
   SELECT pg_size_pretty(pg_database_size('postgres')) FROM gp_dist_random('gp_id');
   ```

   You can view the space occupied by the dbname database in each compute node after the statement is executed. If the space occupied in one or more compute nodes is significantly greater than that of other compute nodes, it indicates the data in this database is skewed.

   ii. Execute the following statement to determine whether the data in a table has been skewed:

   ```
   SELECT pg_size_pretty(pg_relation_size('tblname')) FROM gp_dist_random('gp_id');
   ```

   Using the preceding statement, you can view the space occupied by the tblname table within each compute node after the statement is executed. If the space occupied within one or more compute nodes is significantly greater than that of other compute nodes, it indicates the data in this table is skewed. You must modify the partition key to redistribute the data.

2. You can use the system views to determine whether data has been skewed.
   i. Execute the following statement to check whether the storage space is skewed. The principle of this method is similar to that of the preceding space-viewing method:

      ```
      SELECT * FROM gp_toolkit.gp_skew_coefficients
      ```

      You can use the view to check the data volume of rows in a table. The larger the table, the more time it will take for the check to complete.

   ii. Use the gp_toolkit.gp_skew_idle_fractions view to calculate the percentage of idle system resources during a table scan to check whether the data is skewed:

      ```
      SELECT * FROM gp_toolkit.gp_skew_idle_fractions
      ```

      For more information, see Checking for Uneven Data Distribution.

9.9.3.2. Execute VACUUM and ANALYZE statements

You can execute `VACUUM` and `ANALYZE` statements on a regular basis for frequently updated tables and databases. You can also execute `VACUUM` and `ANALYZE` statements after you have performed a large number of update or write operations to prevent the operations from consuming excessive resources and storage space.

9.9.4. Security maintenance

9.9.4.1. Network security maintenance

Regular maintenance will help ensure the security of networks and devices.

Device security

Check network devices and enable the security management protocols and configurations for the devices you want to secure. Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner. For more information about security maintenance methods, see the product documentation of each device.

Network security

You can select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check public and internal traffic, and defend the network against abnormal behaviors and attacks.

9.9.4.2. Account password maintenance

Account passwords include the superuser password of AnalyticDB for PostgreSQL and the password of the host operating system.

To ensure account security, use complex passwords and periodically change the passwords of systems and devices.

9.10. KVStore for Redis

9.10.1. O&M tool

The Apsara Stack Operation console provides the following operations and maintenance (O&M) features for KVStore for Redis:

- Instance management: allows you to view instance details, instance logs, and user information.
- Host management: allows you to view and manage hosts.

9.10.2. Architecture diagram

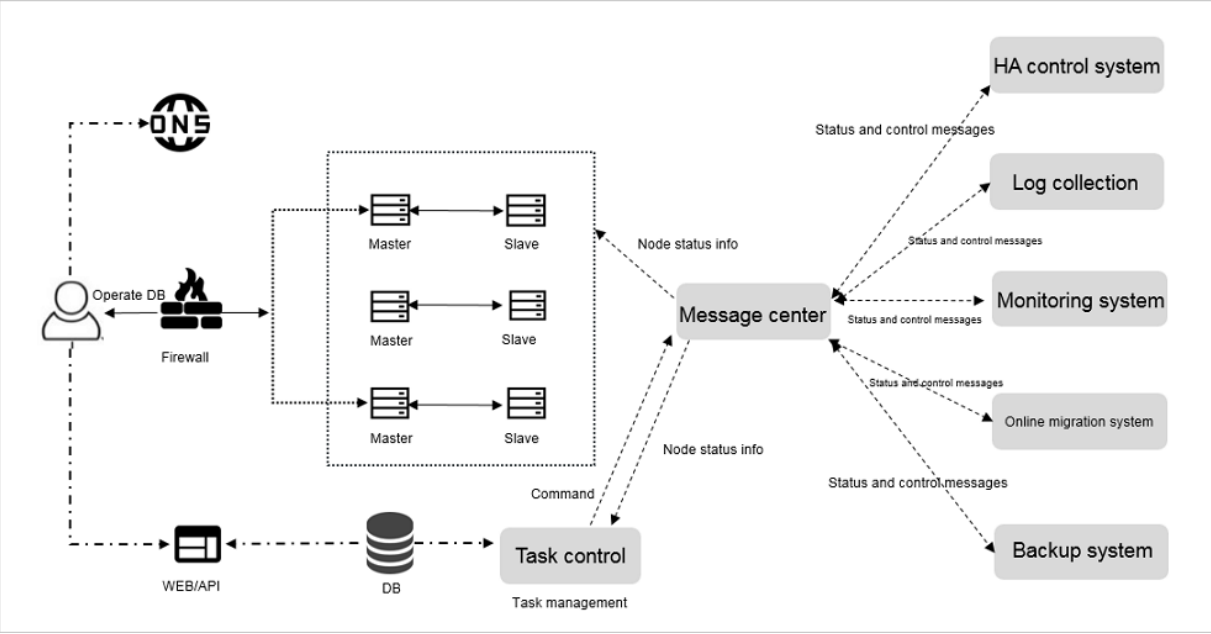

9.10.3. Log on to Apsara Stack Operations

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

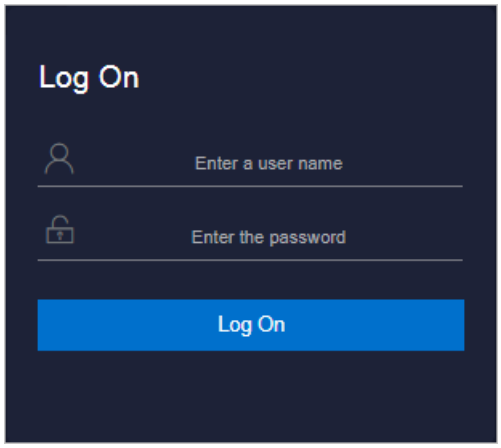

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.

9.10.4. Instance management

You can view instance details, logs, and user information.

Procedure

1. Log on to Apsara Stack Operations.
2. In the left-side navigation pane, choose Products > RDS to go to the **RDS** page. Click the **Instance Management** tab. On the Instance Management tab, you can perform these operations:
   - View the list of instances.

     On the **Instance Management** tab, you can view the instances under your account.
   - View the details of an instance.

     Click the ID of a target instance to view the details of the instance.
   - **View user information.**

Click **User Information** in the **Actions** column.

9.10.5. Host management

Host management allows you to view and manage hosts.

Procedure

1. Log on to Apsara Stack Operations.
2. In the left-side navigation pane, choose Products > RDS to go to the **RDS** page. Click the **Host Management** tab to view the information about all hosts.

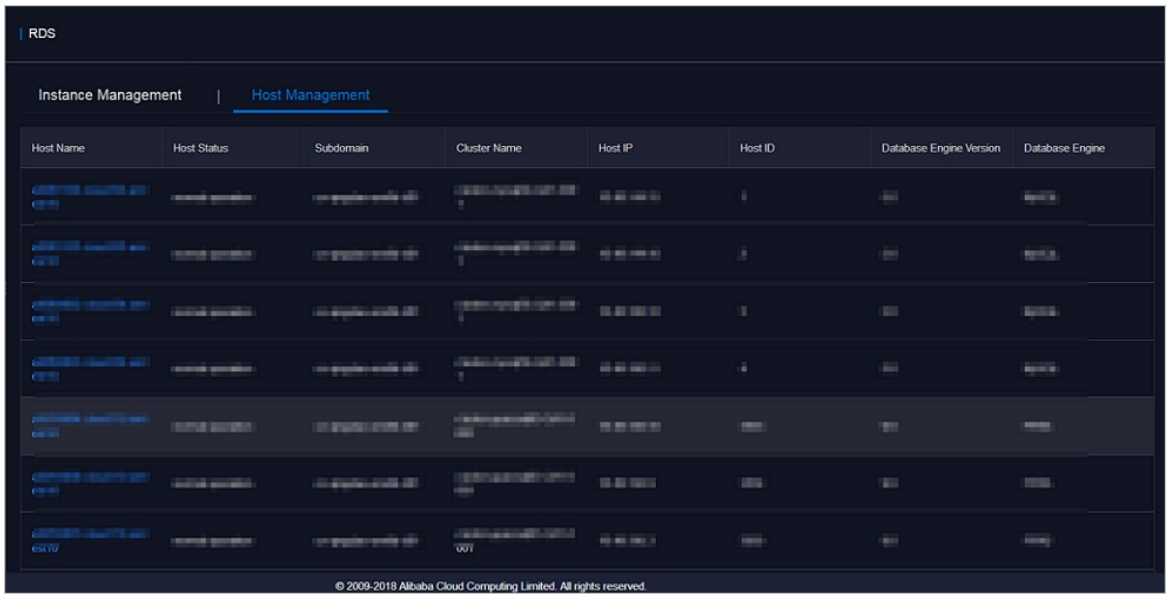

3. Click a host name to go to the **RDS Instance** page. You can view all instances on this host.

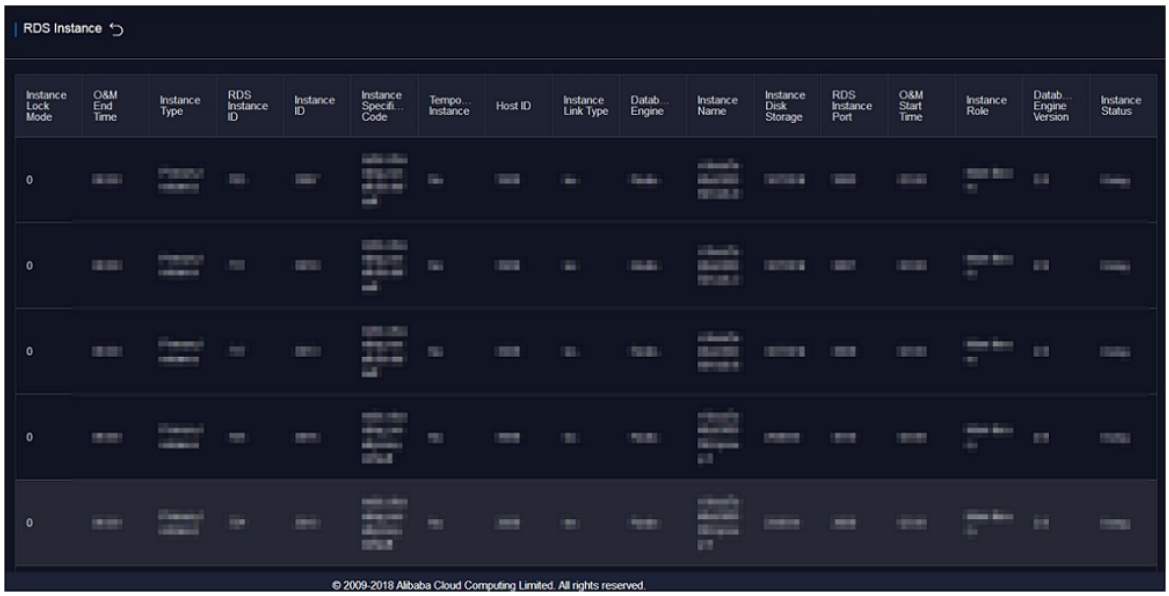

9.10.6. Security maintenance

9.10.6.1. Network security maintenance

Network security maintenance involves device security and network security.

Device security

Check network devices, and enable security management protocols and configurations for these devices.

Check software versions of network devices and update them to more secure versions in time.

For more information about security maintenance methods, see documents of related devices.

Network security

Based on your network conditions, select the intrusion detection system (IDS) or intrusion prevention system (IPS) to detect abnormal Internet and intranet traffic and protect against abnormal behavior and attacks in real time.

9.10.6.2. Password maintenance

Passwords include system passwords and device passwords in KVStore for Redis.

To secure your account, you must periodically change the system and device passwords, and use complex passwords.

9.11. ApsaraDB for MongoDB

9.11.1. Service architecture

9.11.1.1. System architecture

9.11.1.1.1. Backup system

Automatic backup

ApsaraDB for MongoDB supports both physical backup and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can initiate a temporary backup as required. The backup files are retained for seven days.

Log management

ApsaraDB for MongoDB generates operation logs and allows you to download them. You can use the operation logs for local incremental backup.

Data backtracking

ApsaraDB for MongoDB can use backup files and logs to generate a temporary instance for any time point within the past seven days. After verifying that the data in the temporary instance is correct, you can use the temporary instance to restore data to the specified time point.

Creating a temporary instance does not affect the running of the current instance.

Only one temporary instance can be created for each ApsaraDB for MongoDB instance at a time. A temporary instance is valid for 48 hours. You can create a maximum of 10 temporary instances for an ApsaraDB for MongoDB instance each day.

9.11.1.1.2. Data migration system

Database replication between instances

ApsaraDB for MongoDB allows you to easily migrate databases from one instance to another.

Data migration to or from ApsaraDB for MongoDB

ApsaraDB for MongoDB provides a professional tool and a migration wizard to help you migrate data to or from ApsaraDB for MongoDB.

Backup file download

ApsaraDB for MongoDB retains backup files for seven days. During this period, you can log on to the ApsaraDB for MongoDB console to download the backup files.

9.11.1.1.3. Monitoring system

Performance monitoring

ApsaraDB for MongoDB provides nearly 20 metrics for monitoring system performance, such as the disk capacity, IOPS, number of connections, CPU utilization, network traffic, transactions per second (TPS), queries per second (QPS), and cache hit rate. You can obtain such status information for an ApsaraDB for MongoDB instance within the past one year.

SQL auditing

The system records SQL statements and additional information sent to ApsaraDB for MongoDB instances, such as the IP addresses of connections, database names, access accounts, execution time, and number of records returned. You can use SQL auditing to locate problems and check instance security.

Threshold alerting

ApsaraDB for MongoDB provides short message service (SMS) notifications to indicate status or performance exceptions that occur in ApsaraDB for MongoDB instances.

These exceptions include instance locking, disk capacity, IOPS, connection quantity, and CPU exceptions. You can configure alert thresholds and up to 50 alert recipients (of which five are effective at a time). If a metric of an ApsaraDB for MongoDB instance exceeds a specific threshold, an SMS notification is sent to alert the recipients.

Web operation logging

The system logs all modification operations in the ApsaraDB for MongoDB console for administrators to check. These logs are retained for a maximum of 30 days.

9.11.1.1.4. Control system

If a host or an instance crashes, the ApsaraDB for MongoDB high-availability (HA) component fails services over within 30 seconds after the exception is detected. This guarantees that applications run properly and ApsaraDB for MongoDB is highly available.

9.11.1.1.5. Task scheduling system

You can use the ApsaraDB for MongoDB console or APIs to create or delete instances or switch instances between the intranet and Internet. All instance operations are scheduled, traced, and displayed as tasks.

9.11.2. ApsaraDB for MongoDB O&M overview

Apsara Stack Operations Console provides the following O&M features for ApsaraDB for MongoDB:

- Instance management: allows you to view instance details, instance logs, and user information.
- Host management: allows you to view and manage hosts.

9.11.3. Log on to the Apsara Stack Operations console

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   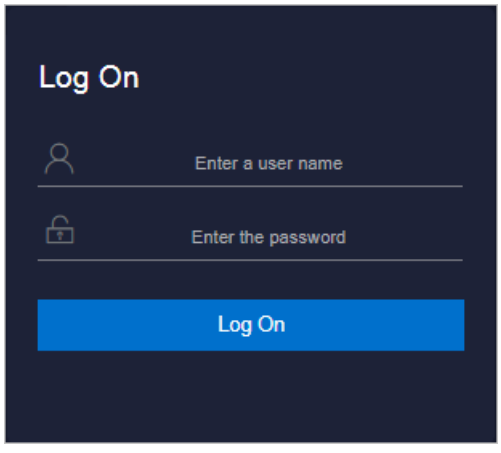

   Note You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.

9.11.4. Manage ApsaraDB for MongoDB instances

This topic describes how to manage ApsaraDB for MongoDB instances. You can view instance details, logs, and user information.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose Products > RDS to go to the **RDS** page. On the **Instance Management** tab, you can perform the following operations:
   - View the list of instances.

     View the instances that belong to the current account on the **Instance Management** tab, as shown in Instance list.

     Instance list

     

   - View the details of an instance.

Click the ID of an instance to view its details, as shown in Instance details.

Instance details

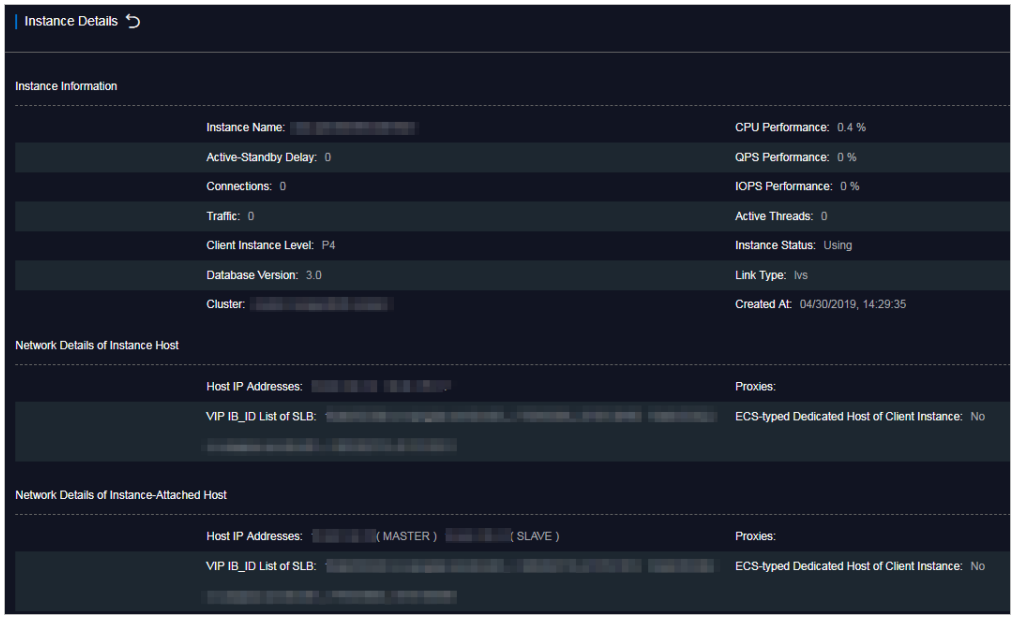

> **Note** On the **Instance Details** page, you can also perform primary/secondary switchovers and query historical operations.

- **View user information.**

  Click **User Information** in the **Actions** column, as shown in User information.

  User information

9.11.5. Host management

Host management allows you to view and manage hosts.

Procedure

1. Log on to the Apsara Stack Operations console.
2. On the **Host Management** tab of the **RDS** page, view information about all hosts.

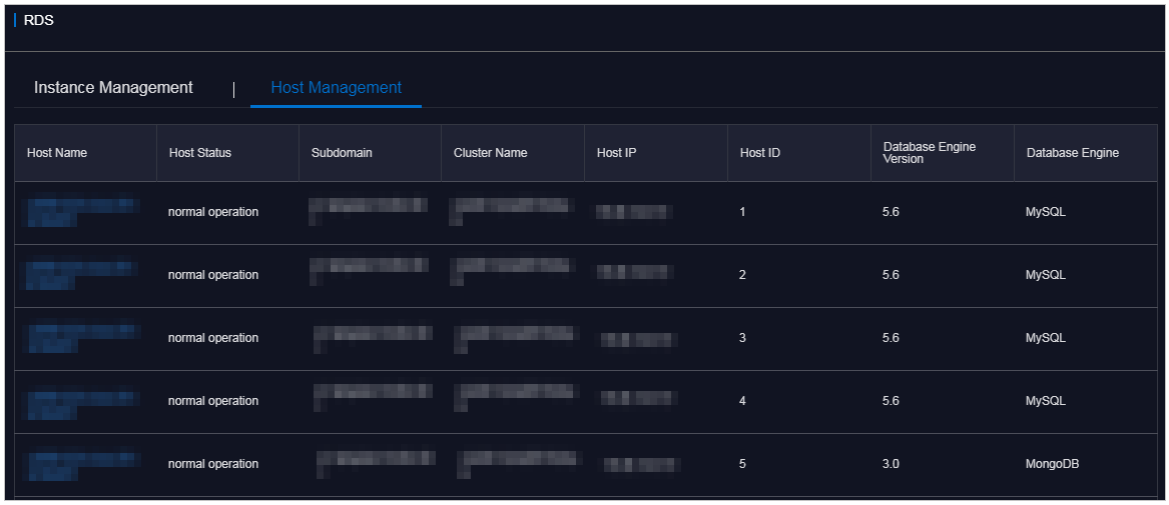

3. Click a host name to go to the **RDS Instance** page. On this page, you can view all instances on this host.

9.11.6. Security maintenance

9.11.6.1. Network security maintenance

Network security maintenance is aimed at ensuring device security and network security.

Device security

Check network devices, and enable security management protocols and configurations of devices.

Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner.

For more information about the security maintenance method, see the product document of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check Internet and intranet traffic and defend the network against abnormal behaviors and attacks.

9.11.6.2. Account password maintenance

Account passwords include the ApsaraDB for MongoDB system and device passwords.

To ensure account security, change the system and device passwords periodically, and use passwords that meet the complexity requirements.

9.12. ApsaraDB for OceanBase

9.12.1. Overview

Purpose

This document is a guide for you to perform operations and maintenance (O&M) tasks such as routine inspections, monitoring, and maintenance on ApsaraDB for OceanBase clusters. These tasks ensure the long-term stable running of the system.

You can follow the instructions in this guide to handle the issues that are identified during the maintenance. If you encounter system issues that are not covered in this guide, contact technical support.

Requirements

You must acquire IT skills, including knowledge for computer networks, knowledge for computer operations, issue analysis, and troubleshooting. You must pass the pre-job training to learn the knowledge for the Apsara Stack system. The required knowledge for the system includes, but is not limited to, system principles, networking, features, and the usage of maintenance tools.

Note that during maintenance, you must comply with operation procedures to ensure personal safety and system security. User data must be kept strictly confidential and must not be copied or disseminated without the written consent of the users.

Introduction to O&M commands and tools

Apsara Stack Operations console for ApsaraDB for OceanBase

The Apsara Stack Operations console for ApsaraDB for OceanBase is a cloud management platform for ApsaraDB for OceanBase. ApsaraDB for OceanBase is a financial-grade distributed relational database service. The console provides a wide range of modules that allow you to implement various features. For example, you can use the modules to manage resources, capacities, clusters, and instance lifecycles. You can also use the modules to monitor performance based on real-time computing and to implement API-related features. These modules provide ApsaraDB for OceanBase cloud services and simplify O&M operations.

Considerations

To ensure a stable system and avoid unexpected events, you must follow these guidelines:

- Hierarchical permission management

  O&M engineers are granted with only the network, device, system, and data permissions that are required to fulfill their duties. This prevents system faults that are caused by unauthorized operations.

- System security

  Before you perform operations on the system, you must be aware of the impacts of the operations on the system. This ensures that the system is not affected by the high-risk operations. You must record the details about the issues that you encounter during the operations for issue analysis and troubleshooting.

- Personal safety and data security
  - You must take safety measures to ensure personal safety based on device manuals when you use electrical equipment.
  - You must use secure devices to access the business network.
  - Unauthorized data replication and dissemination are prohibited.

Technical support

You can contact technical support for help during the maintenance.

9.12.2. Architecture

9.12.2.1. System architecture

Cluster management

You can view the status, the details, and the operation logs of each server in the current ApsaraDB for OceanBase cluster. You can also upgrade, delete, scale out, and restart a cluster or the servers in the cluster.

System monitoring

In the Apsara Stack Operations console for ApsaraDB for OceanBase, you can view monitoring information based on various dimensions, such as clusters, data centers, servers, and tenants.

- Performance monitoring

  In the Apsara Stack Operations console for ApsaraDB for OceanBase, you can view the monitoring information about more than 50 system performance metrics. For example, you can monitor the following performance metrics: transactions per second (TPS), queries per second (QPS), response time (RT), CPU utilization, input/output operations per second (IOPS), memory usage, disk usage, cache hit ratio, and network traffic. You can also view the instance status information that is collected in the last year. Note that one year is the longest allowed time duration during which the status information can be displayed. You can shorten the time duration based on your business requirements.

- SQL auditing

  ApsaraDB for OceanBase records the Structured Query Language (SQL) statements that are executed by each tenant and the statement details. For example, you can view the following details by checking the values of relevant parameters: the number of executions, RT, and execution period. Based on the recorded details, you can locate issues or perform troubleshooting to ensure tenant security.

- Threshold alerting

  If the status of a server in a cluster is abnormal or the tenant status of a cluster is abnormal, the Apsara Stack Operations console for ApsaraDB for OceanBase sends alerts by using SMS messages or phone calls. The exceptions may occur for the following performance metrics: disk capacity, memory usage, CPU utilization, major freeze status, Network Time Protocol (NTP) status, and latency. The exceptions may include the errors that are recorded in error logs. You can customize alert thresholds for clusters or tenants based on your business requirements. If alert conditions are met, the specified alert contacts receive the SMS messages or phone calls from the system.

- Check log records

You can check the log records of different OBServers in clusters for troubleshooting.

System backups

ApsaraDB for OceanBase supports full backups and incremental backups. A full backup is to back up all the data that is stored in disks. An incremental backup is to back up the data that is added, deleted, and modified in memory in real time. The incremental backup data is the data that is recorded in the commit logs of tenants.

You can perform full backups and incremental backups on your clusters. If errors occur on your clusters, you can use the backups to restore your data on other clusters.

Version management

The Apsara Stack Operations console for ApsaraDB for OceanBase allows you to upload the installation packages of multiple ApsaraDB for OceanBase versions. When you create a cluster in the console, you can choose an installation package for the cluster based on your business requirements. You can log on to the console to add or delete the installation packages of ApsaraDB for OceanBase. The console also allows you to manage the installation packages of multiple OBProxy versions. You can use each of the installation package to automatically deploy the OBProxy on the proxy server. You can log on to the console to add or delete OBProxy installation packages.

9.12.2.2. Deployment solutions

9.12.2.2.1. Add ApsaraDB for OceanBase RPM packages

In the Apsara Stack Operations console for ApsaraDB for OceanBase, you can add an RPM Package Manager (RPM) package of the current ApsaraDB for OceanBase version.

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, click **O&M**. In the list that appears, select **RPM Package Management**.
3. On the RPM Package Management page, click the **Database** tab. On the tab, click **Upload RPM Package**.
4. In the Upload RPM Package dialog box, select observer from the File Type drop-down list and click **Upload**. Then, select the RPM package that you want to upload from the local directory.
5. Click **OK**.

9.12.2.2.2. Create clusters

Before you use ApsaraDB for OceanBase, create clusters in the Apsara Stack Operations console for ApsaraDB for OceanBase.

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, click **Clusters**.
3. On the page that appears, click **Create Cluster**. In the **Create Cluster** dialog box, configure the parameters as prompted. For more information about the parameters, see Parameters for creating

a cluster.

Parameters for creating a cluster

Parameter	Description
Cluster Name	The name of the ApsaraDB for OceanBase cluster.
OceanBase Version	The version of the ApsaraDB for OceanBase cluster. We recommend that you select the latest version.
Cluster Group	The group to which the cluster belongs. You can enter the name of the cluster group.
Cluster Model	The cluster model. If the minimal specifications are used for ApsaraDB for OceanBase, the cluster model is 1-1-1. You can set this parameter to 2-2-2 or 3-3-3 based on your cluster model. If five replicas are deployed, set this parameter to 2-2-2-2-2. To obtain the cluster model, contact the owner of the ApsaraDB for OceanBase project.
Data Center & Model Distribution	The data center and the machine model. Select the data center and the machine model based on the specified cluster model.

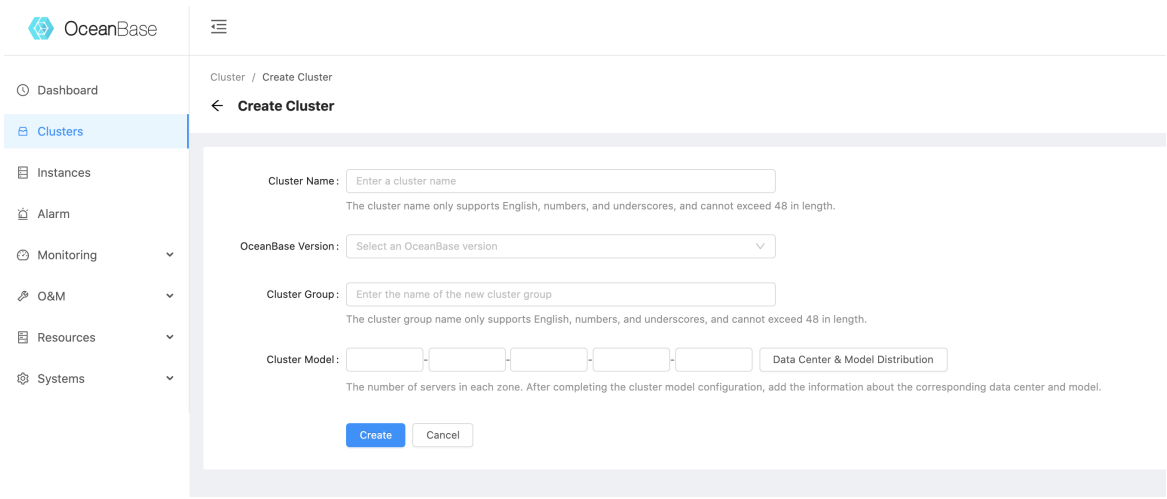

4. After you specify the preceding parameters, click **Create**.

9.12.2.2.3. Add OBProxy RPM packages

In the Apsara Stack Operations console for ApsaraDB for OceanBase, you can add an RPM package of the OBProxy that matches the current ApsaraDB for OceanBase version.

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, click **O&M**. In the list that appears, select **RPM Package Management**.
3. On the RPM Package Management page, click the **Database** tab. On the tab, click **Upload RPM Package**.
4. In the dialog box that appears, select obproxy from the File Type drop-down list, and click

Upload. Then, select the RPM package that you want to upload.

5. Click **OK**.

9.12.2.2.4. Install the OBProxy

After you add an OBProxy RPM package, you can install the OBProxy.

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, click **O&M**. In the list that appears, select **OBProxy**.
3. On the **Servers** tab of the **OBProxy** page, find the IP address of the OBProxy server to be installed. In the Actions column for the IP address, click **Install**.
4. In the dialog box that appears, set OBProxy Name to the name of the OBProxy based on your project name. From the OBProxy Version drop-down list, select an OBProxy version, such as obproxy-1.5.0-1410335.el7.x86_64.rpm. You can select the current time for Start Time.
5. Click **OK**.

9.12.2.3. OCP V2.0 components and their features

9.12.2.3.1. Components and their features

ApsaraDB for OceanBase consists of database nodes and a management node. The database nodes are OBServers that are deployed on multiple physical servers. The management node is OceanBase Cloud Platform (OCP). The management components of OCP V2.0 are deployed as Docker containers on physical servers. You can obtain the information about each server from the Apsara Infrastructure Management Framework portal, such as hosts and IP addresses.

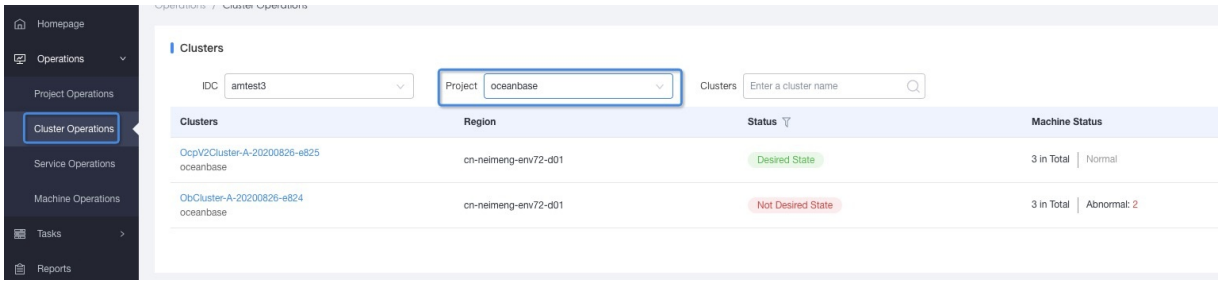

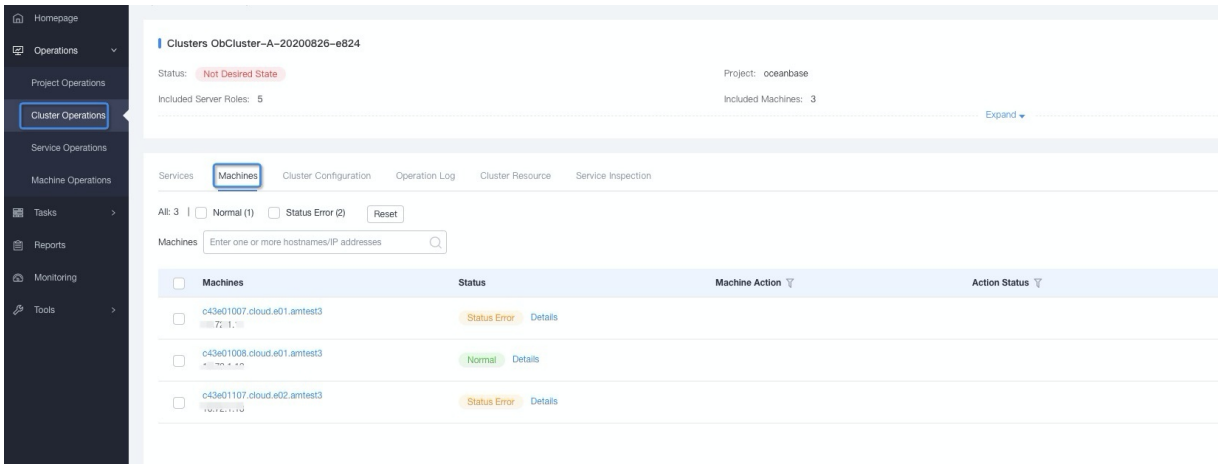

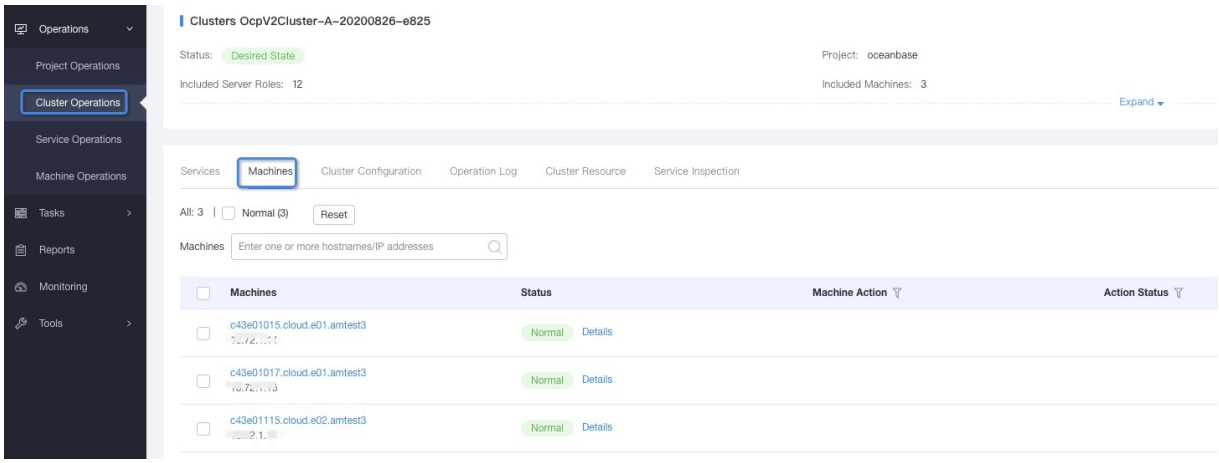

OCP V2.0 consists of the testimage component and the following feature components: OcpMetaServer, OcpMetaInit, OcpObproxy, OcpApiV2, OcpOdc, and OcpTengine. The testimage component is used for automated tests. Each component is deployed as a Docker container.

You can run the commands in the following steps to manage Docker containers:

1. Log on to your server of Apsara Stack.
2. Log on to the physical server of OCP V2.0 over Secure Shell (SSH).
3. Run the following command to view all the processes of the components: docker ps.
4. Run the following command to view the logs of a Docker container: docker logs ${containerID}. The logs record the information about the startup and running of the container.

5. Run the following command to go to a Docker container: docker exec -ti ${containerID} bash.
6. Run the following command to restart a Docker container: docker restart ${containerID}.

```
[root@a010038145178.cloud.h05.amtest70 /root]
#docker ps
CONTAINER ID        IMAGE               COMMAND                  CREATED             STATUS              PORTS               NAMES
73615653d439        bdef4d27a224        "/usr/bin/supervisord"   41 hours ago        Up 22 hours                             ocp-ocpv2.OcpApiV2__.ocp-api-v2.1589385296
eeef9f7121a2        8e63c5c9dd99        "/usr/bin/supervisord"   41 hours ago        Up 29 hours                             ocp-ocpv2.Ocp0dc__.ocp-odc.1589385295
7ea7bef943d8        6e8e09a91016        "sh /home/admin/start"   41 hours ago        Up 29 hours                             ocp-ocpv2.OcpObproxy__.ocp-obproxy.1589384482
007388e5bbd4        93f83158ed70        "sh /usr/local/rds/me"   41 hours ago        Up 29 hours                             ocp-ocpv2.OcpMetaInit__.ocp-meta-init.1589383621
78b150f669c6        f66f982357b3        "/usr/bin/supervisord"   41 hours ago        Up 29 hours                             ocp-ocpv2.OcpMetaServer__.ocp-meta-server.1589383219
30fb547ad645        315f0eab93cc        "/usr/bin/supervisord"   41 hours ago        Up 41 hours                             ocp-ocpv2.OcpTengine__.ocp-tengine.1589382799
```

OCP V2.0 architecture

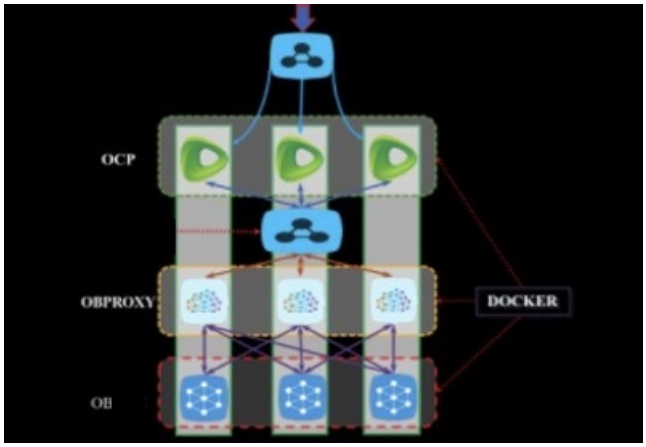

9.12.2.3.2. OcpMetaServer

Feature description

This component functions as a metadatabase of the management components and tools for ApsaraDB for OceanBase. You can create an ApsaraDB for OceanBase cluster where three nodes and three replicas are deployed. This cluster can be used to provide metadatabase services.

Related commands

Command	Description	Impact
ps -ef | grep observer | grep -v grep	Views the OBServer process.	None.
su - admin -c "cd /home/admin/oceanbase; ulimit -s 10240; ulimit -c unlimited; LD_LIBRARY_PATH=/home/admin/oceanbase/lib:/usr/local/lib:/usr/lib:/usr/lib64:/usr/local/lib64: LD_PRELOAD="" /home/admin/oceanbase/bin/observer"	Starts the OBServer process.	An error occurs if the OBServer process already exists.
tail -f /home/admin/oceanbase/log/observer.log	Views the OBServer log.	None.
ps -ef | grep -E "ob_|obstat" | grep -v grep | grep -v observer	Views the OBAgent process.	None.

Command	Description	Impact
/home/admin/obztools_agent/ob_agent.py stop agent -f	Stops the OBAgent service.	Stops monitoring ApsaraDB for OceanBase clusters and generating cluster alerts.
/home/admin/obztools_agent/ob_agent.py start agent	Starts the OBAgent service.	None.
tail -f /home/admin/obztools_agent/log/*.log	Views the OBAgent log.	None.

9.12.2.3.3. OcpMetaInit

Feature description

You can use this ApsaraDB for OceanBase component to initialize the clusters in the metadatabase, modify system parameters, and create metadata tenants that are required to provide services.

9.12.2.3.4. OcpObproxy

Feature description

This component functions as the OBProxy that you can use to access the metadata in the metadatabase clusters of ApsaraDB for OceanBase.

Related commands

Command	Description	Impact
ps -ef | grep obproxy | grep -v grep	Views the OBProxy process.	None
cd /home/admin/obproxy && ./bin/obproxy	Starts the OBProxy process.	None
tail -f /home/admin/obproxy/log/obproxy. *.log	Views the OBProxy log.	None

9.12.2.3.5. OcpApiV2

Feature description

This component is a service node of ApsaraDB for OceanBase OCP V2.0. The component provides features such as service management, O&M, monitoring, alerting, and backup and restoration.

Related commands

Command	Description	Impact
ps -ef | grep ocp-server.jar | grep -v grep	Views the process of the OCP server.	None
tail -f /home/admin/logs/ocp/ocp.*.log	Views the log of the OCP server.	None
ps -ef | grep -E "ob_|monitor" | grep -v grep	Views the process that collects monitoring data in OCP.	None
tail -f /home/admin/obztools_agent/log/*.log	Views the log that records the collected monitoring data in OCP.	None

9.12.2.3.6. OcpTengine

Feature description

This component functions as a frontend reverse proxy for the management components and tools of ApsaraDB for OceanBase. The component provides HTTPS proxy services.

Related commands

Command	Description	Impact
ps -ef | grep nginx | grep -v grep	Views the Tengine process.	None.
/opt/taobao/tengine/bin/tengine -c /opt/taobao/tengine/conf/nginx_https_sm.conf	Starts the Tengine process.	An error occurs if the Tengine process already exists.
tail -f /home/admin/logs/access.*.log tail -f /home/admin/logs/error.*.log	Views the log of the Tengine process.	None.

9.12.3. Routine maintenance

9.12.3.1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase

This topic describes how to use the Apsara Stack O&M system to log on to the Apsara Stack Operations console for ApsaraDB for OceanBase. The Google Chrome browser is used as an example in this topic.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are

obtained from the deployment personnel or an administrator.

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

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

Procedure

1. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

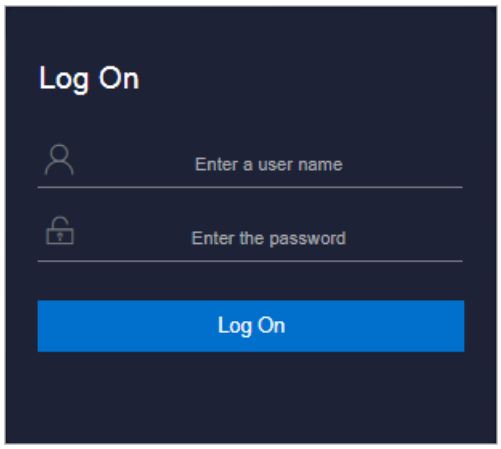

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

2. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

3. Click **Log On** to go to the **ASO** console.
4. In the left-side navigation pane, click **Products**. Then, choose **Product List > Database Services** and click **OceanBase Cloud Platform**. You are directed to the Apsara Stack Operations console for ApsaraDB for OceanBase.

9.12.3.2. Create instances

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.

2. In the left-side navigation pane, click **Instances**. The Instances page appears.
3. On the page that appears, click **Create Instance**. In the **Create Instance** dialog box, configure the parameters as prompted. Parameters for creating an instance describes the parameter configurations.

Parameters for creating an instance

Parameter	Description
OceanBase Version	The ApsaraDB for OceanBase version. You can specify this parameter based on your business requirements.
Cluster Group	The cluster group where the instance is created.
Instance Name	The name of the instance. You can specify this parameter based on your business requirements.
Zone	The zone where the instance is deployed. Select Default Zone from the Zone drop-down list.
Instance Specifications	The type of the instance. You can specify this parameter based on your business requirements.
Tenant Whitelist	The tenant whitelist of the instance. Separate IP addresses or CIDR blocks with commas (,). The parameter value % indicates that all the tenants of the instance are added to the whitelist. We recommend that you use the default value %.
Tenants	The total number of tenants that are bound to the instance.

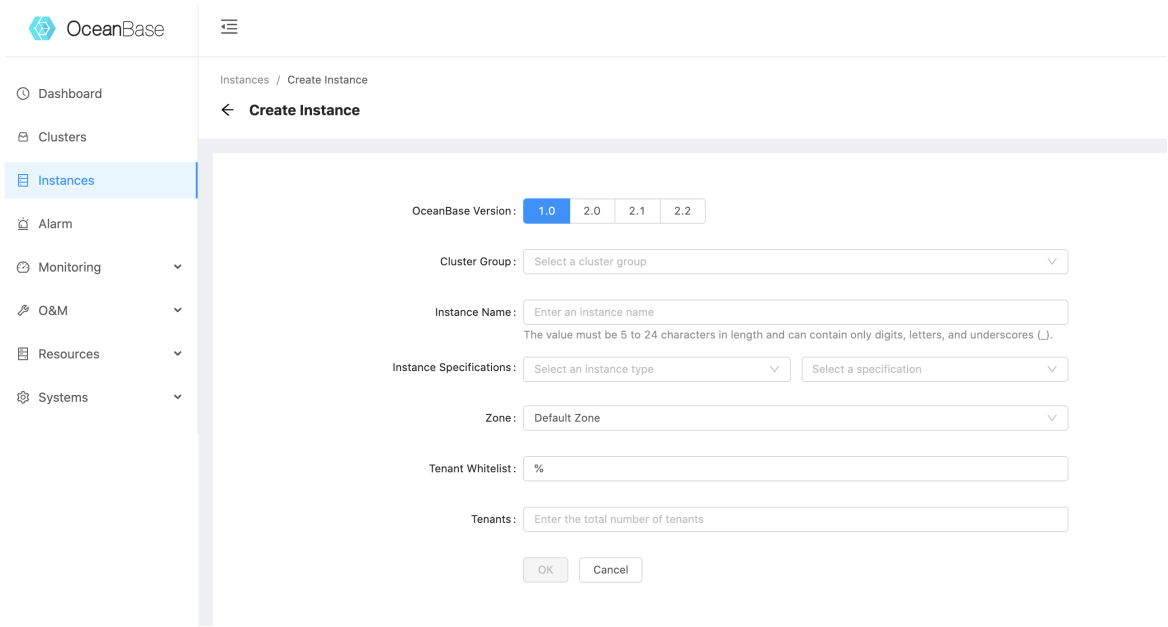

4. After you specify the preceding parameters for the instance, click **OK**. Then, the instance is created.
5. On the **Instances** page, click the name of the specified instance in the **Instance Name** column. On the page that appears, you can check the relevant parameter values and view the basic information and the performance metrics of the instance. You can also reset the password for the

instance, change the instance specifications, and perform other operations.

9.12.3.3. View performance metrics

9.12.3.3.1. Procedure of viewing monitoring metrics

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation page, click Dashboard. On the **Dashboard** page, you can view the following global performance metrics: performance trends, major freeze time, O&M tasks, inspection data, and server usage.

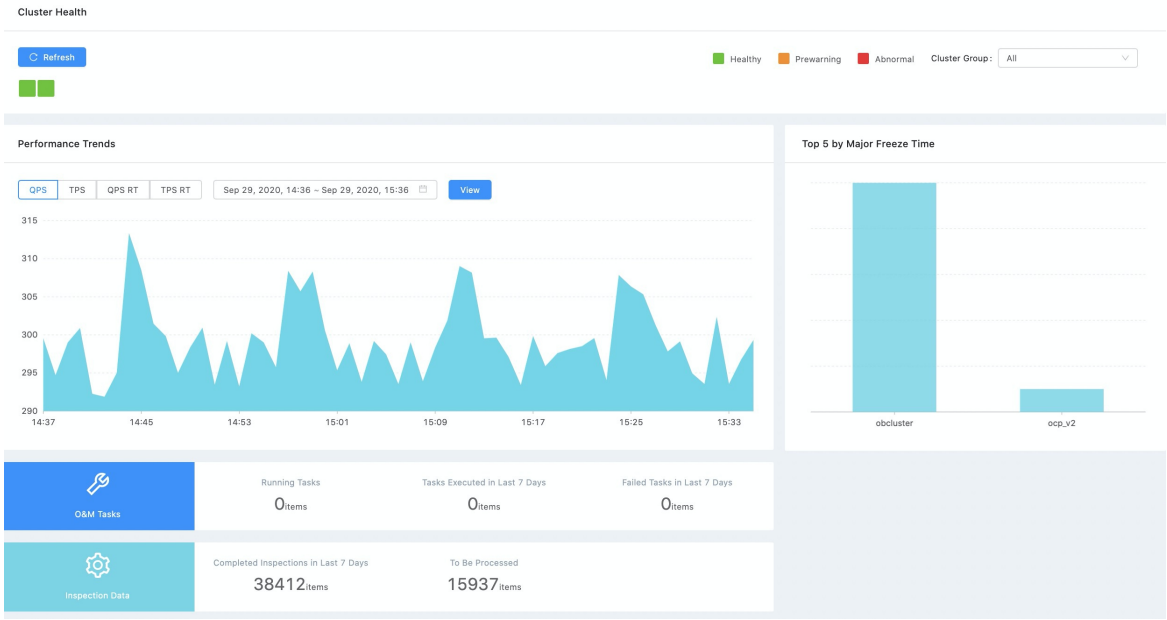

3. In the left-side navigation pane, click **Monitoring**. In the list that appears, click **OceanBase Data Trends**. On this page, you can view the disk usage, CPU utilization, and memory usage of ApsaraDB for OceanBase clusters. You can also view the details about host distribution and the following performance metrics: TPS, QPS, TPS RT, and QPS RT.

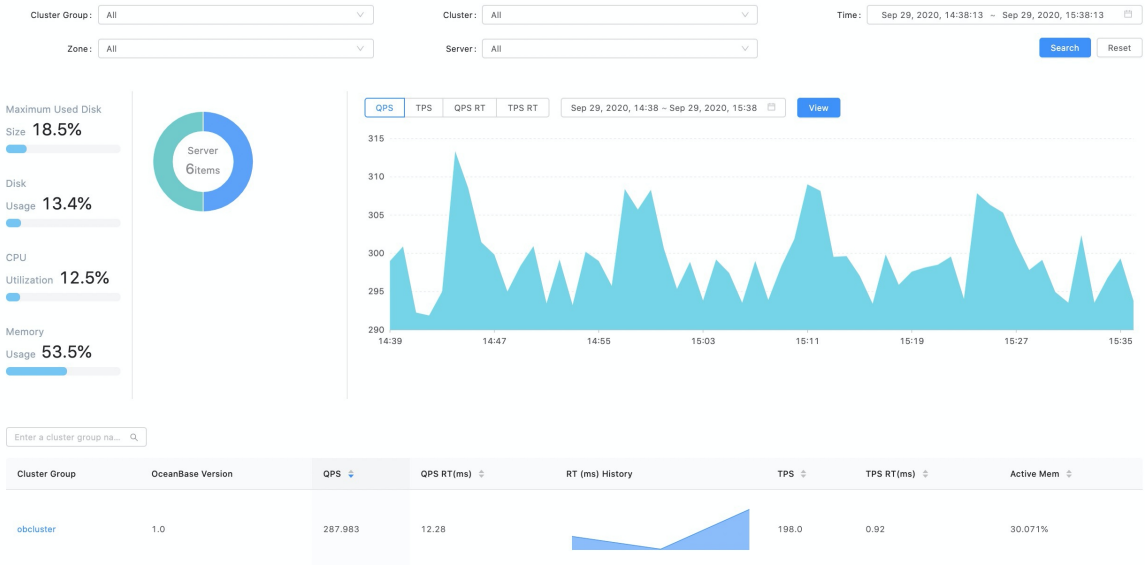

9.12.3.3.2. Description of performance metrics

In ApsaraDB for OceanBase, network packets are divided into two types. The network packets of one type are transferred over a remote procedure call (RPC) protocol. The RPC protocol is used as an internal communication protocol in ApsaraDB for OceanBase. The network packets of the other type are used to send SQL requests and are transferred over the MySQL protocol. The statistics about the two types of network packets are collected in a separate way.

In an OBServer, multiple queues are used to process different requests. After the OBServer receives network packets, the server sends the packets to the relevant queues.

Performance metric	Description
QPS	The number of SELECT executions per second.
TPS	The number of INSERT, REPLACE, UPDATE, and DELETE executions per second.
ACTIVE_SESSION	The number of active connections.
QPS_RT	The execution time of the SELECT statements. The statistics are collected in real time.
TPS_RT	The average execution time of each transaction.
SQL_DISTRIBUTED_COUNT	The number of executions for distributed SQL statements.
SQL_LOCAL_COUNT	The number of executions for the SQL statements that are executed in the on-premises environment.
SQL_REMOTE_COUNT	The number of executions for the SQL statements that are executed in the remote environment.
INNER_SQL_CONNECTION_EXECUTE_COUNT	The number of executions for the SQL requests that are sent over the RPC protocol.

Performance metric	Description
INNER_SQL_ CONNECTION_EXECUTE_RT	The execution time of the SQL requests that are sent over the RPC protocol.
BLOCK_CACHE_HIT_PERCENT	The cache hit percentage for the block cache.
BLOCK_INDEX_CACHE_HIT_PERCENT	The cache hit percentage for the block index cache.
BLOOM_FILTER_CACHE_HIT_PERCENT	The cache hit percentage for the Bloom filter cache.
BLOOM_FILTER_FILT_PERCENT	The cache interception percentage for the Bloom filter cache.
CLOG_CACHE_HIT_PERCENT	The cache hit percentage for the commit log cache.
LOCATION_CACHE_HIT_PERCENT	The cache hit percentage for the location cache.
LOCATION_CACHE_PROXY_HIT_PERCENT	The cache hit percentage for the OBProxy location cache.
PLAN_CACHE_HIT_RATE	The cache hit percentage for the SQL plan cache.
ROW_CACHE_HIT_PERCENT	The cache hit percentage for the row cache.
MYSQL_PACKET_IN	The number of the received MySQL requests.
MYSQL_PACKET_OUT	The number of the delivered MySQL requests.
MYSQL_PACKET_IN_BYTES	The number of bytes in the received MySQL request.
MYSQL_PACKET_OUT_BYTES	The number of bytes in the delivered MySQL request.
MYSQL_DELIVER_FAIL	The number of MySQL requests that failed to be delivered.
RPC_DELIVER_FAIL	The number of RPC requests that failed to be delivered.
RPC_PACKET_IN	The number of the received RPC requests.
RPC_PACKET_OUT	The number of the delivered RPC requests.
RPC_PACKET_IN_BYTES	The number of bytes in the received RPC request.
RPC_PACKET_OUT_BYTES	The number of bytes in the delivered RPC request.
RPC_NET_FRAME_RT	The delay of an RPC network packet at the network layer. The delay is the interval between the time when the packet is received and the time when the packet is distributed to the relevant queue.

Performance metric	Description
RPC_NET_RT	The delay of a network packet. The delay is the interval between the time when the packet is sent and the time when the packet is received. Note that the sending time and the receiving time are recorded on different servers. For example, if a network packet is sent from Server A to Server B, the sending time is recorded on Server A and the receiving time is recorded on Server B.
REQUEST_DEQUEUE_COUNT	The number of requests that are sent to the queue.
REQUEST_ENQUEUE_COUNT	The number of requests that are moved out of the queue.
REQUEST_QUEUE_TIME	The interval between the time when a request is received and the time when the request starts to be processed.
TRANS_COMMIT_LOG_SYNC_COUNT	The number of commit logs for the transaction.
TRANS_COMMT_LOG_SYNC_RT	The interval between the time when the commit logs are submitted and the time when the logs are synchronized to a majority of replicas.
TRANS_COMMT_COUNT	The number of transaction commits.
TRANS_MULTI_PARTITION_COUNT	The number of transactions that are executed across partitions.
TRANS_ROLLBACK_COUNT	The number of transaction rollbacks.
TRANS_SINGLE_PARTITION_COUNT	The number of transactions that are executed in a single partition.
TRANS_START_COUNT	The number of times that transactions are started. The number includes the number of times that the transaction is automatically committed.
TRANS_SYSTEM_TRANS_COUNT	The number of executions for system transactions. This number also indicates the number of SQL statement executions. Note that the SQL statements are executed to retrieve data from ApsaraDB for OceanBase and manage metadata in ApsaraDB for OceanBase.
TRANS_TIMEOUT_COUNT	The number of time-out transactions.
TRANS_USER_TRANS_COUNT	The number of user transactions.
TRANS_COMMIT_RT	The time that is consumed in the commit phase of the transaction.

Performance metric	Description
TRANS_ROLLBACK_RT	The time that is consumed in the rollback phase of the transaction.
TRANS_RT	The total consumption time of the transaction. The time period starts from the time when the transaction starts to be performed and ends at the time when the transaction is completed.
MEMSTORE_READ_LOCK_FAIL_COUNT	The number of times that data failed to be read from the lock in the MemStore.
MEMSTORE_READ_LOCK_SUCC_COUNT	The number of times that data was read from the lock in the MemStore.
MEMSTORE_WRITE_LOCK_FAIL_COUNT	The number of times that data failed to be written to the lock in the MemStore.
MEMSTORE_WRITE_LOCK_SUCC_COUNT	The number of times that data was written to the lock in the MemStore.
MEMSTORE_READ_LOCK_WAIT_TIME	The waiting time to read data from the lock in the MemStore.
MEMSTORE_ WRITE_LOCK_WAIT_TIME	The waiting time to write data to the lock in the MemStore.
MEMSTORE_APPLY_FAIL_COUNT	The number of times that data failed to be written to the MemStore.
MEMSTORE_ APPLY_SUCC_COUNT	The number of times that data were written to the MemStore.
MEMSTORE_GET_FAIL_COUNT	The number of times that GET queries failed to retrieve data from the MemStore.
MEMSTORE_GET_SUCC_COUNT	The number of times that GET queries retrieved data from the MemStore as expected.
MEMSTORE_SCAN_FAIL_COUNT	The number of times that scan queries failed to retrieve data from the MemStore.
MEMSTORE_SCAN_SUCC_COUNT	The number of times that scan queries retrieved data from the MemStore as expected.
MEMSTORE_APPLY_RT	The time that is consumed by writing data to the MemStore.
MEMSTORE_GET_RT	The time that is consumed by GET requests. The requests are sent to retrieve data from the MemStore.

Performance metric	Description
MEMSTORE_SCAN_RT	The time that is consumed by scan query requests. The requests are sent to retrieve data from the MemStore.
MEMSTORE_ROW_COUNT	The number of rows in the MemStore.
IO_READ_COUNT	The number of data reads in the I/O operations.
IO_WRITE_COUNT	The number of data writes in the I/O operations.
IO_READ_RT	The time that is consumed by the data reads in the I/O operations.
IO_WRITE_RT	The time that is consumed by the data writes in the I/O operations.
IO_READ_SIZE	The number of bytes that are read in the I/O operations.
IO_WRITE_SIZE	The number of bytes that are written in the I/O operations.
IO_PREFETCH_MICRO_BLOCK_COUNT	The number of micro-blocks that are pre-read in the I/O operations.
IO_PREFETCH_UNCOMPRESS_MICRO_BLOCK_COUNT	The number of uncompressed micro-blocks that are pre-read in the I/O operations.
IO_READ_MICRO_INDEX_COUNT	The number of times that the micro-block indexes are read in the I/O operations.
IO_PREFETCH_MICRO_BLOCK_SIZE	The number of bytes in the micro-blocks that are pre-read in the I/O operations.
IO_PREFETCH_UNCOMPRESS_MICRO_BLOCK_SIZE	The number of bytes in the uncompressed micro-blocks that are pre-read in the I/O operations.
IO_READ_MICRO_INDEX_SIZE	The number of bytes for the micro-block indexes that are read in the I/O operations.
PARTITION_TABLE_OPERATOR_GET_COUNT	The number of times that the partitioned table receives GET requests.
PARTITION_TABLE_OPERATOR_GET_RT	The response time of the GET requests for the partitioned table.
REFRESH_SCHEMA_COUNT	The number of times that the table schema is updated.
REFRESH_SCHEMA_RT	The time that is consumed to update the table schema.

Performance metric	Description
CLOG_CB_RT	The consumption time of callbacks in the commit logs.
CLOG_COUNT	The number of commit logs.
CLOG_EVENT_RT	The number of events in the commit logs, such as takeover and revoke events.
CLOG_FLUSH_TASK	The number of tasks in the commit logs. In the tasks, the data records are flushed to disks.
CLOG_GROUP_SIZE	The statistics about the group commits in the commit logs.
CLOG_READ	The statistics about the data reads in the commit logs. Note that data is read from disks.
CLOG_RPC_RT	The statistics about the RPC requests in the commit logs.
CLOG_RT	The statistics about the response time in the commit logs.
CLOG_SIZE	The statistics about the sizes in the commit logs.
ACTIVE_MEMSTORE_USED	The memory space that is occupied by the active MemStore.
MAJOR_FREEZE_TRIGGER	The threshold that triggers the major freeze operation.
CPU_USAGE	The current CPU utilization for the tenant.
MIN_CPU_SIZE	The minimum number of CPU cores for the tenant.
MAX_CPU_SIZE	The maximum number of CPU cores for the tenant.
LOCATION_CACHE_PROXY_HIT	The number of cache hits for the OBProxy location cache.
LOCATION_CACHE_PROXY_MISS	The number of cache misses for the OBProxy location cache.
LOCATION_CACHE_RENEW_COUNT	The number of updates in the location cache.
LOCATION_CACHE_RPC_SUCC_COUNT	The number of RPC requests that retrieved data from the location cache as expected.
MEMORY_USAGE	The current memory usage of the tenant.
MIN_MEMORY_SIZE	The minimum memory space of the tenant.

Performance metric	Description
MAX_MEMORY_SIZE	The maximum memory space of the tenant.
MEMSTORE_LIMIT	The maximum memory space for the incremental data that is stored in the MemStore.
TOTAL_MEMSTORE_USED	The total memory space that is occupied by the MemStore.
DATA_SIZE	The total data size of the tenant.
DISK_USAGE	The disk usage.
LEADER_DATA_SIZE	The data size of the server where the leader partition resides.
ABORT_LOG_REPLAY_RT	The duration in which the abort logs are replayed in the memory during the two-phase commit process.
CLEAR_LOG_REPLAY_RT	The duration in which the clear logs are replayed in the memory during the two-phase commit process.
COMMIT_LOG_REPLAY_RT	The duration in which the commit logs are replayed in the memory during the two-phase commit process.
PREPARE_LOG_REPLAY_RT	The duration in which the prepare logs are replayed in the memory during the two-phase commit process.
REDO_LOG_REPLAY_RT	The duration in which the redo logs are replayed in the memory during the two-phase commit process.
PLAN_CACHE_MEM_HOLD	The memory space that can be occupied by the SQL plan cache.
PLAN_CACHE_MEM_USED	The actual memory space that is occupied by the SQL plan cache.
PLAN_CACHE_PLAN_NUM	The number of SQL plans in the SQL plan cache.
PLAN_CACHE_SQL_NUM	The number of executed SQL statements that are recorded the SQL plan cache.
PLAN_CACHE_STMTKEY_NUM	The number of declared keys that are recorded in the SQL plan cache.
RE_SUBMITTED_FREEZE_TASK_COUNT	The number of freeze tasks that are resubmitted.
RE_SUBMITTED_OFFLINE_TASK_COUNT	The number of tasks that are resubmitted to disable objects.

Performance metric	Description
RE_SUBMITTED_TRANS_TASK_COUNT	The number of transaction tasks that are resubmitted.

9.12.3.4. Exception monitoring

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, click **Alarm**. On the **Alarm Event** tab, you can view the list of alerts for the exceptions that occur on the ApsaraDB for OceanBase clusters.

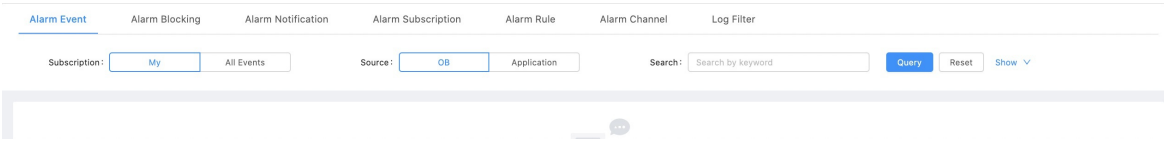

3. Click an alert event. On the page that appears, you can view the **event details**. On this page, you can click **View Alarm Rules** to view the details about the alert rule. You can also click **Block Alarms** to create an alert blocking rule.

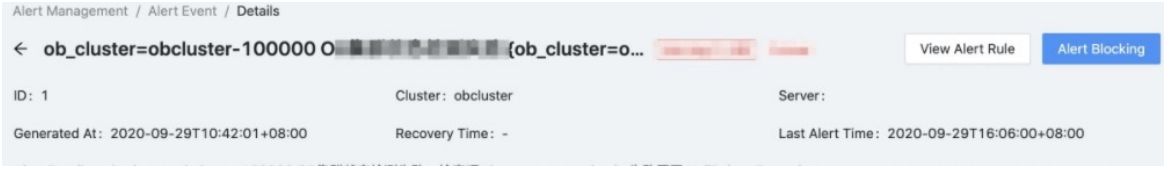

9.12.3.5. Resource management

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, choose **Resources > Resource Status**. On the **Resource Status** page, you can view the usage details of the resources that are managed in the ApsaraDB for OceanBase console.

9.12.3.6. Upgrade and optimization

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, choose **O&M > RPM Package Management**. On the page that appears, click the **Database** tab. On the tab, check whether the RPM package that you want to use for the upgrade is included in the list. If the RPM package is not in the list, click **Upload RPM Package** to add the RPM package.
3. In the left-side navigation pane, click **Clusters**. In the Actions column for the cluster that you want to upgrade, choose **Routine Actions > Cluster Upgrade**.
4. In the **Cluster Upgrade** dialog box, specify the zones for the cluster, the RPM package version, and the start time of the upgrade. Then, click **OK**.

5. In the left-side navigation pane, choose **O&M > Task**. On the **Task** page, you can view the status of the upgrade task.

9.12.4. Security maintenance

9.12.4.1. Network security maintenance

Network security maintenance helps you ensure device and network security.

- Device security
  - Check network devices and enable security management protocols and configurations for these devices.
  - Check up-to-date versions of network device software and update the software to a secure version in a timely manner.
  - For more information about the security maintenance method, see the product documentation of each device.
- Network security

  Based on your network conditions, select the intrusion detection system (IDS) or intrusion prevention system (IPS) to detect abnormal traffic from the Internet or internal networks. This protects services against abnormal behaviors and attacks in real time.

9.12.4.2. Account password maintenance

Account passwords include system tenant passwords, service tenant passwords, and device passwords of ApsaraDB for OceanBase clusters. To ensure account security, you must use complex passwords for your system tenants and devices and change the passwords on a regular basis.

9.12.4.3. Establish a fault response mechanism

Designate the owners for handling various types of faults

The O&M engineers of ApsaraDB for OceanBase must establish a fault emergency response mechanism. This ensures that the services can be resumed in a timely manner after a fault or a security issue occurs.

Stock-up mechanism

A stock-up mechanism must be established for fragile hardware devices to ensure that hardware faults are rectified in a timely manner. This mitigates the negative impacts of hardware faults.

Technical support

After a system fault is detected during routine maintenance, you can use the O&M platform of ApsaraDB for OceanBase to check fault details. Then, you can analyze the fault causes and rectify the fault based on detailed analysis. If the fault cannot be rectified, you can collect related information such as system information and fault symptoms, and contact technical support.

After you rectify a fault, analyze its causes, review the troubleshooting process, and make improvements.

9.12.5. Backup and restoration

9.12.5.1. Overview

The backup and restoration feature of ApsaraDB for OceanBase ensures high-availability and high-performance services. In ApsaraDB for OceanBase, you can use various storage tools such as Object Storage Service (OSS) and disk arrays for backup and restoration. You can use a backup and restoration tool of ApsaraDB for OceanBase to back up data for multiple ApsaraDB for OceanBase clusters at a time. This is known as the 1+N backup. You can use the backup and restoration feature to back up and restore data by cluster or by tenant. You can also restore data to a specific point in time. The backup and restoration feature allows you to back up and restore the data of all the operations that are performed in databases. The backup data supports all the physical data and some logical data. The supported logical data includes user permissions, table definitions, system variables, user information, and view information.

9.12.5.2. Back up data

9.12.5.2.1. Deploy a backup server

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, choose **O&M > Backup and Restoration**. On the Servers tab of the **Backup and Restoration** page, click **Add Server**. In the Add Server dialog box, specify the Server IP Address field as the IP address of the backup server to be deployed. From the Region drop-down list, select a region. We recommend that you select the region of the cluster whose data is to be backed up. From the Server Type drop-down list, select **Backup**. In the Root Password field, enter the password of the root user. In the Administrator Password field, enter the password of the administrator. Then, click **OK**.
3. Refresh the page. You can view the list of backup servers. In the **Actions** column for the added server, click Actions and then click **Deploy**.
4. In the **Deploy Backup Server** dialog box, specify the parameters as prompted, such as **Version to Be Deployed**, **Region**, and **Profile**. Below the selected profile, you can view the configuration data in the selected profile. If you need to modify the profile, modify the profile in the configuration list below the selected profile. After you specify the parameters, click **OK**.
5. In the left-side navigation pane, choose **O&M > Task**. On the Task page, view the status of the task for deploying the backup server. If the backup server is deployed as expected, the task is in the Success state.

9.12.5.2.2. Create a backup task

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, click **Clusters**. On the Clusters page, find the cluster whose data is to be backed up. In the Actions column for the cluster, choose **More > Cluster Backup**.
3. On the **Backup Configuration** tab, specify the backup parameters.

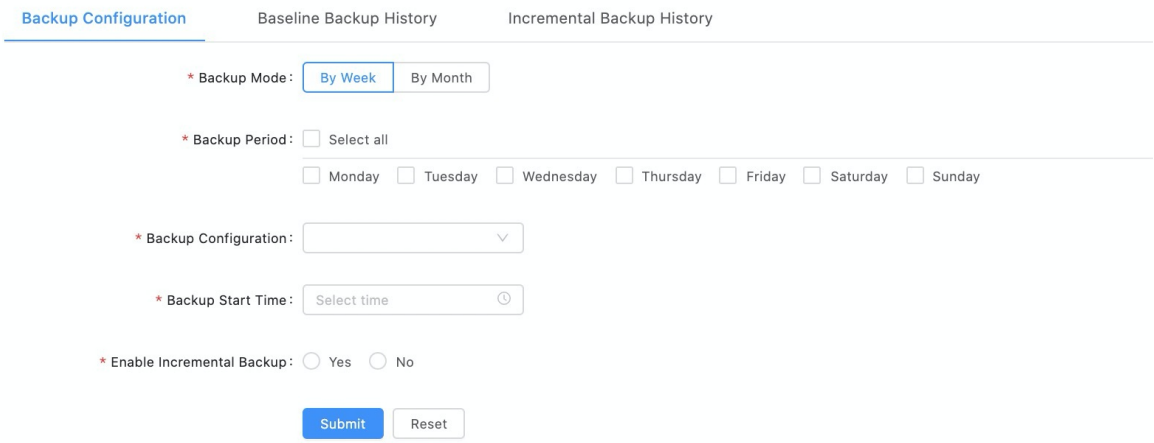

4. Click **Submit** to back up data for the specified cluster.
5. In the left-side navigation pane, choose **O&M > Task**. On the Task page, you can view the status of the backup task that you created.

9.12.5.2.3. View the status of a backup task on a regular basis

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, choose **O&M > Backup and Restoration > Backup Scheduling**. On the Backup Scheduling tab, click the name of the cluster for which you want to view the backup task.

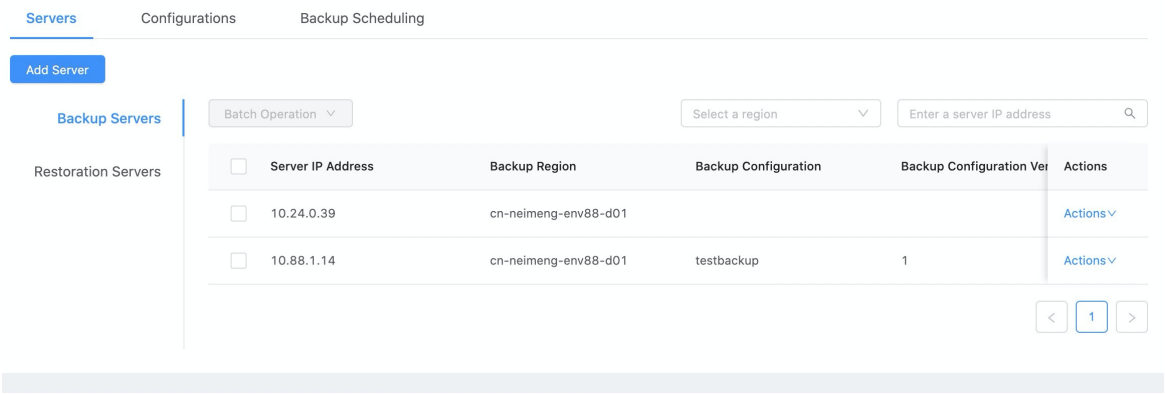

3. On the **Backup Scheduling** tab, you can view the status of the current backup task for the cluster.

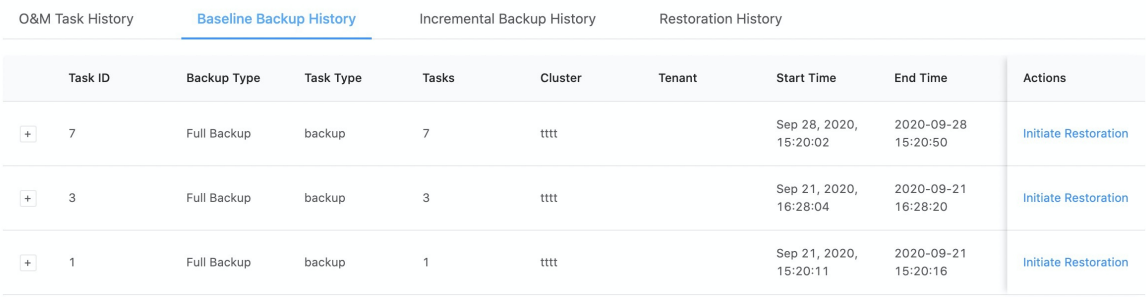

9.12.5.3. Restore data

9.12.5.3.1. Deploy a restoration server

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, choose **O&M > Backup and Restoration**. On the Servers tab of the **Backup and Restoration** page, click **Add Server**. In the Add Server dialog box, specify Server IP Address as the IP address of the restoration server to be deployed. From the Region drop-down list, select a region. We recommend that you select the region of the cluster whose data is to be restored. From the Server Type drop-down list, select **Restore**. If you need to use the same server for backup and restoration, select both Backup and Restore from the Server Type drop-down list. In the Root Password field, enter the password of the root user. In the Administrator Password field, enter the password of the administrator user. Then, click **OK**.
3. Refresh the page. You can view the list of restoration servers. In the **Actions** column for the added server, click Actions and then click **Deploy**.
4. In the **Deploy Server with Restored Service** dialog box, specify the parameters as prompted, such as **Version to Be Deployed**, **Region**, and **Profile**. Below the selected profile, you can view the configuration data in the selected profile. After you specify the parameters, click **OK**.
5. In the left-side navigation pane, choose **O&M > Task**. On the Task page, view the status of the task for deploying the restoration server. If the restoration server is deployed as expected, the task is in the Success state.

9.12.5.3.2. Restore data

Procedure

1. Log on to the Apsara Stack Operations console for ApsaraDB for OceanBase.
2. In the left-side navigation pane, choose **O&M > Backup and Restoration**. On the page that appears, click the Backup Scheduling tab. On the **Backup Scheduling** tab, click name of the cluster whose data needs to be restored. In the dialog box that appears, click the **Baseline Backup History** tab.

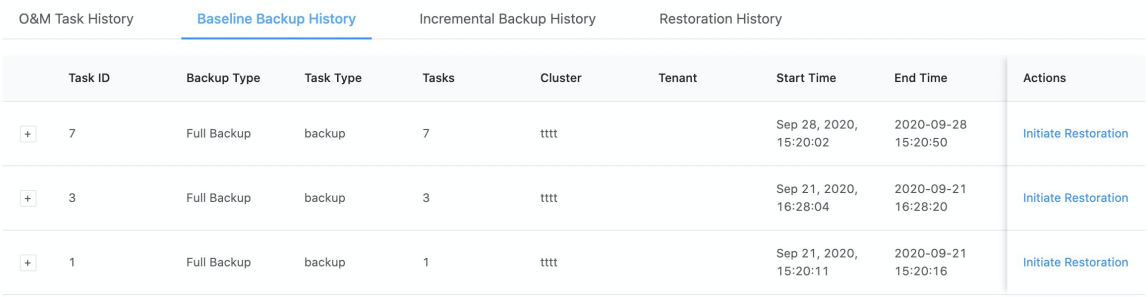

3. On the tab, click **Initiate Restoration**. Then, select the cluster whose data is used for restoration and click **OK**. On the page that appears, select the tenant whose data needs to be restored and keep the default settings for the other parameters. Then, click **OK**.
4. Connect to the cluster whose data is used for restoration and check whether the tenant data is restored as expected.

9.12.6. Troubleshooting

9.12.6.1. Troubleshooting methods

After a system fault is detected during routine maintenance, you can log on to the Apsara Stack Operations console for ApsaraDB for OceanBase to check the fault details. Then, you can analyze the fault causes and rectify the fault based on detailed analysis.

If the fault cannot be rectified, you can collect related information such as system information and fault symptoms, and contact technical support.

After you rectify a fault, analyze its causes, review the troubleshooting process, and make improvements.

9.12.6.2. Troubleshoot common faults

9.12.6.2.1. Insufficient memory

Possible causes

Insufficient memory may occur due to three possible causes. One of the possible causes is that the service data is written to the memory of a tenant at a high rate. Another possible cause is that the tenant is allocated with only a small amount of memory space. The last possible cause is that the cluster processes high loads of service data.

The memory space that is occupied by the MemStore of one tenant or all the tenants may reach the upper limit for the memory space. If this occurs, data cannot be written to ApsaraDB for OceanBase nodes, or the "Over tenant memory limits" error is reported. In this scenario, errors occur if you perform operations in ApsaraDB for OceanBase.

Solutions

- Trigger major freeze operations or configure minor freeze operations to release some memory space.
- Allocate more memory resources to the tenant.
- Scale out the cluster.
- Implement rate limiting and throttling.

9.12.6.2.2. Insufficient disk space

Possible causes

A large number of system log files or commit log files are written into the disks.

Solutions

- If a large number of system log files are written to the disks, delete the earliest system log files in the oceanbase/log directory.
- If a large number of commit log files are written to the disks, delete the earliest files in the clog folder. Then, you can perform two major freeze operations.

9.12.6.2.3. High CPU utilization

Possible causes

High CPU utilization may be caused by high program workloads, slow SQL queries, or inappropriate resource allocation.

Solutions

- If high CPU utilization is detected on some servers, increase the value of the unit_number parameter to increase the CPU resources that are allocated to the tenants of the servers. This allows the servers to occupy more resources and helps you balance the server loads.
- If high CPU utilization is caused by slow SQL queries, contact the developers to check the issue and optimize the SQL statements that result in the issue.

9.12.6.2.4. High loads

Possible causes

One of the possible causes is that a large number of programs are concurrently running. Another possible cause is that SQL statements are executed at a low efficiency.

Solutions

- If the loads of some servers are high, increase the value of the unit_number parameter to increase the resources that are allocated to the tenants of the servers. This allows the servers to occupy more resources and helps you balance the server loads.
- If a large number of programs are concurrently running, contact the developers to reduce the number of concurrent programs.
- If high loads are caused by low efficiency of executing SQL statements, execute the EXPLAIN statement to check how the system processes the SQL statements and to find performance bottlenecks. If you cannot use indexes to improve the execution efficiency, contact the developers.

9.13. Log Service

9.13.1. O&M methods

This topic describes two O&M methods of Log Service.

Log Service is deployed, operated, and maintained by using the Apsara Infrastructure Management Framework console. Log Service supports the following two O&M methods:

- Terminal: In the Apsara Infrastructure Management Framework console, you can use the terminal service to log on to the server where Log Service resides and view logs.
- Portal: The portal service provides a user interface for you to perform operations in Log Service. The portal service complies with the standard Java applications of Alibaba Cloud.

Terminal

1. Log on to the Apsara Stack Operations (ASO) system.For more information about how to log on to the ASO system, see **Log on to the ASO console** in *Operations and Maintenance Guide*.
2. In the left-side navigation pane, choose **Products > Product List**.
3. On the page that appears, click **Apsara Infrastructure Management Framework** to go to the Apsara Infrastructure Management Framework console.
4. In the left-side navigation pane, choose **Operations > Service Operations**.
5. On the page that appears, find `sls-backend-server` in the Services column, and then click **Operations** in the Actions column.
6. On the **Clusters** tab, find the target cluster in the Clusters column, and then click **Operations** in the Actions column.
7. On the **Services** tab, select the target server role, for example, sls-backend-server.ServiceTest#, and then click **Terminal** in the Actions column.

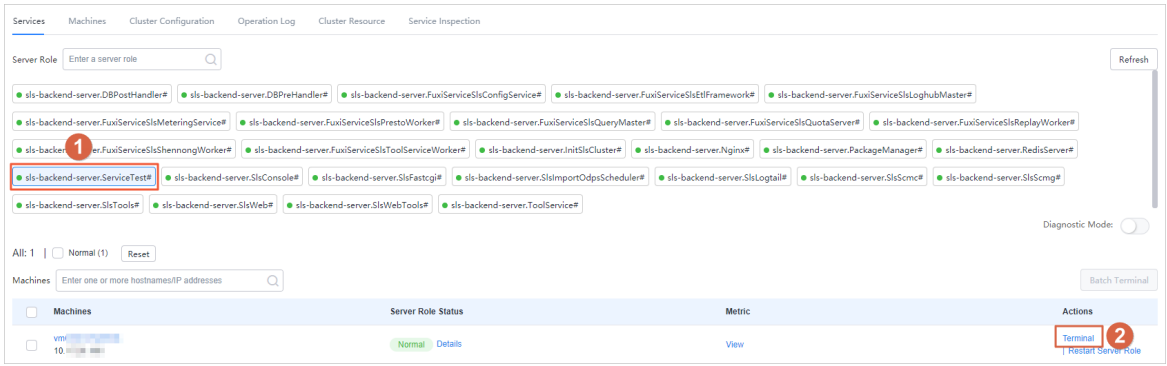

8. Log on to the server by using the terminal service and go to the corresponding directory to view logs.

Portal

You can collect logs of your server and send the logs to the portal service. Then, you can search for, retrieve, and analyze these logs by using the portal service.

1. Log on to the Apsara Infrastructure Management Framework console to obtain the endpoint of the portal service.
   i. Log on to the Apsara Infrastructure Management Framework console. For more information, see Terminal.
   ii. In the left-side navigation pane, click **Reports**. You are redirected to the **All Reports** page.
   iii. In the Report column, click **Registration Vars of Services**.

iv. In the dialog box that appears, click the icon in the column, enter **sls** in the search box, and then click **Apply Filter**.

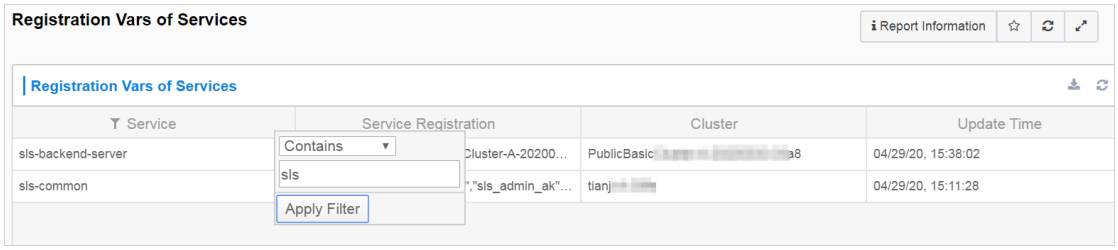

v. Right-click the **Service Registration** column of the **sls-backend-server** service, and then select **Show More**.

vi. On the **Details** page, find the endpoint of the portal service.

2. Log on to the Apsara Infrastructure Management Framework console to obtain the AccessKey pair of the portal service.

   i. Log on to the Apsara Infrastructure Management Framework console. For more information, see Terminal.

   ii. In the left-side navigation pane, click **Reports**. You are redirected to the **All Reports** page.

iii. On the left side of the page, click the **C** tab, find the target cluster, and then choose > **Cluster Configuration File**.

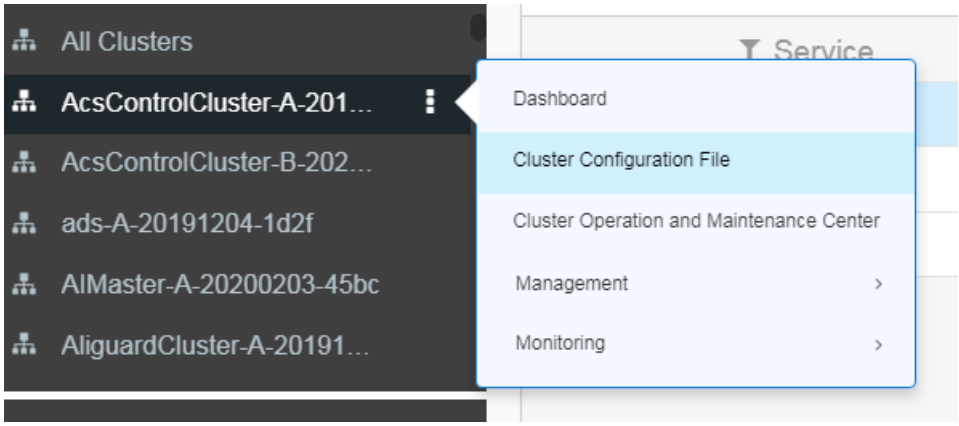

iv. Click **kv.conf** to obtain the AccessKey pair of the portal service.

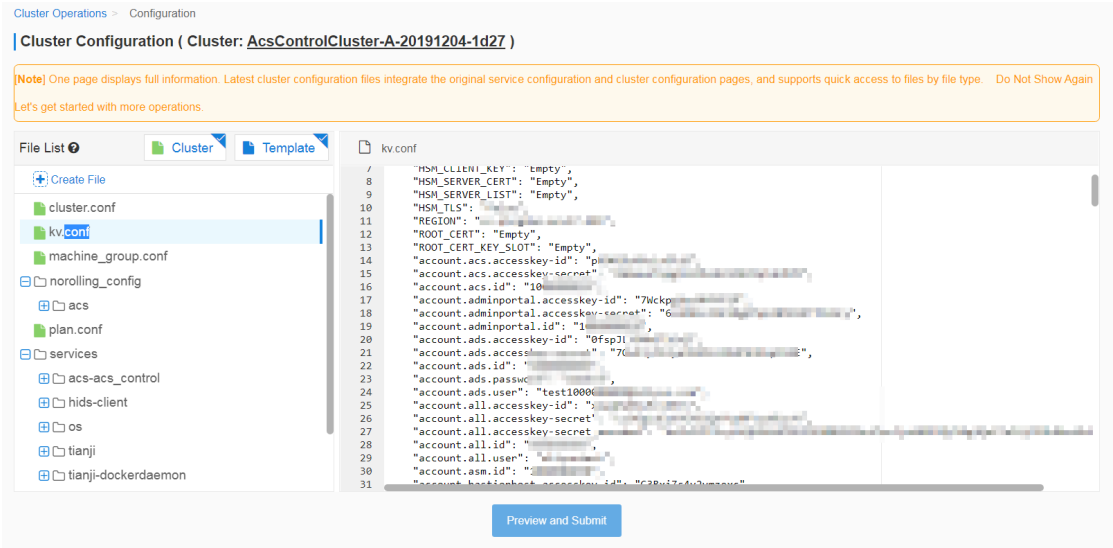

3. Log on to the Apsara Infrastructure Management Framework console.Log on to the portal service by using the endpoint obtained in Step 1 and the AccessKey pair obtained in Step 2.
4. Find the target project and Logstore, and then query logs.

9.13.2. O&M

9.13.2.1. View logs on machines

InitSlsCluster#

- Startup log: /cloud/app/sls-backend-server/InitSlsCluster#/init_sls_cluster/current/log/start.log
- Service log: none

Nginx#

- Startup log: /cloud/app/sls-backend-server/Nginx#/nginx/current/log/start.log
- Service logs:

- /apsara/nginx/logs/access.log
- /apsara/nginx/logs/error.log
- /apsara/nginx/logs/fastcgi_agent_access.log
- /apsara/nginx/logs/offline_access.log
- /apsara/nginx/logs/scmc_access.log
- /apsara/nginx/logs/scmc_err_log
- /apsara/nginx/logs/scmc_op_log
- /apsara/nginx/logs/scmg_access.log
- /apsara/nginx/logs/scmg_err_log
- /apsara/nginx/logs/scmg_op_log
- /apsara/nginx/logs/sls_console.log
- /apsara/nginx/logs/web_access.log

PackageManager#

- Startup log: /cloud/app/sls-backend-server/PackageManager#/package_manager/current/log/start.log
- Service log: none

RedisServer#

- Startup log: /cloud/app/sls-backend-server/RedisServer#/sls_redis/current/log/start.log
- Service log: /var/log/redis/redis.log

SlsConsole#

- Startup log: /cloud/app/sls-backend-server/SlsConsole#/sls_console/current/log/start.log
- Service logs: /alidata/www/logs/
  - /alidata/www/logs/java/sls/
    - /alidata/www/logs/java/sls/dashboard.log
    - /alidata/www/logs/java/sls/debug.log
    - /alidata/www/logs/java/sls/error.log
    - /alidata/www/logs/java/sls/info.log
    - /alidata/www/logs/java/sls/reasons.log
    - /alidata/www/logs/java/sls/tairSave.log
  - /alidata/www/logs/java/sls-service/applog
    - /alidata/www/logs/java/sls-service/applog/error.log
    - /alidata/www/logs/java/sls-service/applog/info.log
    - /alidata/www/logs/java/sls-service/applog/warn.log
  - /usr/share/jetty/logs/
    - /usr/share/jetty/logs/request.log
    - /usr/share/jetty/logs/stderrout.log

SlsFastcgi#

- Startup log: /cloud/app/sls-backend-server/SlsFastcgi#/sls_fastcgi/current/log/start.log
- Service logs:
  - /apsara/fcgi_agent/FastcgiAgent.LOG
  - /apsara/fcgi_agent/metering.LOG
  - /apsara/fcgi_agent/monitor.LOG
  - /apsara/fcgi_agent/ols_operation.LOG

SlsLogtail#

- Startup log: /cloud/app/sls-backend-server/SlsLogtail#/sls_ilogtail/current/log/start.log
- Service logs
  - Service log on Apsara Stack: /usr/local/ilogtail_private/ilogtail.LOG
  - Service log on on-premises machines: /usr/local/ilogtail/ilogtail.LOG

SlsScmc#

- Startup log: /cloud/app/sls-backend-server/SlsScmc#/sls_scmc/current/log/start.log
- Service logs:
  - /var/www/html/SCMC/logs/scm_op_log
  - /var/www/html/SCMC/logs/scm_err_log

SlsScmg#

- Startup log: /cloud/app/sls-backend-server/SlsScmg#/sls_scmg/current/log/start.log
- Service logs:
  - /var/www/html/SCMG/logs/scm_err_log
  - /var/www/html/SCMG/logs/scm_op_log

SlsTools#

- Startup log: /cloud/app/sls-backend-server/SlsTools#/aliyun_log_cli/current/log/start.log
- Service log: none

SlsWeb#

- Startup log: /cloud/app/sls-backend-server/SlsWeb#/sls_web/current/log/start.log
- Service logs:
  - /apsara/sls/web/logs/access.log
  - /apsara/sls/web/logs/apidetail.log
  - /apsara/sls/web/logs/httpclient.log
  - /apsara/sls/web/logs/normal.log
  - /apsara/sls/web/logs/sysinfo.log
  - /apsara/sls/web/logs/worker.log

SlsWebTools#

- Startup log: /cloud/app/sls-backend-server/SlsWebTools#/sls_web_tools/current/log/start.log
- Service log: none

ToolService#

- Startup logs:
  - /cloud/app/sls-backend-server/ToolService#/init_db/current/log/start.log
  - /cloud/app/sls-backend-server/ToolService#/init_diamond/current/log/start.log
  - /cloud/app/sls-backend-server/ToolService#/init_odps/current/log/start.log
  - /cloud/app/sls-backend-server/ToolService#/init_pop/current/log/start.log
  - /cloud/app/sls-backend-server/ToolService#/jdk_uploader/current/log/start.log
- Service log: none

SlsImportOdpsScheduler#

- Startup log: /cloud/app/sls-backend-server/SlsImportOdpsScheduler#/sls_import_odps_scheduler/current/log/start.log
- Service Logs: Job Scheduler service

FuxiServiceSlsConfigService#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsConfigService#/sls_config_service/current/log/start.log
- Service log: none

FuxiServiceSlsEtlFramework#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsEtlFramework#/sls_etl_framework/current/log/start.log
- Service log: none

FuxiServiceSlsLoghubMaster#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsLoghubMaster#/sls_loghub_master/current/log/start.log
- Service log: none

FuxiServiceSlsMeteringService#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsMeteringService#/sls_metering_service/current/log/start.log
- Service log: none

FuxiServiceSlsPrestoWorker#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsPrestoWorker#/sls_presto_worker/current/log/start.log
- Service log: none

FuxiServiceSlsQueryMaster#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsQueryMaster#/sls_query_master/current/log/start.log
- Service log: none

FuxiServiceSlsQuotaServer#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsQuotaServer#/sls_quota_server/current/log/start.log
- Service log: none

FuxiServiceSlsReplayWorker#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsReplayWorker#/sls_replay_worker/current/log/start.log
- Service log: none

FuxiServiceSlsShennongWorker#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsShennongWorker#/sls_shennong_worker/current/log/start.log
- Service log: none

FuxiServiceSlsToolServiceWorker#

- Startup log: /cloud/app/sls-backend-server/FuxiServiceSlsToolServiceWorker#/sls_tool_service_worker/current/log/start.log
- Service log: none

NGINX

Error log: */apsara/nginx/log/error.log*

Error	Action
Bind Address Failed	Check the port listening information in */etc/init.d/nginx.conf*.
open() ... failed	Check whether the item that you want to open exists in the static resource file.

Console

Error log: */alidata/www/logs/java/sls/error.log*

Error	Action
SLS SDK Exception	No action is required.
Create Bean Failed	Check the dubbo settings in the console configurations of SlsConsole.

Service

Error log: */alidata/www/logs/java/sls-service/applog/error.log*

Error	Action

Error	Action
Create Bean Failed	Check the dubbo settings in the service configurations of SlsConsole.
Invoke failed	Check the scmg settings in the service configurations of SlsConsole.

Query Job Scheduler service logs

1. In the startup log, find the **rpc sql** command.

   For example, if the command is **/apsara/deploy/pc_wrapper/rpc.sh spl EtlFramework**, EtlFramework is the name of the Job Scheduler service.

   ```
   [root@vm010025018250 /cloud/app/sls-backend-server/FuxiServiceSlsEtlFramework#/sls_etl_framework/current]
   #tail -n 10 /cloud/app/sls-backend-server/FuxiServiceSlsEtlFramework#/sls_etl_framework/current/log/start.log
   2020-01-07 15:06:55,213 - 83648 - root - tianji_starter.handle_check_alive:353 - INFO - Enter the check alive phase, deploy_flag=True
   2020-01-07 15:06:55,213 - 83648 - root - command_executor.exec_cmd:12 - INFO - Prepare to execute cmd, cmd=[/apsara/deploy/rpc_wrapper/rpc.sh spl EtlFramework]
   2020-01-07 15:06:55,414 - 83648 - root - tianji_proxy_client.report_status:23 - INFO - Prepare to report status, monitor=sls_etl_framework_monitor_app, level=good, description=, hostname=vm010025018250, server_role=sls-backend-server.FuxiServiceSlsEtlFramework#
   2020-01-07 15:06:55,854 - 83648 - root - tianji_starter.do_check_conf_notify:214 - INFO - Check conf_notify, last_check_time=1576942460.83, cur_check_time=1576942460.83
   2020-01-07 15:07:05,357 - 83648 - root - tianji_starter.handle_check_alive:353 - INFO - Enter the check alive phase, deploy_flag=True
   2020-01-07 15:07:05,358 - 83648 - root - command_executor.exec_cmd:12 - INFO - Prepare to execute cmd, cmd=[/apsara/deploy/rpc_wrapper/rpc.sh spl EtlFramework]
   2020-01-07 15:07:05,426 - 83648 - root - tianji_starter.handle_check_alive:353 - INFO - Enter the check alive phase, deploy_flag=True
   2020-01-07 15:07:05,427 - 83648 - root - command_executor.exec_cmd:12 - INFO - Prepare to execute cmd, cmd=[/apsara/deploy/rpc_wrapper/rpc.sh spl EtlFramework]
   2020-01-07 15:07:05,580 - 83648 - root - tianji_proxy_client.report_status:23 - INFO - Prepare to report status, monitor=sls_etl_framework_monitor_app, level=good, description=, hostname=vm010025018250, server_role=sls-backend-server.FuxiServiceSlsEtlFramework#
   2020-01-07 15:07:05,856 - 83648 - root - tianji_starter.do_check_conf_notify:214 - INFO - Check conf_notify, last_check_time=1576942460.83, cur_check_time=1576942460.83
   ```

2. Find the Job Scheduler machine.

```
/apsara/deploy/rpc_wrapper/rpc.sh spl EtlFramework
Partition | WorkerName                                | LastUpdateTime       | status
66    | EtlFrameworkPartitionRole@a34h11080.cloud.h11.amtest87 | Sun Jan  5 16:03:01 2020 | loaded
62    | EtlFrameworkPartitionRole@a34h11080.cloud.h11.amtest87 | Sun Jan  5 16:03:01 2020 | loaded
111   | EtlFrameworkPartitionRole@a34h11080.cloud.h11.amtest87 | Sun Jan  5 16:03:01 2020 | loaded
113   | EtlFrameworkPartitionRole@a34h11080.cloud.h11.amtest87 | Sun Jan  5 16:03:01 2020 | loaded
```

3. Log on to the Job Scheduler machine without using a password.

```
ssh a34h11080.cloud.h11.amtest87
```

4. View the logs.

```
[root@a34h11078.cloud.h11.amtest87 /root]
#ls /apsara/tubo/TempRoot/sys/EtlFramework/EtlFrameworkPartitionRole@a34h11078.cloud.h11.amtest87/etl_worker.LOG
/apsara/tubo/TempRoot/sys/EtlFramework/EtlFrameworkPartitionRole@a34h11078.cloud.h11.amtest87/etl_worker.LOG
```

   - /apsara/tubo/TempRoot/sys/: fixed directory
   - EtlFramework: the service name obtained in Step 1.
   - EtlFrameworkPartitionRole@a34h11078.cloud.h11.amtest87: the Job Scheduler machine name obtained in Step 2.
   - etl_worker.LOG: the log name.

9.13.2.2. Use Log Service Portal to view logs

Project admin

Logstore	Log directory
metering	/tmp/metering_*.LOG metering.log
sls_service_error_log	/alidata/www/logs/java/sls-service/applog/error.log
sls_service_info_log	/alidata/www/logs/java/sls-service/applog/info.log
sls_console_error_log	/alidata/www/logs/java/slserror.log
sls_console_info_log	/alidata/www/logs/java/slsinfo.log
scmc_access_log	/apsara/nginx/logsscmc_access.log
scmc_err_log	/apsara/nginx/logs/scmc_err_log
scmc_op_log	/apsara/nginx/logs/scmc_op_log

Logstore	Log directory
sls_operation_agg_log	/apsara/fcgi_agent/metering_*.LOG
sls_operation_log	/apsara/fcgi_agent/ols_operation*.LOG
offline_scheduler_log	/apsara/sls/import_odps/scheduler/*.[Ll][Oo][Gg]
sls_fastcgi_log	/apsara/fcgi_agent/FastcgiAgent*.LOG
trace_log	/apsara/shennong_agent/tracer/index_worker_trace.LOG
dispatch_worker_log	/apsara/tubo/TempRoot/sys/DispatchWorker/[[user@ip]]/log_dispatch_worker.LOG
etl_framework_log	/apsara/tubo/TempRoot/sys/EtlFramework/[[user@ip]]/etl_worker.LOG
etl_golang_worker_log	/apsara/tubo/TempRoot/sys/EtlFramework/[[user@ip]]/etl_golang_worker.LOG
fc_trigger_log	/apsara/tubo/TempRoot/sys/FcTriggerWorker/[[user@ip]]/fc_trigger.log
query_master_log	/apsara/tubo/TempRoot/sys/QueryMaster/[[user@ip]]/query_master.LOG
sls_configservice_log	/apsara/tubo/TempRoot/sys/ConfigService/[[user@ip]]/sls_config_service.LOG
sls_configservice_query_log	/apsara/tubo/TempRoot/sys/ConfigService/[[user@ip]]/config_service_query.LOG
sls_consumergroup_log	/apsara/tubo/TempRoot/sys/QuotaServer/[[user@ip]]/monitor.LOG
sls_index_status_log	/apsara/tubo/TempRoot/sys/ShennongWorker/[[user@ip]]/project_index_size.LOG
sls_indexworker_log	/apsara/tubo/TempRoot/sys/OlsIndexWorker/[[user@ip]]/ols_index_worker.LOG
sls_loghub_shard_status_log	/apsara/tubo/TempRoot/sys/LoghubMaster/[[user@ip]]/loghub_master_meta.LOG
sls_loghubmaster_log	/apsara/tubo/TempRoot/sys/LoghubMaster/[[user@ip]]/sls_loghub_master.LOG
sls_quotaserver_log	/apsara/tubo/TempRoot/sys/QuotaServer/[[user@ip]]/quota_server.LOG
sls_quotausage_log	/apsara/tubo/TempRoot/sys/QuotaServer/[[user@ip]]/charge.LOG

Logstore	Log directory
sls_replayworker_log	/apsara/tubo/TempRoot/sys/ShennongReplayWorker/[[user@ip]]/shennong_replay_worker.LOG
sls_shennongworker_log	/apsara/tubo/TempRoot/sys/ShennongWorker/[[user@ip]]/shennong_worker.LOG
worker_input_log	/apsara/tubo/TempRoot/sys/ShennongWorker/[[user@ip]]/shennong_worker_input.LOG

Project scmg

Logstore	Log directory
scmg_access_log	/apsara/nginx/logs/scmg_access.log
nginx_error_log	/apsara/nginx/logs/error.log
scmg_err_log	/apsara/nginx/logs/scmg_err_log
scmg_op_log	/apsara/nginx/logs/scmg_op_log
sls_portal_access_log	/apsara/sls/web/logsaccess.log
sls_portal_http_req	/apsara/sls/web/logshttpclient.log
sls_portal_sys_info	/apsara/sls/web/logssysinfo.log
sls_portal_normal	/apsara/sls/web/logsnormal.log
sls_portal_api_audit	/apsara/sls/web/logsapidetail.log

9.14. Apsara Stack Security

9.14.1. Log on to the Apsara Infrastructure Management Framework console

This section describes how to log on to the Apsara Infrastructure Management Framework console.

Prerequisites

You have obtained the URL of the Apsara Stack Operations console and the username and password to log on to the console from your system administrator.

Procedure

1. In the browser address bar, enter *https://Apsara Stack Operations URL*, and press Enter.
2. On the logon page, enter the username and password, and click **Log On**.
3. In the left-side navigation pane, choose **Products** .

4. In the product list, click **Apsara Infrastructure Management Framework** to go to the Apsara Infrastructure Management Framework console.

9.14.2. Routine operations and maintenance of Server Guard

9.14.2.1. Check the service status

9.14.2.1.1. Check the client status

Check the following status information about the Server Guard client to verify that the client is running properly:

Client logs

Client logs are stored in the data directory under the directory of the Server Guard process file, for example, */usr/local/aegis/aegis_client/aegis_xx_xx/data*.

Client logs are saved by day, for example, *data.1 to data.7*

Client's online status

Run the following command to check the client's online status:

```
ps -aux | grep AliYunDun
```

Network connectivity

Run the following command to check whether the client has set up a TCP connection with the server:

```
netstat -tunpe |grep AliYunDun
```

Client UUID

Open the client log file data.x and check the character string following `Currentuid Ret`. This character string is the UUID of the current client.

Client processes

The Server Guard client has three resident processes: AliYunDun, AliYunDunUpdate, and AliHids.

When the client runs properly, all of the three processes run normally.

> **Note** On a Windows OS client, the AliYunDun and AliYunDunUpdate processes exist in the form of services. The service names are Server Guard Detect Service and Server Guard Update Service, respectively.

9.14.2.1.2. Check the status of Aegiserver

Context

To check the running status of Aegiserver, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server of Aegiserver.
2. Run the following command to find the Aegiserver image ID:

   ```
   docker ps -a |grep aegiserver
   ```

   The following message is displayed:

   ```
   b9e59994df41
   reg.docker.alibaba-inc.com/aqs/aegiserverlite@sha256:f9d292f54c58646b672a8533a0d78fba534d26d3
   76a194034e8840c70d9aa0b3 "/bin/bash /startApp." 2 hours ago Up 2 hours 80/tcp, 7001/tcp, 8005/tcp, 8
   009/tcp yundun-aegis.Aegiserverlite__.aegiserverlite. 1484712802
   ```

3. Run the following command to go to the Docker container:

   ```
   docker exec -it [imageId] /bin/bash
   ```

4. Run the following command to check whether the Java process is normal:

   ```
   ps aux |grep aegiserver
   ```

   The following message is displayed:

   ```
   root 153 0.6 25.8 2983812 1084588 ? Sl 12:13 1:01 /opt/taobao/java/bin/java -Djava.util.logging.config.fil
   e=/home/admin/aegiserverlite/.default/conf/logging.properties -Djava.util.logging.manager=org.apach
   e.juli.ClassLoaderLogManager -server -Xms2g -Xmx2g -XX:PermSize=96m -XX:MaxPermSize=384m -Xmn1
   g -XX:+UseConcMarkSweepGC -XX:+UseCMSCompactAtFullCollection -XX:CMSMaxAbortablePrecleanTim
   e=5000 -XX:+CMSClassUnloadingEnabled -XX:+UseCMSInitiatingOccupancyOnly -XX:CMSInitiatingOccup
   ancyFraction=80 -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/home/admin/logs/java.h
   prof -verbose:gc -Xloggc:/home/admin/logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -Djava.
   awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun.net.client.defaultReadTimeo
   ut=30000 -XX:+DisableExplicitGC -Dfile.encoding=UTF-8 -Ddruid.filters=mergeStat -Ddruid.useGloalData
   SourceStat=true -Dproject.name=aegiserverlite -Dcatalina.vendor=alibaba -Djava.security.egd=file:/de
   v/./urandom -Dlog4j.defaultInitOverride=true -Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_EQ
   UALS_IN_VALUE=true -Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_HTTP_SEPARATORS_IN_V0=
   true -Djava.endorsed.dirs=/opt/taobao/tomcat/endorsed -classpath /opt/taobao/tomcat/bin/bootstra
   p.jar:/opt/taobao/tomcat/bin/tomcat-juli.jar -Dcatalina.logs=/home/admin/aegiserverlite/.default/logs
   -Dcatalina.base=/home/admin/aegiserverlite/.default -Dcatalina.home=/opt/taobao/tomcat -Djava.io.t
   mpdir=/home/admin/aegiserverlite/.default/temp org.apache.catalina.startup.Bootstrap -Djboss.serve
   r.home.dir=/home/admin/aegiserverlite/.default -Djboss.server.home.url=file:/home/admin/aegiserverl
   ite/.default start
   ```

5. Run the following command to perform the health check:

   ```
   curl 127.0.0.1:7001/checkpreload.htm
   ```

   If the response is "success", the service is normal.

6. View related logs.
   - **Protocol logs**: View logs about upstream and downstream protocol messages between the server and client in */home/admin/aegiserver/logs/AEGIS_MESSAGE.log*.
   - **Operation logs**: View abnormal stack information during operation in */home/admin/aegiserver/logs/aegis-default.log*.
   - **Offline logs**: View the logs about client disconnection caused by time-out in */home/admin/aegiserver/logs/AEGIS_OFFLINE_MESSAGE.log*.

9.14.2.1.3. Check the Server Guard Update Service status

Context

To check the status of Server Guard Update Service, follow the following steps:

Procedure

1. Run the `ssh host IP address` command to log on to the server of Aegiserver.
2. Run the following command to find the Aegiserver image ID:

   ```
   docker ps -a |grep aegiserver
   ```

3. Run the following command to go to the Docker container:

   ```
   docker exec -it [imageId] /bin/bash
   ```

4. Run the following command to check whether the Java process is normal:

   ```
   ps aux |grep aegisupdate
   ```

5. Run the following command to perform the health check:

   ```
   curl 127.0.0.1:7001/checkpreload.htm
   ```

   If the response is "success", the service is normal.

9.14.2.1.4. Check the Defender module status

Context

To check the status of the Defender module of Server Guard, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Defender module of Server Guard.
2. Run the following command to find the image ID of the Defender module of Server Guard:

   ```
   docker ps -a |grep defender
   ```

3. Run the following command to go to the Docker container:

   ```
   docker exec -it [imageId] /bin/bash
   ```

4. Run the following command to check whether the Java process is normal:

   ```
   ps aux |grep defender
   ```

`ps aux |grep defender`

5. Run the following command to perform health check:

   `curl 127.0.0.1:7001/checkpreload.htm`

   If the response is "success", the service is normal.

9.14.2.2. Restart Server Guard

Context

To restart Server Guard when a fault occurs, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts Server Guard.
2. Run the following command to find the image ID of Server Guard:

   `docker ps -a |grep application name`
3. Run the following command to go to the Docker container:

   `docker exec -it [imageId] /bin/bash`
4. Restart related services.
   - Restart the Server Guard client service.
     - For a server running a Windows OS, go to the service manager, locate *Server Guard Detect Service*, and restart this service.
     - For a server running a Linux OS, use either of the following methods to restart the Server Guard client service:
       - Run the `service aegis restart` command to restart the service.
       - Run the `killall AliYunDun` command as the root user to stop the current process, and then restart the */usr/local/aegis/aegis_client/aegis_xx_xx/AliYunDun* process.
   - Restart the Aegiserver service.
     a. Run the following command to view the Java process ID:

        `ps aux |grep aegiserver`

     b. Run the following command to stop the current process:

        `kill -9 process`

     c. Run the following command to restart the process:

        `sudo -u admin /home/admin/aegiserever/bin/jbossctl restart`

     d. Run the following command to check whether the process has been successfully restarted:

        `curl 127.0.0.1:7001/checkpreload.htm`
   - Restart Server Guard Update Service:

a. Run the following command to view the Java process ID:

```
ps aux |grep aegisupdate
```

b. Run the following command to stop the current process:

```
kill -9 process
```

c. Run the following command to restart the process:

```
sudo -u admin /home/admin/aegisupdate/bin/jbossctl restart
```

d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/checkpreload.htm
```

- Restart the Defender service of Server Guard.
  a. Run the following command to view the Java process ID:

```
ps aux |grep secure-service
```

  b. Run the following command to stop the current process:

```
kill -9 process
```

  c. Run the following command to restart the process:

```
sudo -u admin /home/admin/secure-service/bin/jbossctl restart
```

  d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/checkpreload.htm
```

9.14.3. Routine operations and maintenance of Network Traffic Monitoring System

9.14.3.1. Check the service status

9.14.3.1.1. Basic inspection

The basic inspection of Network Traffic Monitoring System checks whether the service status is normal.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. In the **Clusters** search box, enter **BeaverCluster**.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. On the **Services** tab, enter yundun-beaver-advance in the **Services** search box. Then, check whether the service status is normal.

9.14.3.1.2. Advanced inspection

The advanced inspection of Network Traffic Monitoring System checks the service status and features.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical machines of Network Traffic Monitoring System.
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **BeaverCluster**. Then, click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-beaver-advance** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **BeaverAdvance#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two physical machines of Network Traffic Monitoring System.
3. Check the log status of Network Traffic Monitoring System. Run the `sudo cat /var/log/messages` command. If a record is returned, the log status is normal.
4. Check the status of mirrored traffic. Run the `sudo cat /proc/ixgbe_debug_info` command. If the value of **speed** in the second-to-last row is not 0, the mirrored traffic is normal.
5. Check the protected CIDR block in the log file. Run the `tail -f /dev/shm/banff-2018-xx.log` command. In the command, set *xx* to the month. For example, the log file for May in 2018 is named *banff-2018-05.log*. The CIDR block in the command output is an SLB or EIP CIDR block in the classic network. However, if the CIDR block is connected to Network Traffic Monitoring System through CSWs, a CIDR block in a VPC is returned.
6. Check the network connectivity between Network Traffic Monitoring System and a VM. Run the `ping VM IP address` command to check the network connectivity. In the command, set *VM IP address* to an IP address in the CIDR block returned in the previous step.
7. Check the tcp_decode process status. Run the `ps -ef | grep tcp_decode` command. If a record is returned, the tcp_decode process is normal.
8. Check configurations of the traffic scrubbing server. Run the `cat /home/admin/beaver-dj-schedule/conf/dj.conf` command. Check whether the value of the ip parameter in the aliguard_smart field that is not commented out is set to the DNS virtual IP address mapped to the aliguard.${global:internet-domain} domain name.
9. View the following logs:
   - DDoS alert logs

     Run the `grep -A 10 -B 10 LIDS /var/log/messages` command to view the DDoS alert logs.
   - TCP intercept command logs

     Run the `grep add_to_blacklist.htm /var/log/messages` command to view the TCP intercept command logs.
   - Outbound attack logs

     Run the `grep zombie_new /var/log/messages` command to view the outbound attack logs.

9.14.3.2. Common operations and maintenance

9.14.3.2.1. Restart the Network Traffic Monitoring System process

Context

To restart the Network Traffic Monitoring System process, follow the following steps:

Procedure

1. Log on to the physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to restart the Network Traffic Monitoring System process: `rm -rf /dev/shm/drv_setup_path`

9.14.3.2.2. Uninstall Network Traffic Monitoring System

Context

To uninstall Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to uninstall Network Traffic Monitoring System:

   ```
   bash /opt/beaver/bin/uninstall.sh
   ```

9.14.3.2.3. Disable TCP blocking

Context

To disable TCP blocking for Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Open the */beaver_client.sh* file on each server of Network Traffic Monitoring System, and add a number sign (`#`) to the start of the `./tcp_reset` line to comment out the line.
4. Run the following command on each server of Network Traffic Monitoring System to disable TCP blocking:

   ```
   killall tcp_reset
   ```

9.14.3.2.4. Enable TCPDump

Context

To enable TCPDump for Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to enable TCPDump:

   ```
   echo 1 > /proc/ixgbe_debug_dispatch
   ```

   > **Note**
   >
   > When TCPDump is enabled, the performance of Network Traffic Monitoring System may be affected. We recommend that you run the following command to disable TCPDump after packet capture is complete.
   >
   > ```
   > echo 0 > /proc/ixgbe_debug_dispatch
   > ```

9.14.4. Routine operations and maintenance of Anti-DDoS Service

9.14.4.1. Check the service status

9.14.4.1.1. Basic inspection

The basic inspection of Traffic Scrubbing checks whether the service status is normal.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. In the **Clusters** search box, enter **AliguardCluster**.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. On the **Services** tab, enter yundun-aliguard in the **Services** search box. Then, check whether the service status is normal.

9.14.4.1.2. Advanced inspection

The advanced inspection of Traffic Scrubbing checks the service status and features.

Procedure

1. Log on to the two physical machines of Traffic Scrubbing.

i. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose **Operations > Cluster Operations**.

ii. In the **Clusters** search box, enter **AliguardCluster**. Then, click the cluster name. The **Cluster Details** page appears.

iii. On the **Services** tab, enter **yundun-aliguard** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.

iv. In the **Service Role** search box, enter **AliguardConsole#**.

v. View the machine information, and click **Terminal** in the Actions column to log on to the two physical machines of Traffic Scrubbing.

2. Check the deployment status of Traffic Scrubbing. Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard_defender_check` command and check the command output.

> **Note** If the host of Traffic Scrubbing has just restarted, wait for 3 to 5 minutes before you run the command to check the deployment status.

- If `aliguard status check OK!` is returned, Traffic Scrubbing is properly deployed, and its service status is normal, as shown in the Traffic Scrubbing status figure.

  Traffic Scrubbing status

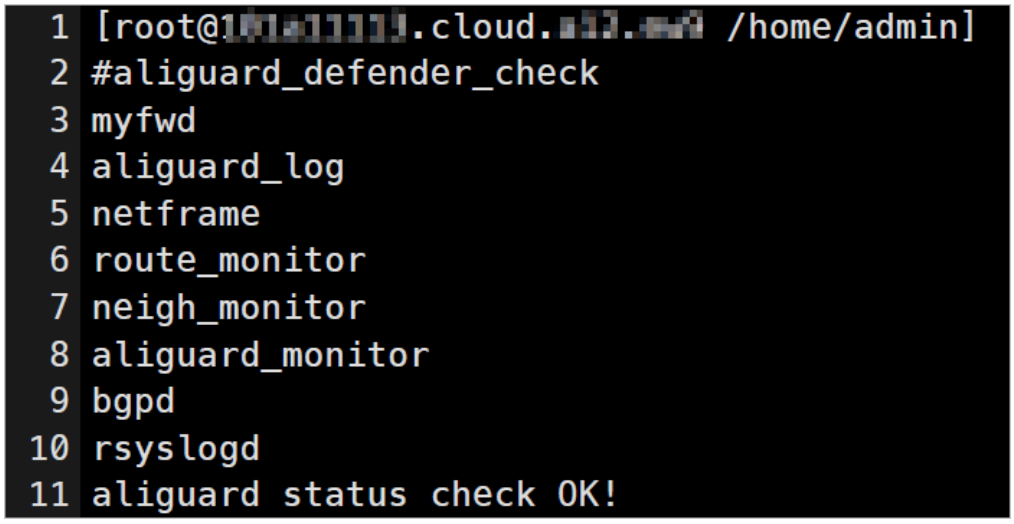

- If the error message shown in Reinjection route error message is returned, the reinjection route is incorrect.

  Reinjection route error message

```
Error: route status error, we need two default routes to reinject the net flow!
Error: route error, can't get to the target ip.
```

  Troubleshooting: The reinjection route is a default route generated by Traffic Scrubbing. Its next hop is the ISW interface that is bound to the VPN. If an error occurs, check whether this route is generated by Traffic Scrubbing. If the route is generated, check ISW configurations to determine whether the route to downstream devices is available.

- If the error message shown in BGP route error message is returned, the BGP route is incorrect.

  BGP route error message

```
1 Error: bgp status error!
```

Troubleshooting: If the BGP route is incorrect, troubleshoot the error based on the following operations:

a. Check whether the BGP neighbor is normal on the ISW.

b. Check whether the destination of the BGP route is an attacked IP address with a 32-bit subnet mask and the next hop of the BGP route is the Traffic Scrubbing address.

c. Check whether the BGP routing policy on the ISW is correct.

- If other errors are reported, the core process is faulty. Contact Alibaba Cloud technical support.

3. Check the status of the NICs or optical modules of Traffic Scrubbing.

> **Note** Traffic Scrubbing must use NICs or optical modules equipped with Intel X520 or Intel 82599.

Run the `lspci | grep Eth` command. Information containing Intel X520 or Intel 82599 is returned.

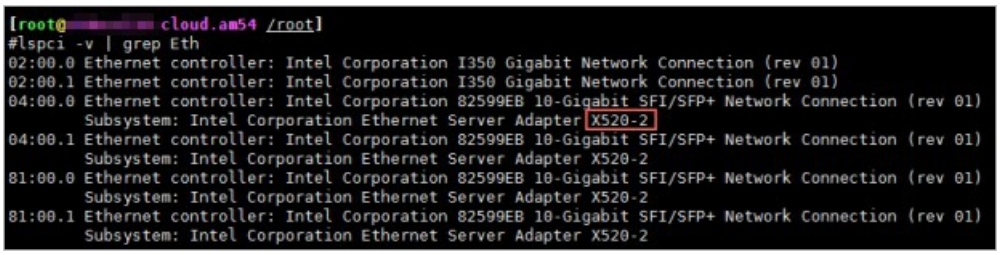

9.14.4.2. Common operations and maintenance

9.14.4.2.1. Restart Anti-DDoS Service

Context

To restart Anti-DDoS Service when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts Anti-DDoS Service.
2. Run the following command to stop Anti-DDoS Service: `/home/admin/aliguard/target/AliguardDefender/bin/aliguard stop`

> **Note** If the `ERROR: Module net_msg is in use` message is displayed, run the command again later. If Anti-DDoS Service cannot be stopped after several attempts, restart the server of Anti-DDoS Service.

3. Run the following command to restart Anti-DDoS Service: `/home/admin/aliguard/target/AliguardDefender/bin/aliguard start`
4. Run the service status check command five minutes after Anti-DDoS Service is restarted.

9.14.4.2.2. Troubleshoot common faults

Context

When an error occurs in Anti-DDoS Service, follow the following troubleshooting steps:

Procedure

1. Restart Anti-DDoS Service.
   - If Anti-DDoS Service is in the normal status after being restarted but an error message is returned during the health check performed later, non-standard NICs or optical modules are used. To check whether standard NICs or optical modules are used, see Check the status of the NICs or optical modules of Anti-DDoS Service. If non-standard NICs or optical modules are used, change the NICs or optical modules.
   - If Anti-DDoS Service is in an unusual status after being restarted, go to the next step.
2. View the `aliguard_dynamic_config` file. Carefully check whether each configuration item in the file is exactly the same as that in the plan.

   > **Note** Ensure that the AS number specified in `aliguard local` is 65515 and that the BGP password is correct.

3. Check the wiring and switch configuration.

   > **Note** If any incorrect configuration is found, the current fault is caused by incorrect wiring or switch IP address configuration, rather than incorrect deployment of Anti-DDoS Service. In this case, contact the network engineer.

   Assume that the Anti-DDoS Service configurations to be checked are listed in the following figure, among which the server IP address is 10.1.4.12. To check whether the four ports of Anti-DDoS Service can ping the ports of the switch, follow the following steps:

   Anti-DDoS Service configuration example

aliguard_host_ip	port	aliguard_port_ip	csr_port_ip
10.1.4.12	T0	10.1.0.34	10.1.0.33
10.1.4.12	T1	10.1.0.38	10.1.0.37
10.1.4.12	T2	10.1.0.50	10.1.0.49
10.1.4.12	T3	10.1.0.54	10.1.0.53
10.1.4.28	T0	10.1.0.42	10.1.0.41
10.1.4.28	T1	10.1.0.46	10.1.0.45
10.1.4.28	T2	10.1.0.58	10.1.0.57
10.1.4.28	T3	10.1.0.62	10.1.0.61

   i. Run the following commands to check the NIC PCI IDs of Anti-DDoS Service:

   ```
   cd /sys/bus/pci/drivers/igb_uio
   ls
   ```

   Record the PCI IDs of the four NICs, for example, 0000:01:00.0, 0000:01:00.1, 0000:82:00.0, and 0000:82:00.1.

ii. Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard stop` command to stop Anti-DDoS Service.

iii. In the /sys/bus/pci/drivers/igb_uio directory, unbind the four NICs recorded in the first step from the igb_uio driver, as shown in Unbind NICs.

Unbind NICs

```
1 echo "0000:01:00.0"  >> unbind
2 echo "0000:01:00.1"  >> unbind
3 echo "0000:82:00.0"  >> unbind
4 echo "0000:82:00.1"  >> unbind
```

iv. In the /sys/bus/pci/drivers/ixgbe directory, bind the four NICs to the ixgbe driver for Linux, as shown in Bind NICs.

Bind NICs

```
1 echo "0000:01:00.0"  >> bind
2 echo "0000:01:00.1"  >> bind
3 echo "0000:82:00.0"  >> bind
4 echo "0000:82:00.1"  >> bind
```

v. Set Anti-DDoS Service IP addresses for the NICs.

The local server IP address is 10.1.4.12, and the NIC IP addresses are set to 10.1.0.34, 10.1.0.38, 10.1.0.50, and 10.1.0.54, as shown in Anti-DDoS Service configuration example.

a. Run the `ifconfig-a` command to display all NICs, and run the `ethtool-i` command to view the PCI ID of each NIC. Find the four NICs of which the IDs are the same as those recorded in the first step, for example, eth0, eth1, eth2, and eth3.

b. Run the following commands to move these NICs to the top of the queue:

```
ifconfig eth0 up
ifconfig eth1 up
ifconfig eth2 up
ifconfig eth3 up
```

c. Set Anti-DDoS Service IP addresses for the NICs. Run the following commands to set Anti-DDoS Service IP addresses for the NICs based on their PCI IDs in an ascending order:

```
ifconfig eth0 10.1.0.34 netmask 255.255.255.252
ifconfig eth1 10.1.0.38 netmask 255.255.255.252
ifconfig eth2 10.1.0.50 netmask 255.255.255.252
ifconfig eth3 10.1.0.54 netmask 255.255.255.252
```

vi. Try to ping the peer IP addresses configured. If the peer IP addresses cannot be pinged, the switch configuration or wiring is incorrect.

```
ping 10.1.0.33
ping 10.1.0.37
ping 10.1.0.49
ping 10.1.0.53
```

vii. If these four IP addresses can all be pinged, you can directly start Anti-DDoS Service without unbinding the NICs.

Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard start` command to start Anti-DDoS Service.

After Anti-DDoS Service has been started for a while, run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard_rule -v 0.0.0.0 -d drop_icmp` command to disable the drop_icmp policy.

viii. Ping the peer IP addresses again.

```
ping 10.1.0.33
ping 10.1.0.37
ping 10.1.0.49
ping 10.1.0.53
```

If the peer IP addresses cannot be pinged, non-standard NICs or optical modules are used or the configuration is incorrect.

4. If these four peer IP addresses can be pinged after Anti-DDoS Service is started but an error is reported during a status check of Anti-DDoS Service, contact Alibaba Cloud technical support.

9.14.5. Routine operations and maintenance of Threat Detection Service

9.14.5.1. Check the service status

9.14.5.1.1. Basic inspection

The basic inspection of Threat Detection Service (TDS) checks whether the service status is normal.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. In the **Clusters** search box, enter **BasicThinCluster**.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. On the **Services** tab, enter yundun-sas in the **Services** search box. Then, check whether the service

status is normal.

9.14.5.1.2. Advanced inspection

The advanced inspection of Threat Detection Service (TDS) checks the service status and features.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical machines of TDS.
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **BasicThinCluster**. Then, click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-sas** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SasApp#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two physical machines of TDS.
3. Log on to the two Docker containers of TDS. Run the `sudo docker exec -it $(sudo docker ps | grep sas | awk '{print $1}') bash` command.
4. Check the service process status. Run the `ps aux |grep sas` command. If a record is returned, the process is normal.
5. Check the health status. Run the `curl 127.0.0.1:3008/check.htm` command. If ok is returned, the service is normal.
6. View logs.
   - View all logs in the */home/admin/sas/logs/sas-default.log* file, including metaq message logs, execution logs of scheduled tasks, and error logs. You can locate TDS faults based on these logs.
   - View info logs generated when TDS is running in the */home/admin/sas/logs/common-default.log* file.
   - View TDS error logs in the */home/admin/sas/logs/common-error.log* file.
   - View logs about metaq messages received by TDS in the */home/admin/sas/logs/SAS_LOG.log* file.

     > **Note** Asset verification is performed on messages in this log file. Therefore, the number of messages in this log file is less than that in the sas-default.log file.

   - View logs generated when the alert contact sends alert notifications in the */home/admin/sas/logs/notify.log* file.

9.14.5.2. Restart Threat Detection Service

Context

When a fault occurs, you can restart Threat Detection Service (TDS).

Procedure

1. Run the `ssh Host IP address` command to log on to the host of TDS.
2. Run the following command to find the image ID of TDS:

   `docker ps -a |grep sas`
3. Run the following command to enter the Docker container:

   `docker exec -it [imageId] /bin/bash`
4. Run the following command to find the Java process:

   `ps aux |grep sas`
5. Run the following command to stop the process:

   `kill -9 Process ID`
6. Run the following command to restart the process:

   `sudo -u admin /home/admin/sas/bin/jbossctl restart`
7. Run the following command to check whether the process is restarted:

   `curl 127.0.0.1:7001/check.htm`

9.14.6. Routine operations and maintenance of Cloud Firewall

9.14.6.1. Check the service status

To check the running status of the Cloud Firewall server, follow the following steps:

Procedure

1. Run the **ssh server IP address** command to log on to the server that hosts Cloud Firewall.
2. Run the following command to find the image ID of the Cloud Firewall server: `sudo docker ps | grep cloudfirewall`

   The following message is displayed:

   ```
   af8a1a182a17  reg.docker.aliyun-inc.com/yundun-advance/cloudfirewall:696dcf568512deceb7199dc4c
   9aa70855e66aca71244296cf13bfaf4a0897ebe  "/bin/bash /home/admi"  23 hours ago  Up 23 hours  yu
   ndun-cloudfirewall.CloudFirewallApp__.cloudfirewall-app. 1523453835
   ```
3. Run the following command to go to the Docker container: `docker exec –it [imageId] /bin/bash`
4. Run the following command to check whether the Java process is normal: `ps aux |grep cloudfirewall`

   The following message is displayed:

```
admin 118 0.5 28.2 6261808 2368828 ? Sl Apr11 8:06 /opt/taobao/java/bin/java -Djava.util.logging.config.file=/home/admin/cloud-firewall/.default/conf/logging.properties -Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager -server -Xms4g -Xmx4g -XX:PermSize=96m -XX:MaxPermSize=384m -Xmn2g -XX:+UseConcMarkSweepGC -XX:+UseCMSCompactAtFullCollection -XX:CMSMaxAbortablePrecleanTime=5000 -XX:+CMSClassUnloadingEnabled -XX:+UseCMSInitiatingOccupancyOnly -XX:CMSInitiatingOccupancyFraction=80 -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/home/admin/logs/java.hprof -verbose:gc -Xloggc:/home/admin/logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -Djava.awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun.net.client.defaultReadTimeout=30000 -XX:+DisableExplicitGC -Dfile.encoding=UTF-8 -Ddruid.filters=mergeStat -Ddruid.useGloalDataSourceStat=true -Dproject.name=cloud-firewall -Dhsf.server.port=21015 -Djdk.tls.ephemeralDHKeySize=2048 -Dcatalina.vendor=alibaba -Djava.security.egd=file:/dev/./urandom -Dlog4j.defaultInitOverride=true -Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_EQUALS_IN_VALUE=true -Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_HTTP_SEPARATORS_IN_V0=true -Dcatalina.logs=/home/admin/cloud-firewall/.default/logs -Dignore.endorsed.dirs= -classpath /opt/taobao/tomcat/bin/bootstrap.jar:/opt/taobao/tomcat/bin/tomcat-juli.jar -Dcatalina.base=/home/admin/cloud-firewall/.default -Dcatalina.home=/opt/taobao/tomcat -Djava.io.tmpdir=/home/admin/cloud-firewall/.default/temp org.apache.catalina.startup.Bootstrap -Djboss.server.home.dir=/home/admin/cloud-firewall/.default -Djboss.server.home.url=file:/home/admin/cloud-firewall/.default startroot 4931 0.0 0.0 61208 764 ? S+ 21:32 0:00 grep cloud-firewall
```

5. Run the following command to perform the health check: `curl 127.0.0.1:2015/cloud-firewall/check_health`

   If OK is returned, the service is normal.

6. View related logs.
   - View the routine printing logs of the Cloud Firewall server in */home/admin/cloud-firewall/logs/cloud-firewall-info.log*.
   - View the error printing logs of the Cloud Firewall server in */home/admin/cloud-firewall/logs/cloud-firewall-error.log*.
   - View the logs returned after Cloud Firewall calls OpenAPI in */home/admin/cloud-firewall/logs/cloud-firewall-openApi.log*.

9.14.6.2. Restart Cloud Firewall

To restart Cloud Firewall when an error occurs, follow the following steps:

Procedure

1. Run the **ssh server IP address** command to log on to the server that hosts Cloud Firewall.
2. Run the following command to find the image ID of Cloud Firewall: `sudo docker ps | grep cloudfirewall`
3. Run the following command to go to the Docker container: `docker exec –it [imageId] /bin/bash`
4. Restart Cloud Firewall.

i. Run the following command to view the Java process ID: `ps -aux | grep cloud-firewall`

ii. Run the following command to stop the current process: `kill -9 process`

iii. Run the following command to restart the process: `sudo –u admin /home/admin/cloud-firewall/bin/jbossctl restart`

iv. Run the following command to check whether the process has been successfully restarted: `curl 127.0.0.1:2015/cloud-firewall/check_health`

If OK is returned, the service is normal.

9.14.7. Routine operations and maintenance of WAF

9.14.7.1. Check the service status

9.14.7.1.1. Basic inspection

The basic inspection of Web Application Firewall (WAF) checks whether the service status is normal.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. In the **Clusters** search box, enter **SemaWaf Cluster**.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. On the **Services** tab, enter yundun-semawaf in the **Services**. Then, check whether the service status is normal.

9.14.7.1.2. Advanced inspection

The advanced inspection of Web Application Firewall (WAF) checks the system status and service status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical machines of WAF.
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **SemaWaf Cluster**. Then, click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-semawaf** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **YundunSemawafApp#**.

v. View the machine information, and click **Terminal** in the Actions column to log on to the two physical machines of WAF.

3. Check the system status.
   i. View system logs. Run the `dmesg -T |tail -30` command to check for exception logs.
   ii. Check the system loads.
      - Run the `free -h` command to check whether the memory usage is normal.
      - Run the `df -h` command to check whether the disk usage is normal.
      - Run the `uptime` command to check whether the average system load is normal.
      - Run the `top` command to check whether the CPU utilization is normal.
4. Check the service status.

   > **Note** Perform the following operations in the WAF installation directory, which is */home/safeline* by default.

   i. Run the `cd /home/safeline` command to open the installation directory.
   ii. Check the minion service.
      a. Run the `systemctl status minion` command to check the execution time and status of the Minion service.
      b. Run the `tail -100 logs/minion/minion.log` command to check for exception logs.
   iii. Check the mgt-api service.
      a. Run the `docker logs --tail 50 mgt-api` command to check for exception logs.
      b. Run the `docker exec -it mgt-api supervisorctl status` command to check whether the service properly runs and whether uptime is normal.
      c. Run the `tail -50 logs/management/gunicorn.log` command to check for exception logs.
      d. Run the `tail -50 logs/management/daphne.log` command to check for exception logs.
      e. Run the `tail -50 logs/management/scheduler.log` command to check for exception logs.
      f. Run the `tail -50 logs/management/dramatiq.log` command to check for exception logs.
   iv. Check the redis service. Run the `docker logs --tail 50 mgt-redis` command to check for exception logs.
   v. Check the detector service.
      a. Run the `docker logs --tail 50 detector-srv` command to check for exception logs.
      b. Run the `tail -50 logs/detector/snserver.log` command to check for exception logs.
      c. Run the `curl 127.0.0.1:8001/stat | grep num` command to check whether the service responds and whether the real-time request processing data is normal. For example, check the req_num_total parameter, which indicates the number of requests that were processed within the last 5 seconds.

vi. Check the tengine service.

a. Run the `docker logs --tail 50 tengine` command to check for exception logs.

b. Run the `tail -50 logs/nginx/error.log` command to check for exception logs.

vii. Check the mario service.

a. Run the `docker logs --tail 50 mario` command to check for exception logs.

b. Run the `tail -50 logs/mario/mario.log` command to check for exception logs.

c. Run the `curl 127.0.0.1:3335/api/v1/state` command to check whether the service responds and whether the real-time request processing data is normal. For example, check whether the num_pending parameter remains at a high value of nearly 10,000 or whether the num_processed_last_10s parameter, which indicates the number of requests that were processed within the last 10 seconds, is normal.

9.14.8. Routine operations and maintenance of Sensitive Data Discovery and Protection

9.14.8.1. Check the service status

9.14.8.1.1. Basic inspection

During the basic inspection of Sensitive Data Discovery and Protection (SDDP), check whether the service has reached the final status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Choose **Operations > Project Operations**.
3. In the **Fuzzy Search** field, enter *yundun-sddp*.
4. Click **Details** in the **Actions** column of the yundun-sddp project to go to the **Cluster Operations** page.
5. In the cluster list, click the cluster name that starts with **SddpCluster**.
6. In the **Service Instances** section of the **Cluster Dashboard** page, check whether the **yundun-sddp** service instance is in the final status.

9.14.8.1.2. Advanced inspection: Check the status of the SddpService service

This topic describes how to check the running status of the SddpService service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.

i. Choose **Operations > Project Operations**.

ii. In the **Fuzzy Search** field, enter *yundun-sddp*. Click **Details** in the **Actions** column of the yundun-sddp project to go to the **Cluster Operations** page.

iii. In the cluster list, click the cluster name that starts with **SddpCluster**.

iv. In the Service Instances section, find **yundun-sddp** and click **Details** in the **Actions** column to go to the **Service Instance Information Dashboard** page.

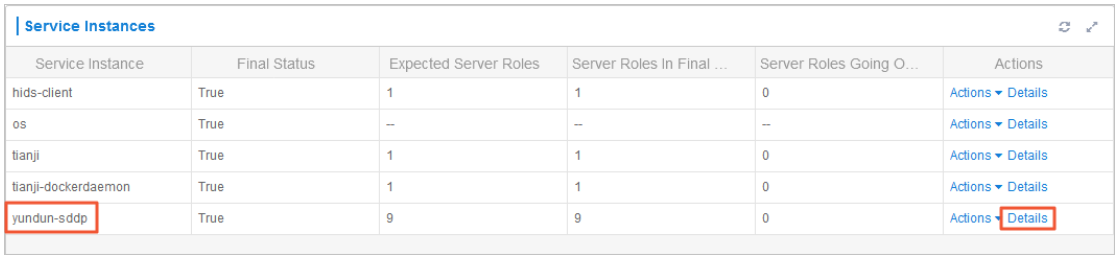

v. In the **Server Role List** section, find **SddpService#** and click **Details** in the **Actions** column to go to the **Server Role Dashboard** page.

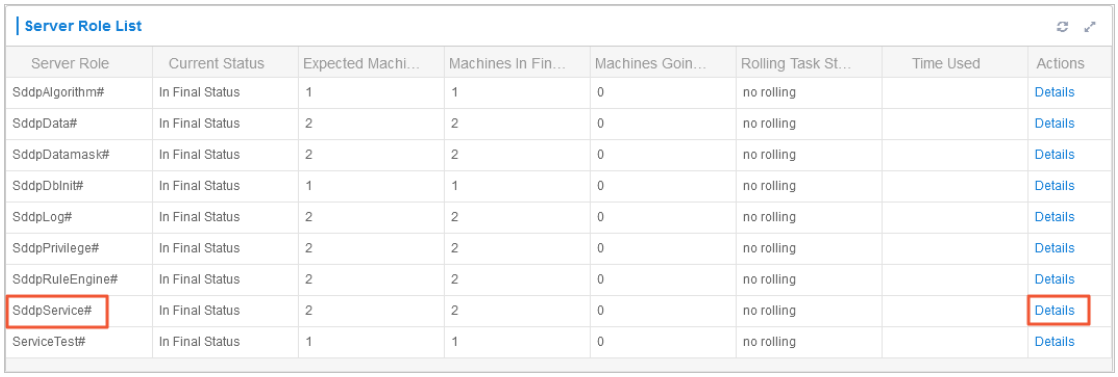

vi. In the **Machine Information** section, click **Terminal** in the **Actions** column to log on to the two physical servers of SDDP, respectively.

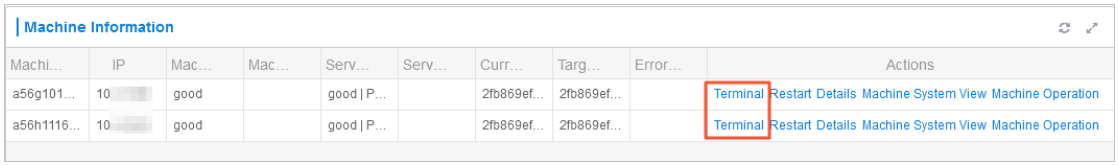

3. Log on to two Docker containers of the SddpService service, respectively. Run the `sudo docker exec -it $(sudo docker ps | grep SddpService | awk '{print $1}') bash` command.

4. Check the process status of the SddpService service. Run the `ps aux | grep java | grep yundun-sddp-service` command. If any record is returned, the service is normal.

```
#ps aux | grep java | grep yundun-sddp-service
root       162  0.1 30.7 7224188 2579604 ?     Sl   May31  26:35 /opt/taobao/java/bin/java -Dspring.profiles.active=cloud -server -Xms4g -Xmx4g -Xmn2g -XX:MetaspaceSize=256m -XX:MaxMetaspaceSize=512m -XX:MaxDirectMemorySize=1g -XX:SurvivorRatio=10 -XX:+UseConcMarkSweepGC -XX:CMSMaxAbortablePrecleanTime=5000 -XX:+CMSClassUnloadingEnabled -XX:CMSInitiatingOccupancyFraction=80 -XX:+UseCMSInitiatingOccupancyOnly -XX:+ExplicitGCInvokesConcurrent -Dsun.rmi.dgc.server.gcInterval=2592000000 -Dsun.rmi.dgc.client.gcInterval=2592000000 -XX:ParallelGCThreads=4 -Xloggc:/root/logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/root/logs/java.hprof -Djava.awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun.net.client.defaultReadTimeout=30000 -DJM.LOG.PATH=/root/logs -DJM.SNAPSHOT.PATH=/root/snapshots -Dfile.encoding=UTF-8 -Dhsf.publish.delayed=true -Dproject.name=yundun-sddp-service -Dpandora.boot.wait=true -Dlog4j.defaultInitOverride=true -Dserver.port=7001 -Dmanagement.port=7002 -Dmanagement.server.port=7002 -Dpandora.location=/home/admin/yundun-sddp-service/target/taobao-hsf.sar -classpath /home/admin/yundun-sddp-service/target/yundun-sddp-service -Dapp.location=/home/admin/yundun-sddp-service/target/yundun-sddp-service -Djava.endorsed.dirs= -Djava.io.tmpdir=/home/admin/yundun-sddp-service/.default/temp com.taobao.pandora.boot.loader.SarLauncher
```

5. Check the health status. Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If the response is **success**, the service is normal.

```
#curl 127.0.0.1:7001/checkpreload.htm
"success"
```

6. View related logs.
   - View common logs in the */home/admin/yundun-sddp-service/logs/common-log.log* file.
   - View application logs in the */home/admin/yundun-sddp-service/logs/application.log* file.
   - View front-end request logs in the */home/admin/yundun-sddp-service/logs/common-request.log* file.
   - View system logs in the */home/admin/yundun-sddp-service/logs/service-stdout.log* file.

9.14.8.1.3. Advanced inspection: Check the status of the SddpData service

This topic describes how to check the running status of the SddpData service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
   i. Choose **Operations > Project Operations**.
   ii. In the **Fuzzy Search** field, enter *yundun-sddp*. Click **Details** in the **Actions** column of the yundun-sddp project to go to the **Cluster Operations** page.
   iii. In the cluster list, click the cluster name that starts with **SddpCluster**.
   iv. In the **Service Instances** section, find **yundun-sddp** and click **Details** in the **Actions** column to go to the **Service Instance Information Dashboard** page.
   v. In the **Server Role List** section, find **SddpData#** and click **Details** in the **Actions** column to go to the **Server Role Dashboard** page.
   vi. In the **Machine Information** section, click **Terminal** in the **Actions** column to log on to the two physical servers of SDDP, respectively.
3. Log on to two Docker containers of the SddpData service, respectively. Run the `sudo docker exec -it $(sudo docker ps | grep SddpData | awk '{print $1}') bash` command.
4. Check the process status of the SddpData service. Run the `ps aux | grep yundun-sddp-data` command. If any record is returned, the service is normal.
5. View related logs. View logs in the */home/admin/yundun-sddp-data/logs/sddp.log* file.

9.14.8.1.4. Advanced inspection: Check the status of the SddpPrivilege service

This topic describes how to check the running status of the SddpPrivilege service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
   i. Choose **Operations > Project Operations**.
   ii. In the **Fuzzy Search** field, enter *yundun-sddp*. Click **Details** in the **Actions** column of the yundun-sddp project to go to the **Cluster Operations** page.
   iii. In the cluster list, click the cluster name that starts with **SddpCluster**.
   iv. In the **Service Instances** section, find **yundun-sddp** and click **Details** in the **Actions** column to go to the **Service Instance Information Dashboard** page.
   v. In the **Server Role List** section, find **SddpPrivilege#** and click **Details** in the **Actions** column to go to the **Server Role Dashboard** page.
   vi. In the **Machine Information** section, click **Terminal** in the **Actions** column to log on to the two physical servers of SDDP, respectively.
3. Log on to two Docker containers of the SddpPrivilege service, respectively. Run the `sudo docker exec -it $(sudo docker ps | grep SddpPrivilege | awk '{print $1}') bash` command.
4. Check the process status of the SddpPrivilege service. Run the `ps aux | grep java | grep yundun-sddp-privilege` command. If any record is returned, the service is normal.
5. Check the health status. Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If the response is **success**, the service is normal.
6. View related logs.
   - View exception logs in the */home/admin/yundun-sddp-privilege/logs/exception.log* file.
   - View application logs in the */home/admin/yundun-sddp-privilege/logs/application.log* file.
   - View task logs in the */home/admin/yundun-sddp-privilege/logs/task.log* file.
   - View system logs in the */home/admin/yundun-sddp-privilege/logs/service-stdout.log* file.

9.14.8.1.5. Advanced inspection: Check the status of the SddpLog service

This topic describes how to check the running status of the SddpLog service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
   i. Choose **Operations > Project Operations**.
   ii. In the **Fuzzy Search** field, enter *yundun-sddp*. Click **Details** in the **Actions** column of the yundun-sddp project to go to the **Cluster Operations** page.
   iii. In the cluster list, click the cluster name that starts with **SddpCluster**.
   iv. In the **Service Instances** section, find **yundun-sddp** and click **Details** in the **Actions** column to go to the **Service Instance Information Dashboard** page.

v. In the **Server Role List** section, find **SddpLog#** and click **Details** in the **Actions** column to go to the **Server Role Dashboard** page.

vi. In the **Machine Information** section, click **Terminal** in the **Actions** column to log on to the two physical servers of SDDP, respectively.

3. Log on to two Docker containers of the SddpLog service, respectively. Run the `sudo docker exec -it $(sudo docker ps | grep SddpLog | awk '{print $1}') bash` command.
4. Check the process status of the SddpLog service. Run the `ps aux | grep java | grep yundun-sddp-log` . If any record is returned, the service is normal.
5. Check the health status. Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If the response is **success**, the service is normal.
6. View related logs.
   - View exception logs in the */home/admin/yundun-sddp-log/logs/exception.log* file.
   - View application logs in the */home/admin/yundun-sddp-log/logs/application.log* file.
   - View debug logs in the */home/admin/yundun-sddp-log/logs/debug.log* file.
   - View system logs in the */home/admin/yundun-sddp-log/logs/service-stdout.log* file.

9.14.8.2. Restart SDDP

This topic describes how to restart Sensitive Data Discovery and Protection (SDDP) when a fault occurs.

Procedure

1. Run the `ssh Server IP address` command to log on to the server that hosts SDDP.
2. Run the following command to find the image ID of the service:

   ```
   docker ps -a |grep service name
   ```

3. Run the following command to log on to the Docker container:

   ```
   docker exec -it [imageId] /bin/bash
   ```

4. Restart related services.
   - Restart the yundun-sddp-service service.
     a. Run the following command to stop the current process:

        ```
        kill -9 $(ps -ef | grep java | grep yundun-sddp-service | grep -v grep | awk '{print$2}')
        ```

     b. Run the following command to restart the process:

        ```
        /bin/bash /home/admin/start.sh
        ```

     c. Run the following command to check whether the process is restarted:

        ```
        curl 127.0.0.1:7001/check.htm
        ```

        If the response is **success**, the service is normal.
   - Restart the yundun-sddp-log service.

a. Run the following command to stop the current process:

```
kill -9 $(ps -ef | grep java | grep yundun-sddp-log | grep -v grep | awk '{print $2}')
```

b. Run the following command to restart the process:

```
/bin/bash /home/admin/start.sh
```

c. Run the following command to check whether the process is restarted:

```
curl 127.0.0.1:7001/check.htm
```

If the response is **success**, the service is normal.

- Restart the yundun-sddp-privilege service.
  a. Run the following command to stop the current process:

```
kill -9 $(ps -ef | grep java | grep yundun-sddp-privilege | grep -v grep | awk '{print $2}')
```

  b. Run the following command to restart the process:

```
/bin/bash /home/admin/start.sh
```

  c. Run the following command to check whether the process is restarted:

```
curl 127.0.0.1:7001/check.htm
```

  If the response is **success**, the service is normal.

- Restart the yundun-sddp-data service.
  a. Run the following command to stop the current process:

```
kill -9 $(ps -ef | grep yundun-sddp-data | grep -v grep | awk '{print $2}')
```

  b. Run the following command to restart the process:

```
/bin/bash /home/admin/yundun-sddp-data/start.sh
```

  c. Check whether the process is restarted.

  Run the `ps aux | grep yundun-sddp-data` command. If any record is returned, the service is normal.

9.14.9. Routine operations and maintenance of Apsara Stack Security Center

9.14.9.1. Check service status

9.14.9.1.1. Basic inspection

During the basic inspection of Apsara Stack Security Center, check whether the service has reached the final status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.

2. Choose **Operations > Project Operations**. Enter *yundun-advance*, and click Details to go to the Cluster Operations page.
3. Select **BasicCluster**.
4. Check whether yundun-secureconsole has reached the final status in **Service Instances List**.

9.14.9.1.2. Advanced inspection

Check the running status of Apsara Stack Security Center.

Context

To check the running status of Apsara Stack Security Center, follow the following steps:

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to two physical machines, respectively.
   i. Choose **Operations > Project Operations**.
   ii. Enter *yundun-advance*, and click **Details** to go to the Cluster Operations page.
   iii. Select **BasicCluster**.
   iv. Select **yundun-secureconsole** from Service Instances List, and click **Details** to go to the **Service Instance Dashboard** page.
   v. Select **SecureConsoleApp#** from Service Role List, and click **Details** to go to the **Service Role Dashboard** page.
   vi. View Server Information, and use TerminalService to log on to two physical machines, respectively.
3. Log on to two secure-console Docker containers, respectively. Run `sudo docker exec -it $(sudo docker ps | grep secureconsole | awk '{print $1}') bash`.
4. Check the console progress status. Run `ps aux |grep console`. If any record is returned, the console progress is normal.
5. Check the health status. Run `curl 127.0.0.1:3014/check.htm`. If `OK` is returned, the service is normal.
6. View related logs.
   - View the Tomcat logs in */home/admin/console/logs/jboss_stdout.log*.

9.14.9.2. Restart the secure-console service

Context

To restart the secure-console service when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the secure-console service.
2. Run the following command to find the image ID of the secure-console service:

   `sudo docker ps -a |grep console`

```
sudo docker ps -a |grep console
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to locate the Java process:

```
ps aux |grep console
```

5. Run the following command to stop the current process:

```
kill -9 process
```

6. Run the following command to restart the process:

```
sudo -u admin /home/admin/console/bin/jbossctl restart
```

7. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/check.htm
```

9.14.10. Routine operations and maintenance of secure-service

9.14.10.1. Check the service status

9.14.10.1.1. Basic inspection

During the basic inspection of secure-service, check whether the service has reached the final status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Choose **Operations > Project Operations**. On the page that appears, enter *yundun-advance*, and click Details to go to the Cluster Operations page.
3. Select **BasicCluster**.
4. Check whether yundun-secureservice has reached the final status in **Service Instances List**.

9.14.10.1.2. Advanced inspection: Check the secure-service status

This topic describes how to check the secure-service running status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to two physical machines, respectively.
   i. Choose **Operations > Project Operations**.
   ii. Enter *yundun-advance*, and click **Details** to go to the Cluster Operations page.
   iii. Select **BasicCluster**.

iv. Select **yundun-secureservice** from Service Instances List, and click **Details** to go to the **Service Instance Dashboard** page.

v. Select **SecureServiceApp#** from Service Role List, and click **Details** to go to the **Service Role Dashboard** page.

vi. View Server Information, and click Terminal to log on to two physical machines, respectively.

3. Log on to two secure-service Docker containers, respectively. Run `sudo docker exec -it $(sudo docker ps | grep secureservice | awk '{print $1}') bash` .
4. Check the secure-service process status. Run `ps aux |grep secure-service` . If any record is returned, the secure-service process is normal.
5. Check the health status. Run `curl 127.0.0.1:3010` . If `OK` is returned, the service is normal.
6. Run the following command to go to the Docker container:

   ```
   sudo docker exec -it [imageId] /bin/bash
   ```

7. View related logs.
   - View the Server Guard logs in */home/admin/secure-service/logs/aegis-info.log*.
   - View the error logs in */home/admin/secure-service/logs/Error*.
   - View the vulnerability analysis and scanning logs in */home/admin/secure-service/logs/leakage-info.log*.
   - View the cloud intelligence logs in */home/admin/secure-service/logs/threat-info.log*.
   - View the web attack logs in */home/admin/secure-service/logs/web-info.log*.

9.14.10.1.3. Check the Dolphin service status

Context

To check the running status of the Dolphin service, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Dolphin service.
2. Run the following command to find the image ID of the Dolphin service:

   ```
   sudo docker ps -a |grep dolphin
   ```

3. Run the following command to go to the Docker container:

   ```
   sudo docker exec -it [imageId] /bin/bash
   ```

4. Run the following command to check whether the Java process is normal:

   ```
   ps aux |grep dolphin
   ```

5. Run the following command to perform the health check:

   ```
   curl 127.0.0.1:7001/checkpreload.htm
   ```

   If the response is "success", the service is normal.

6. View related logs.
   - View the info logs generated when the Dolphin service is running in */home/admin/dolphin/logs/*

- View the info logs generated when the Dolphin service is running in */home/admin/dolphin/logs/common-default.log*.
- View the Dolphin service error logs in */home/admin/dolphin/logs/common-error.log*.
- View the metaq messages received by the Dolphin service in */home/admin/dolphin/logs/dolphin-message-consumer.log*.

  > **Note** Currently, only Threat Detection Service (TDS) sends messages to the Dolphin service.

- View the metaq messages sent by the Dolphin service in */home/admin/dolphin/logs/dolphin-message-producer.log*.

  > **Note** Currently, the Dolphin service sends messages only to TDS.

9.14.10.1.4. Check the data-sync service status

Context

To check the running status of the data-sync service, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the data-sync service.
2. Run the following command to find the image ID of the data-sync service:

   ```
   sudo docker ps -a |grep data-sync
   ```

3. Run the following command to go to the Docker container:

   ```
   sudo docker exec -it [imageId] /bin/bash
   ```

4. Run the following command to check whether the Java process is normal:

   ```
   ps aux |grep data-sync
   ```

5. Run the following command to perform health check:

   ```
   curl 127.0.0.1:7001/check_health
   ```

   If OK is returned, the service is normal.

6. View related logs.

   View the data-sync service logs in *data-sync.log*.

9.14.10.2. Restart secure-service

Context

To restart secure-service when a fault occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server of the service.

2. Run the following command to find the image ID of the service:

   ```
   docker ps -a |grep application name
   ```

3. Run the following command to go to the Docker container:

   ```
   docker exec -it [imageId] /bin/bash
   ```

4. Restart related services.
   - Restart secure-service.
     a. Run the following command to view the Java process ID:

        ```
        ps aux |grep secure-service
        ```

     b. Run the following command to stop the current process:

        ```
        kill -9 process
        ```

     c. Run the following command to restart the process:

        ```
        sudo -u admin /home/admin/secure-service/bin/jbossctl restart
        ```

     d. Run the following command to check whether the process has been successfully restarted:

        ```
        curl 127.0.0.1:7001
        ```

   - Restart the Dolphin service.
     a. Run the following command to view the Java process ID:

        ```
        ps aux |grep dolphin
        ```

     b. Run the following command to stop the current process:

        ```
        kill -9 process
        ```

     c. Run the following command to restart the process:

        ```
        sudo -u admin /home/admin/dolphin/bin/jbossctl restart
        ```

     d. Run the following command to check whether the process has been successfully restarted:

        ```
        curl 127.0.0.1:7001/checkpreload.htm
        ```

   - Restart the data-sync service.
     a. Run the following command to view the Java process ID:

        ```
        ps aux |grep data-sync
        ```

     b. Run the following command to stop the current process:

        ```
        kill -9 process
        ```

     c. Run the following command to restart the process:

        ```
        sudo -u admin /home/admin/data-sync/bin/jbossctl restart
        ```

     d. Run the following command to check whether the process has been successfully restarted:

        ```
        curl 127.0.0.1:7001/check_health
        ```

9.15. Key Management Service (KMS)

9.15.1. O&M of KMS components

9.15.1.1. Overview

You can deploy KMS and perform O&M on KMS components in the Apsara Infrastructure Management Framework console.

You can log on to the machine where KMS resides from **Machine Operations** in the Apsara Infrastructure Management Framework console.

9.15.1.2. Log on to the Apsara Infrastructure Management Framework console

This topic describes how to log on to the Apsara Infrastructure Management Framework console.

Prerequisites

- 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.

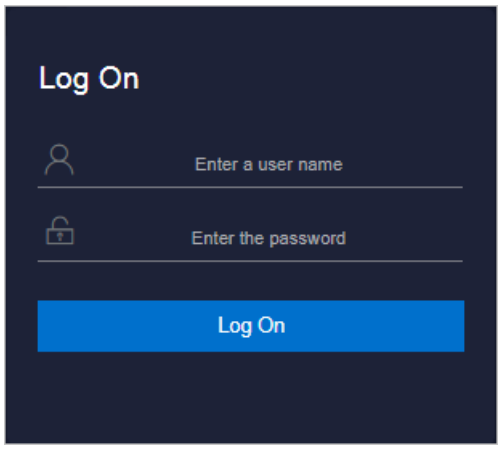

> **Note** You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

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 > Product List > Apsara Stack O&M**.
6. Click **Apsara Infrastructure Management Framework**.

9.15.1.3. KMS_HOST

This topic describes how to check the running status of the KMS_HOST service.

Check whether the server role is normal

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Operations** and then **Cluster Operations**.
3. On the **Cluster Operations** page, search for the KMS cluster.

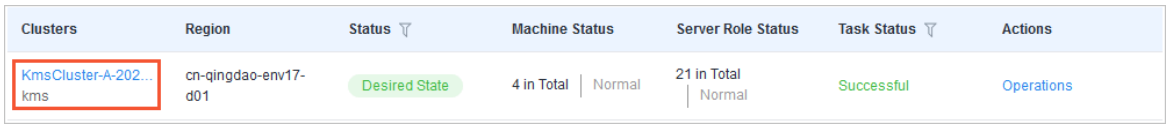

4. Click the link in the Clusters column to go to the **Cluster Details** page.
5. On the **Services** tab, click **kms** in the Services column.

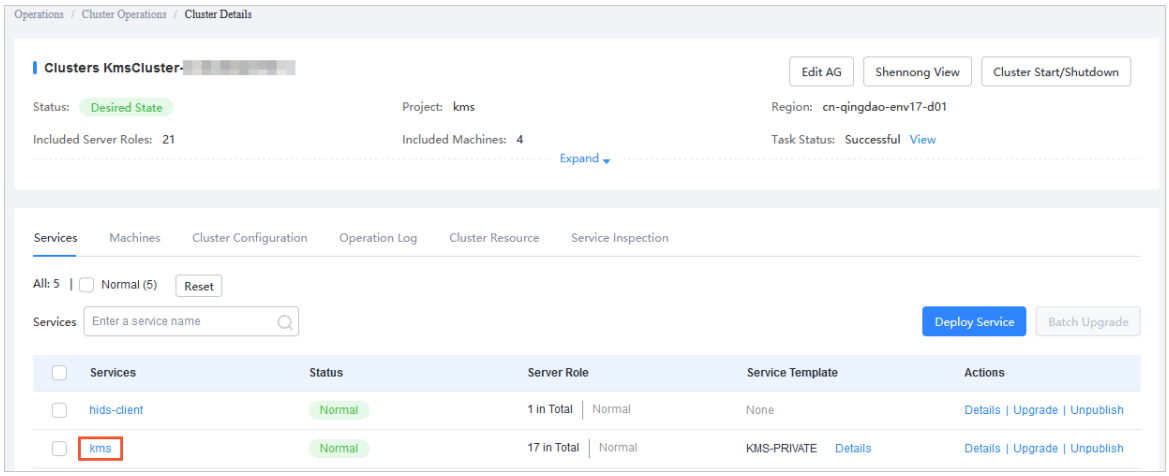

6. On the **Service Details** page, check whether the KMS_HOST server role has reached the desired state.If the indicator for kms.KmsHost# is green, the server role has reached the desired state.
7. On the **Service Details** page, click **kms.KmsHost#**. The information of the machines on which the KMS_HOST service is deployed appears in the lower part of the page. The IP addresses of the machines are required in subsequent steps.

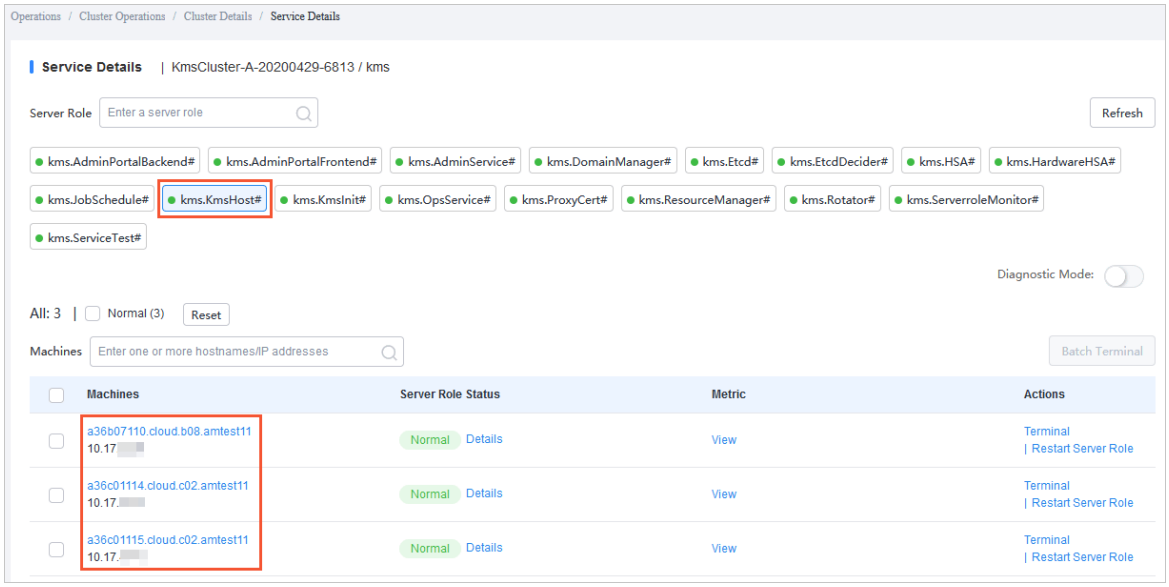

8. Select a machine and click **Terminal** in the Actions column to log on to this machine.
9. Run the `curl http://ip:5555/status.html` command and check whether success is returned.

```
$curl http://[redacted]:5555/status.html
success
```

> **Note**
> - Replace ip in the command with the IP address of this machine you obtained in Step 7.
> - Use this method to verify all machines on which the KMS_HOST service is deployed.

Identify exceptions

1. View logs in the */cloud/log/kms/KmsHost#/kms_host* directory.
2. Check whether the KMS_HOST service is running normally.
   - If KMS_HOST abnormally exits after it starts, check *debug.log* to identify the specific exception.
   - If KMS_HOST is running but cannot function normally, check *status.log* to identify the specific exception.

Possible errors

Error	Troubleshooting
`xxx selfCheck error` **Note** *xxx* refers to a dependency service.	• Check whether the dependency configuration is correct. You can check *debug.log* to identify the specific exception. • Check whether the xxx service runs normally.
`exit code 1`	View *debug.log* to identify the cause of the abnormal exit.

9.15.1.4. HSA

This topic describes how to check the running status of the HSA service.

Check whether the server role is normal

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Operations** and then **Cluster Operations**.
3. On the **Cluster Operations** page, search for the KMS cluster.

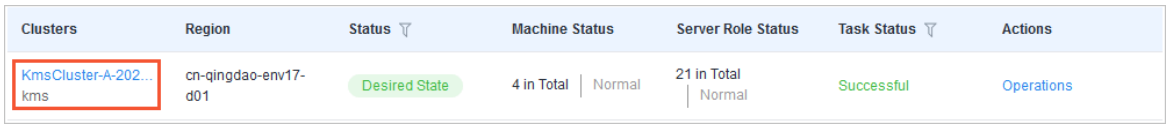

4. Click the link in the Clusters column to go to the **Cluster Details** page.
5. On the **Services** tab, click **kms** in the Services column.

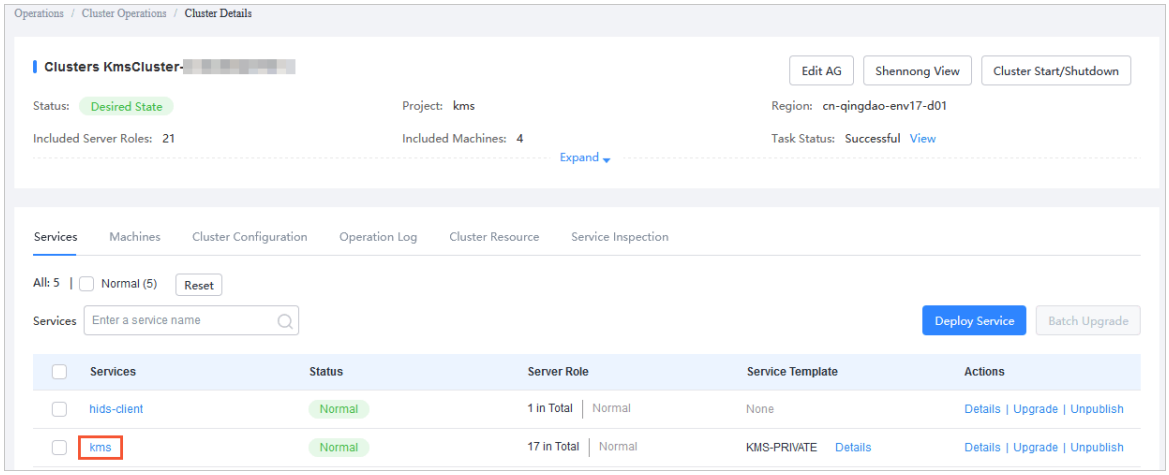

6. On the **Service Details** page, check whether the HSA server role has reached the desired state.If the indicator for kms.HSA# is green, the server role has reached the desired state.
7. On the **Service Details** page, click **kms.HSA#**. The information of the machines on which the HSA service is deployed appears in the lower part of the page. The IP addresses of the machines are required in subsequent steps.

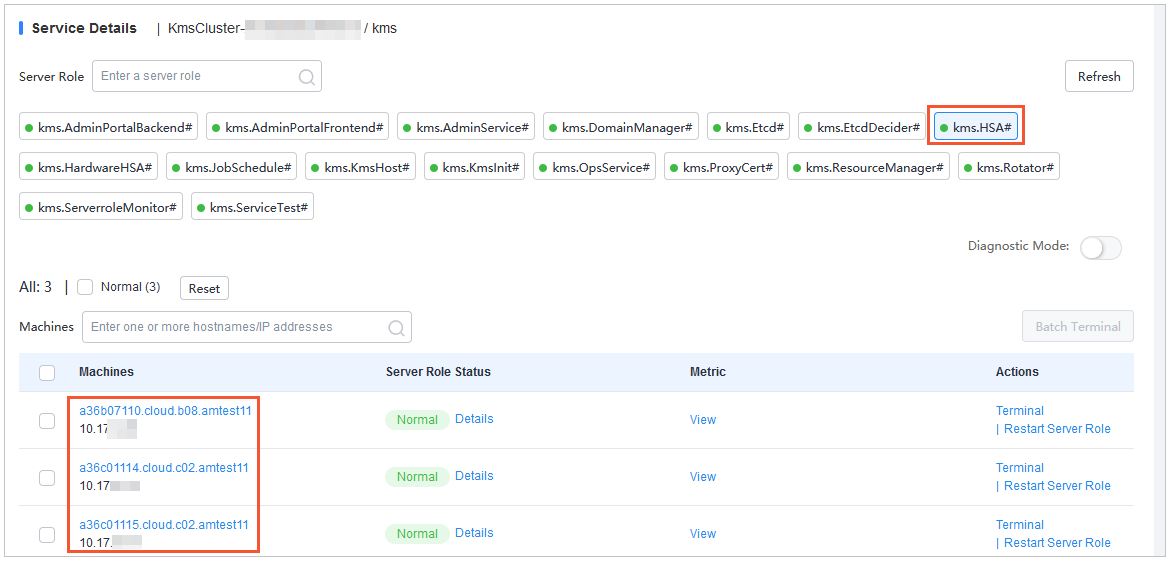

8. Select a machine and click **Terminal** in the Actions column to log on to this machine.
9. Run the `curl http://ip:5555/status.html` command and check whether success is returned.

> **Note**
> - Replace ip in the command with the IP address of this machine you obtained in Step 7.
> - Use this method to verify all machines on which the HSA service is deployed.

Identify exceptions

1. View logs in the */cloud/log/kms/HSA#/hsa* directory.
2. Check whether the HSA service is running normally.
   - If HSA abnormally exits after it starts, check *debug.log* to identify the specific exception.
   - If HSA is running but cannot function normally, check *status.log* to identify the specific exception.

Possible errors

Error: `exit code 1`

Troubleshooting: View debug.log to identify the cause of the abnormal exit. This error can be caused by one of the following reasons:

- etcd is not started.
- etcd is started, but its data is invalid.

> **Note** During disaster recovery, synchronization errors in the secondary cluster may cause this error.

9.15.1.5. etcd

This topic describes how to check the running status of the etcd service.

Check whether the server roles are normal

In the Apsara Infrastructure Management Framework console, check whether the Etcd and EtcdDecider server roles have reached the desired state.

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Operations** and then **Cluster Operations**.
3. On the **Cluster Operations** page, search for the KMS cluster.

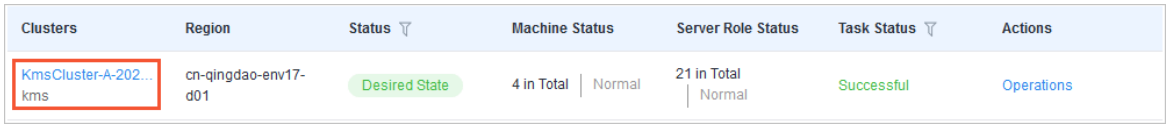

4. Click the link in the Clusters column to go to the **Cluster Details** page.
5. On the **Services** tab, click **kms** in the Services column.

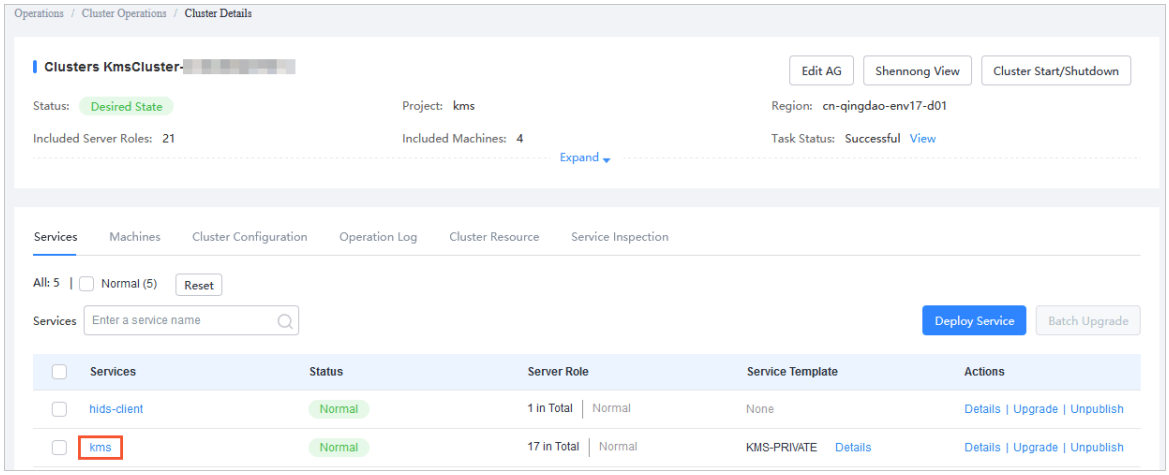

6. On the **Service Details** page, check whether the Etcd and EtcdDecider server roles have reached the desired state. If the indicators for kms.Etcd# and kms.EtcdDecider# are green, they have reached the desired state.

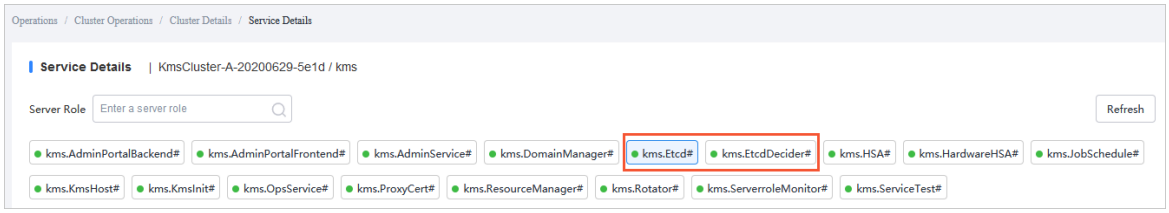

Identify exceptions

View logs under the */cloud/log/kms/Etcd#/etcd* and */cloud/log/kms/EtcdDecider#/decider* directories to identify exceptions.

Possible errors

Error	Troubleshooting
The startup parameters of etcd are invalid.	Find the correct settings of the startup parameters in the historical records of *debug.log*. Then, manually start etcd. **Note** Retain the error log and request the technical support team to identify the cause.
Errors in EtcdDecider during service upgrades cause errors in etcd.	In most scenarios, this issue occurs when there is a rolling task. You can analyze the issue and identify the cause based on *debug.log* of EtcdDecider.
The data directory of etcd is missing and etcd cannot start.	Solution: Use the Apsara Infrastructure Management Framework console to remove the abnormal etcd node from its server role group, and then add it back.

9.15.1.6. Rotator

9.15.1.6.1. Primary data center

This topic describes how to check the running status of the rotator of the primary data center.

The rotator is a special component. Even if the server role is in the desired state in the Apsara Infrastructure Management Framework console, the rotator is not necessarily working normally.

Rotator exceptions do not have any impact on the API logic of KMS.

In most scenarios, you need to identify the cause of a rotator exception only when unexpected results are found, for example, when the data on ApsaraDB for RDS does not meet expectations.

Check whether the rotator starts in primary data center mode

Check the value of `current idc master` in */cloud/log/kms/Rotator#/rotator/debug.log*, as shown in the following figure.

If the value of `current idc master` is true, the rotator starts in primary data center mode. If the value of `current idc master` is false, the rotator starts in secondary data center mode.

Check whether the rotator is in the working state

The rotator of the primary data center is deployed on all nodes. The nodes are in distributed lock mode. Only one node can work at a time. All the other nodes remain in the standby state.

View */cloud/log/kms/Rotator#/rotator/status.log* and check the status of each node.

Check the value of RotatorState in status.log. Valid values:

- ExecuteWorker: The node is in the working state.
- TryLock: The node is in the standby state.

Working state

```
[2017-10-23 16:51:51.554310]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  module:Rotator  host:a27d
05007.cloud.d05.ew9-5  RotatorState:ExecuteWorker
[2017-10-23 16:52:51.554415]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  module:Rotator  host:a27d
05007.cloud.d05.ew9-5  RotatorState:ExecuteWorker
```

Standby state

```
[2017-10-17 18:35:20.618575]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  module:Rotator  host:a27d
11007.cloud.d11.ew9-5  RotatorState:TryLock
[2017-10-17 18:36:11.867967]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  module:Rotator  host:a27d
11007.cloud.d11.ew9-5  RotatorState:TryLock
[2017-10-17 18:36:20.620963]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  module:Rotator  host:a27d
11007.cloud.d11.ew9-5  RotatorState:TryLock
```

Possible errors

- Abnormal RDS database access. Statistics collection and key deletion tasks cannot be executed.
- Abnormal HSA. Key rotation tasks cannot be executed.
- Abnormal Log Service. Metering tasks cannot be executed.
- Abnormal etcd. Distributed locks are unavailable and tasks cannot be executed.
- If one of the tasks on the rotator is abnormal, the rotator may be unable to reach the desired state. However, when this occurs, the rotator may be displayed to be in the desired state in the Apsara Infrastructure Management Framework console.

9.15.1.6.2. Secondary data center

This topic describes how to check the running status of the rotator of the secondary data center.

The rotator of the secondary data center is deployed on all nodes, which are all in the working state. The work scopes of the nodes are idempotent in a certain time range.

Check whether the rotator starts in secondary data center mode

Check the value of `current idc master` in */cloud/log/kms/Rotator#/rotator/debug.log*, as shown in the following figure.

```
[2017-10-21 16:34:34.412535]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  PKIVersion:pssl
[2017-10-21 16:34:34.446620]    [INFO]  [logger.(*LoggerWrapper).Infof] [logwrapper.go:37]  current idc master: false
```

If the value of `current idc master` is false, the rotator starts in secondary data center mode. If the value of `current idc master` is true, the rotator starts in primary data center mode.

Possible errors

- Abnormal network of the primary data center. The etcd of the primary data center is inaccessible.
- Abnormal etcd of the primary data center. The etcd of the primary data center is inaccessible.
- Abnormal etcd of the secondary data center. Data cannot be written into etcd.
- Incorrect etcd information of the primary data center. Data synchronization errors occur.

> **Notice** Rotator exceptions of the secondary data center has a severe impact on KMS in the secondary data center. You must fix the exceptions in a timely manner.

9.15.2. Log analysis

9.15.2.1. Overview

Logtail is a log collection client provided by Log Service to facilitate your access to logs. After installing Logtail on a host that has KMS deployed, you can monitor a specified log. The newly written log entries are automatically uploaded to a specified log library.

Logtail is used to transmit the logs of KMS to Log Service. Then the portal or API of Log Service analyzes the logs. If Log Service has no portals, you have to log on to the hosts that have KMS deployed individually and check the hosts one by one.

9.15.2.2. View logs by using request IDs

After you send a request to KMS, KMS sends you a message that contains a request ID.

Request IDs can be used in the following scenarios:

- You can view the KMS audit log in */cloud/log/kms/KmsHost#/kms_host/audit.log*.

  You can view the audit log information of the current access based on the value of request_id.

- For log entries whose expected_code values are not 200, you can view error information in debug.log based on the value of request_id.

  Path to the local log: */cloud/log/kms/KmsHost#/kms_host/debug.log*

  > **Note** */cloud/log/kms/KmsHost#/kms_host/debug.log* and *audit.log* are stored on the same machine.

- If you need all details of a request, you can view detailed information in the trace log.

  Path to the local log: */cloud/log/kms/KmsHost#/kms_host/trace.log*

  > **Note** */cloud/log/kms/KmsHost#/kms_host/trace.log* and *audit.log* may be stored on different machines.

- You can associate a cryptographic API operation with the trace log of HSA by using the value of request_id.

  Path to the local log: */cloud/log/kms/HSA#/hsa/trace.log*

  > **Note** */cloud/log/kms/KmsHost#/kms_host/trace.log* and *audit.log* may be stored on different machines.

- You can retrieve log information based on other information.

  You can retrieve the information in the audit log of KMS by using information other than request_id. If you want to associate the audit log with other logs, you must use request_id.

9.15.2.3. Common KMS errors

9.15.2.3.1. Overview

KMS has two HTTP status codes in audit.log: expected_code and status_code.

Typically, the expected code and status code of an error are the same. (`expected_code = status_code`). However, there are exceptions.

status_code is the HTTP status code that is actually returned to a user.

9.15.2.3.2. Errors with HTTP status code 4XX

Most errors with HTTP status code 4XX are included in the business logic of KMS. For example, HTTP status code 403 indicates a user request authentication failure, and HTTP status code 400 indicates that an input parameter is invalid.

You can view the details of an error in the debug log by using the value of request_id.

9.15.2.3.3. Errors with HTTP status code 500

This type of error is not included in the business logic of KMS. They are severe errors and must be fixed immediately.

In most scenarios, if the status code of an error is 500, the expected code of this error is also 500.

Such an error may be caused by an unexpected exception in a dependency service. We recommend that you contact the technical support personnel of the dependency service for further assistance.

You can view the details of an error in the debug log by using the value of request_id.

9.15.2.3.4. Errors with HTTP status code 503

The status code and expected code of such an error may be different or consistent.

- The status code is 503 but the expected code is not 503.

  Possible error causes:

  - The user (client) interrupts the connection in advance.
  - The client times out because the response of KMS (server) is too slow.

  You can check the trace log by using the value of request_id to determine whether the error is caused by a slow response of the server and identify the specific module.

- Both the status code and expected code are 503.

  Such an error is an expected error in a dependency service of KMS. It may occur when the performance of the dependency service is unstable.

  You can view the details of an error in the debug log by using the value of request_id. We recommend that you contact the technical support personnel of the dependency service for further assistance.

9.15.2.3.5. Degradation of dependency on a service

KMS stores the data of its dependency services in the local cache. If a dependency service is unavailable, KMS uses the obsolete data stored in the cache.

In this scenario, the status code in the audit log of KMS is 200, but an additional debug log will be generated.

When this situation occurs, users with cached data can access KMS. However, users without cached data encounter a 503 error when they try to access KMS.

9.16. Apsara Stack DNS

9.16.1. Introduction to Apsara Stack DNS

This topic describes Apsara Stack DNS and the features of its modules.

Database management system

The database management system compares the versions in the baseline configuration with those in the database to better manage databases. This allows you to validate the database version in each update.

API system

The API system determines the business logic of all calls and manages all data and tasks. This system is written in Java.

DNS

The DNS system consists of BIND and Agent. Agent receives and processes task information passed from the API system. Agent parses the tasks into commands, and then delivers the commands to the BIND system.

9.16.2. Maintenance

9.16.2.1. View operational logs

During operations and maintenance, you can query and view logs that are stored at specific locations in different systems to troubleshoot errors.

The operational logs of the API service are stored in the */home/admin/gdns/logs/* directory. You can query logs as needed.

The operational logs of the Agent service are stored in the */var/log/dns/* directory of the DNS server. Each log contains log entries of a specific day.

The operational logs of the BIND service are stored in the */var/named/chroot/var/log/* directory of the DNS server.

9.16.2.2. Enable and disable a service

You can log on to the API server as an administrator and run the `/home/admin/gdns/bin/appctl.sh restart` command to restart the API service. We recommend that you run the command on one server at a time to ensure that another server can provide services. You can specify the start, stop, and restart parameters in the preceding command.

Apsara Stack DNS provides services by using anycast IP addresses. You must run the `service ospfd stop` command to disable the OSPF service before you run the `service named stop` command to disable the DNS service.

You must run the `service named start` command to enable the DNS service before you run the `service ospfd start` command to enable the OSPF service.

You can run the `/usr/local/AgentService/agent -s start` command to enable the Agent service. If you receive a message that indicates the PID file already exists, delete the `/var/dns/dns.pid` file and run the command again.

You can run the `/usr/local/AgentService/agent -s stop` command to disable the Agent service.

9.16.2.3. Data backup

If you need to back up data before updating the service, copy the */var/named/* and */etc/named/* directories to a backup location. When you need to restore your data, copy the backup data to the original directories. Do not trigger automatic update during a data restoration process. Otherwise, data inconsistency may occur.

9.16.3. DNS API

9.16.3.1. Manage the API system

You can manage the API system in the Apsara Infrastructure Management Framework console. To log on to the server in which the API system resides, choose **Operations > Machine Operations** in the Apsara Infrastructure Management Framework console.

Context

To determine whether a service role is running as expected, follow these steps:

Procedure

1. In the Apsara Infrastructure Management Framework console, check whether the API is at desired state.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the top navigation bar, choose **Tasks > Deployment Summary** to open the **Deployment Summary** page.
   iii. Click **Deployment Details**.
   iv. On the **Deployment Details** page, find the dnsProduct project.

v. Find the dnsServerRole# service role, and click **Details** in the Deployment Progress column to check whether the service role is at desired state. If a green check mark is displayed after dnsServerRole#, then dnsServerRole# is at desired state.

View API status

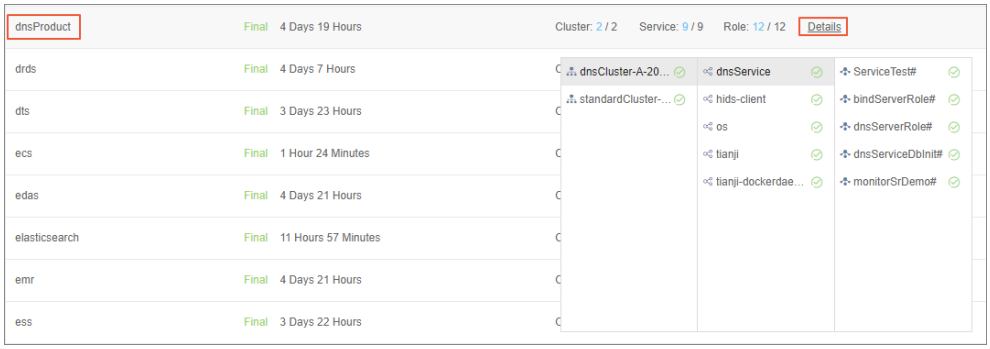

2. Obtain the IP addresses of servers where the API services are deployed.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the top navigation bar, choose **Operations > Cluster Operations**.
   iii. Click a cluster URL to open the **Cluster Dashboard** page.
   iv. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**.

Cluster Operation and Maintenance Center

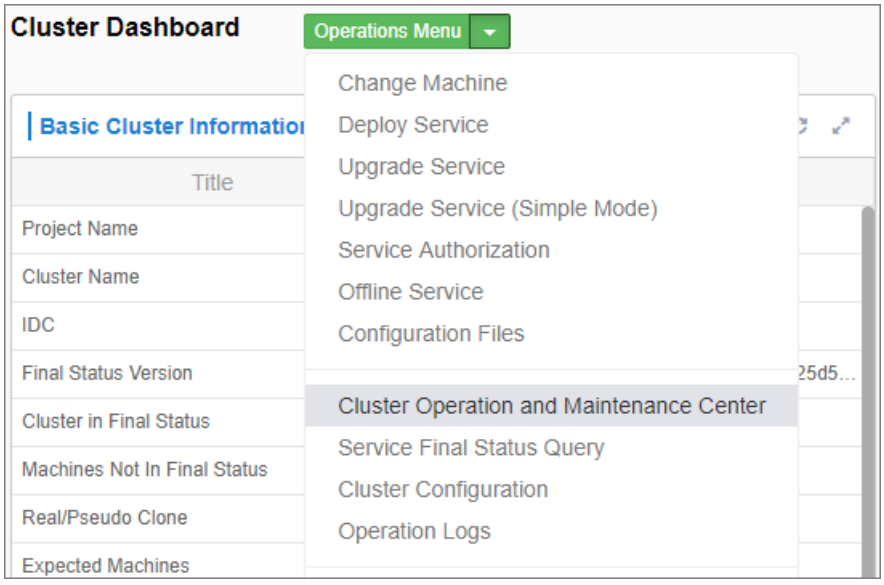

v. On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of servers that are deployed with the API service.

View the IP addresses of servers

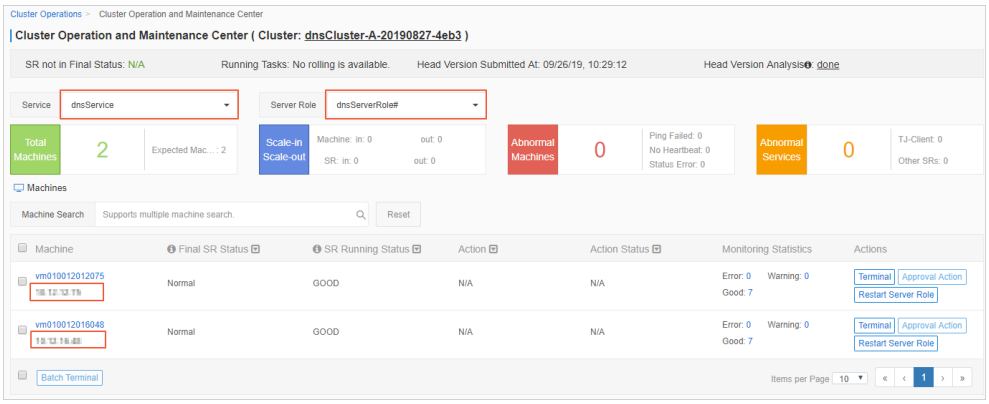

3. Log on to the DNS API server. Run the `curl http://localhost/checkpreload.htm` command, and check whether the command output is "success".
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the top navigation bar, choose **Operations > Machine Operations**.
   iii. Click **Terminal** in the Actions column of a server to log on to the server.
   iv. Run the `curl http://localhost/checkpreload.htm` command on the server where the API service is deployed and check whether the command output is "success".

Verify the server

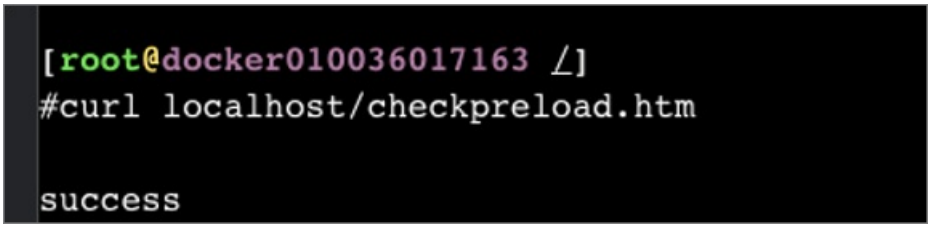

9.16.3.2. Troubleshooting

Procedure

1. View logs stored in */home/admin/gdns/logs/*.
2. Check whether the API service is running. If an error occurs when you call an API operation, check the log to troubleshoot the error.
3. If the API service is running, but its features do not function as expected, check the application.log file.

9.16.4. DNS system

9.16.4.1. Check whether a server role is normal

Procedure

1. In the Apsara Infrastructure Management Framework console, check whether the Apsara Stack DNS system is in its final state.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the top navigation bar, choose **Tasks > Deployment Summary**.
   iii. On the **Deployment Summary** page, click **Deployment Details**.
   iv. On the **Deployment Details** page, find dnsProduct.
   v. Click **Details** in the **Deployment Progress** column to check whether the bindServerRole# role is in its final state.

   Checking whether the bindServerRole# server role is in its final state

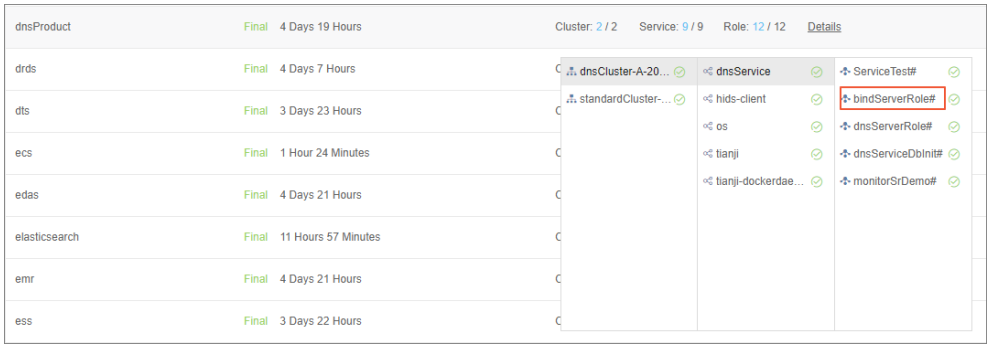

2. Obtain the IP addresses of the servers where DNS services are deployed.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the top navigation bar, choose **Operations > Cluster Operations**.
   iii. Click a cluster URL to go to the Cluster Dashboard page.
   iv. On the Cluster Dashboard page, choose **Operations Menu > Cluster Operation and Maintenance Center**.

   Cluster Operation and Maintenance Center

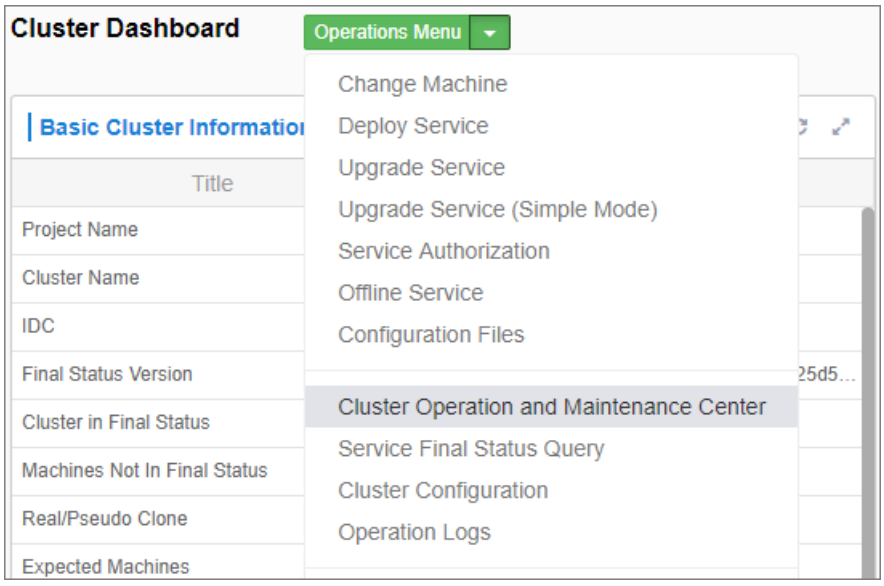

   v. On the Cluster Operation and Maintenance Center page, view and obtain IP addresses of all the servers that are assigned with the bindServerRole# role.

3. Log on to the DNS server, run the **python /bind/hello/check_health.py|echo $?** command, and check whether the command output is 0.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. Choose **Operations > Machine Operations**.
   iii. Select a server and click **Terminal** to log on to the server.
   iv. Run the **python /bind/hello/check_health.py|echo $?** command on each server that is assigned with the bindServerRole# role and check whether the command output is 0.

   Verifying the server

```
[root@101h08207.cloud.h10.amtest1284 /bind/hello]
#python check_health.py|echo $?
0
```

9.16.4.2. Troubleshooting

Procedure

1. Check the operational logs of the BIND service that are stored in the */var/named/chroot/var/log/* directory, and determine whether errors have occurred.
2. Check the operational logs of the Agent service that are stored in the */var/log/dns/* directory, and determine whether errors have occurred.
3. Run the **named-checkconf** command to check whether errors have occurred in the configuration file.

9.16.4.3. Errors and exceptions

Error: exit code 1

Run the health check script to view the cause of this error.

Common causes include:

- The DNS service is not running.
- The Agent service is not running.
- The OSPF service is not running, or anycast and public IP addresses cannot be advertised because of a network information retrieval error.
- Failed to run the task.

9.16.5. Log analysis

Query log entries by request ID

After you send a request, you will receive a response that contains the request ID. The request ID can be used in the following scenarios:

1. Query the tasks that are associated with the current request from the database.
2. Retrieve the execution results and error messages of the current request from the API system log.
3. Retrieve the results of the current request from the log of `bindServerRole#`, and verify the results with information that is retrieved from multiple other systems.

9.16.6. View and process data

Context

You can view task records and execution results.

Procedure

1. Log on to the API server to view database connection details.
2. Run the **use genesisdns** command of MySQL to log on to the database and then run the **select * from task** command to retrieve the progress and status of each task.

9.17. API Gateway

9.17.1. API Gateway introduction

This topic describes Apsara Stack API Gateway and the features of its modules.

API Gateway console

The API Gateway console is used to configure and manage your APIs and related policies. With the API management system, you can query, update, edit, and delete APIs. You can also create, associate, disassociate, and delete API management policies. API Gateway also provides a full range of API lifecycle management functions, including creating, testing, publishing, and unpublishing APIs. It improves API management and iteration efficiency. All your data will eventually be used as the API metadata for API Gateway.

API Gateway

API Gateway is a complete API hosting service. It helps you use APIs to provide capabilities, services, and data to your partners. API Gateway is initialized based on the API metadata generated by the API management system, and ultimately acts as the agent to send API requests. API Gateway provides a range of mechanisms to enhance security and reduce risks arising from APIs. These mechanisms include attack prevention, replay prevention, request encryption, identity authentication, permission management, and throttling.

9.17.2. Routine maintenance

9.17.2.1. View operational logs

During O&M, you can query and view logs that are stored in specific directories of different systems to troubleshoot issues.

API Gateway pop logs: The operational log files are stored in the */apsara/alidata/www/logs/java/cloudapi-openapi/* directory. You can query the files as required.

API Gateway logs: The operational log files are stored in the */apsara/alidata/logs/* directory. Each log file contains log entries that are generated over a single day. You can query the files as required.

9.17.2.2. Enable and disable a service

Perform the following operations to enable a service: Log on to the Apsara Infrastructure Management Framework console. Find the apigateway service instance in the Service Instances section of the Cluster Dashboard page and click Details in the Actions column.

On the Service Instance Information Dashboard page, find the target SR in the Server Role List section and click Details in the Actions column.

On the Server Role Dashboard page, find the target machine in the Machine Information section and click `Restart` in the Actions column.

In the message that appears, click OK. To disable a service, click Terminal in the Actions column and run the docker stop [containerId] command.

9.17.3. API Gateway O&M

9.17.3.1. System O&M

9.17.3.1.1. Check the desired state of API Gateway

You can use Apsara Infrastructure Management Framework to operate and maintain API Gateway. To log on to the machines in which the API Gateway console resides, choose Operations > Server Operations in the Apsara Infrastructure Management Framework console.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the top navigation bar, choose **Tasks > Deployment Summary**.
3. On the Deployment Summary page that appears, click **Deployment Details**.
4. On the **Deployment Details** page, find the apigateway project.
5. Click **Details** in the **Deployment Progress** column corresponding to the apigateway project. Check whether the ApigatwayLite# server role is in the desired state.If a green tick appears for the server role item, the server role has reached the desired state.

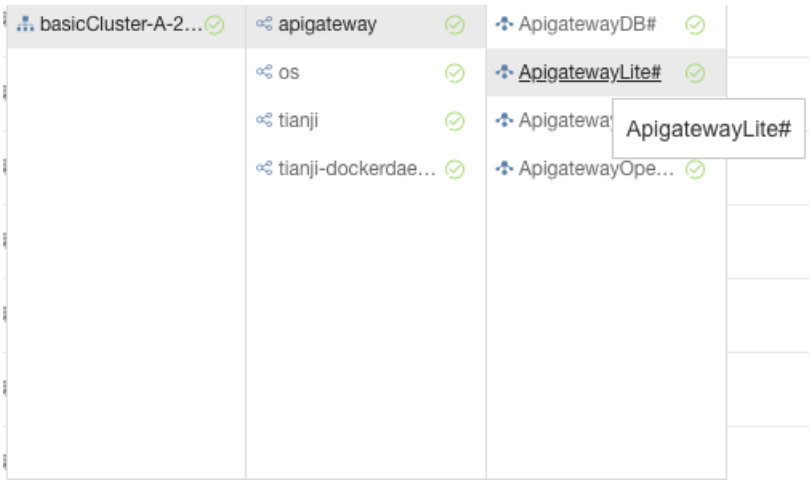

9.17.3.1.2. Check the service status of OpenAPI

Procedure

1. Find machines in the ApigatewayOpenAPI# server role.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. Click the **C** tab in the left-side navigation pane.
   iii. Select apigateway from the Project drop-down list.

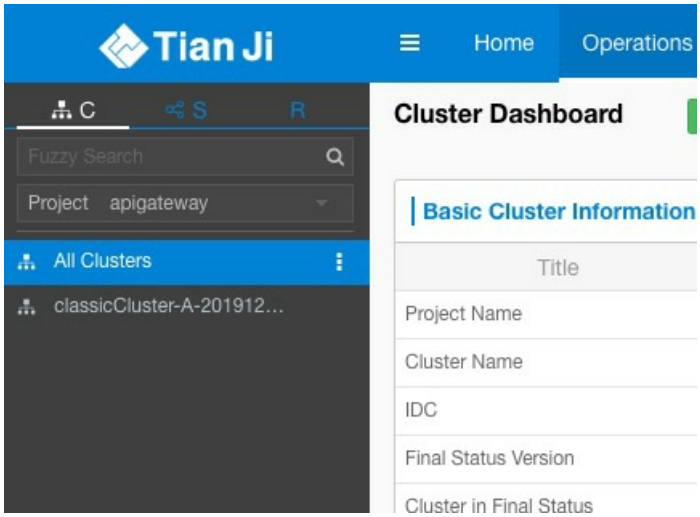

   iv. Place the pointer over the ⋮ icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.
   v. In the **Service Instance List** section, click Details in the Actions column corresponding to the apigateway service instance.
   vi. In the **Server Role List** section, you can view the deployment status of each role.

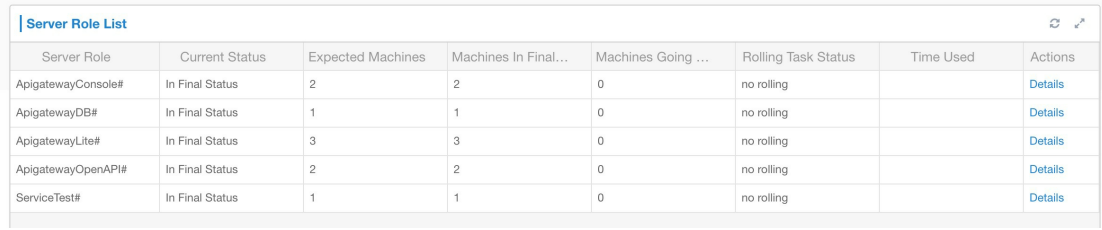

   vii. Click Details in the Actions column corresponding to the ApigatewayOpenAPI# role and view machine information of the role in the **Machine Information** section.

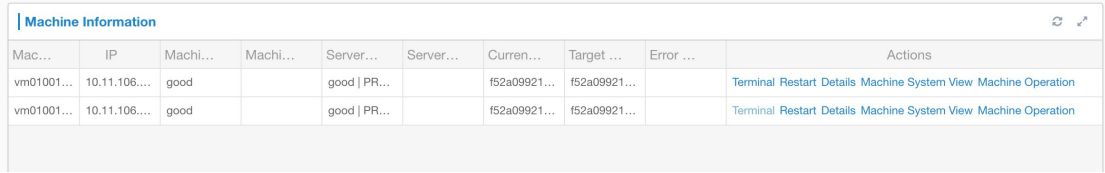

2. Click **Terminal** in the Actions column corresponding to a machine to log on to the machine.
3. Run the following command to find the container: **docker ps|grep cloudapi-openapi**
4. Run the following command to find the container IP address: **docker inspect [container ID] | grep IPAddress**
5. Run the following command to check whether OK is returned: **curl -i http://localhost:18080/cloudapi-openapi/check_health**

```
[admin@vm010148065157 /home/admin]
$docker inspect 81e002d83e7b |grep IPAddress
            "SecondaryIPAddresses": null,
            "IPAddress": "",
                    "IPAddress": "10.148.65.158",

[admin@vm010148065157 /home/admin]
$curl http://10.148.65.158:18080/cloudapi-openapi/check_health
ok
```

If OK is returned, the service status of the OpenAPI component is normal.

9.17.3.1.3. Check the service status of the API Gateway console

Procedure

1. Find machines in the ApigatewayConsole# server role.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. Click the **C** tab in the left-side navigation pane.
   iii. Select apigateway from the Project drop-down list.

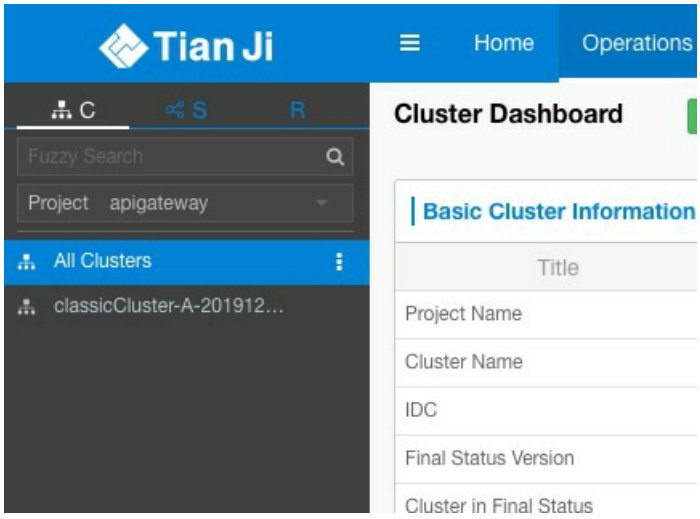

   iv. Place the pointer over the icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.
   v. In the **Service Instance List** section, click Details in the Actions column corresponding to the apigateway service instance.

vi. In the **Server Role List** section, you can view the deployment status of each role.

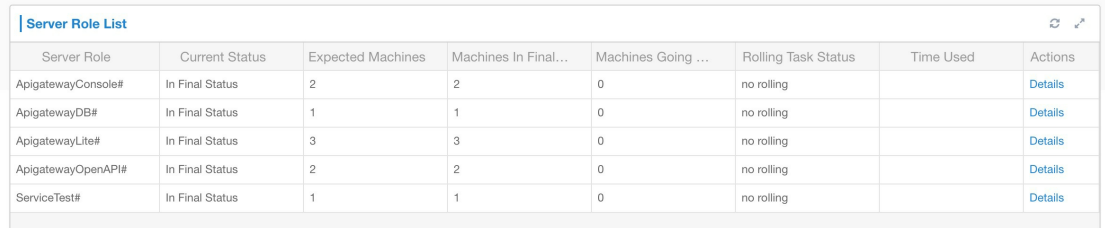

Server Role List

Server Role	Current Status	Expected Machines	Machines In Final...	Machines Going ...	Rolling Task Status	Time Used	Actions
ApigatewayConsole#	In Final Status	2	2	0	no rolling		Details
ApigatewayDB#	In Final Status	1	1	0	no rolling		Details
ApigatewayLite#	In Final Status	3	3	0	no rolling		Details
ApigatewayOpenAPI#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

vii. Click Details in the Actions column corresponding to the ApigatewayConsole# role and view machine information of the role in the **Machine Information** section.

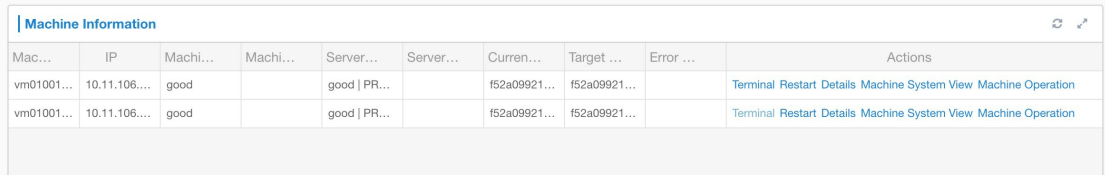

Machine Information

Mac...	IP	Machi...	Machi...	Server...	Server...	Curren...	Target ...	Error ...	Actions
vm01001...	10.11.106....	good		good | PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation
vm01001...	10.11.106....	good		good | PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation

2. Click **Terminal** in the Actions column corresponding to a machine to log on to the machine.
3. Run the following command to find the container: **docker ps|grep cloudapi-openapi**
4. Run the following command to find the container IP address: **docker inspect [container ID] | grep IPAddress**
5. Run the following command to check whether OK is returned: **curl -i http://localhost:18080/cag-console-aliyun-com/check_health**

```
[admin@vm010148065157 /home/admin]
$docker ps|grep console-backend
bc0d1d8295ea        696fa22ae150        "/bin/sh -c /alidata/"   3 days ago          Up 3 days                               apigateway.ApigatewayConsole__
.console-backend.1566551509

[admin@vm010148065157 /home/admin]
$docker inspect bc0d1d8295ea|grep IPAddress
            "SecondaryIPAddresses": null,
            "IPAddress": "",
                    "IPAddress": "10.148.65.159",

[admin@vm010148065157 /home/admin]
$curl http://10.148.65.159:18080/cag-console-aliyun-com/check_health
ok
[admin@vm010148065157 /home/admin]
```

If OK is returned, the service status of the API Gateway console is normal.

9.17.3.1.4. Check the service status of API Gateway

Procedure

1. Find machines in the ApigatewayLite# server role.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. Click the **C** tab in the left-side navigation pane.

iii. Select apigateway from the Project drop-down list.

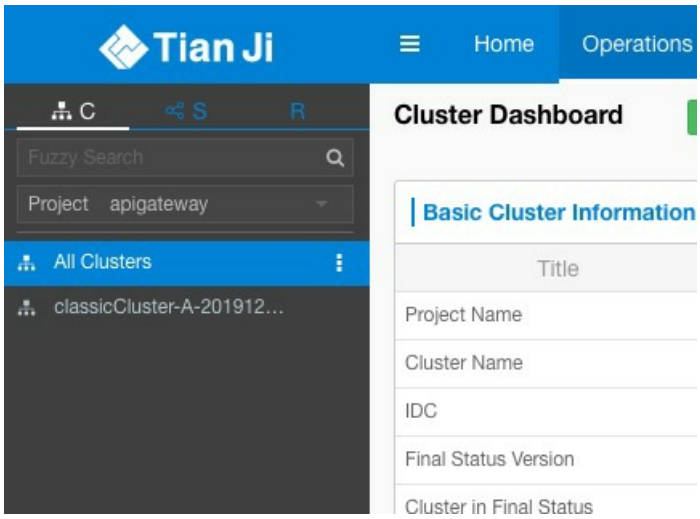

iv. Place the pointer over the icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.

v. In the **Service Instance List** section, click Details in the Actions column corresponding to the apigateway service instance.

vi. In the **Server Role List** section, you can view the deployment status of each role.

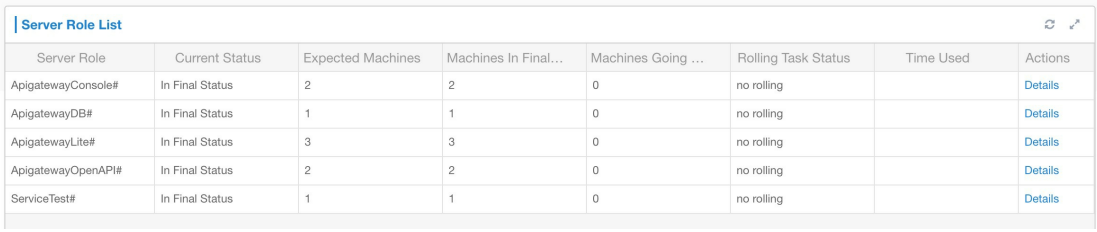

vii. Click Details in the Actions column corresponding to the ApigatewayLite# role and view machine information of the role in the **Machine Information** section.

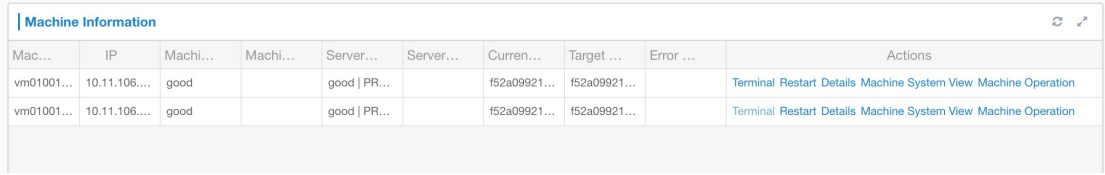

2. Click **Terminal** in the Actions column corresponding to a machine to log on to the machine.
3. Run the following command to check whether the `I'm fine, thank you, and you?` message is returned: **curl -i http://localhost/status -H Host:status.taobao.com**

9.17.3.1.5. View results of automated test cases

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Click the **C** tab in the left-side navigation pane.
3. Select **apigateway** from the **Project** drop-down list.

4. Place the pointer over the icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.
5. In the **Service Instance List** section, click **Details** in the Actions column corresponding to the apigateway service instance.
6. In the **Service Monitoring Information** section, click **Details** in the Actions column to view the automated test case report.

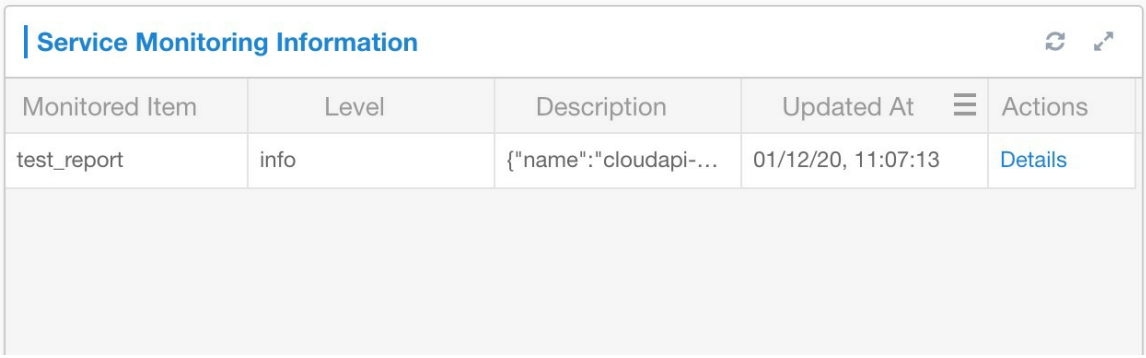

9.17.3.2. Troubleshooting

Context

> **Note**
> - /alidata/logs/system.log: API Gateway logs.
> - /usr/share/jetty/logs/stderrout.log: API Gateway console and OpenAPI logs.

Procedure

1. Start the application and check whether any errors have occurred. Check whether the system is operating normally.
   - If the system is operating but does not function properly, check the logs to troubleshoot errors.
   - If the system quits shortly after being started up, check the logs to troubleshoot errors.

9.17.4. Log analysis

You can perform log analysis based on the ID of an individual API request.

After you send a request, you will receive a response that contains the request ID from API Gateway.

You can use the request ID to perform the following operations:

- All API Gateway logs are uploaded to Log Service, where you can view the request ID.
- You can use the request ID to query the response to or error message for the current request in the API system logs.

10.Operations of middleware products

10.1. Enterprise Distributed Application Service (EDAS)

10.1.1. O&M overview

This topic describes the system architecture, component architecture, and O&M architecture of EDAS.

10.1.1.1. Architecture

This topic describes the system architecture and component architecture of EDAS. Familiarize yourself with this knowledge when performing O&M for EDAS.

System architecture

The system architecture of EDAS consists of the console, data collection center, configuration registry, authentication center, and file system. EDAS architecture shows the overall architecture of EDAS.

EDAS architecture

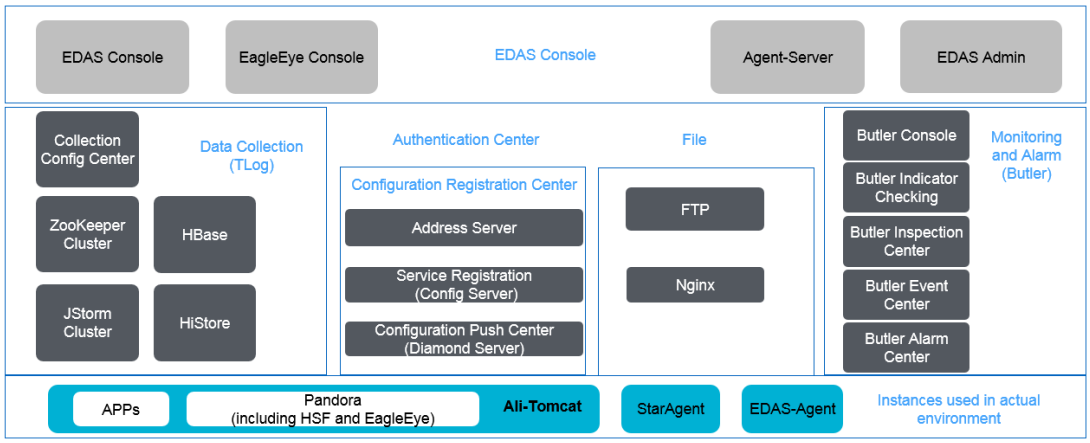

- EDAS console

  The EDAS console is the only EDAS system component that you can use directly. You can implement resource management, application lifecycle management, maintenance control and service governance, three-dimensional monitoring, and digital operation in the console.

- Data collection system

  It allows you to collect, compute, and store the runtime status, trace logs, and other information about clusters and instances where applications are deployed in EDAS in real time.

- Configuration registry

  This is a central server that is used to publish and subscribe to HSF services (RPC framework) and to push distributed configurations.

- Authentication center

This system component controls permissions for user data to ensure data security.

- O&M system

  EDAS uses Butler as its O&M system. Butler monitors the data of EDAS and triggers alarms when specified criteria are met. Butler provides routine inspection and alarm functions for all EDAS components.

- File system

  This system component stores WAR packages and required components, such as JDK and Ali-Tomcat, uploaded by users.

Component architecture

Each system component of EDAS consists of one or more components. The following figure shows the component architecture of EDAS.

Component	Node type	Node quantity	Description
EDAS console	Control node	2	The console of EDAS. It provides the core functions of the PaaS platform, including resource management, application lifecycle management, service governance, and auto scaling.
EDAS admin	Control node	2	A background task service. It provides the instance synchronization and application health check functions.
EDAS server	Control node	2	The EDAS server synchronizes status information with EDAS Agent.
Cai-fs	Control node	2	A file server. It stores the EDAS Agent installation package and EDAS application packages.
EagleEye console	Control node	2	You can query and view service traces in the EagleEye console.
Cai-address	Control node	3	An address discovery service. It provides the address lists for DiamondServer and ConfigServer.

Component	Node type	Node quantity	Description
DiamondServer	Control node	3	A configuration management service. It provides configuration storage, query, and notification functions, and mainly stores database metadata and EDAS function switch configurations in EDAS.
ConfigServer	Control node	3	An RPC service registry. It is used to query and store the publishing and subscription data of services.

For information about other external components, such as Butler, DAuth, and TLog, see the corresponding O&M documents.

10.1.1.2. O&M architecture

This topic describes the O&M architecture of Enterprise Distributed Application Service (EDAS). Before you use this topic, familiarize yourself with the system architecture of EDAS.

You can perform O&M for EDAS by mainly using the command-line interface (CLI).

O&M category	Description	O&M tool
Routine maintenance	Perform inspection and monitoring.	CLI: You can use the CLI to manually inspect the containers and components of EDAS.
Power-off maintenance	• Check and determine the statuses of containers and components. • Stop and start containers and components.	CLI
Troubleshooting	Handle component availability and service continuity faults of EDAS.	CLI

10.1.2. Overview of critical operations

Routine O&M for EDAS must be performed in strict accordance with the O&M guide. Failure to follow the O&M guide may cause risks to components and services.

O&M operations are classified into three levels: G1, G2, and G3. Operations vary by level. See the following table.

Definitions of operation levels

Level	Description
G1	L1|L2: Operations can be performed safely based on documented instructions, without having to apply for changes. Such operations will not affect the service.
G2	L1|L2: The onsite personnel must obtain confirmation from the product personnel before performing operations, which require applying for changes and following the documented instructions. Such operations will not affect the service.
G3	L1|L2: The onsite personnel must obtain confirmation from the product personnel and the customer before performing operations, which require applying for changes and following the documented instructions. Such operations may affect the service.

G3 is the highest level, which involves critical operations. See the following table.

A list of critical operations

Operation	Operation or Command
Check the AccessKeyId and AccessKeySecret	**cat /home/admin/.spas_key/default**
Clear logs	• For EagleEye: **find /home/admin/eagleeye/logs/ -name "*log.*" -exec rm {}** • For EDAS: **find /home/admin/edas/logs/ -name "*log.*" -exec rm {}**
Restart containers	**docker start {containerId}**

10.1.3. O&M preparation

This topic describes the logon portal, account, permissions, and tools required for O&M.

O&M preparation

Item	Purpose	Description
Remote Secure Shell (SSH) logon tool (such as *MobaXterm* or *PuTTY*)	Log on to the instances where components are located.	The logon account must be assigned the corresponding permissions. We recommend that you do not use the **root** or **admin** account for logon.

Item	Purpose	Description
Account	Log on to the console or an instance.	• Obtain the account and password for console logon from EDAS Customer Services. • The account used to log on to the instances where EDAS components are located must be assigned the corresponding permissions. We recommend that you do not use the **root** or **admin** account for logon.

10.1.4. Routine maintenance

EDAS routine maintenance includes inspection and monitoring.

- Inspection is the process where a periodic dialing test is performed on URLs or ports to determine whether EDAS services are normal. Currently, inspections in HTTP, TCP, ping, and JDBC modes are supported.
- Monitoring is the process where logs are collected from clients through TLog to summarize key metrics for measuring the system runtime status. Monitoring includes infrastructure monitoring and JVM monitoring.

10.1.4.1. Log on to Apsara Infrastructure Management Framework

This topic describes how to log on to Apsara Infrastructure Management Framework.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

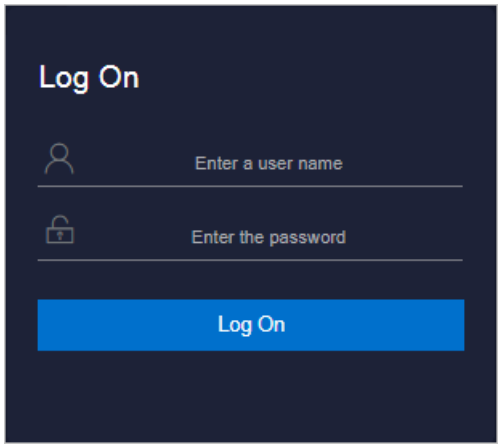

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. On the **Product List** page, choose **Apsara Stack O&M > Apsara Infrastructure Management Framework**.

10.1.4.2. Inspection

You can configure the inspection rules of the HTTP, TCP, ping, and database types to inspect the components and services of Enterprise Distributed Application Service (EDAS). An inspection rule provides a response code that is used to check the configured alert rule. The alert content is configured in Alert Description. You can also log on to the instances where components and services are deployed and run commands for inspection.

EDAS inspection items lists the default inspection items of EDAS. You can create an inspection configuration as needed.

EDAS inspection items

Inspected object	Description	Inspection method
ConfigServer	Request /configserver/serverlist to check whether ConfigServer is normal.	Check /configserver/serverlist.
Diamond	Check whether the API operation for querying the DiamondServer status is normal.	Check /diamond-server/diamond.
	Check whether the DiamondServer database is connected.	Check the database connection.
TLog	Check whether the TLog service is normal.	Check /api/StageHealthCheck.
	Check whether the TLog listening port is normal.	Check port 8080.
	Check whether the TLog database is connected.	Check the database connection.
edas-console	Check whether the edas-console service is normal.	Check /checkpreload.htm.
	Check whether the edas-console port is normal.	Check port 8080.
	Check whether the EDAS database is connected.	Check the database connection.
HiStore	Check whether the HiStore listening port is normal.	Check port 5029.
	Check whether HiStore is connected.	Check the database connection.
Redis	Check whether the Redis listening port is normal.	Check port 6379.
edas-admin	Check whether the edas-admin service is running properly.	Check /index.
	Check whether the edas-admin listening port is normal.	Check port 8080.

10.1.4.2.1. Component inspection

You can inspect Enterprise Distributed Application Service (EDAS) components by using the CLI.

10.1.4.2.1.1. Manual inspection

Enterprise Distributed Application Service (EDAS) allows you to manually perform complex logic inspection by using the CLI.

Manual inspection uses commands typical of Linux operating systems. Set specific parameters based on the actual environment.

10.1.4.3. Monitoring

You can monitor the containers, system components, and services of Enterprise Distributed Application Service (EDAS) by using monitoring metrics.

Container monitoring

By default, the container status of each EDAS component is checked based on the monitoring script that is configured in an environment variable.

System monitoring

System monitoring includes infrastructure monitoring and Java Virtual Machine (JVM) monitoring.

System metrics

Monitoring type	Metric	Threshold
Infrastructure monitoring	CPU utilization	70%
	Memory usage	70%
	Disk usage	90%
JVM monitoring	Young garbage collection (GC) times	60
	FullGC	5
	Old generation usage	90%
	Permanent generation usage	90%

Service monitoring

Service monitoring is configured by EDAS and reported by API operations for monitoring and alerts. The following table lists the service metrics.

Monitored object	Monitoring type	Monitoring description
HTTP service	QPS	Monitors the QPS of the HTTP service.
	Time consumption	Monitors the input/output (I/O) time consumed by the HTTP service.
	Service provisioning QPS	Monitors the QPS of HSF service provisioning.

Monitored object	Monitoring type	Monitoring description
HSF service	Response time (RT) of service provisioning	Monitors the RT of HSF service provisioning.
	Service consumption QPS	Monitors the QPS of HSF service consumption.
	RT of service consumption	Monitors the RT of HSF service consumption.
Container	Heap memory usage	Measures how much heap memory is used by services.
	Off-heap memory usage	Measures how much off-heap memory is used by services.

10.1.4.3.1. Monitoring logs

Logs are critical for O&M of Enterprise Distributed Application Service (EDAS). You can monitor logs to promptly locate runtime faults.

The following table lists the EDAS logs that can be monitored.

Component	Log	Path
edas-console	console.log	/home/admin/edas/logs
edas-admin	admin.log	/home/admin/edas/logs
edas-server	agent-server.log	/home/admin/edas/logs

EDAS also provides other component logs, including the infrastructure monitoring log, service monitoring log, container monitoring log, and JVM monitoring log. For more information, see Log reference.

10.1.5. Troubleshooting

Faults may occur during EDAS usage. This topic describes the typical faults that may occur during O&M as well as their handling methods.

Fault classification

Currently, EDAS-related faults are classified into two categories:

- Component unavailability
- Service discontinuity

Fault locating

You can locate faults through inspection, monitoring, logs, and alarms.

10.1.5.1. Alert handling

During O&M of Enterprise Distributed Application Service (EDAS), you can inspect and monitor the status and metrics of each EDAS component. Alerts are triggered when inspection and monitoring are abnormal. This topic describes how to handle inspection and monitoring alerts.

10.1.5.1.1. CPU utilization alerts

The CPU utilization threshold is 70%. The CPU utilization is abnormal and an alert is triggered if it exceeds the threshold.

Possible causes

- High access concurrency
- Insufficient application instances

Impact on the system

Service performance is compromised.

Procedure

1. Open the Secure Sockets Layer (SSL) tool (MobaXterm Personal Edition).
2. Run **ssh <Username>@<IP address of your Ark client>** and enter your *password* to log on to the client.
3. Go to the target container and run **top** to check the CPU utilization of components.
   - If the CPU utilization is normal, no further action is required.
   - If the CPU utilization is abnormal, check the number of calls.
4. Run **netstat -tnlp | grep -E "80|8080" | wc -l** to check the call status of components.
   - If the number of calls is relatively large and meets the service status, scale out instances.
   - If the number of calls is not large, identify the cause of high CPU utilization by completing the following steps.

     > **Note** You can use the open-source script to locate and print the processes with high CPU utilization.

     a. Run **top -Hp <Component process ID>** to locate the processes with high CPU utilization or memory usage and convert them into the hexadecimal format.

     b. Go to the *jstack process id* path, open **ps.txt**, and identify the specific process class based on the hexadecimal thread ID.

     c. If Full GC occurs, check **gc.log** or run **jstat -gcutil [pid]** to check the corresponding GC log. Record *top/jstack file/full GC* and send the record to EDAS Customer Services.

10.1.5.1.2. Memory usage alerts

The memory usage threshold is 90%. Memory usage is abnormal if it exceeds the threshold.

Possible causes

- High access concurrency
- Insufficient application instances

Impact on the system

Service performance is compromised.

Procedure

1. Open the SSL tool (MobaXterm Personal Edition).
2. Run **ssh <Username>@<IP address of your Ark client>** and enter your *password* to log on to the client.
3. Go to the target container and run **top** to check the memory usage of components.
   - If the memory usage is normal, check the disk usage and JVM metrics.
   - If the memory usage is abnormal, check whether this is caused by JVM.
4. Run **jmap** to check memory usage.
5. Run **jstat -gcutil [pid]** to check memory usage.
6. Run **vmstat** to analyze and collect statistics on virtual memory.

Result

Log on to the target container by using the SSH tool and run **top** to check whether the memory usage is normal (less than 90%).

10.1.5.1.3. Disk usage alerts

The disk usage threshold is 90%. The disk usage is abnormal if it exceeds the threshold.

Possible causes

- High access concurrency
- Insufficient application instances
- Insufficient disk space or no periodic disk cleanup

Impact on the system

Service performance is compromised.

Procedure

1. Open the SSL tool (MobaXterm Personal Edition).
2. Run **ssh <Username>@<IP address of your Ark client>** and enter your *password* to log on to the client.
3. Go to the target container and run **df -lh** to check the disk usage of each directory and identify the directories with excessive and fast disk usage.
4. Run **iostat** to check the data write status. Perform disk cleanup if the *logs* directory occupies excessive disk space.
5. Run **netstat** to check the number of calls, view logs for call errors, and take relevant measures such as scale-out.

Result

Go to the target container by using the SSH tool and run **df -lh** to check whether the disk usage is normal (less than 90%).

10.1.5.1.4. JVM alarms

Butler monitors the JVM usage in containers for the past five minutes. JVM is abnormal if this predefined threshold is exceeded.

Context

JVM metrics include:

- Young GC times (threshold: 60)
- Full GC times (threshold: 5)
- Old generation usage (threshold: 90%)
- Permanent generation usage (threshold: 90%)

Possible causes

- High access concurrency
- Insufficient application instances
- Insufficient disk space or no periodic disk cleanup

Impact on the system

Service performance is compromised.

Procedure

1. Log on to the Butler console. In the left-side navigation pane, choose **Service Monitoring > EDAS**.
2. On the **EDAS Monitoring** page, select the **edasServer** service and the **edas-admin** component in the **Monitoring Status** section, and click the **JVM Monitoring** tab.
3. Check memory usage, including the memory change in each area.
4. Open the SSL tool (MobaXterm Personal Edition).
5. Run the command **ssh <Username>@<IP address of the client>** and enter your *password* to log on to the client.
6. Go to the target container and run **jmap -heap <pid>** to check the settings of each memory area.
7. Log on to the target container and run **jmap-histo:live <pid>** to view the number of instances, memory usage, and full name of each class. Run the preceding command every minute five times in a row.
8. Save the collected JVM-related data and information, and contact EDAS Customer Services.

Result

- Log on to the Butler console to check whether the related alarms are cleared.
- Log on to the Butler console to check whether JVM metrics are normal.
- Go to the target container by using the SSH tool and run **jmap-histo:live <pid>** to check whether JVM metrics are normal.

10.1.5.1.5. Inspection alarms

Butler periodically inspects EDAS over HTTP, TCP, ping, or JDBC to check the conditions of EDAS components. An alarm is triggered when any component is abnormal.

Possible causes

1. Missing check script
2. Component process breakdown
3. Container breakdown

Impact on the system

- The downtime of a single container does not affect the service because the component cluster is highly available.
- When all containers malfunction, console logon and task execution may fail, but the service is not affected.

Procedure

1. In the left-side navigation pane of the Butler console, choose **System Monitoring > Inspection Management** .
2. Locate the component inspection items, click **Dialing Test Now** and check the inspection results.
   - The troubleshooting process ends if the inspection results are normal.
   - Perform if the inspection results are abnormal.
3. Run **docker ps | grep <component name>** to check whether a specific component process exists.
   - If the process does not exist, run **docker restart ${container_id}** to restart the container and perform a dialing test again.
   - If the process exists, perform .
4. Check OVS, VLAN, and other components based on the *Basic Component O&M Guide* to troubleshoot network faults.

Result

- Log on to the Butler console to check whether the related alarms are cleared.
- Log on to the Butler console to check whether components and containers are in the normal state (highlighted in green).

10.1.5.2. Service continuity exceptions

10.1.5.2.1. EDAS monitoring exceptions

This topic describes how to troubleshoot Enterprise Distributed Application Service (EDAS) monitoring exceptions.

Problem description

- No application data can be monitored in the EDAS console.

- Data monitoring in the EDAS console has a significant lag.
- The monitoring and alerts features are ineffective.
- Traces cannot be queried.

Possible causes

The EDAS components and dependent components are abnormal.

Impact on the system

EDAS cannot monitor applications or services, or monitoring is inefficient.

Procedure

1. Check whether related components, such as Tlog, JStorm, and HBase, are normal.
   - If TLog is abnormal, log on to the instance where TLog is located, go to */home/admin/logs/*, and check the *tlogconsole.log* file. Then, troubleshoot the problem and restart TLog.
   - If JStorm is abnormal, log on to the instance where JStorm is located and check the log for the data collection task, such as */home/admin/logs/tlog_eagleeye-worker-6801.log*. Troubleshoot the problem and restart JStorm.

     > **Note** You need to restart each JStorm process. Otherwise, data cannot be written to HBase due to a connection error.

   - If HBase is abnormal, log on to the instance where HBase is located and check the related error log. Troubleshoot the problem and restart HBase.

Result

Check whether EDAS monitoring becomes normal.

10.1.5.2.2. Excessive node logs

This topic describes how to troubleshoot the problem of excessive node logs for EDAS.

Problem description

- Trace queries and system responses slow down, and disk usage alerts are reported.
- The service instance generates excessive log files.

Possible causes

A large amount of log files are not cleared from the disk in a timely manner, which affects system performance.

Impact on the system

Service nodes become less responsive.

Procedure

1. Log on to the EDAS component node to check logs.
   - Path to EagleEye logs: /home/admin/logs/eagleeye

- Path to EDAS logs: /home/admin/edas/logs

2. Ensure that service logs are not printed on the preceding paths or that service logs have been backed up.
3. Clear backup logs by running **find /home/admin/logs/ -name "*log.*" -exec rm {};** or **find /home/admin/edas/logs/ -name "*log.*" -exec rm {};**.

Result

Check whether the service nodes become normal.

10.1.5.2.3. Console access failure

This topic describes how to troubleshoot EDAS console access failures.

Symptoms

The EDAS console cannot be accessed.

Possible causes

- The edas-console node is abnormal.
- An error occurs during DNS resolution.

Impact on the system

The EDAS console is unavailable.

Procedure

1. Troubleshoot the edas-console node errors.
   - If the EDAS console becomes accessible again, no further action is required.
   - If the EDAS console remains inaccessible, proceed with the next step.
2. Log on to the ECS instance where the edas-console node is located, go to */home/admin/edas/logs*, and check console.log for the problem.

Result

The EDAS console becomes accessible again.

10.1.5.2.4. Failure to import an ECS instance

This topic describes how to troubleshoot the failure to import an Elastic Compute Service (ECS) instance.

Problem description

An ECS instance fails to be imported.

Possible causes

- An Alibaba Cloud API operation fails to be called.
- An image fails to be replaced.
- The ECS instance fails to be registered.

Impact on the system

The service availability and reliability are compromised.

Procedure

1. Log on to the EDAS console, and manually import the ECS instance.
2. If the ECS instance failed to be registered, register it again by running edas init.
3. If the image failed to be registered, Log on to the ECS console, and check the specific status.

10.1.5.2.5. TLog data collection errors

This topic describes how to fix TLog data collection errors.

Symptoms

- The application monitoring dashboard is inaccessible.
- Application and service monitoring is inaccessible, and alarms cannot be triggered.
- Infrastructure monitoring is inaccessible, auto scaling is ineffective, and alarms cannot be triggered.
- Traces cannot be queried.

Possible causes

Collection point: Each collection job of TLog is called a collection point.

Collection points are the basic units for task processing by TLog. The monitoring function of EDAS is provided by one or more collection points in TLog. When working properly, collection points are in the activated or running state.

If the collection point for a product encounters an error, the corresponding EDAS monitoring data or page shows an exception.

- The basic data of the monitoring dashboard corresponds to the collection point service group TLog and the collection point infrastructure.
- The service data of the monitoring dashboard corresponds to the collection point service group EagleEye and the collection point stats_logger_agg.
- The zoom-in (more than 30 minutes) function in infrastructure monitoring corresponds to the collection point service group TLog and the collection point infrastructure.
- Service monitoring corresponds to the collection point service group EagleEye and the collection point stats_logger_agg.
- Trace analysis and query corresponds to the collection point service group TLog and the collection point EagleEye.

Impact on the system

An application change fails.

Procedure

1. Identify the corresponding TLog collection point based on the abnormal function.
2. Check whether the collection point has been started properly.

i. On the Collection Points page, locate the row that contains the collection point, and click **More** > **Manually Assign Task** in the Actions column. View the dialog box that appears.

- If the number in the dialog box is greater than 0, the collection point has been started properly. You can go to the next step.
- If the collection point is not started properly, return to the Collection Points page and click **Edit/Deployment Process**. On the page shown in the following figure, click **Start** and wait until "Operation successful" appears in the result. If "Operation successful" does not appear in the result, contact EDAS Customer Services and give feedback like "The xxx collection point does not start properly."
- If "Operation successful" appears in the result, return to the **Collection Points** page, wait for three to five minutes, and click Manually Assign Task again. If the number in the dialog box that appears is greater than 0, the collection point has been started properly. Go to the next step.

3. If the collection point has been started, check whether the collection rules are correctly distributed (operation risk level: G1).
   - On the Collection Points page, check whether the distribution status is Active. If it is Inactive, click **Activate** and **OK** in sequence.
   - Click **Collection Point Details** to go to the **Collector Status** tab. If the status list is not empty, collection rules are distributed properly. Go to the next step.
   - If the status list on the **Collector Status** tab is empty, return to the Collection Points page to manually distribute collection rules as follows: Click **Create Task by Rule** under **Collection Rule**. If the number of created tasks in the dialog box shown in the following figure is greater than 0, click **OK**. Return to the Collection Points page and click **Manually Assign Task** and **OK** in sequence. If manual distribution is successful, a dialog box appears.
   - If the number of created tasks in **Create Task by Rule** is 0, contact EDAS Customer Services for troubleshooting and give feedback like "The xxx collection point has 0 created tasks in Create Task by Rule."
4. If the collection point has been started and collection rules are distributed properly but no data exists, perform troubleshooting as follows (operation risk level: G1):
   - Click **Collection Point Details** to go to the **Collector Status** tab. Check the data in the **Last Collection Attempt** column. Normally, the time in this column is less than 1 minute. If the time in this column is generally greater than 1 minute, contact EDAS Customer Services and give feedback like "On the Collector Status tab for the xxx collection point, the time in the Last Collection Attempt column is generally greater than 1 minute. The collection point must be scaled up."
   - On the **Collector Status** tab, check the data in the Status column. Normally, the states in this column are Normal or File Not Modified. If states such as File Not Found, No Permission, Connection Timeout, and SProxy Not Found appear in this column, contact EDAS Customer Services and give feedback like "On the Collector Status tab for the xxx collection point, the yyy state appears in the Status Column."
   - On the **Collection Points** page, click **More > Perform Health Check**. Then, contact EDAS Customer Services and provide the JSON content on the health check page to help engineers quickly locate the problem.

10.1.6. Log reference

You can check logs to view the status of each EDAS component or locate faults during O&M.

EDAS provides logs for the following components:

- EDAS console
- EDAS admin
- EDAS server
- Cai-fs
- DiamondServer
- ConfigServer
- Cai-address
- EagleEye console

EDAS archives and clears the logs for these components based on predefined policies.

10.1.6.1. EDAS console logs

The EDAS console is the console component of EDAS. It provides the core functions of the PaaS platform, including resource management, application lifecycle management, service governance, and auto scaling.

Log files

EDAS console logs

File	Description
console.log	The EDAS console log.
changeorder.log	The change order log.
openapi.log	The API log.
tengine.log	The TEngine log.
debug.log	The log that records the internal API calls of the EDAS console.

Path

- Logging path: *${user.home}/edas/logs*
- Archive path: *${user.home}/edas/logs/bak*

Format

openapi.log: %msg%n (print log information directly) others: %d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %-5level %logger{50}:%line - %msg%n (date and time, thread name, log level, class name: number of lines - specific log information)

Archiving policies

EDAS console log archiving policies

Log	Archiving policy
console.log	• Maximum size: 100 MB • The name of the new file takes the format console.{d}.log. Seven logs are retained.
changeorder.log	• A file is created every day. • The name of the new file takes the format changeorder.{yyyy-MM-dd}.log. Logs from the last seven days are retained.
openapi.log	• A file is created every day. • The name of the new file takes the format opanapi.{yyyy-MM-dd}.log. Logs from the last seven days are retained.
tengine.log	• A file is created every day. • The name of the new file takes the format tengine.{yyyy-MM-dd}.log. Logs from the last seven days are retained.
debug.log	• Maximum size: 100 MB • The name of the new file takes the format debug.{d}.log. Three logs are retained.

10.1.6.2. EDAS admin logs

The EDAS admin is a background task service that provides the instance synchronization and application health check functions.

File

EDAS admin logs

File	Description
admin.log	The scheduling task log.
tengine.log	The Tengine log.

Path

- Logging path: *${user.home}/edas/logs*
- Archive path: *${user.home}/edas/logs/bak*

Format

%d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %-5level %logger{50}:%line - %msg%n (date and time, thread name, log level, category name: number of lines - specific log information)

Archiving policy

EDAS admin log archiving policies

Log	Archiving policy
admin.log	• Maximum size: 100 MB • The name of the new file takes the format admin.{d}.log. Three logs are retained.
tengine.log	• A file is created every day. • The name of the new file takes the format tengine.{yyyy-MM-dd}.log. Logs from the last seven days are retained.

10.1.6.3. EDAS server logs

The EDAS server synchronizes status information with EDAS Agent.

File

EDAS server logs

File	Description
agent-server.log	The log for the instance where EDAS Agent is installed.
changeorder.log	The change order log.
tengine.log	The Tengine log.

Path

- Logging path: *${user.home}/edas/logs*
- Archive path: *${user.home}/edas/logs/bak*

Format

agent-server: %d{HH:mm:ss.SSS} [%thread] %-5level %logger{36}:%line - %msg%n(time, thread name, log level, category name: number of lines - specific log information) others: %d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %-5level %logger{50}:%line - %msg%n (date and time, thread name, log level, category name: number of lines - specific log information)

Archiving policy

Archiving policy for EDAS server logs

Log	Archiving policy

Log	Archiving policy
agent-server.log	• A file is created every day. • The name of the new file takes the format agent-server-.{yyyy-MM-dd}.log. Logs from the last seven days are retained.
tengine.log	• A file is created every day. • The name of the new file takes the format tengine.{yyyy-MM-dd}.log. Logs from the last seven days are retained.
changeorder.log	• A file is created every day. • The name of the new file takes the format changeorder.{yyyy-MM-dd}.log. Logs from the last seven days are retained.

10.1.6.4. DiamondServer logs

A configuration management service. It provides configuration storage, query, and notification functions, and primarily stores database metadata and EDAS function switch configurations in EDAS.

File

DiamondServer logs

File	Description
diamondServer.log	The DiamondServer log.
fata.log	The most important system log, which records database service errors, "master db not found" messages, and other information.
dump.log	The log that records the dumping of configurations to the local device.

Path

- Logging path: *${user.home}/admin/diamond/logs*
- Archive path: *${user.home}/admin/diamond/logs*

Format

[%p] [%t] %d{MM-dd HH:mm:ss,SSS} [%c{1}] - %m%n (log information priority, log event thread name, logging time, log information category, and specific log information)

Archiving policy

The archiving policies vary depending on the version. For example, in the latest version 3.8.8, a log is 15 MB in size and 10 logs are retained.

10.1.6.5. Cai-fs logs

A file server. It stores the EDAS Agent installation package and EDAS application packages.

Log files

EDAS console logs

File	Description
efs-server.log	Cai-fs logs

Path

- Logging path: *${user.home}/efs/logs*
- Archive path: *${user.home}/efs/logs*

Format

%d{HH:mm:ss} [%thread] %-5level %logger{36} - %msg%n (time, thread name, log level, class name: number of lines - specific log information)

Archiving policies

EDAS admin log archiving policies

Log	Archiving policy
efs-server.log	• Maximum size: 100 MB • The name of the new file takes the format efs-server.log.{d}. Three logs are retained.

10.1.6.6. ConfigServer logs

An RPC service registry. It is used to query and store the publishing and subscription data of services.

Log files

EDAS console logs

File	Description
cluster.log	The log that records cluster operations, such as merging tasks and connecting to or disconnecting from other instances in the cluster.
memory.log	The log that records memory statuses, including the total number of subscriptions and the amount of persistent data of an instance.
persistent.log	The data persistence log.
push.log	The data push log.

File	Description
http.log	The log that records the instance commands called over HTTP.
monitor.log	The warning code log.

Path

- Logging path: *${user.home}/admin/configserver/log*
- Archive path: *${user.home}/admin/configserver/log*

Format

%date %level %msg%n%n (logging time, log level, and specific log information)

Archiving policies

- A file is created every day.
- The name of the new file takes the format {module}.log.%d{yyyy-MM-dd}.log. Logs from the last 15 days are retained.

10.1.6.7. Cai-address logs

An address discovery service. It provides the address lists for DiamondServer and ConfigServer.

Log files

EDAS console logs

File	Description
access.log	All access logs
error.log	Error logs

Path

- Logging path: *${user.home}/admin/cai/logs*
- Archive path: *${user.home}/admin/cai/logs*

Format

"$remote_addr $request_time_usec $http_x_readtime [$time_local] \"$request_method http://$host$request_uri\" $status $body_bytes_sent \"$http_referer\" \"$http_user_agent\" \"$md5_encode_cookie_unb\" \"$md5_encode_$cookie_cookie2\" \"$eagleeye_traceid\""; records the client IP address - request elapsed time - request header - request status - the number of bytes sent to the client - records the link from which the access request is received - records information about the web browser of the client - performs MD5 on cookies to obtain fixed-length cookies - eagleeye trace id

Archiving policies

Log splitting is not performed.

10.1.6.8. EagleEye console logs

You can query and view service traces.

Log files

EDAS console logs

File	Description
eagleeye-console.log	All console access logs
eagleeye-sql.log	Trace query logs

Path

- Logging path: *${user.home}/admin/logs*
- Archive path: *${user.home}/admin/logs*

Format

%d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %msg%n (time, thread name, and specific log information)

Archiving policies

Archiving policies of EagleEye console access logs

Log	Archiving policy
eagleeye-console.log	• Maximum size: 500 MB • Retention period: 30 days
eagleeye-sql.log	• Maximum size: 200 MB • Retention period: 15 days

10.1.7. Configuration reference

You need to complete basic configuration and optimization configuration during the EDAS O&M process.

The configuration during EDAS O&M is divided into component configuration and JVM configuration.

10.1.7.1. Component configuration

You can configure the basic settings of components by using configuration files.

Parameters

Component	Configuration file	Path	Configuration item	Description	Value

EDAS console	config.properties	/home/admin/edas/conf/	dataSource.config.URL	Database connection string	Standard database connection string, such as *jdbc:mysql://edastest.mysql.rds.aliyuncs.com/edas?rewriteBatchedStatements=true*
			dataSource.config.user	Username	Standard database username
			dataSource.config.password	Password	Standard database password
EDAS admin	config.properties	/home/admin/edas/conf/	dataSource.config.URL	Database connection string	Standard database connection string, such as *jdbc:mysql://edastest.mysql.rds.aliyuncs.com/edas?rewriteBatchedStatements=true*
			dataSource.config.user	Username	Standard database username
			dataSource.config.password	Password	Standard database password
Redis console	redis.conf	/home/admin/redis-2.8.17/src/	dataSource.config.URL	Database connection string	Standard database connection string
			dataSource.config.user	Username	Standard database username
			dataSource.config.password	Password	Standard database password

TLog console	tlog-cloud.properties	/home/admin/taobao-tomcat-production-7.0.59.3/lib/	config.tlog.zk.servers	ZooKeeper connection string	Standard ZooKeeper address information, such as *192.168.1.2:2181, 192.168.1.3:2181, and 192.168.1.4:2181*
			config.tlog.hbase.zkServers	ZooKeeper used by HBase	Standard ZooKeeper connection string, which is shared by default and is consistent with the preceding ZooKeeper
			config.tlog.hbase.zkRootNode	HBase root node	Default value: /hbase
			config.nimbus.host	JStorm nimbus node	IP address of the primary node
			config.edas.console.url	EDAS admin address	EDAS admin domain name
HBase	hbase-site.xml	/home/admin/hbase{-Version}/conf/	hbase.rootdir	Host name of the primary instance	Note that the host name cannot be an IP address and must be bound in */etc/hosts* if it cannot be resolved. All HBase-based applications must be bound to the host names of all HBase instances.
			hbase.zookeeper.quorum	ZooKeeper connection string	ZooKeeper connection string

			hbase.zookeeper.property.clientPort	ZooKeeper port	ZooKeeper port
Jstorm	storm.yaml	/home/admin/jstorm/conf/	storm.zookeeper.servers	IP addresses of all storm nodes	IP addresses of all storm nodes
			nimbus.host	Primary node of JStorm nimbus	Primary node of JStorm nimbus
			supervisor.slots.port.cpu.weight	CPU weight	Number of CPUs occupied by each task
ConfigServer	confsrv.conf	/home/admin/configserver/conf/	serverlist	Server list	A list of IP addresses separated with commas (,)
			unitserverlist	Modular server list	The content is the same as that of serverlist.
DiamondServer	config.properties	/home/admin/diamond/target/diamond.war/WEB-INF/classes/	openInnerInterfaceFilter	Indicates whether to enable internal interface access verification.	The default value is false. Enter false to avoid failed verification because Address-Server is typically configured with a virtual IP address rather than a real one.
			OPEN_SPAS	Indicates whether to enable authentication.	The value is true, which indicates that authentication is enabled.

10.1.7.2. JVM configuration

You can optimize system performance through a JVM configuration.

The parameters vary slightly depending on the JDK versions.

Parameters

Name	Description	Applicable JDK version	Reference value
-Xms	Specifies the initial heap memory size for the JVM.	All JDK versions	4 GB
-Xmx	Specifies the maximum heap memory size for the JVM.	All JDK versions	4 GB
-Xmn	Specifies the size of the young generation.	All JDK versions	2 GB
-Xss	Specifies the stack size of each thread.	All JDK versions	2 MB
-XX:+UseCompressedOops	Compresses common object pointers.	All JDK versions	-
-XX:SurvivorRatio	Specifies the ratio of Survivor to Eden.	All JDK versions	10
-XX:+UseConcMarkSweepGC	Uses the Concurrent Mark Sweep (CMS) collector for memory collection.	All JDK versions	-
-XX:+UseCMSCompactAtFullCollection	Instructs the CMS collector to compress the old generation upon full GC.	All JDK versions	-
-XX:CMSMaxAbortablePrecleanTime		All JDK versions	5000
-XX:+CMSClassUnloadingEnabled	Specifies that CMS GC is triggered after class unloading.	All JDK versions	-
-XX:CMSInitiatingOccupancyFraction	Sets the threshold size of the old generation that triggers CMS GC.	All JDK versions	80
-XX:PermSize	Specifies the initial value of the permanent generation.	1.7 and earlier versions	196 MB
-XX:MaxPermSize	Specifies the maximum value of the permanent generation.	1.7 and earlier versions	256 MB

Name	Description	Applicable JDK version	Reference value
MetaspaceSize	Sets the threshold size of the allocated metadata space that triggers full GC.	1.8 and later versions	196 MB
MaxMetaspaceSize	Sets the maximum size of the allocated metadata space that triggers full GC.	1.8 and later versions	256 MB
-XX:+DisableExplicitGC	Disables System.gc().	All JDK versions	-
-XX:+HeapDumpOnOutOfMemoryError		All JDK versions	-
-XX:HeapDumpPath	Specifies the heap dump path.	All JDK versions	/home/admin/logs/oomDump.log

11.Operations of big data products

11.1. Apsara Big Data Manager (ABM) platform

11.1.1. What is Apsara Big Data Manager?

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

ABM supports the following services:

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

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

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

11.1.2. Common operations

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

Search for a project

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

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

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

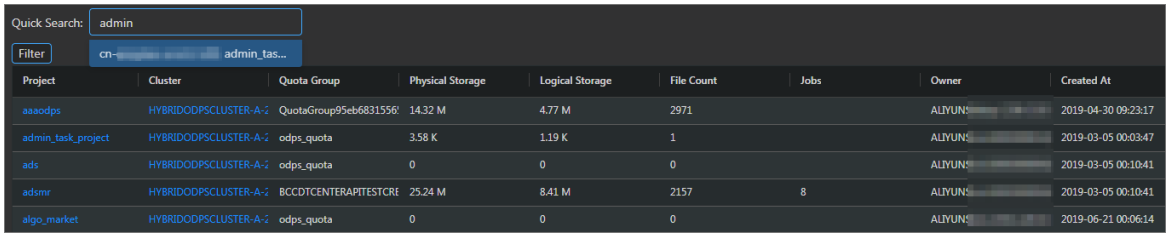

The following figure shows the search result.

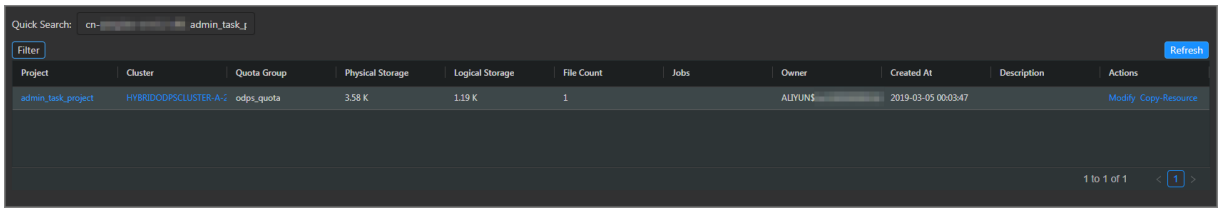

Filter projects

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

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

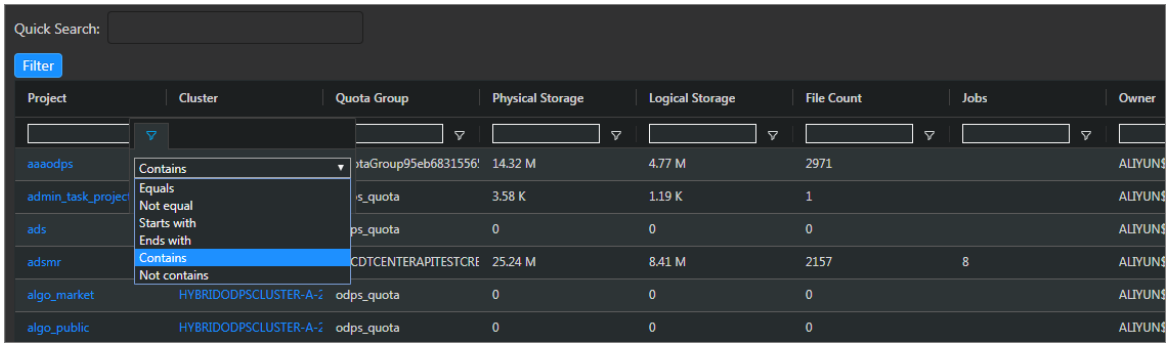

You can select one of the following filtering methods:

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

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

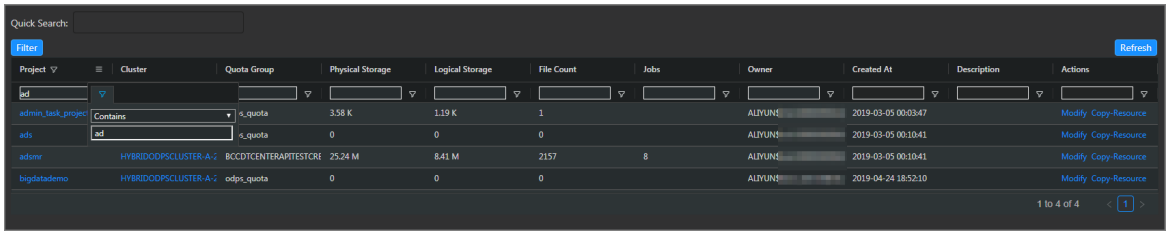

5. If the filtering result is not accurate, you can continue performing this operation on other columns.

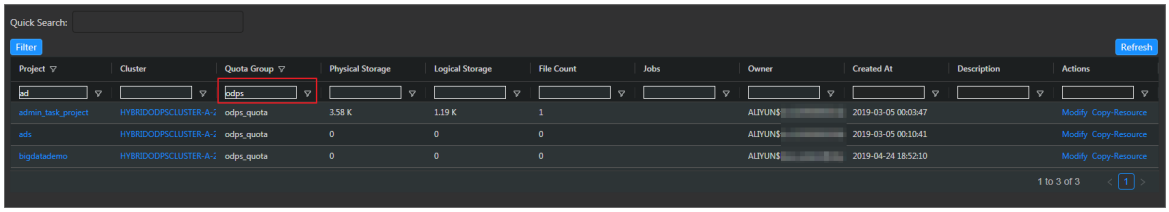

After you set the filter conditions for the projects, the **Filter** button is highlighted. If you need to cancel filtering, click the highlighted **Filter** button.

Search for an item

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

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

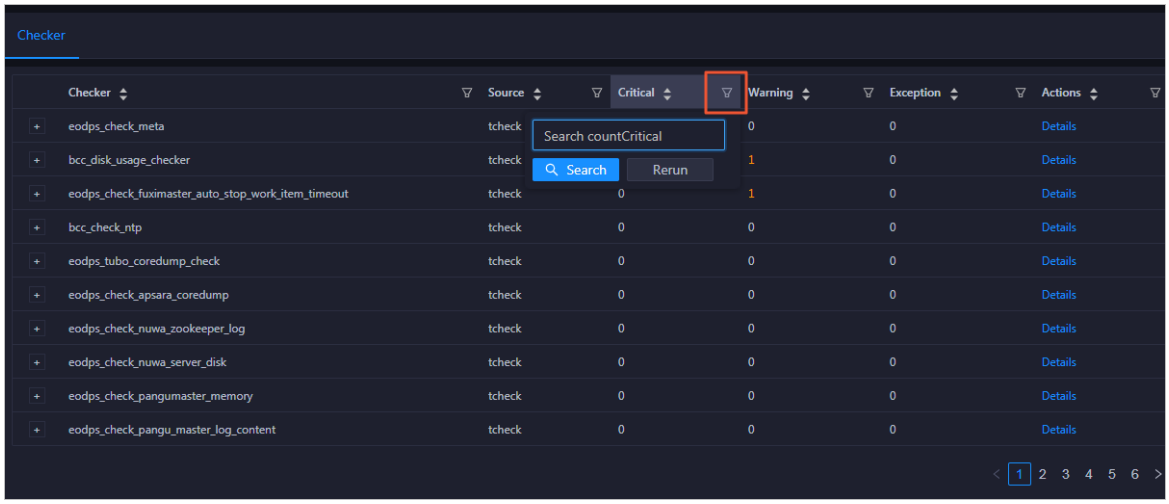

3. Click **Search**. The checkers that meet the requirements appear.
4. If the search result is not accurate, you can continue performing this operation on other columns.

Customize a column

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

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

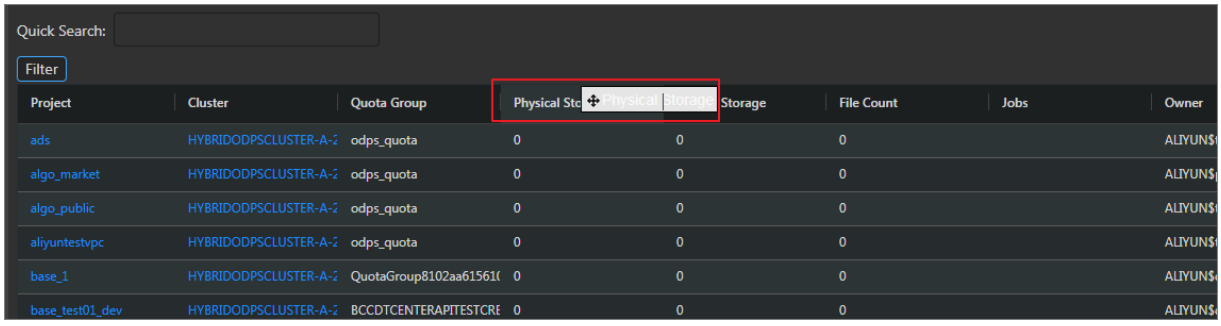

You can click ☰ in a column heading to customize the column.

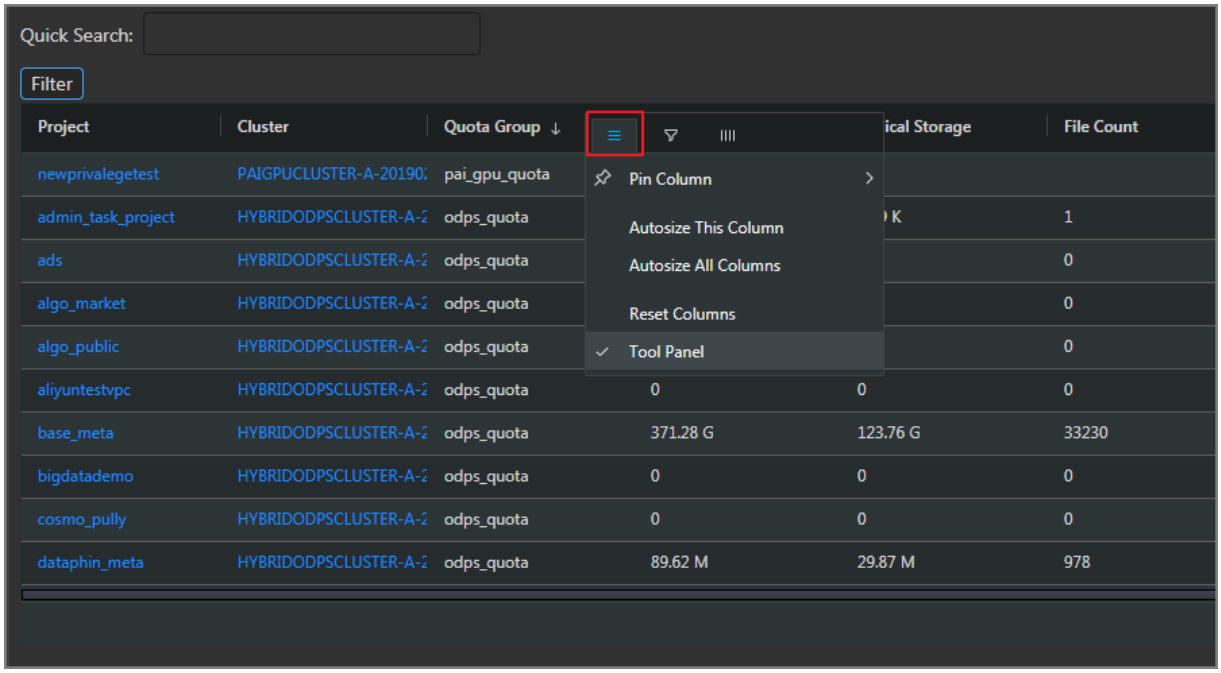

- **Pin Column**: allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.
- **Autosize This Column**: allows you to adjust the width of a column automatically.
- **Autosize All Columns**: allows you to adjust the width of all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

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

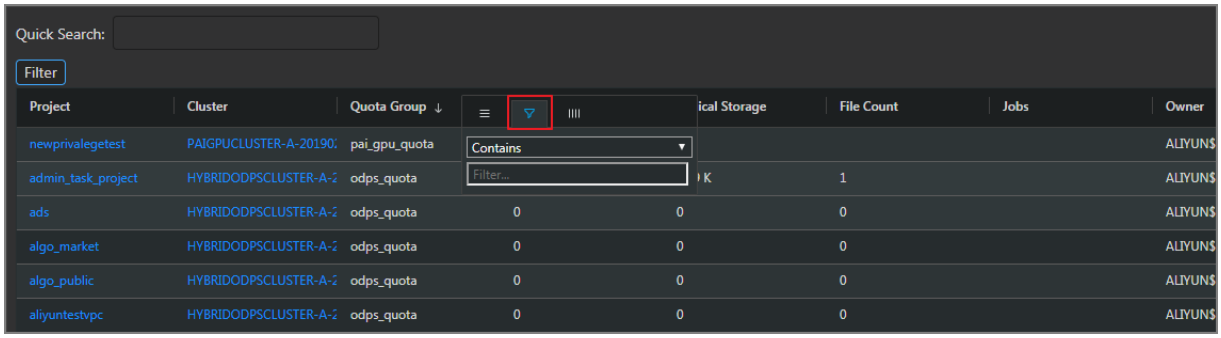

Click in a column heading and select the columns to display.

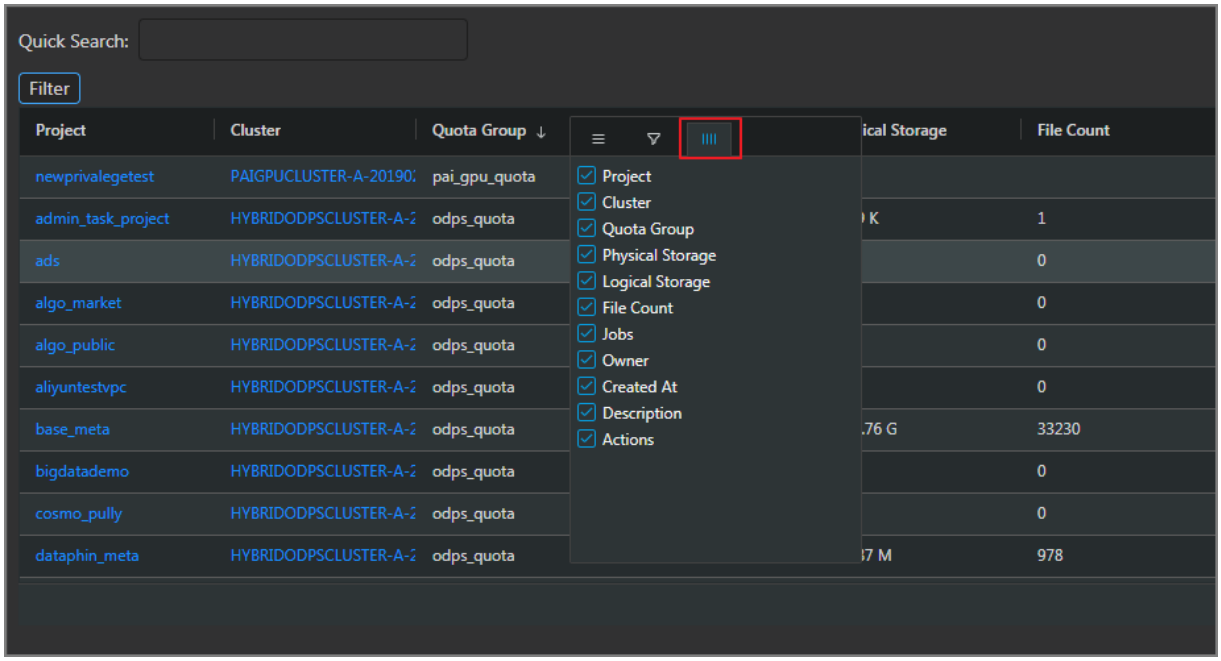

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

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

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

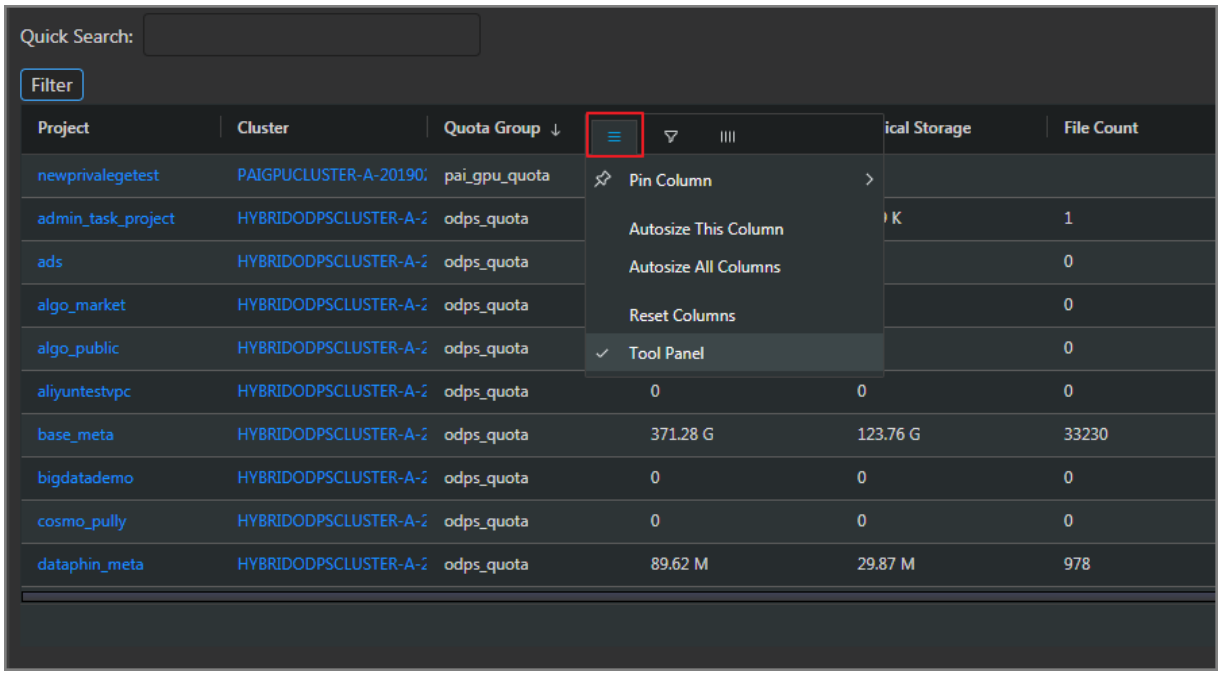

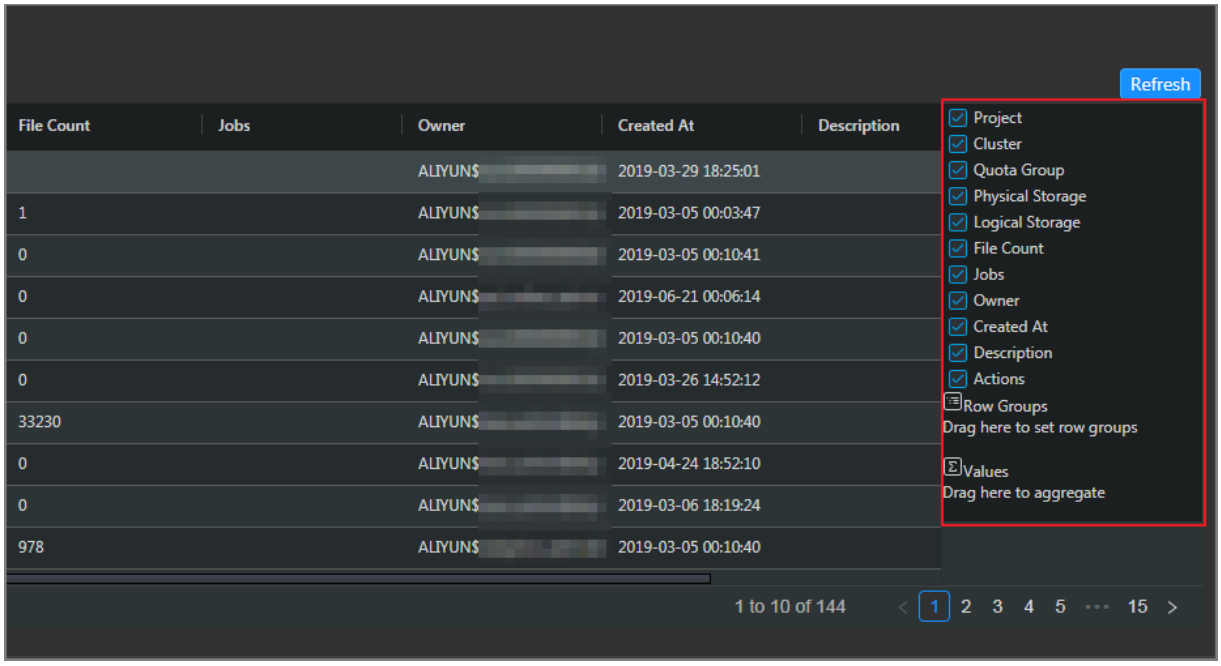

Sort projects based on a column

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

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

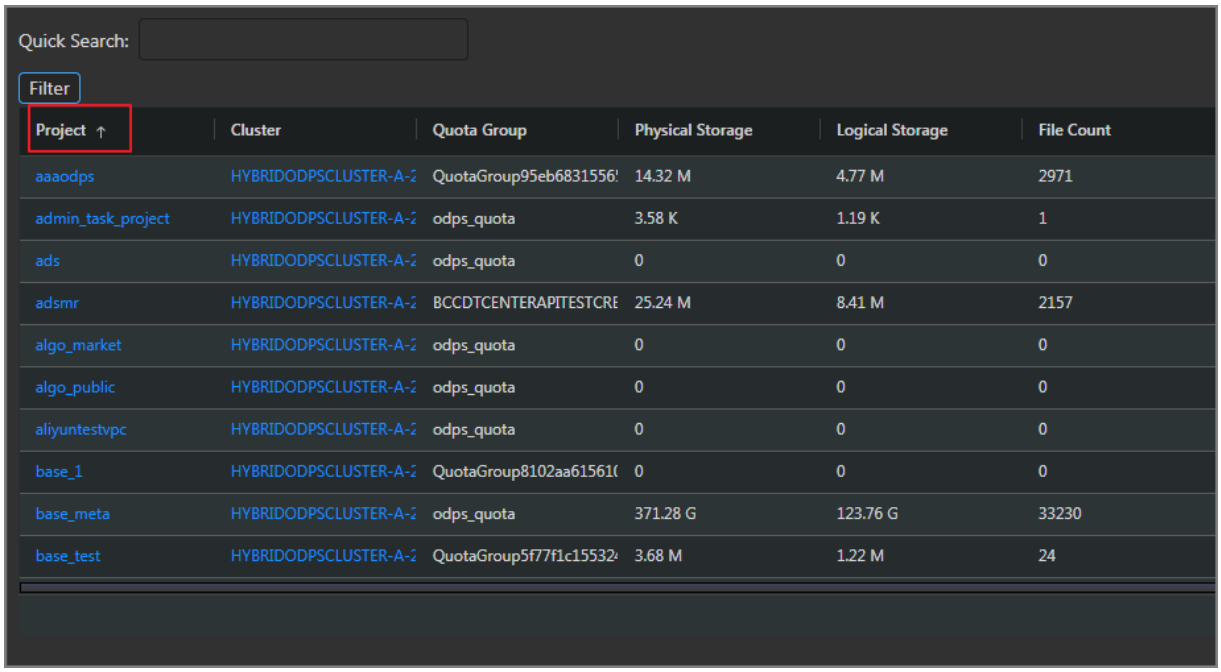

Sort items based on a column

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

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

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

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

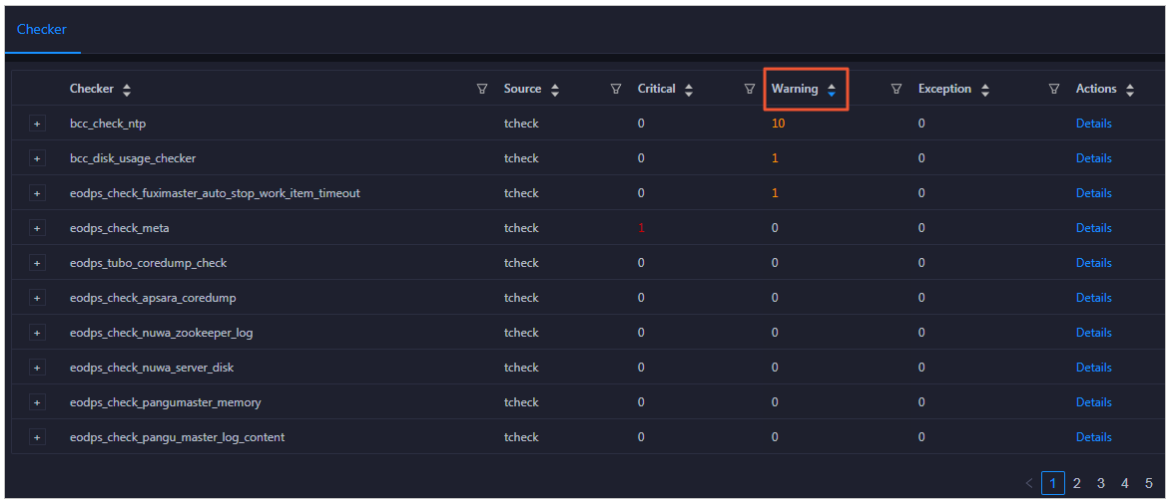

The highlighted up arrow indicates that the checkers are sorted in ascending order. The highlighted down arrow indicates that the checkers are sorted in descending order.

View the trend charts for a MaxCompute cluster

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

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

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

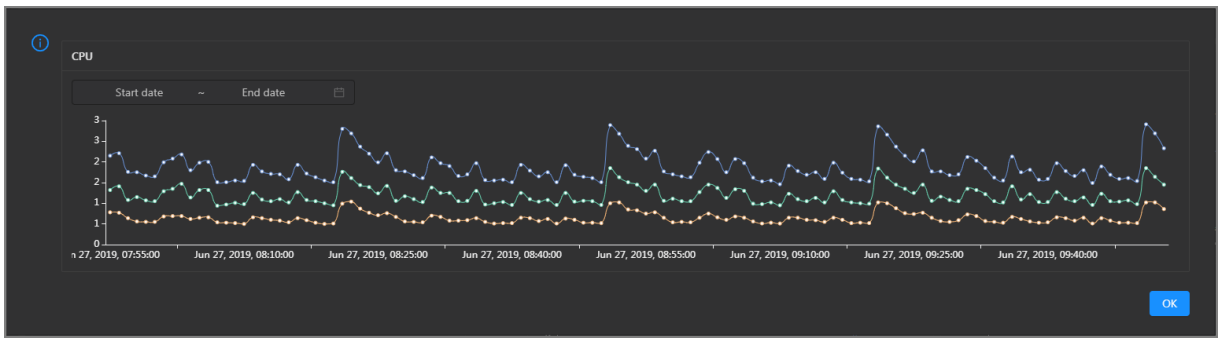

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

View the trend charts for a DataHub cluster

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

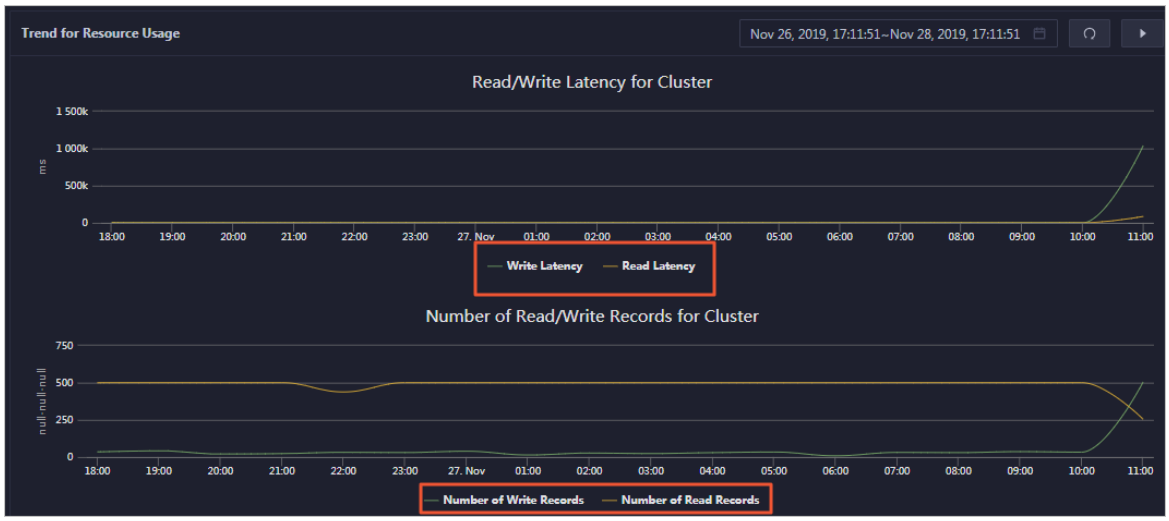

The trend charts, such as the trend charts of the read/write latency and the number of read/write records, appear in the Trend for Resource Usage section. Each chart displays the trend lines of the metrics over time in different colors. You can customize the metrics to display. You can click the name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

11.1.3. Quick start

11.1.3.1. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

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

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

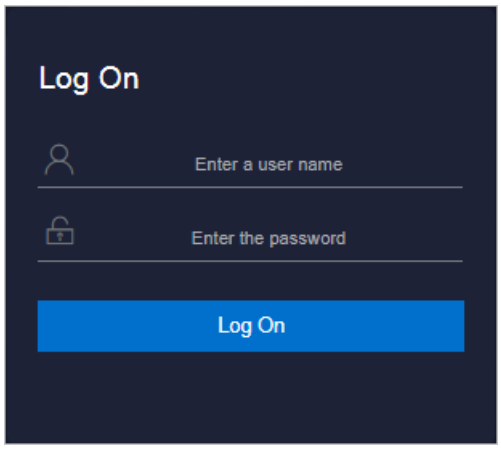

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.1.3.2. Set the theme of the console

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

Prerequisites

An ABM account and the corresponding password are obtained.

Procedure

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

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

11.1.3.3. View the dashboard

The Apsara Big Data Manager (ABM) dashboard shows key operation information about MaxCompute, DataWorks, StreamCompute, and DataHub. It also provides more information about alerts for all big data services. This helps you understand the overall running status of the big data services.

Prerequisites

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

Background information

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

Procedure

1. Log on to the ABM console.

   The **Dashboard** tab appears. To return to the **Dashboard** tab, click the icon in the upper-left corner and then click **ABM**.

2. View and clear service alerts.

   You can view the number of alerts for all big data services. The **Critical** and **Warning** alerts must be cleared in a timely manner.

i. On the **Dashboard** tab, click the number of **Critical** or **Warning** alerts of a service. The **Health Status** tab of the **Clusters** tab for the service appears.

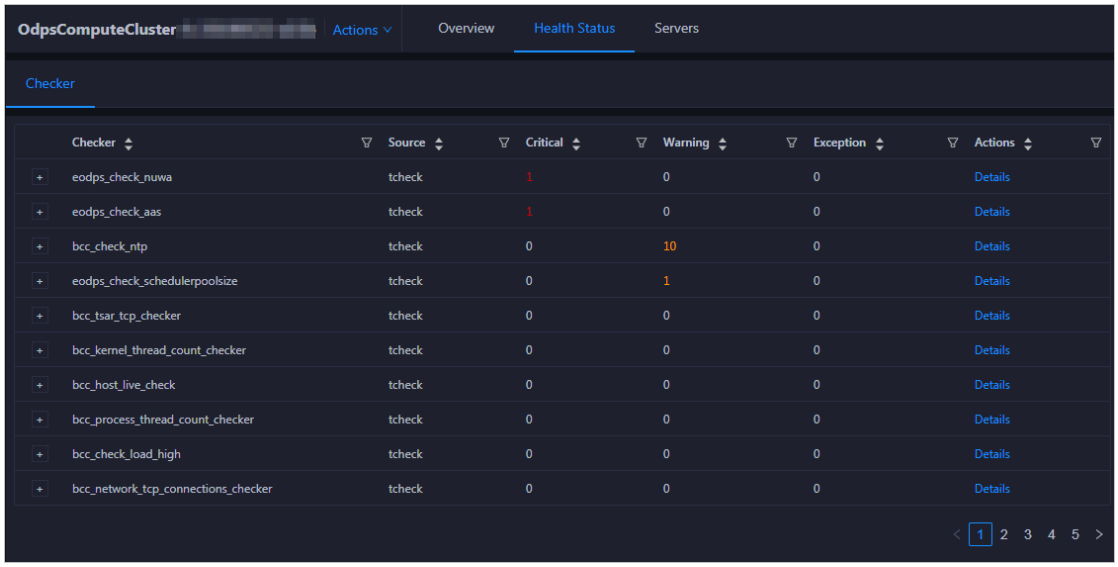

On the **Health Status** tab, you can view all the checkers for the service.

ii. Click **Details** in the Actions column of a checker for which alerts are reported. In the Details dialog box, view the details of the checker and the scheme to clear the alerts. Perform the steps provided in the **Description** section to clear the alerts.

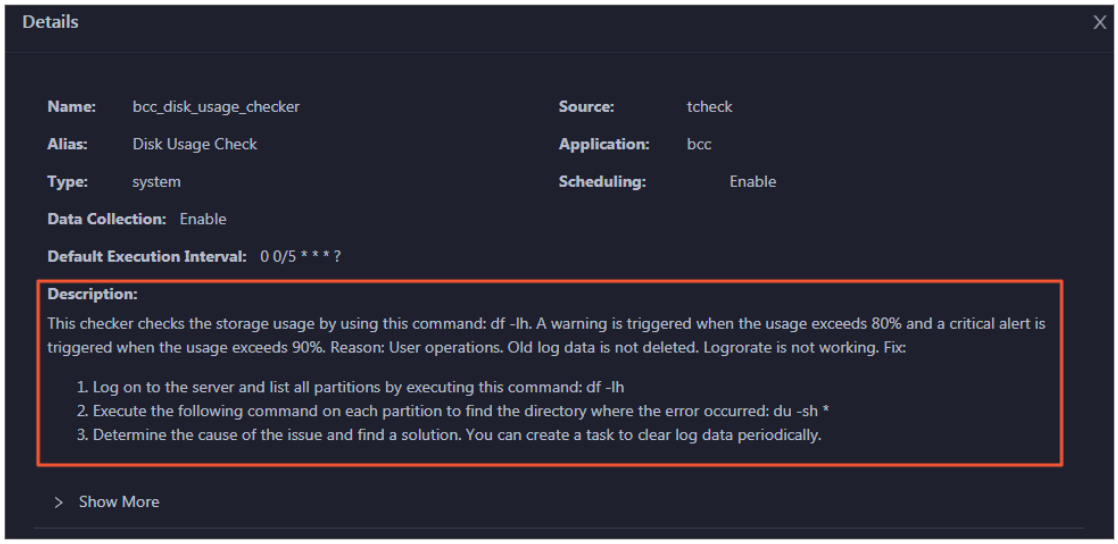

iii. Log on to the hosts on which alerts are detected and clear the alerts.

Click + to expand the checker for which alerts are reported. Then, click the **Log On** icon next to the name of the host on which you want to clear the alerts.

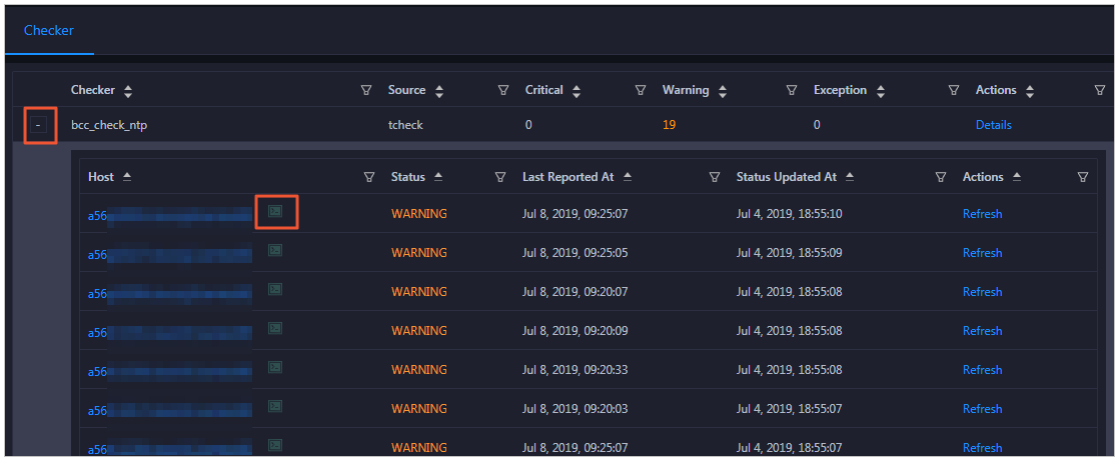

iv. In the left-side navigation pane of the **TerminalService** page, click the hostname to log on to the host.

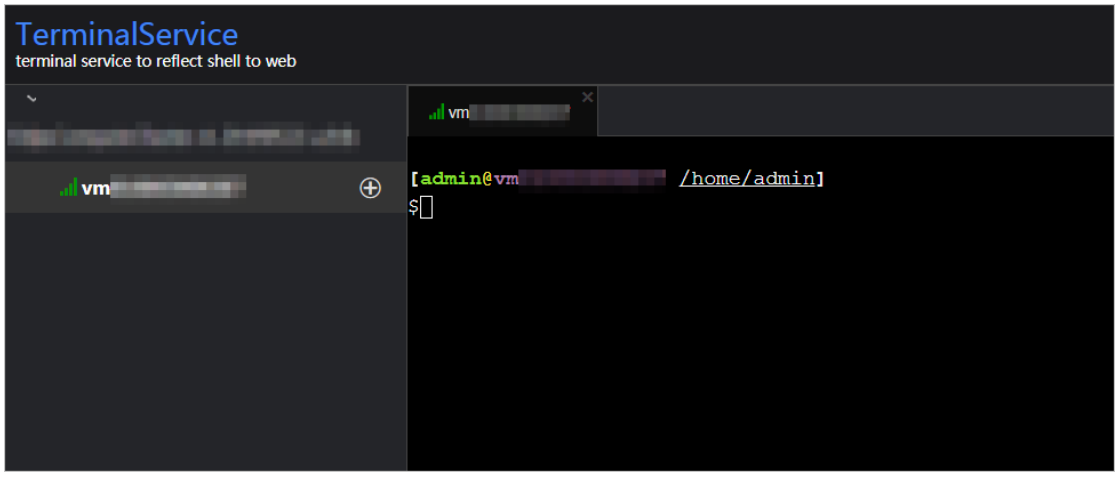

3. On the **Dashboard** tab, click **MaxCompute** to view operation information about **MaxCompute**.

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 utilization, and physical CPU utilization.

4. On the **Dashboard** tab, click **DataWorks** to view operation information about **DataWorks**.

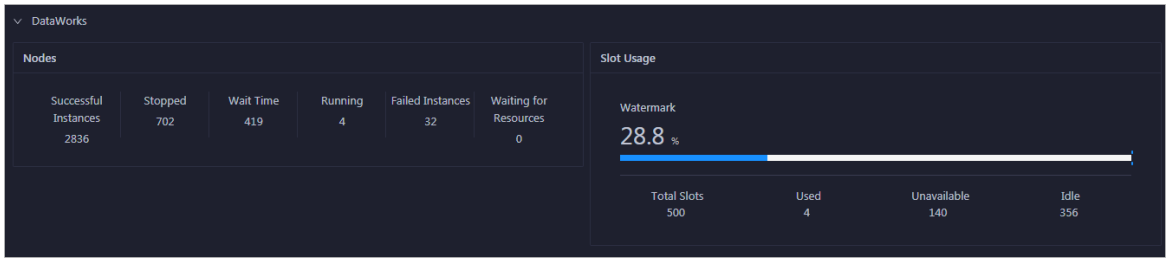

In the **DataWorks** section, you can view the node scheduling and slot usage of a DataWorks cluster. You can also view the trend chart of the total number of finished tasks.

5. On the **Dashboard** tab, click **StreamCompute** to view operation information about **StreamCompute**.

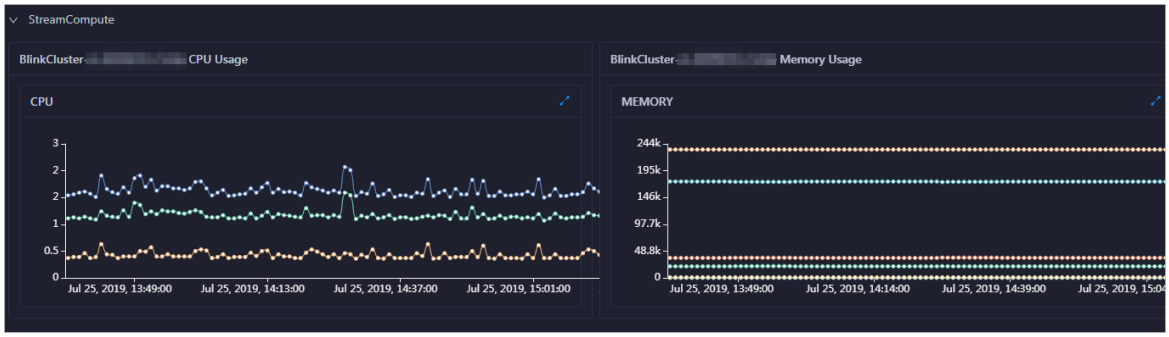

In the **StreamCompute** section, you can view the trend charts of the transactions per second (TPS), failover rate, CPU utilization, and memory usage of a StreamCompute cluster.

6. On the **Dashboard** tab, click **DataHub** to view operation information about **DataHub**.

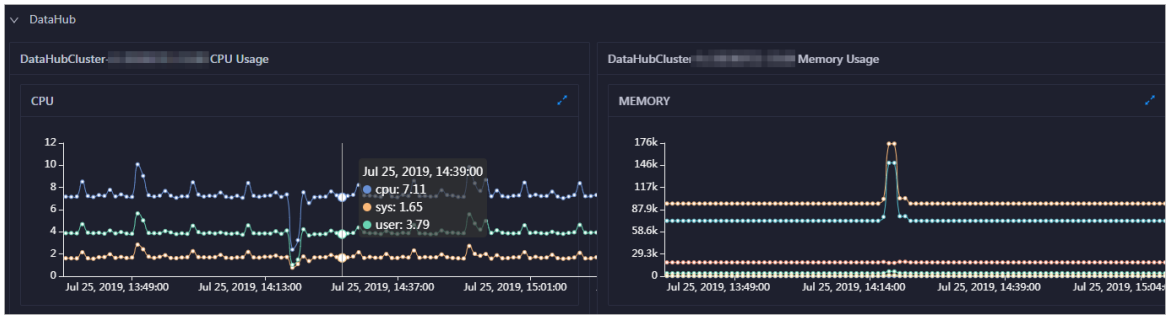

In the **DataHub** section, you can view the trend charts of the read/write latency, read/write records, read/write queries per second (QPS), and read/write throughput. You can also view the trend charts of CPU utilization and memory usage of a DataHub cluster.

11.1.3.4. View the cluster running status

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

Prerequisites

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

Context

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

Procedure

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

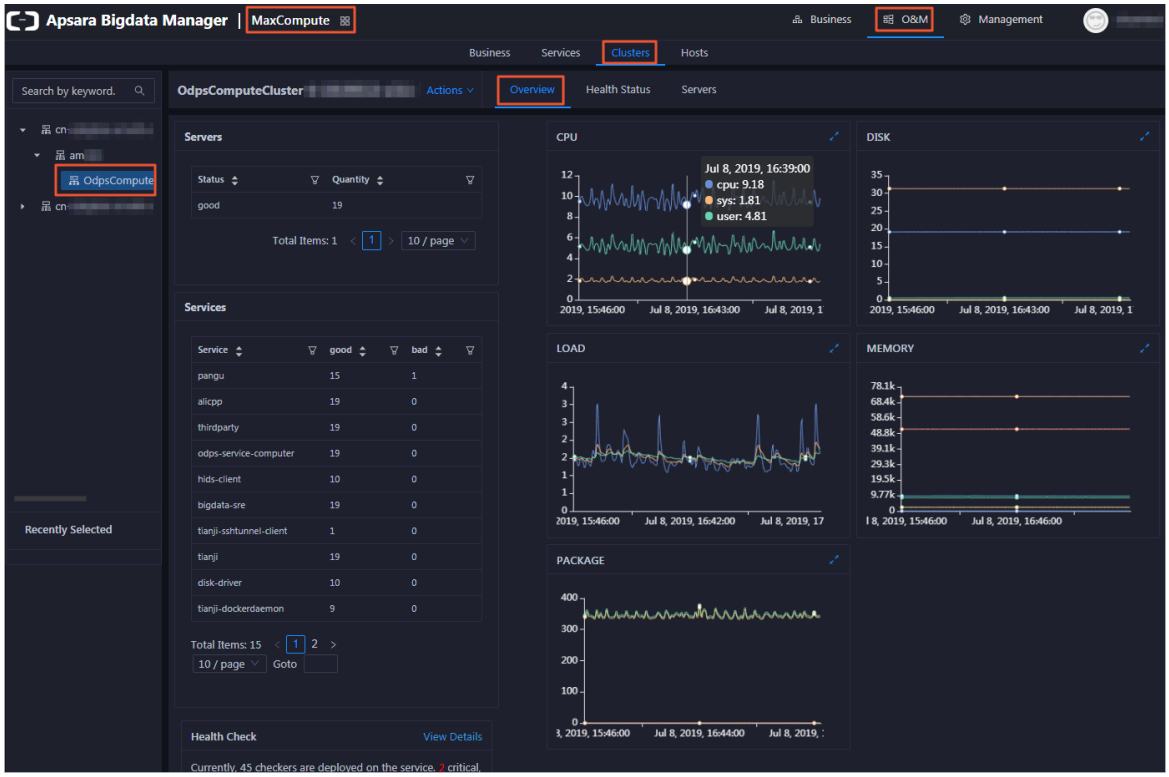

On the **Overview** page, you can view the host status, service status, health check result, and health check history of the selected cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster.

What's next

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

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

11.1.3.5. 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 Big Data Manager (ABM) service as an example to describe how to view and clear alerts.

Prerequisites

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

Context

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

Procedure

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

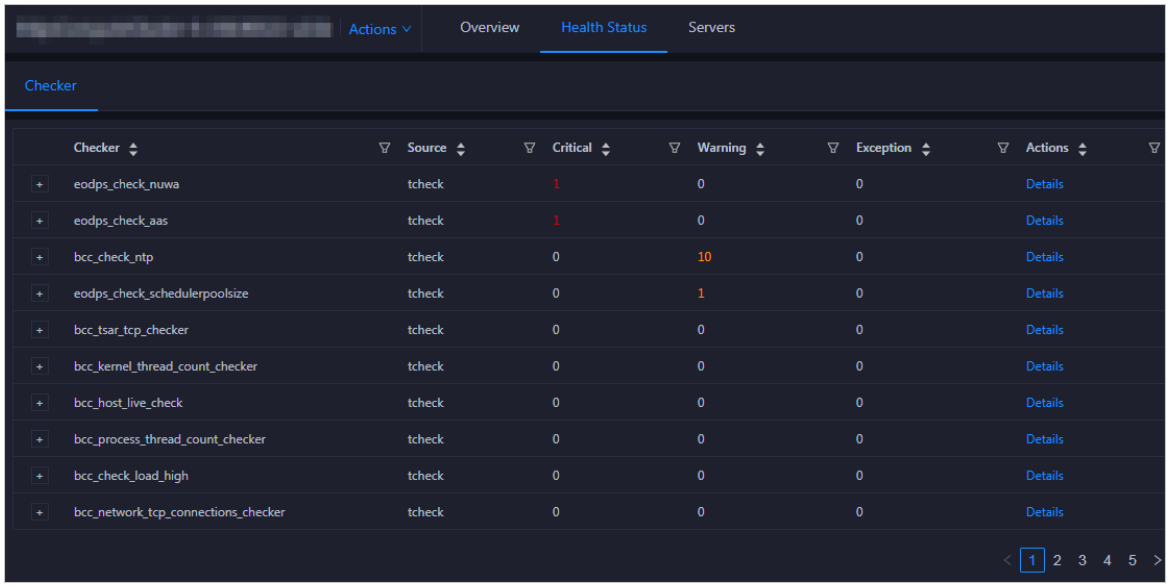

4. On the **Health Status** page, click **+** to expand a checker with alerts. You can view all hosts where the checker is run.

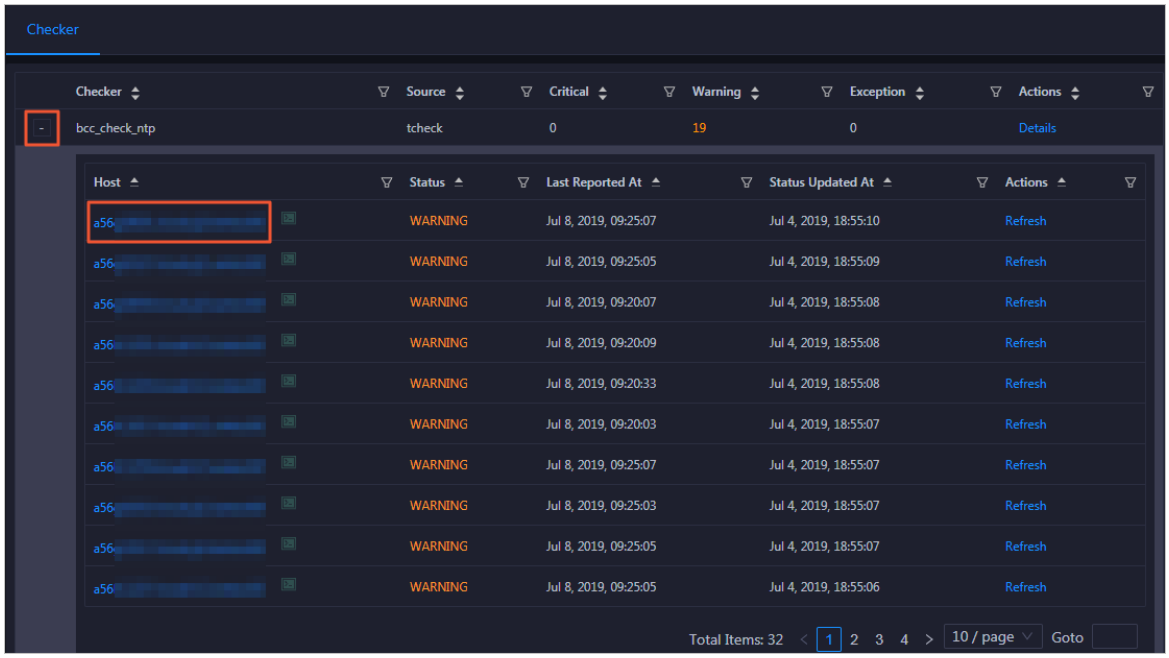

5. Click a hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

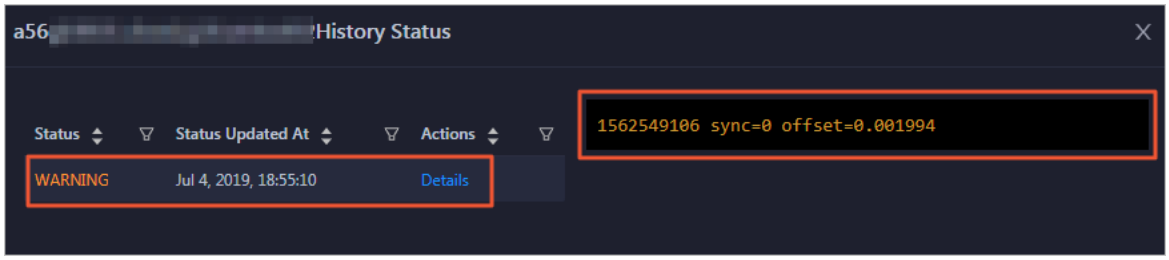

6. On the **Health Status** page, click **Details** in the Actions column of the checker to view the schemes to clear the alerts.

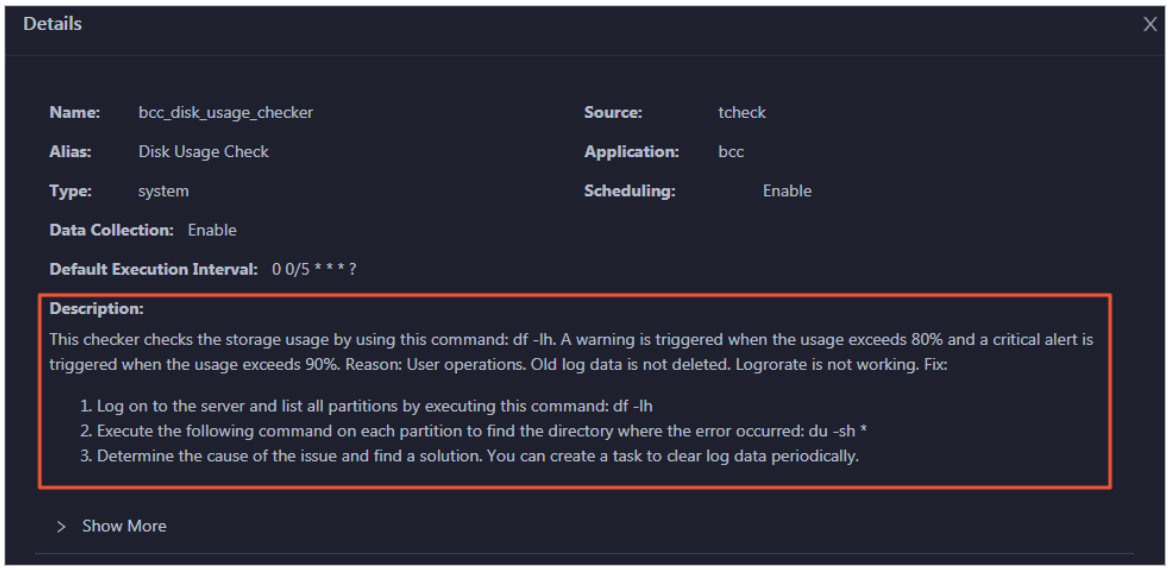

7. Clear the alerts according to the schemes.

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

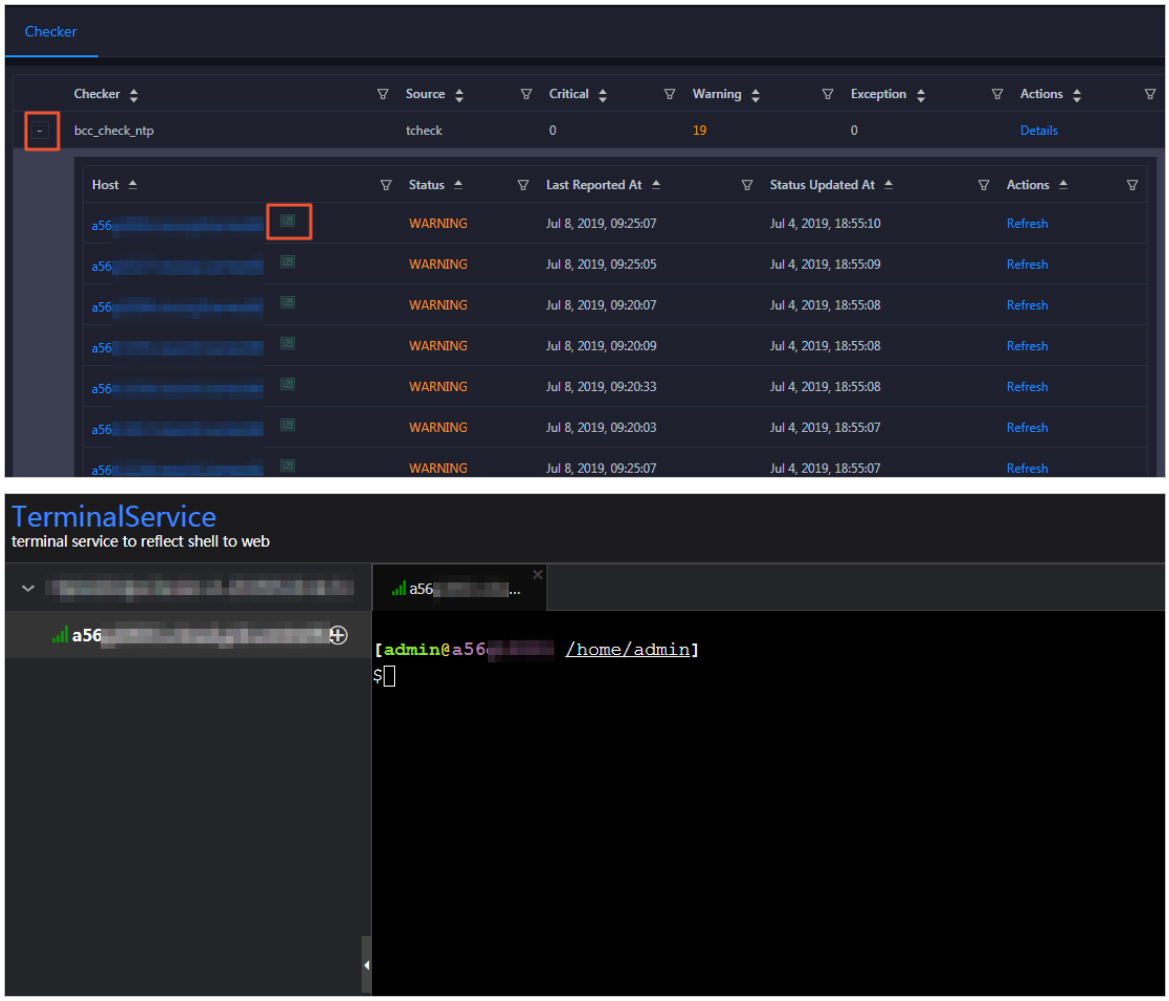

8. After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

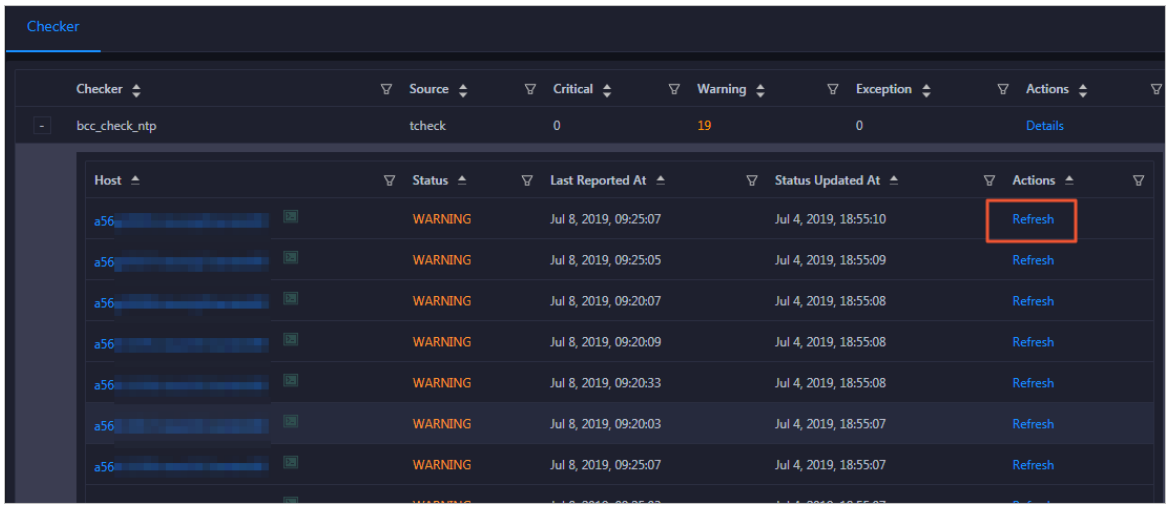

11.1.4. ABM

11.1.4.1. ABM dashboard

The Apsara Big Data Manager (ABM) dashboard shows key operation information about MaxCompute, DataWorks, StreamCompute, and DataHub. It also provides information about alerts for all big data services. This helps you understand the overall running status of the big data services. The dashboard also supports automatic data refresh and full screen display.

Go to the Dashboard tab

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

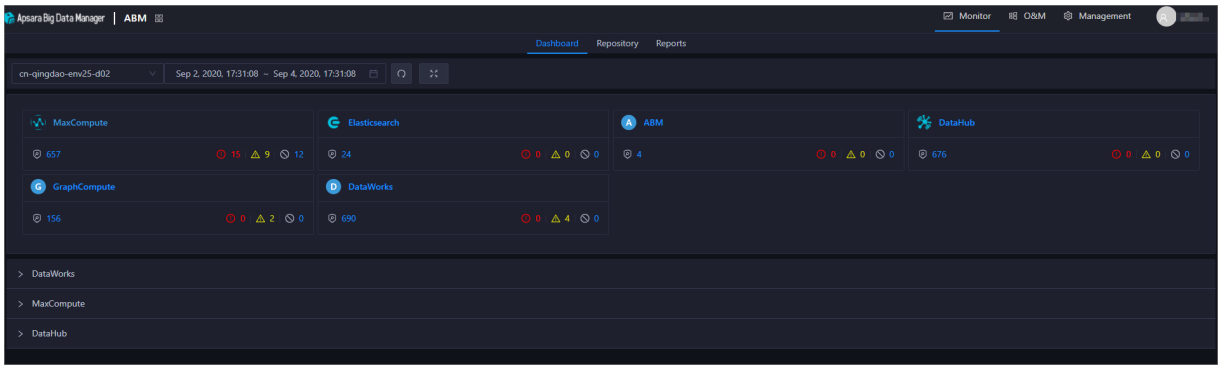

In the upper-left corner of the **Dashboard** tab, you can select a region from the drop-down list to view the cluster running status of each big data service in that region.

View and clear the alerts of various services

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

1. On the **Dashboard** tab, click the number of **Critical** or **Warning** alerts of a service. The **Health Status** tab of the **Clusters** tab for the service appears.

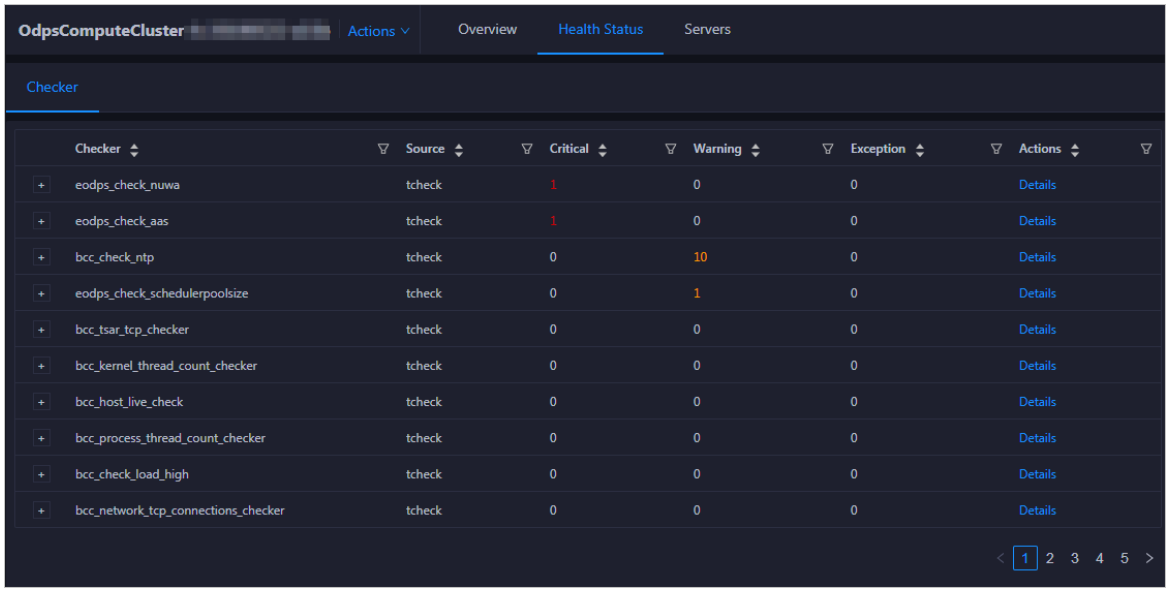

   On the **Health Status** tab, you can view all the checkers of the service.

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

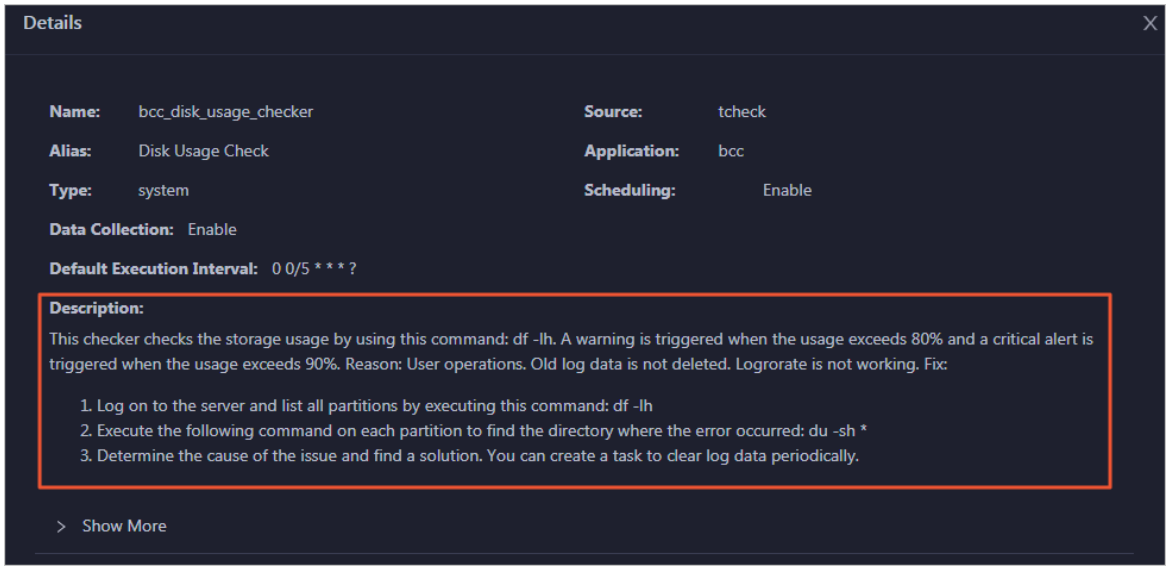

3. Log on to the hosts on which the alerts are detected to clear the alerts.

   Click + to expand a checker. Then, click the **Log On** icon next to the name of the host on which the alerts are detected. On the page that appears, click the hostname on the left to log on to the host.

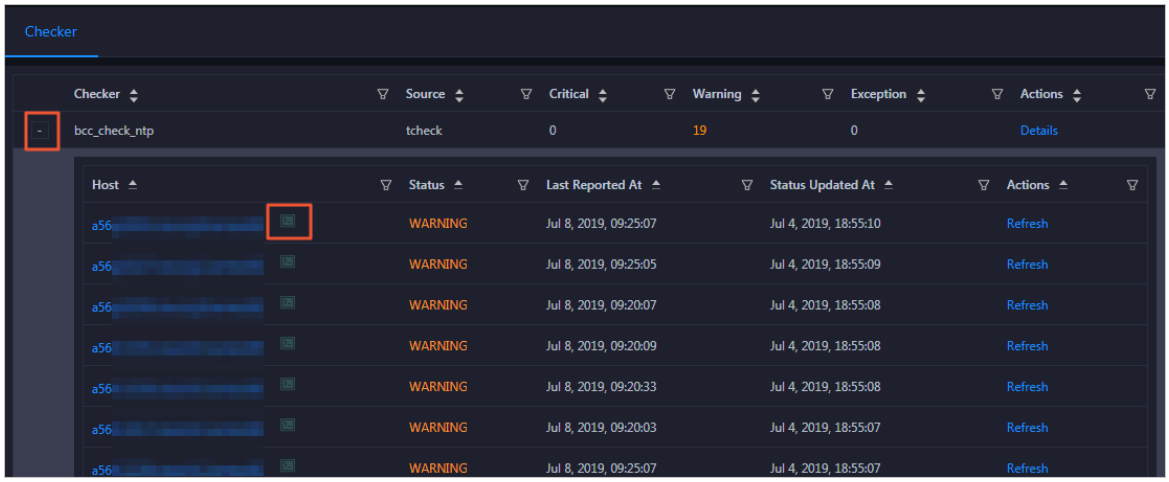

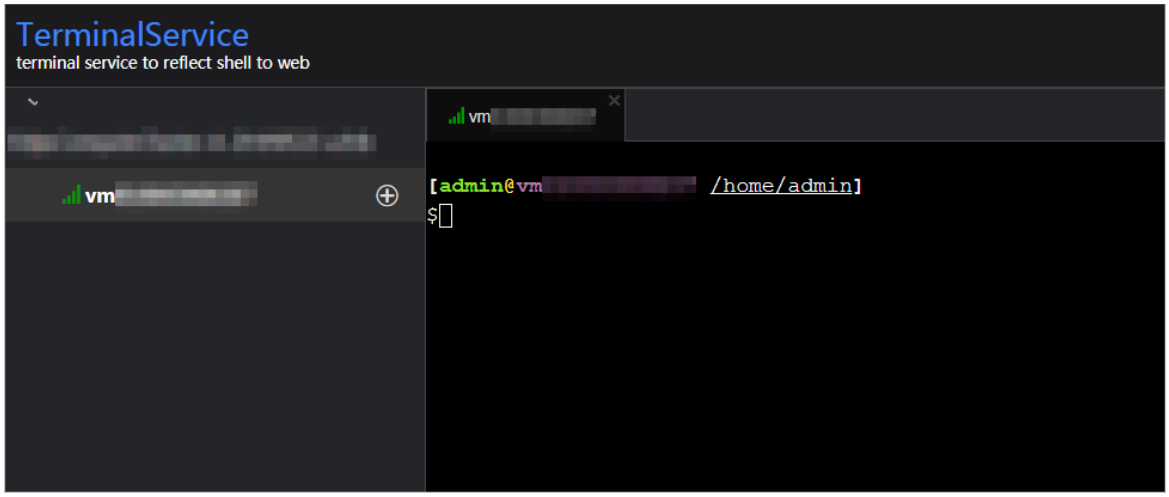

View key operation information about MaxCompute

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

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 utilization, and physical CPU utilization.

View key operation information about DataWorks

The ABM dashboard shows key operation information about DataWorks. On the **Dashboard** tab, click **DataWorks** to view the information.

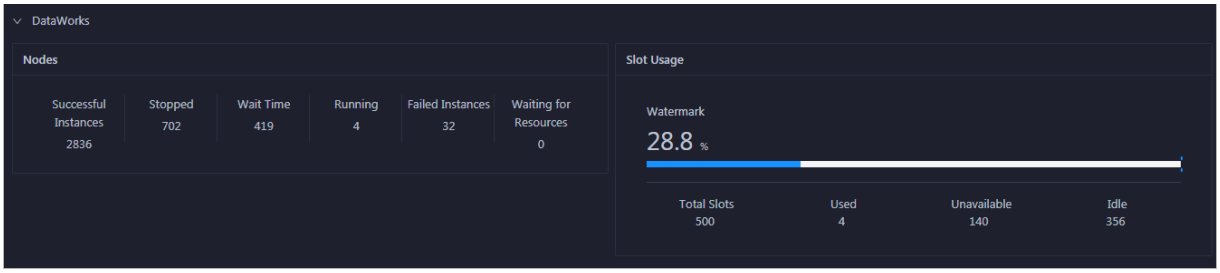

In the **DataWorks** section, you can view the node scheduling and slot usage of a DataWorks cluster. You can also view the trend chart of the total number of finished tasks.

View key operation information about StreamCompute

The ABM dashboard shows key operation information about StreamCompute. On the **Dashboard** tab, click **StreamCompute** to view the information.

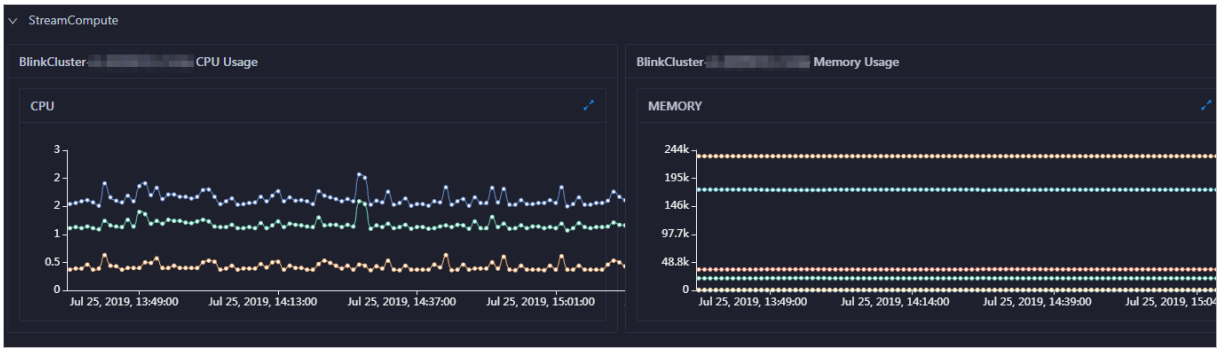

In the **StreamCompute** section, you can view the trend charts of the transactions per second (TPS), failover rate, CPU utilization, and memory usage for a StreamCompute cluster.

View key operation information about DataHub

The ABM dashboard shows key operation information about DataHub. On the **Dashboard** tab, click **DataHub** to view the information.

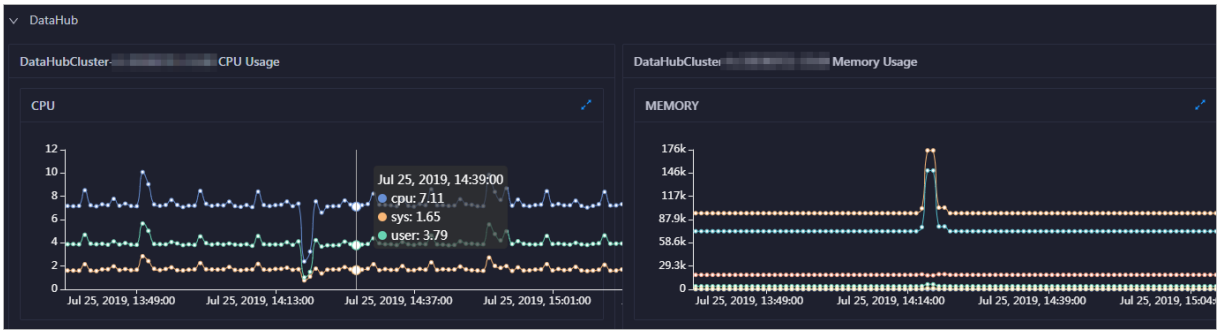

In the **DataHub** section, you can view the trend charts of the read/write latency, read/write records, read/write queries per second (QPS), and read/write throughput. You can also view the trend charts of CPU utilization and memory usage of a DataHub cluster.

Display the dashboard in full screen mode

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

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

11.1.4.2. ABM repository

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

Entry

1. Log on to the ABM console.

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

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

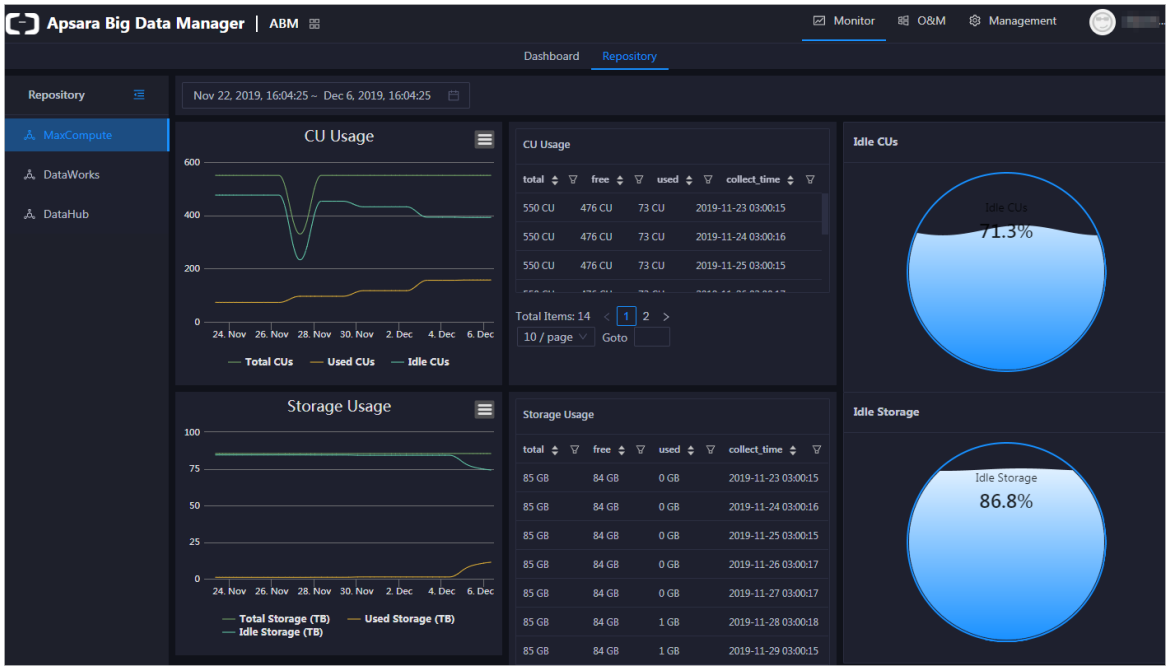

View the resource usage in MaxCompute

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

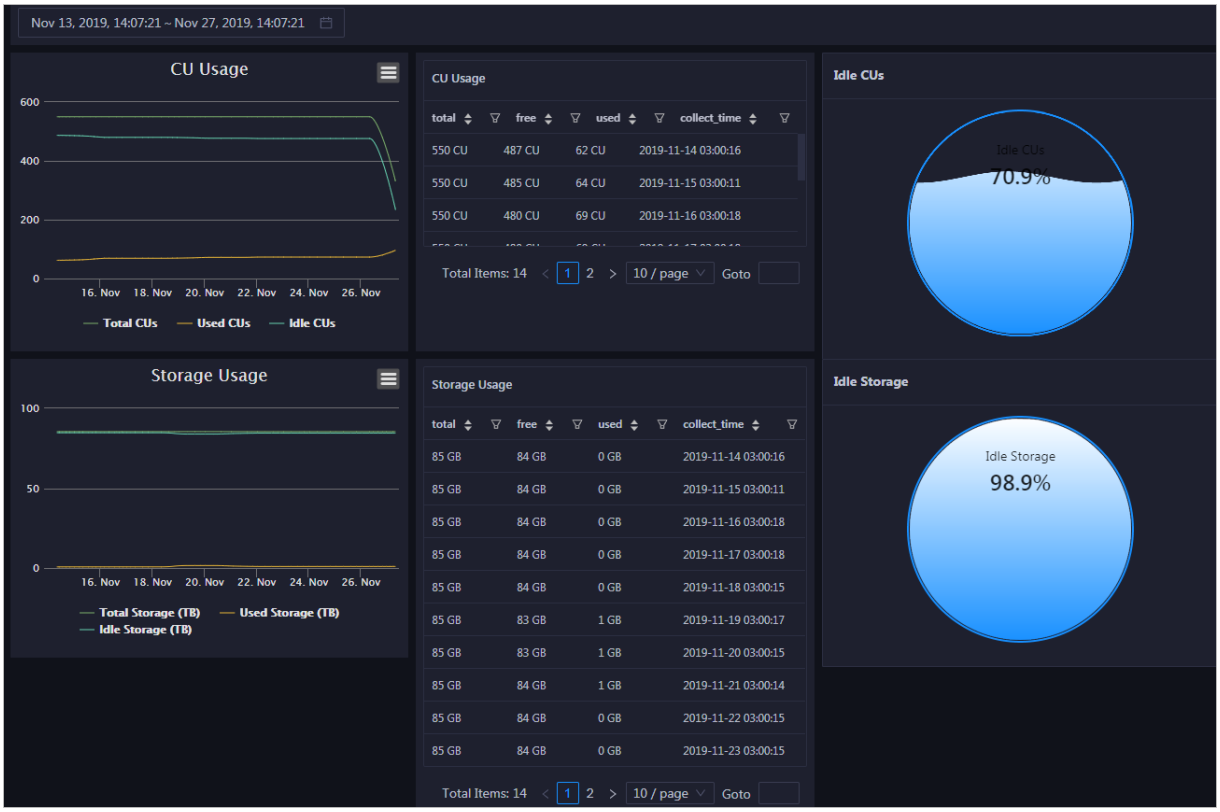

For MaxCompute, the Repository page displays the trend charts of CU and storage usage, records of CU and storage usage, and proportions of idle CUs and storage.

View the resource usage in DataWorks

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

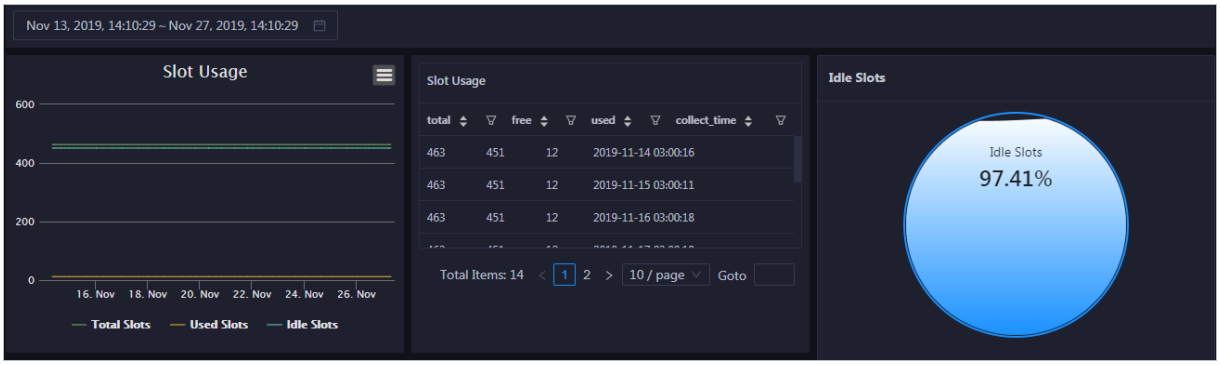

For DataWorks, the Repository page displays the trend chart of slot usage, records of slot usage, and proportion of idle slots.

View the resource usage in DataHub

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

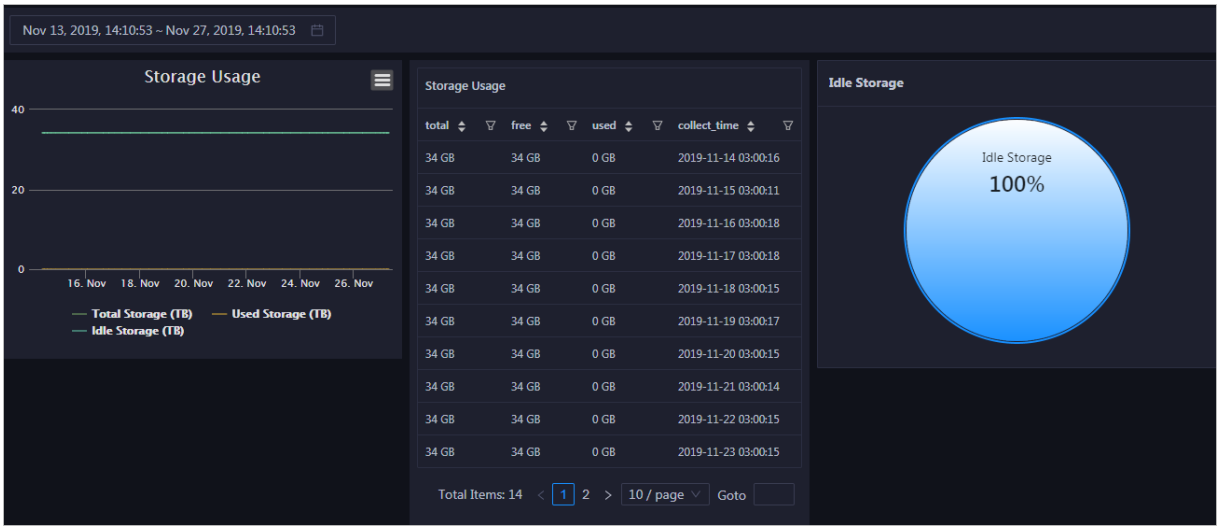

For DataHub, the Repository page displays the trend chart of storage usage, records of storage usage, and proportion of idle storage.

Other operations

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

11.1.4.3. ABM O&M overview

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

Modules

ABM O&M includes the following modules: services, clusters, and hosts. The following table describes these modules.

Module	Feature	Description
Services	Overview	Shows the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each server role in a cluster.
	Server	Shows the host list of each server role in a cluster so that you can understand the deployment of server roles on hosts.
Clusters	Overview	Shows the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Health Status	Shows all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.

Module	Feature	Description
Hosts	Overview	Shows the overall running and health check information about a host. You can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check results, and health check history of the host. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.
	Health Status	Shows the checkers of the selected host, including the checker details, health check results, health check history, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.

Go to the ABM O&M page

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

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

11.1.4.4. Service O&M

11.1.4.4.1. Service overview

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

Entry

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

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

CPU

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

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

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.

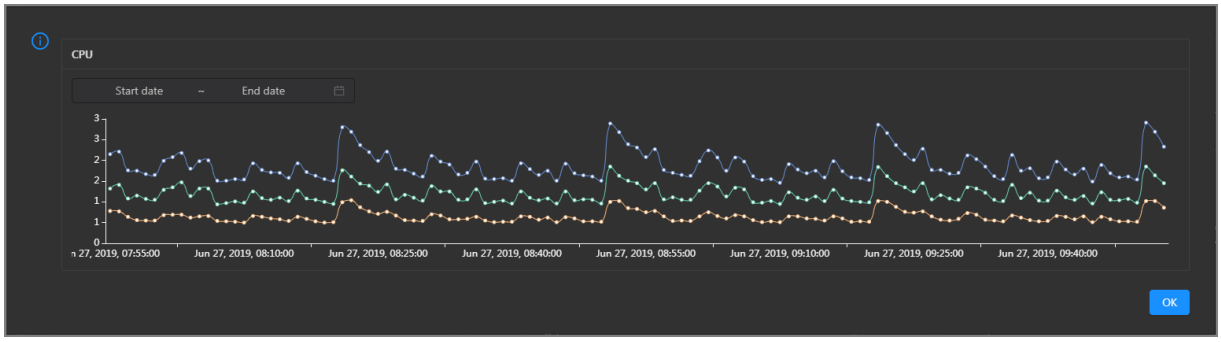

DISK

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 it.

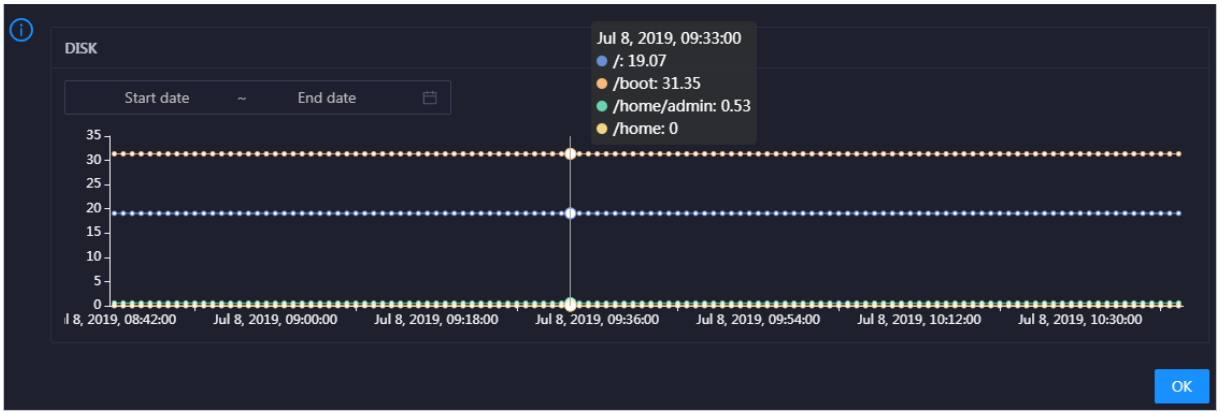

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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 it.

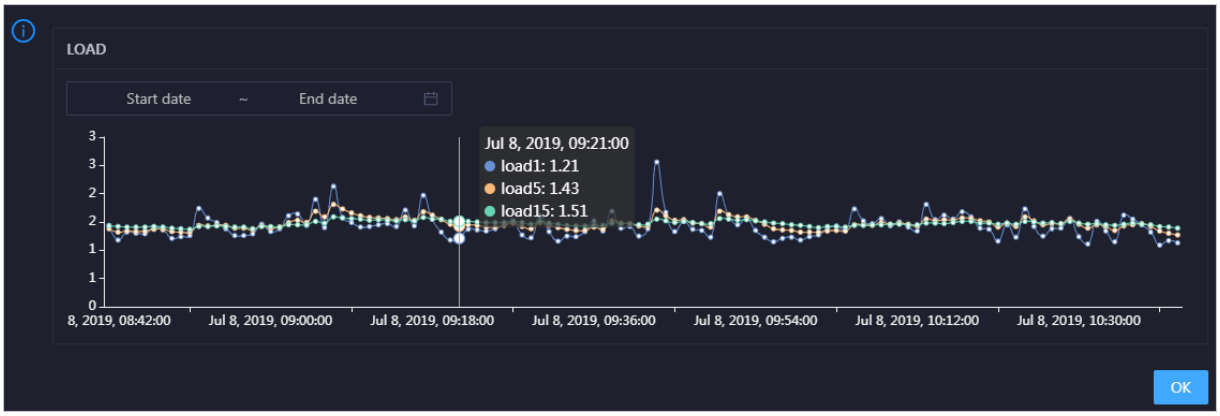

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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 it.

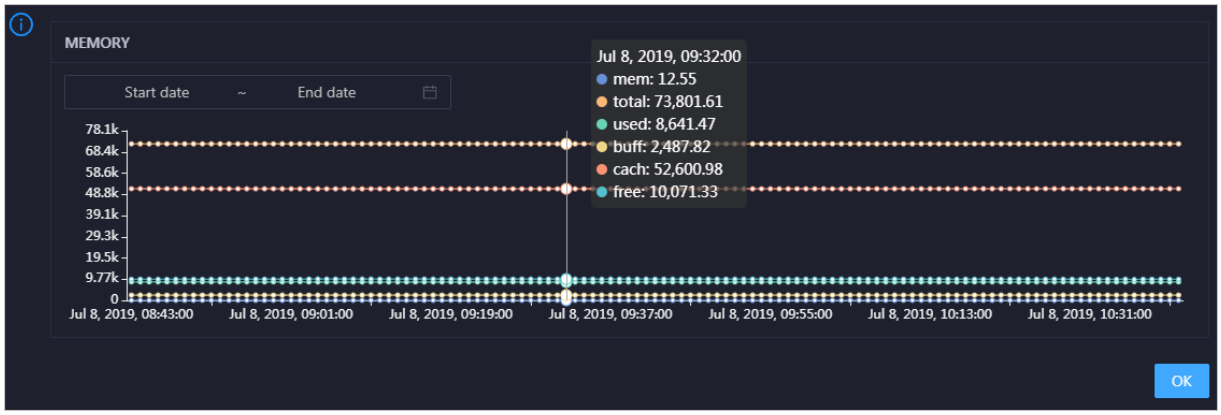

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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 it.

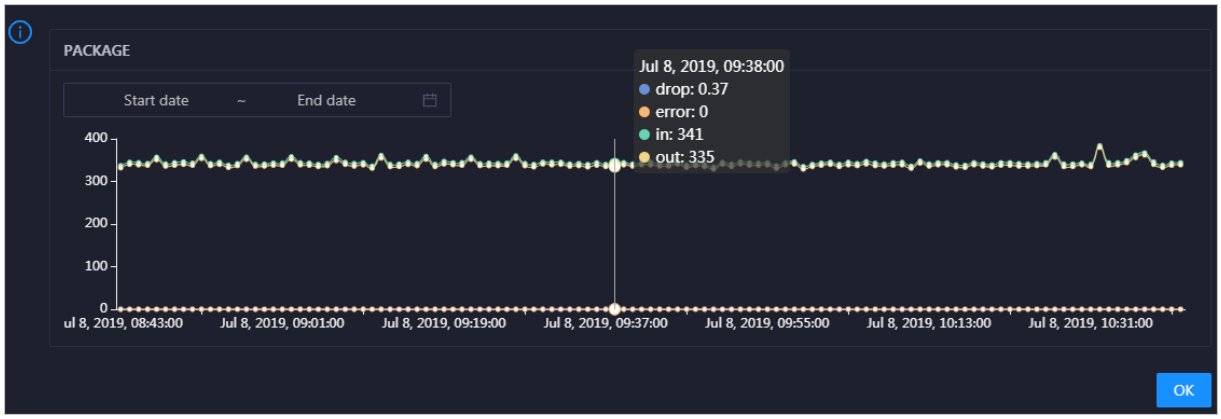

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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 it.

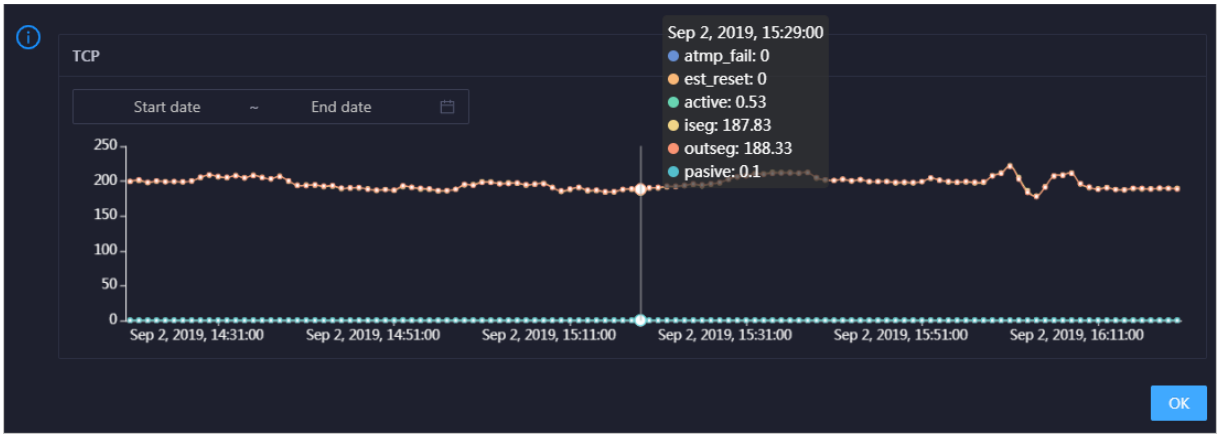

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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 it.

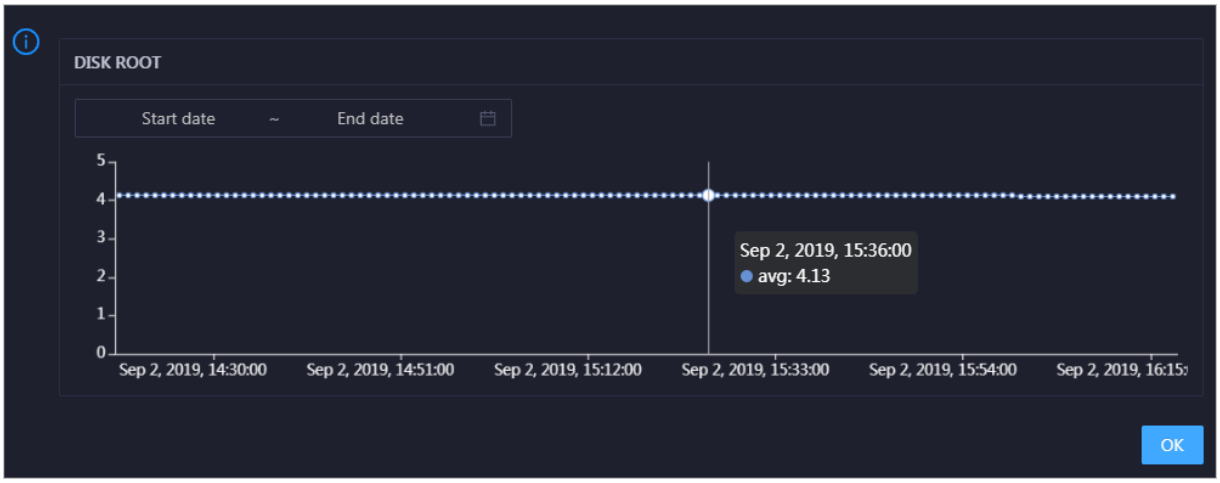

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

11.1.4.4.2. Service hosts

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

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

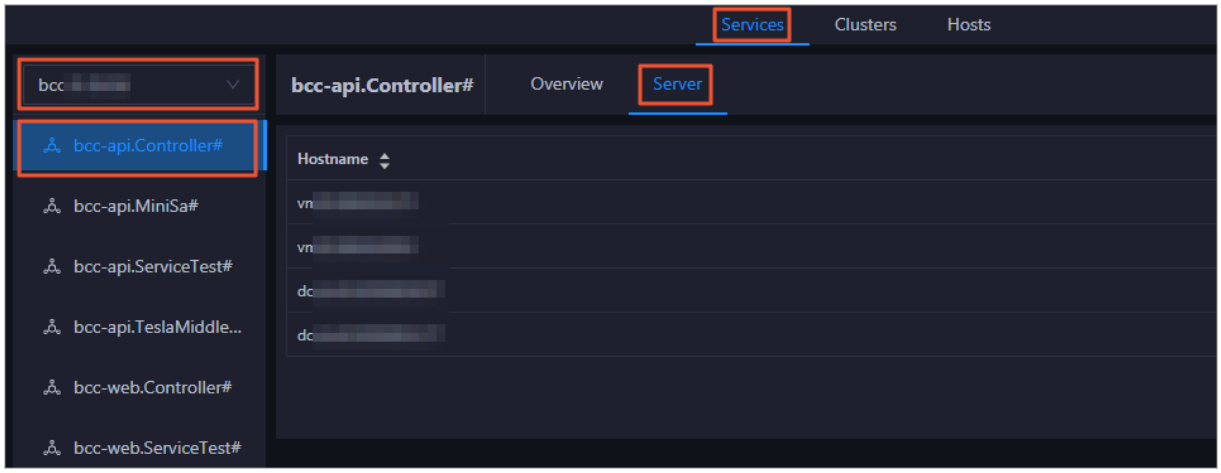

On the **Server** page, you can view the hosts where the selected service is run.

11.1.4.5. Cluster O&M

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

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

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

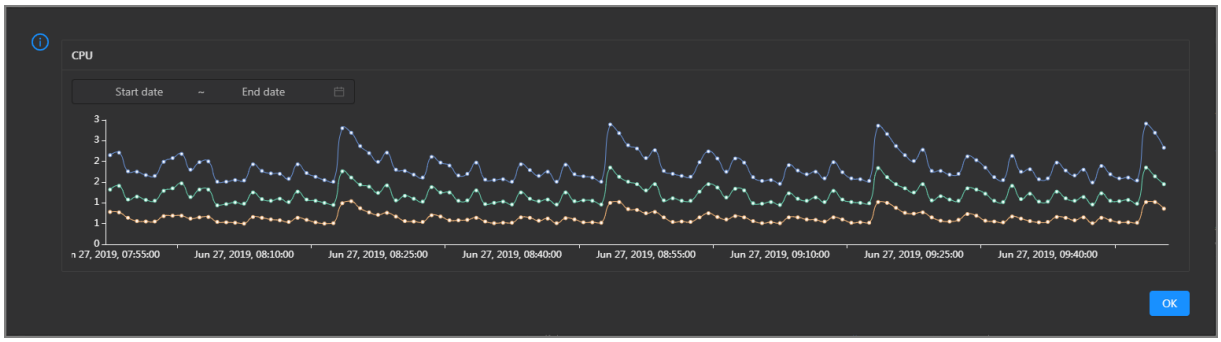

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

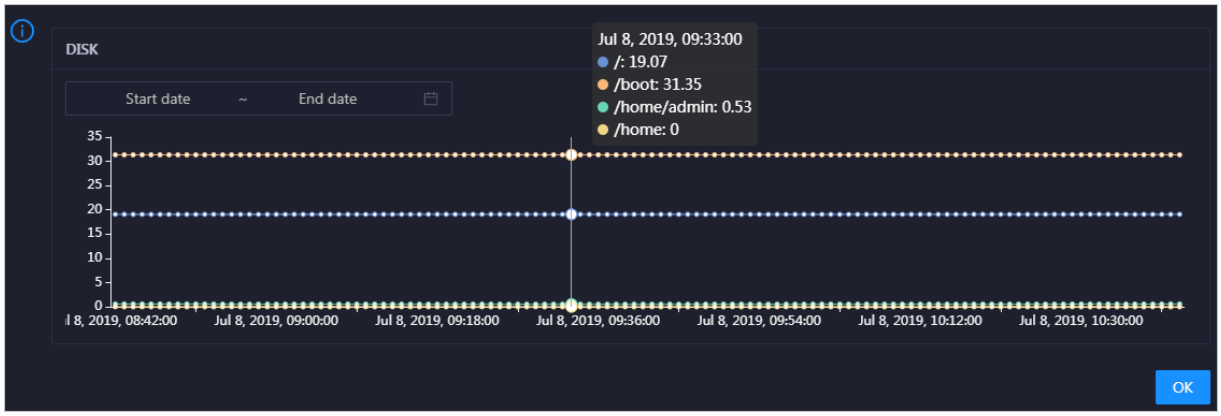

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

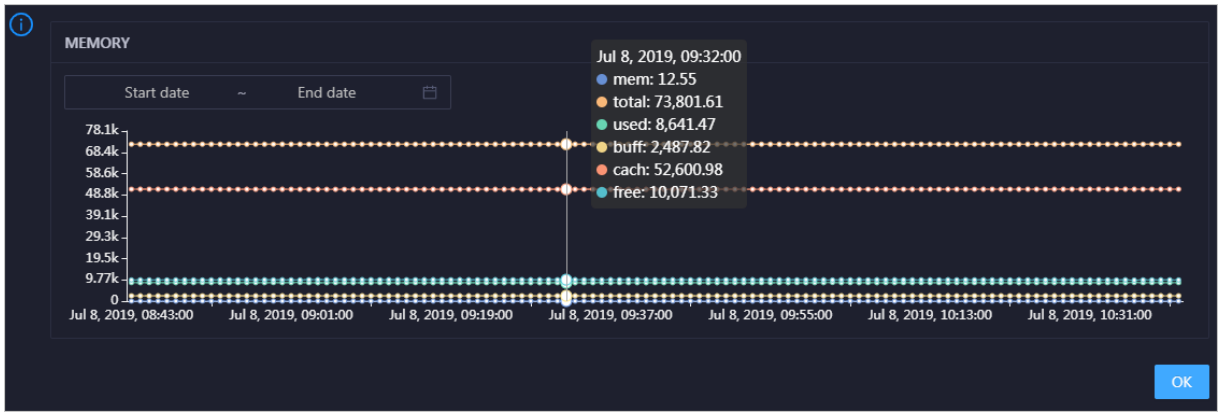

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

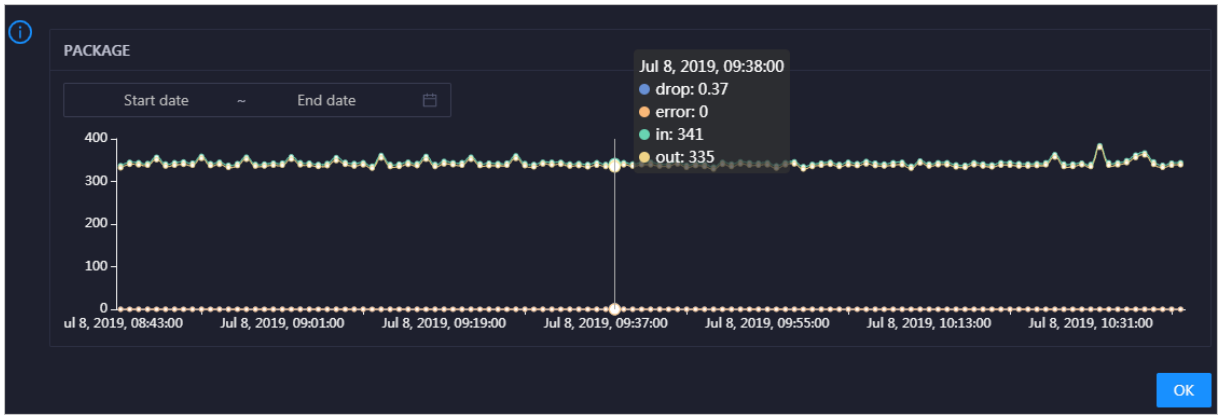

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

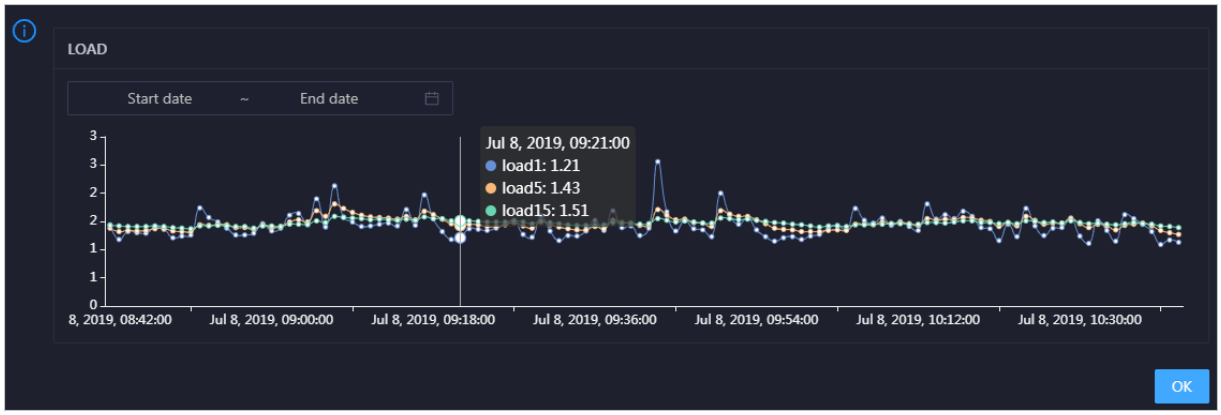

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.1.4.5.2. Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for 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.

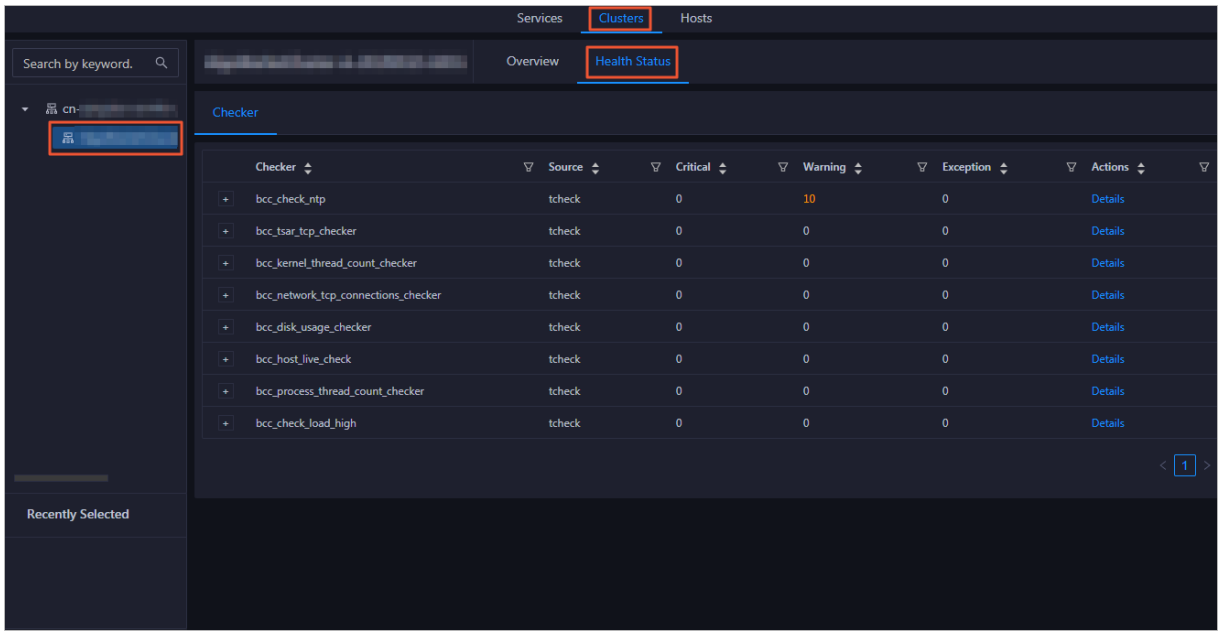

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, the **Critical**, **Warning**, and **Exception** events are alerts. You need to pay attention to them, especially the **Critical** and **Warning** events.

View checker details

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

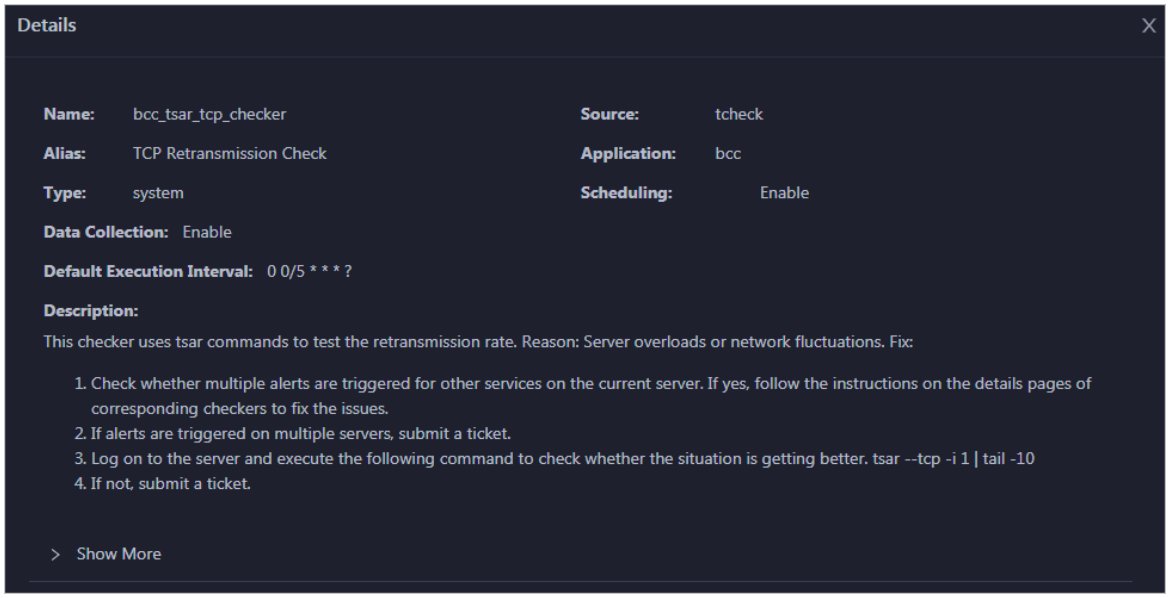

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

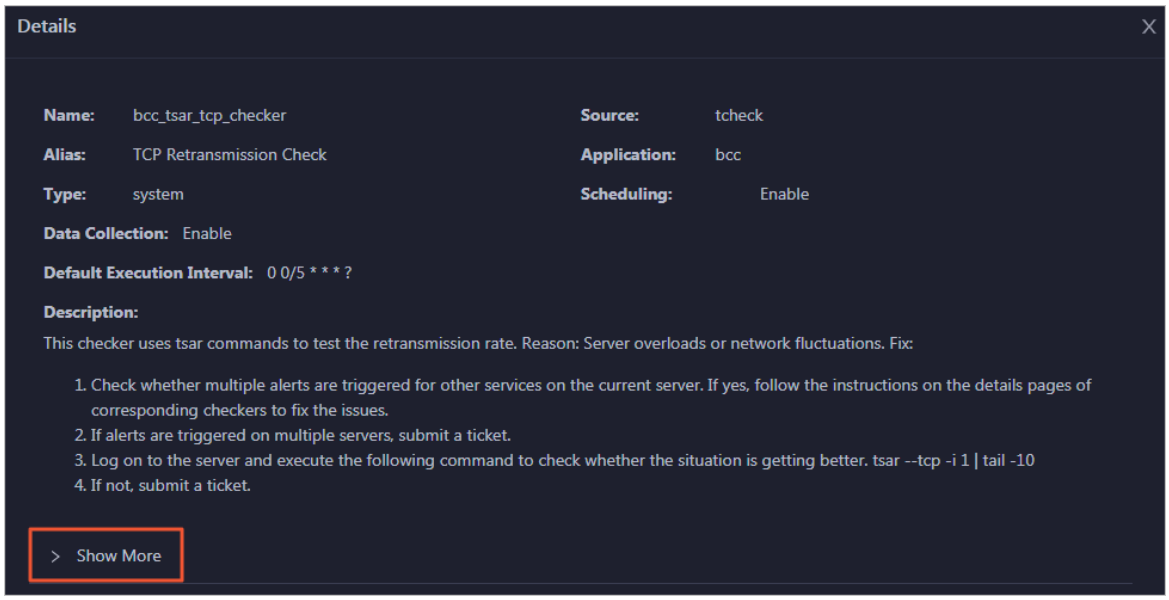

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

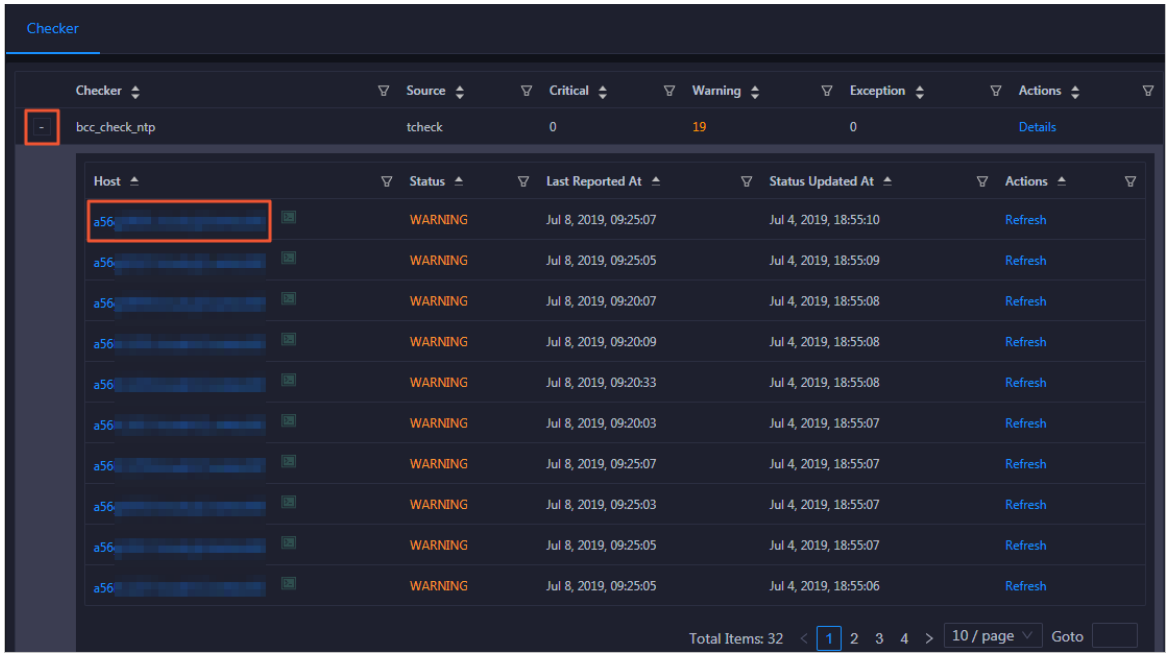

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

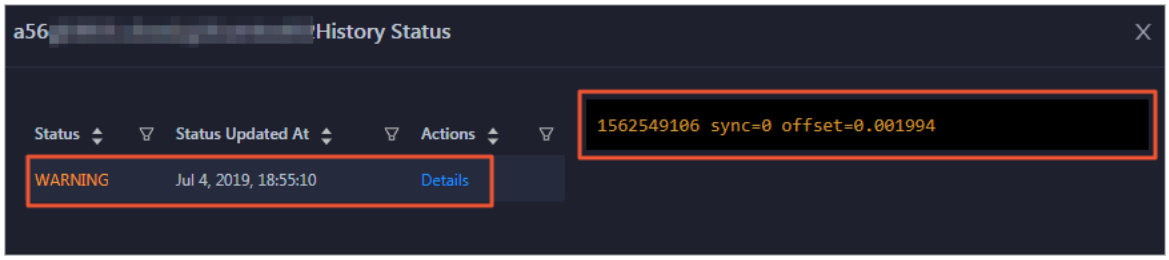

Clear alerts

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

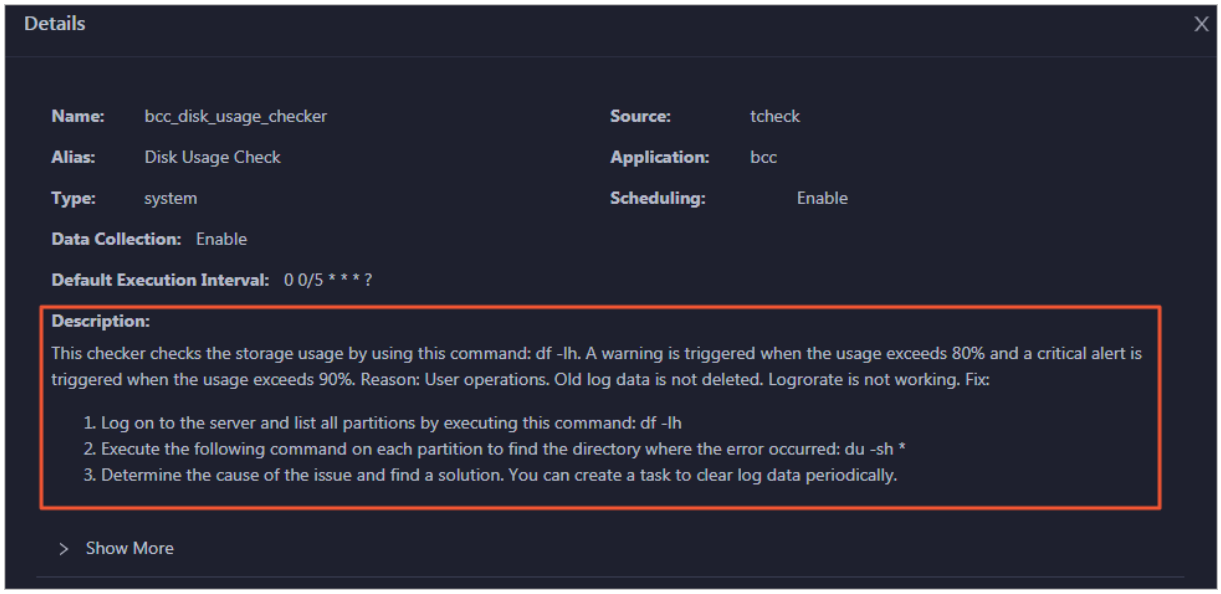

Log on to a host

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

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

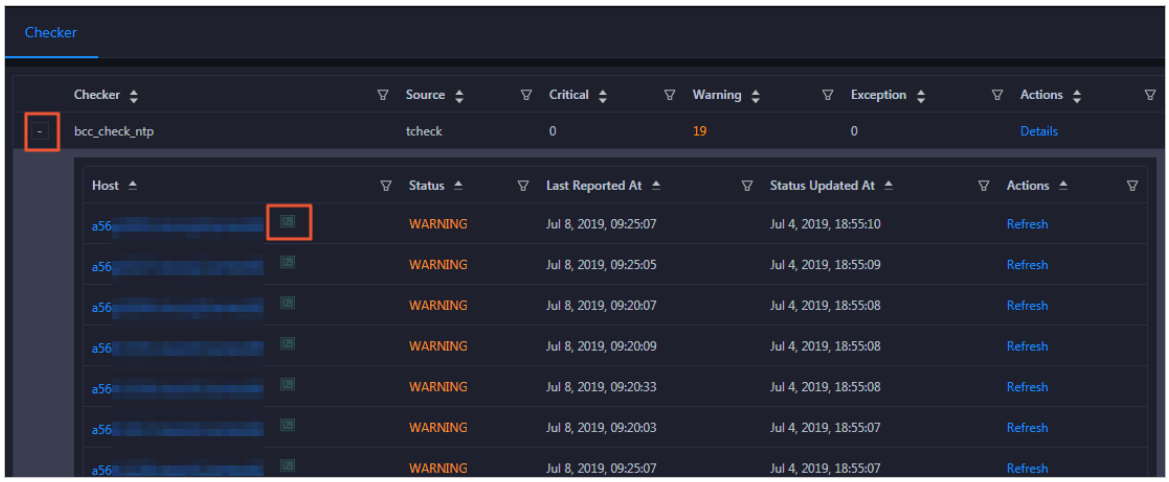

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

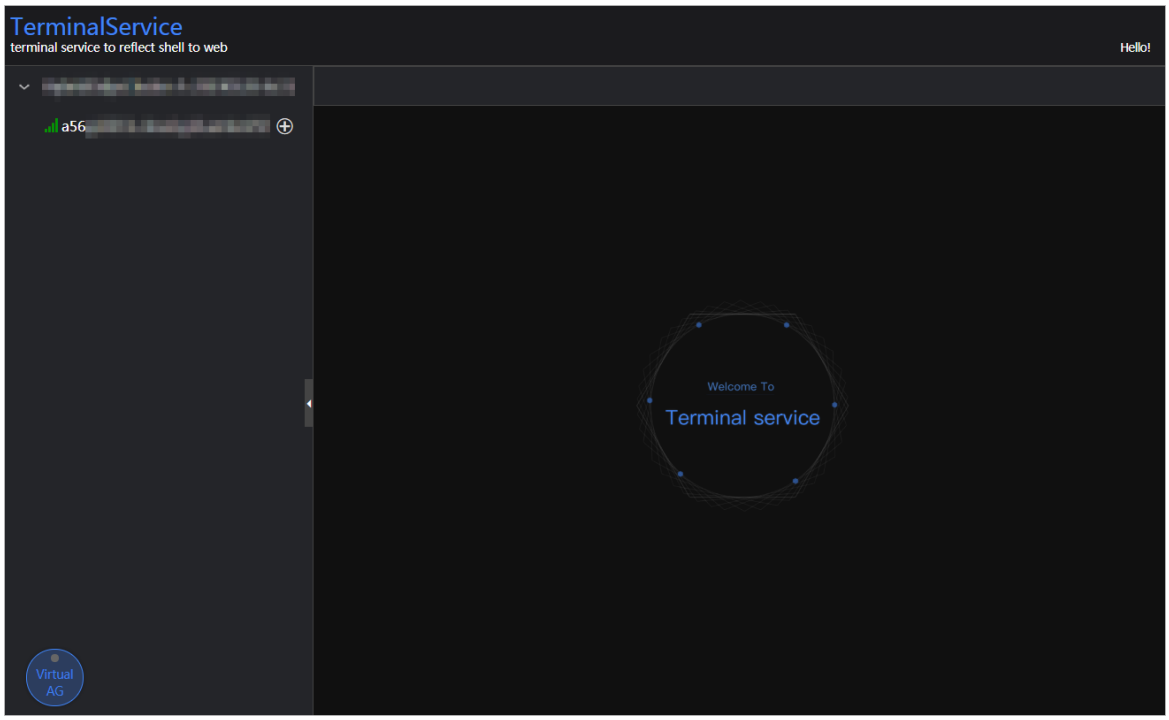

3. On the **TerminalService** page, click the hostname to log on to the host.

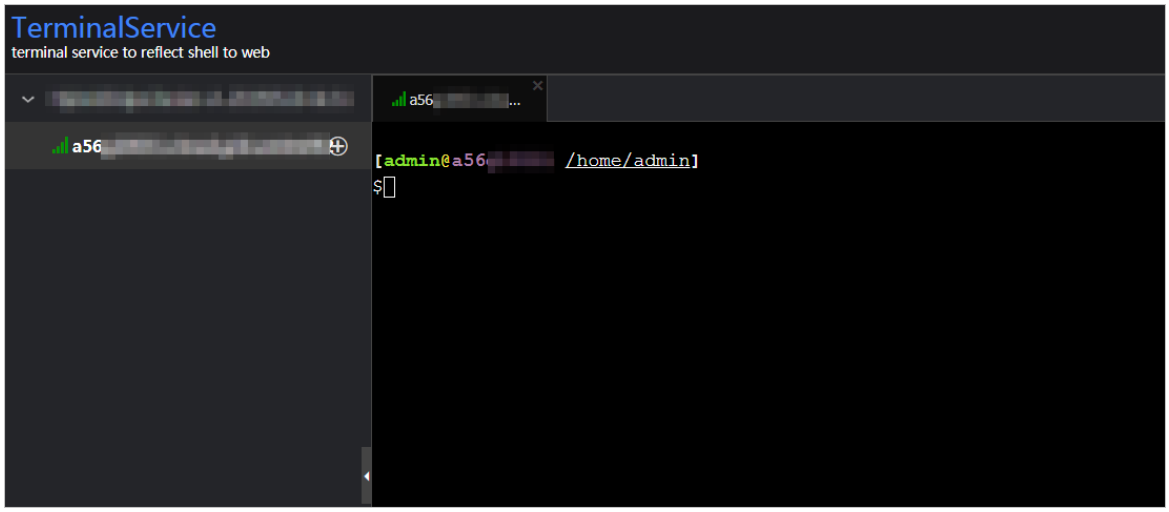

Run a checker again

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

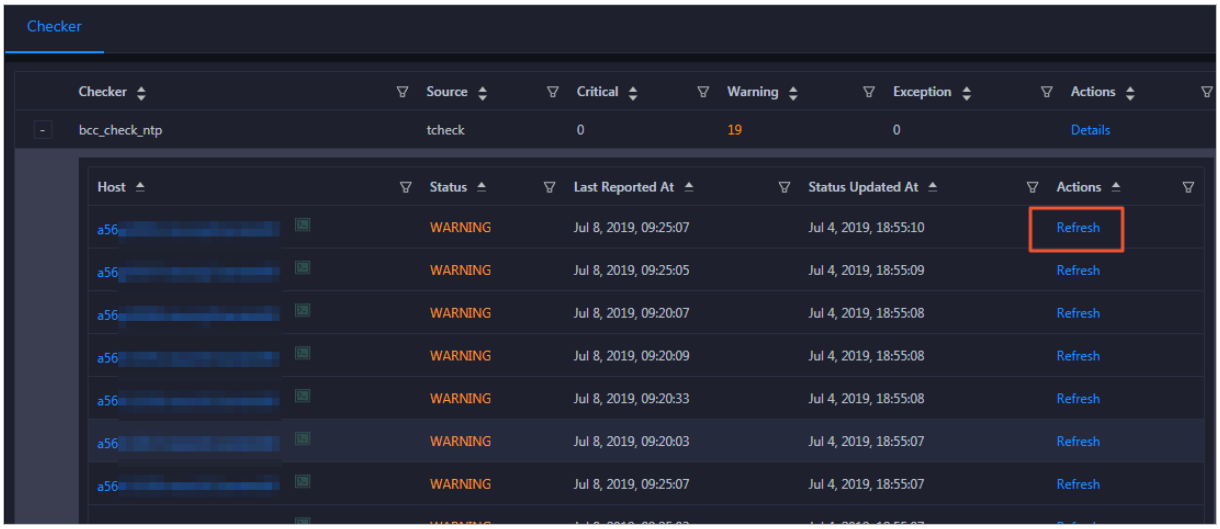

11.1.4.5.3. Restore environment settings

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

Prerequisites

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

Restore environment settings

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

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

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

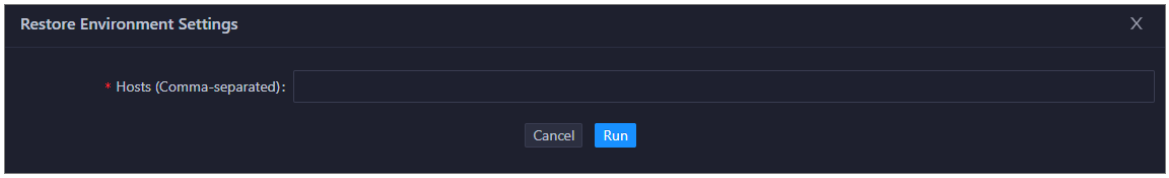

6. Click **Run**.
7. Check the execution status.

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

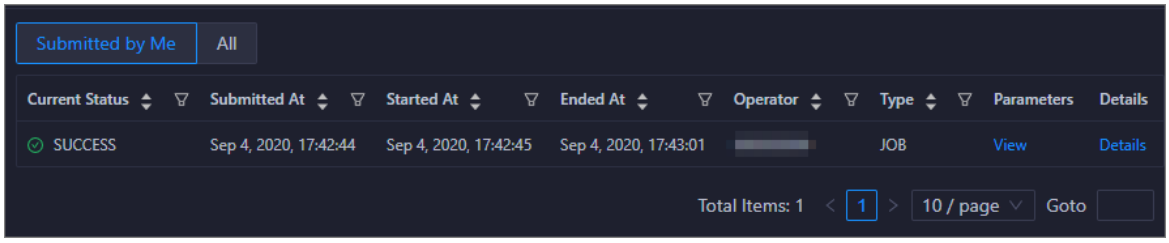

   It requires a long time to restore environment settings. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeded. **FAILED** indicates that the execution failed.

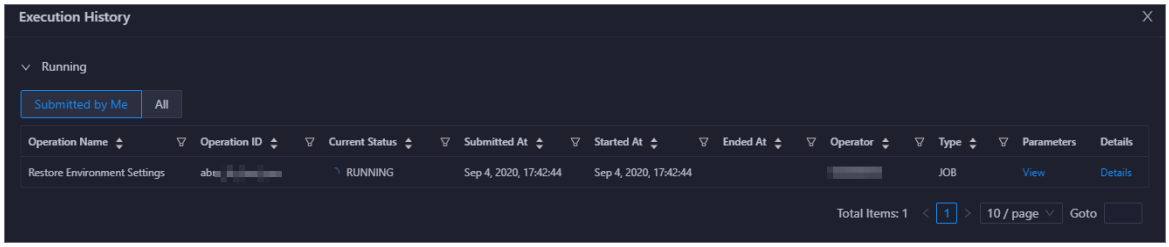

8. If the status is RUNNING, click **Details** in the Details column to view the steps and progress of restoration.

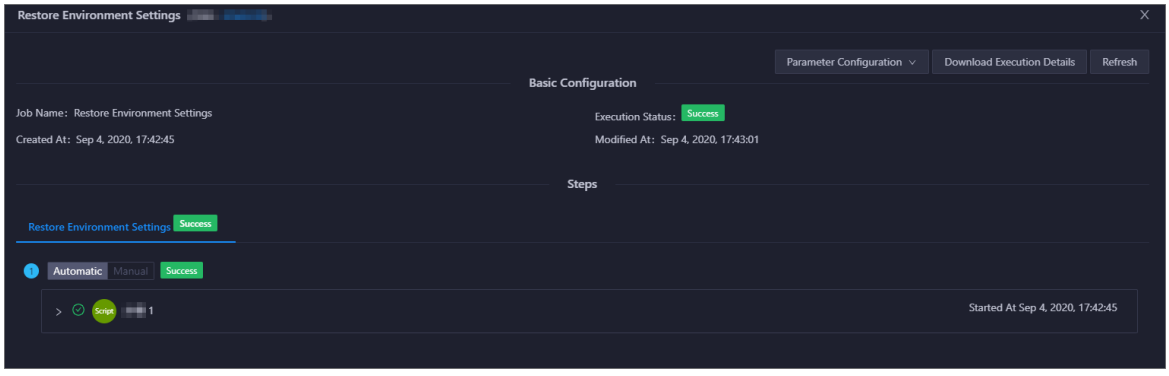

9. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of the failure to restore environment settings.

Identify the cause of the failure to restore environment settings

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

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

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

11.1.4.6. Host O&M

11.1.4.6.1. Host overview

The host overview page displays the overall running information about a host in an 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, 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

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

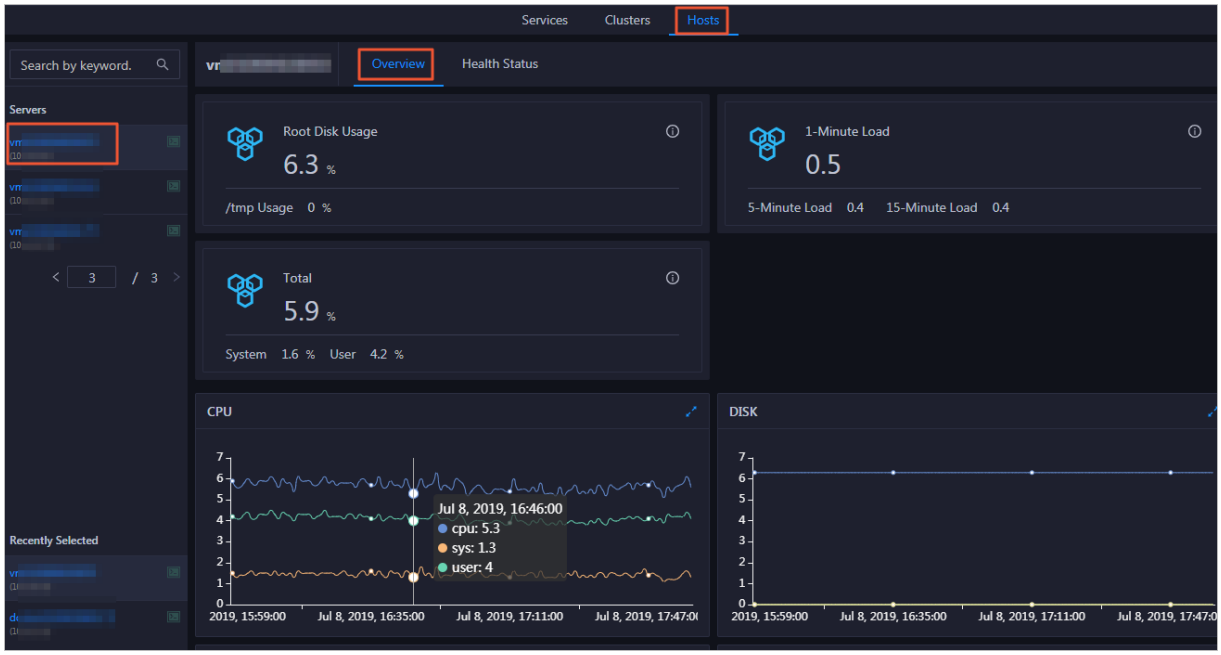

The **Overview** page consists of the following areas:

- Left-side navigation pane: displays a navigation tree of hosts.
- Recently Selected section: displays the recently selected hosts, which allows you to quickly switch between commonly used hosts.
- Right pane: displays the root disk usage, total usage, load, health check result, health check history of the host. It also displays 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.

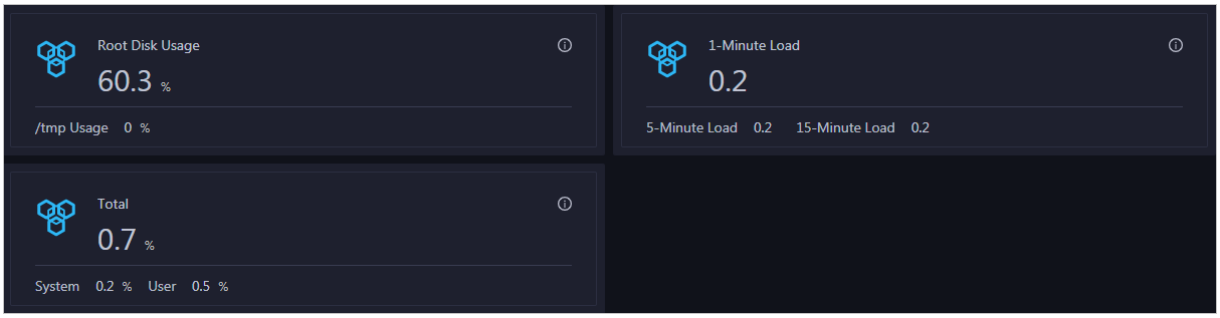

CPU

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.

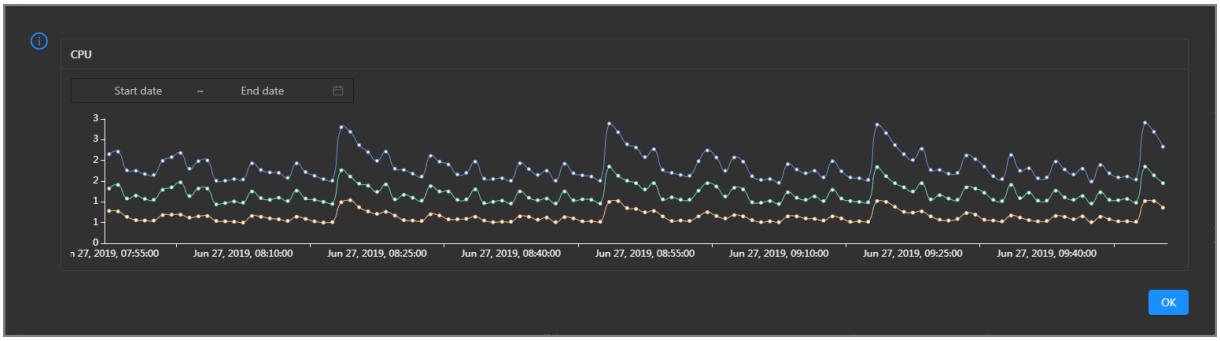

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

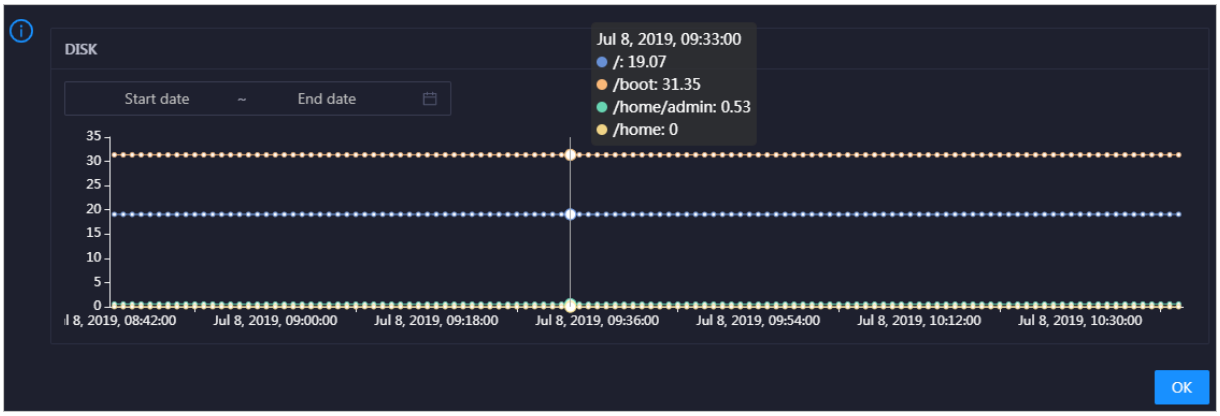

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

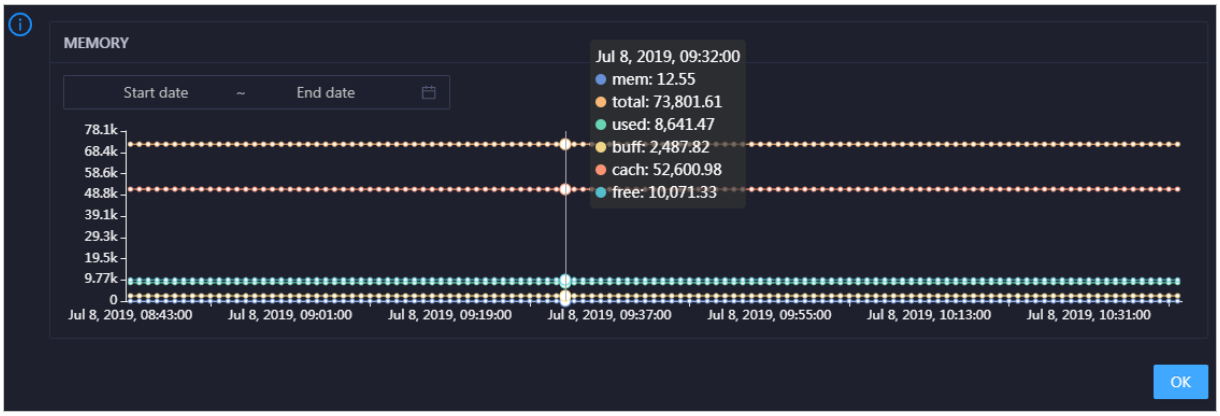

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

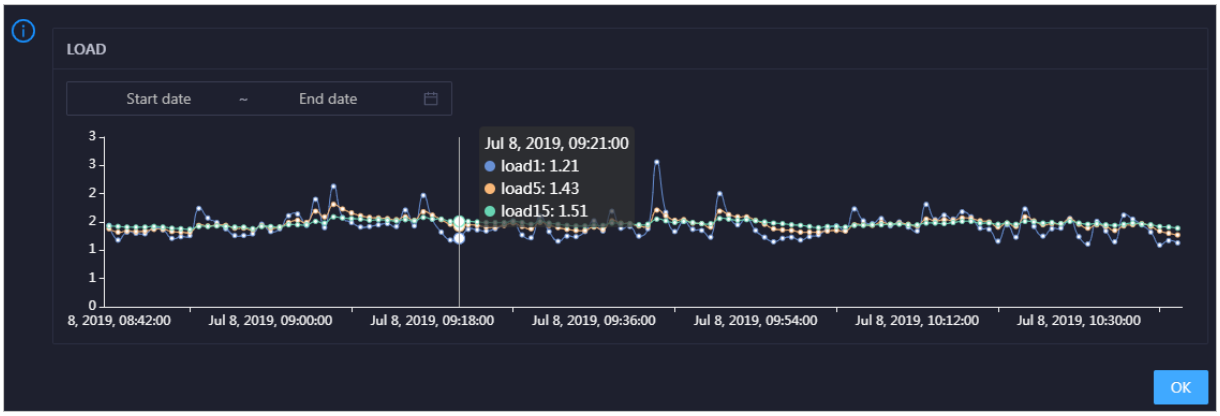

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

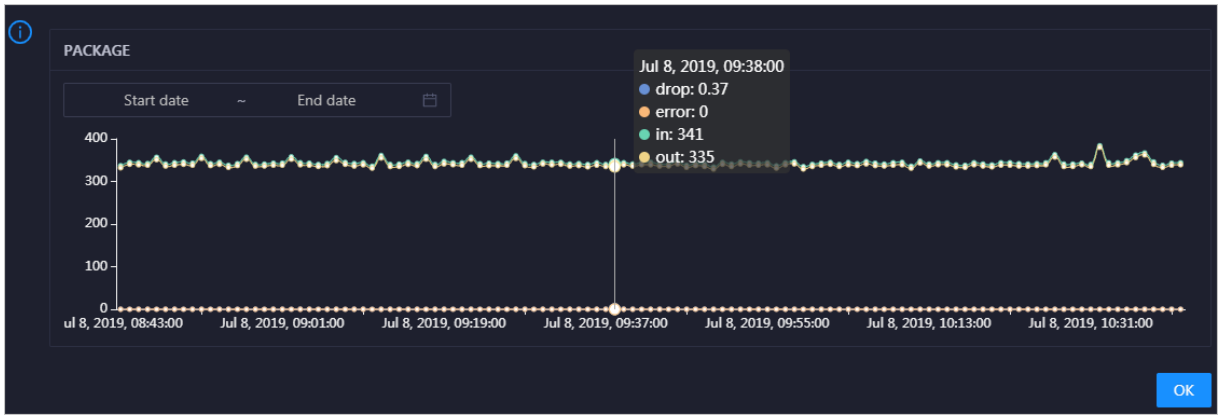

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

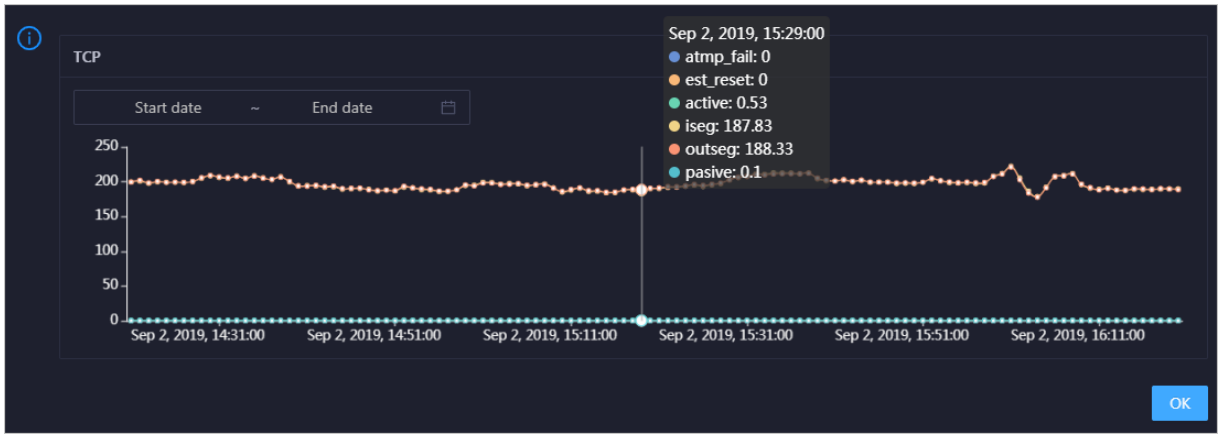

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

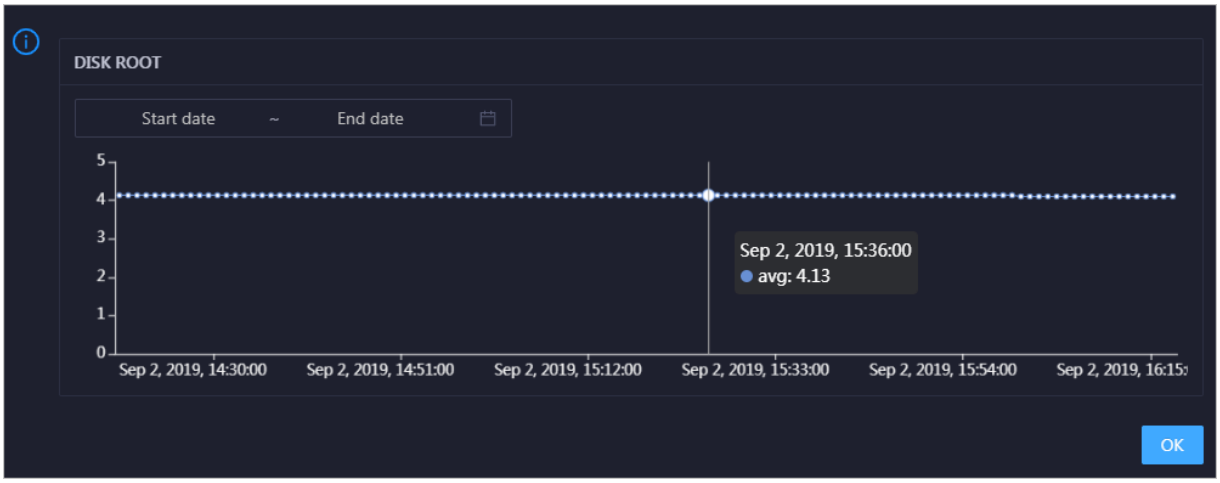

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

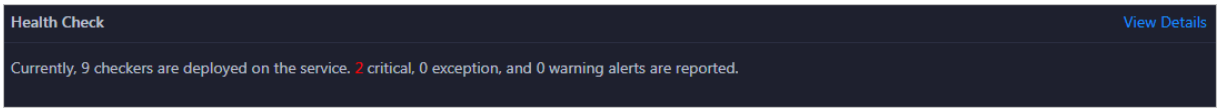

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

Health Check History

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

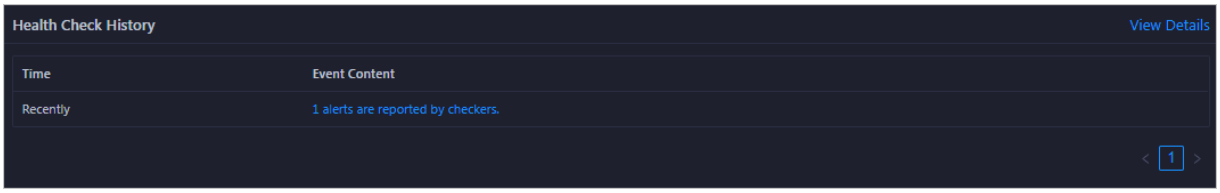

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

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

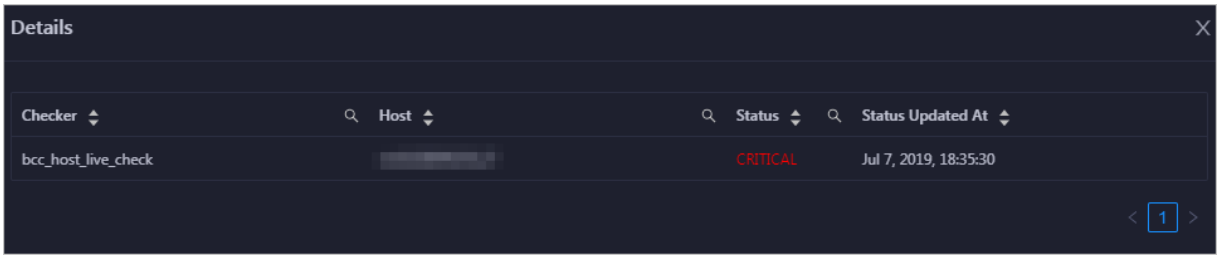

11.1.4.6.2. Host health

On the host health status page, you can view the checkers of all hosts, 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

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

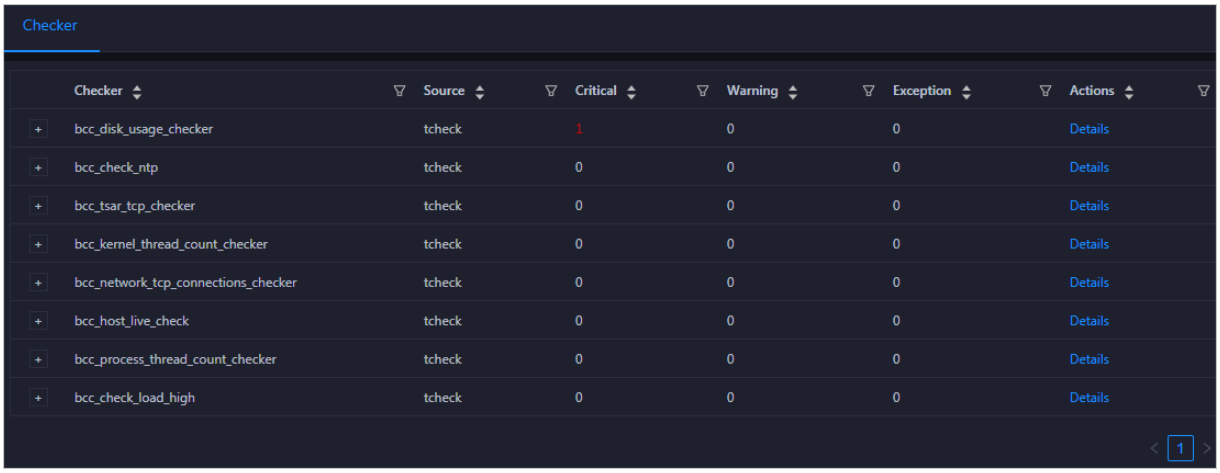

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, the **Critical**, **Warning**, and **Exception** events are alerts. You need to pay attention to them, especially the **Critical** and **Warning** events.

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.

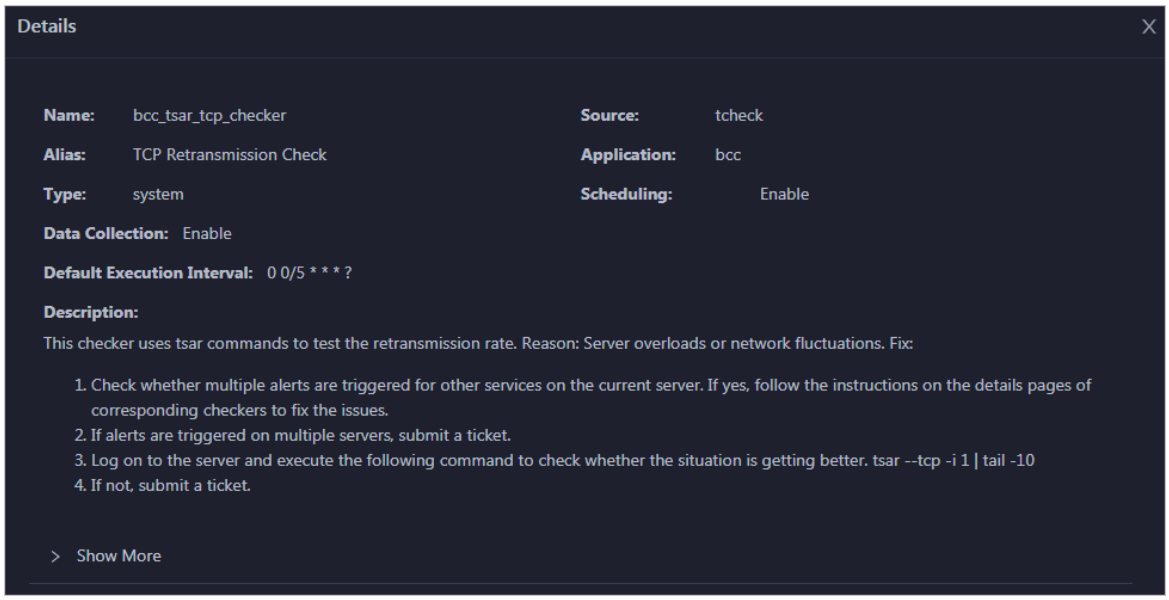

   The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

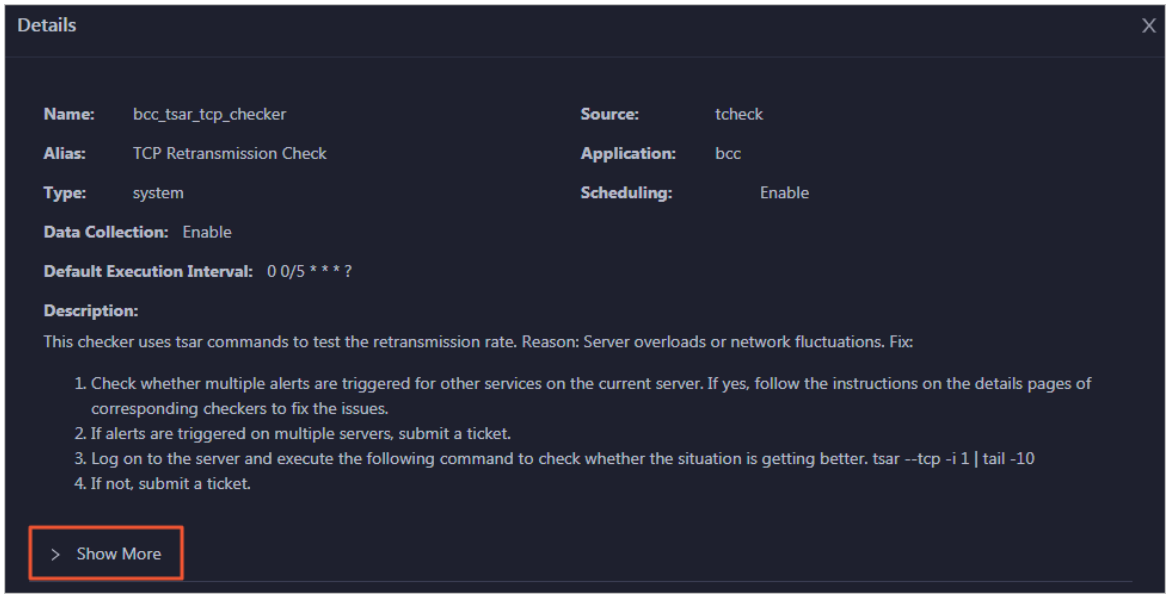

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

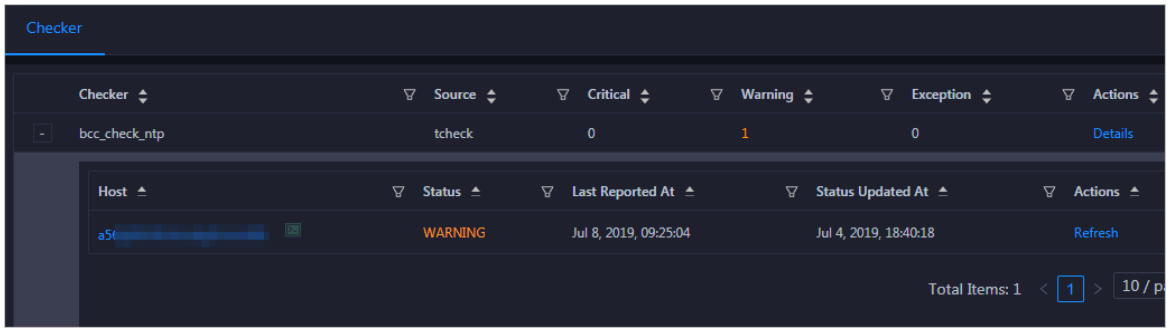

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

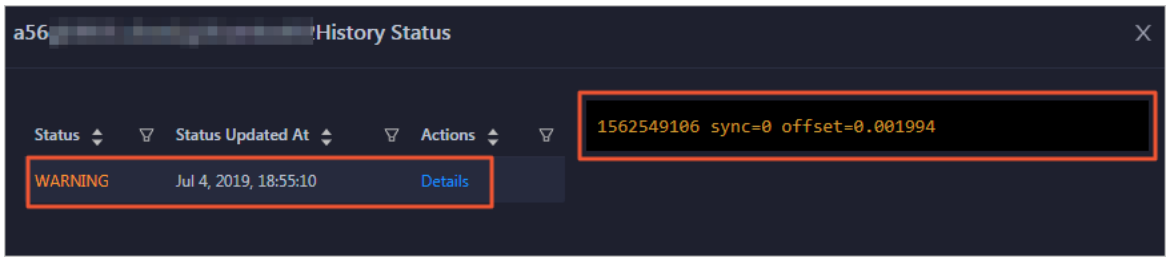

Clear alerts

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.

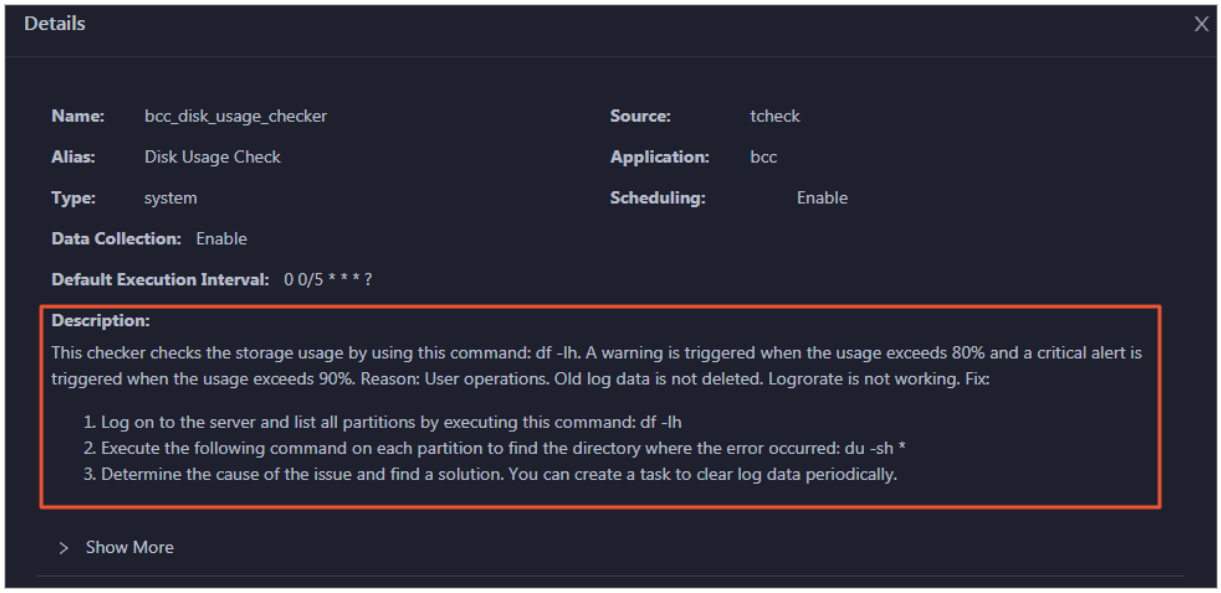

Log on to a host

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.

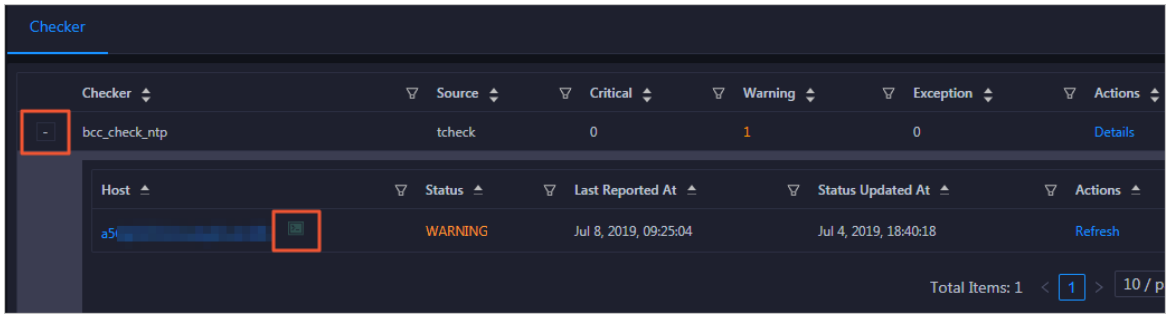

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

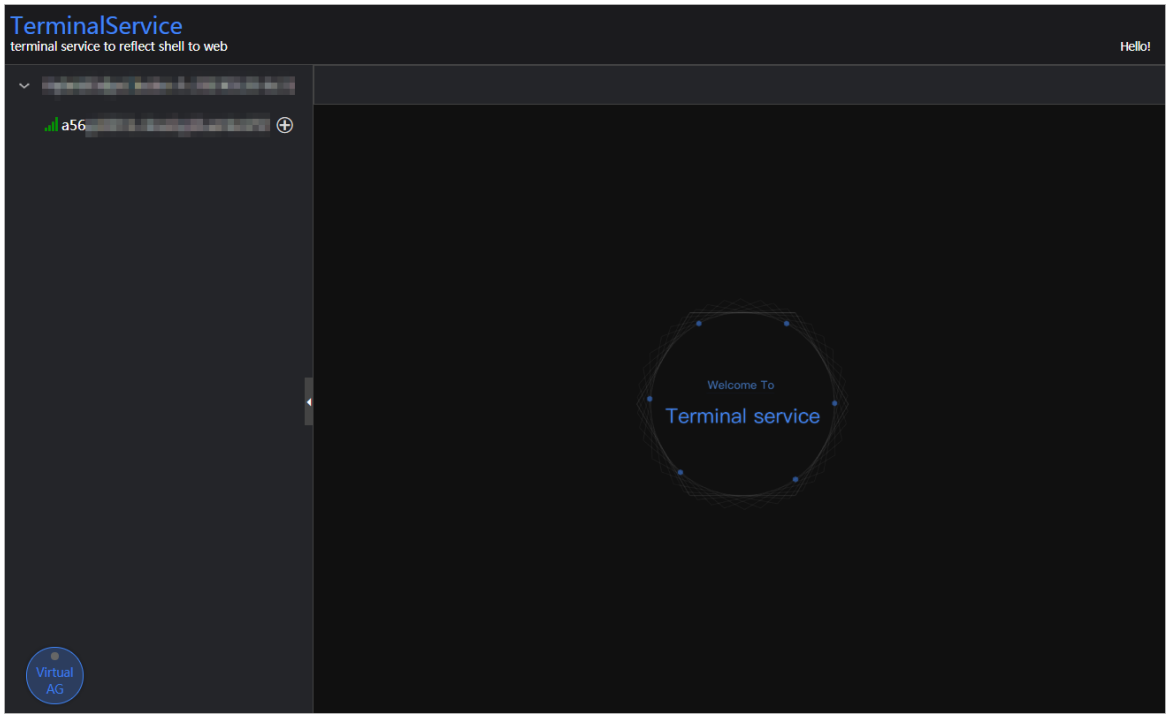

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

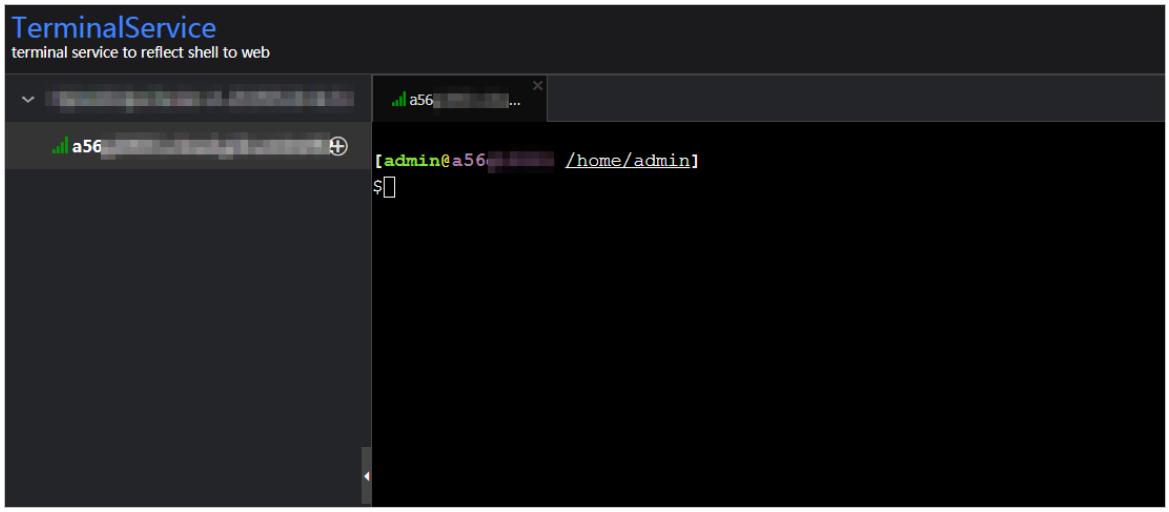

Run a checker again

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.

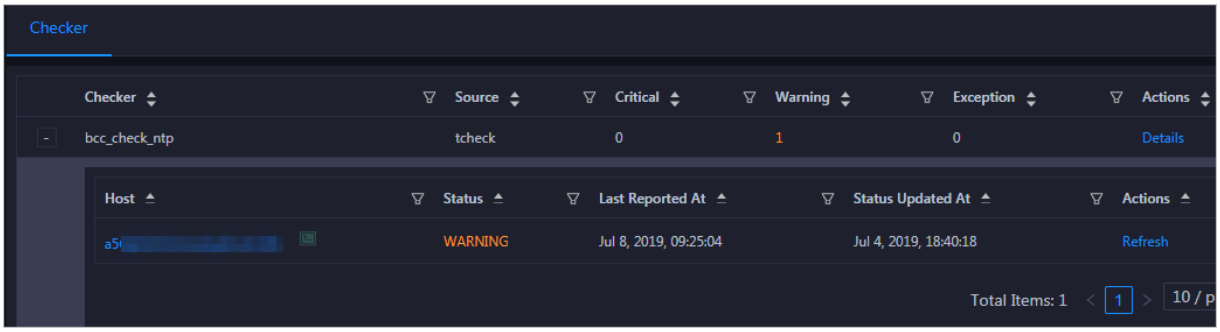

11.1.5. Management

11.1.5.1. Overview

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

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 products.
- Patch management: You can deploy upgrade patches for various products.
- 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 product hosts.
- Operation audit: You can view the records of job execution and other product O&M operations in ABM.

11.1.5.2. Jobs

11.1.5.2.1. Overview

This topic describes the job management interface and concepts related to jobs.

Apsara Big Data Manager (ABM) runs jobs to perform O&M operations on big data products. Jobs, also known as product O&M tasks, are O&M operations performed on physical devices in the cluster. The job management interface consists of two pages: **Job Execution** and **Job Management**.

Concepts

Concepts related to jobs include:

- Ordinary job: jobs that can only be manually run.
- Cron job: jobs that are automatically run based on timer settings.
- Scheme: job templates provided by ABM. You can use schemes to generate jobs.
- Atom: step templates provided by ABM. You can use atoms as steps when generating jobs.
- Ordinary step: steps 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: steps 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

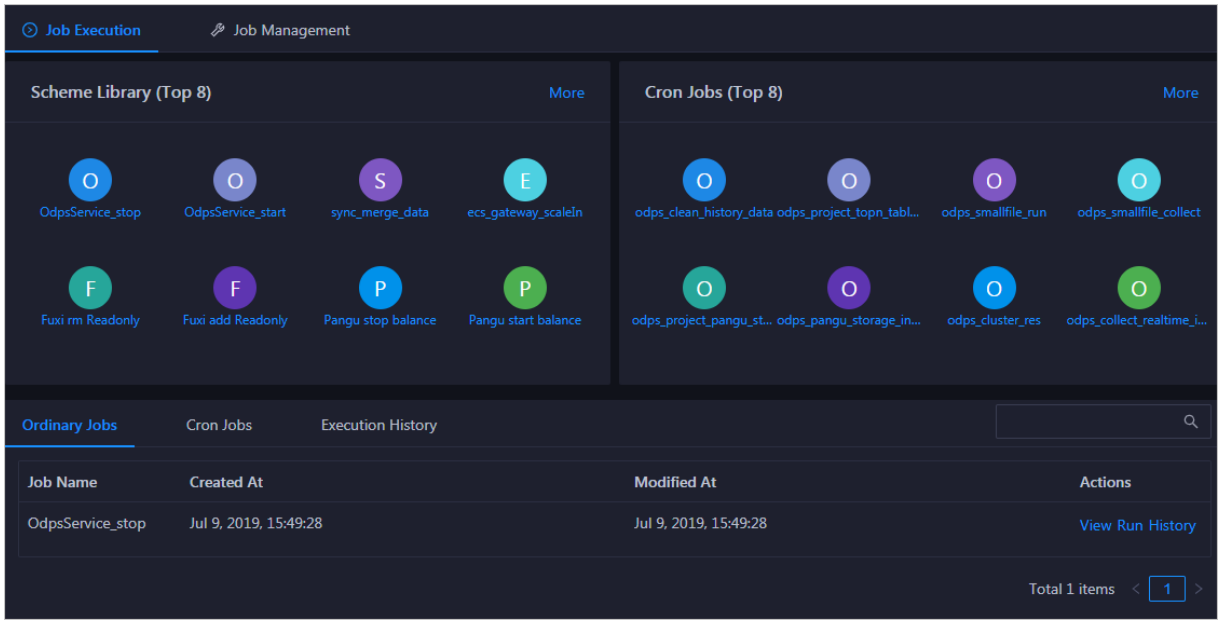

The **Job Execution** page contains the following modules:

- 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, and 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

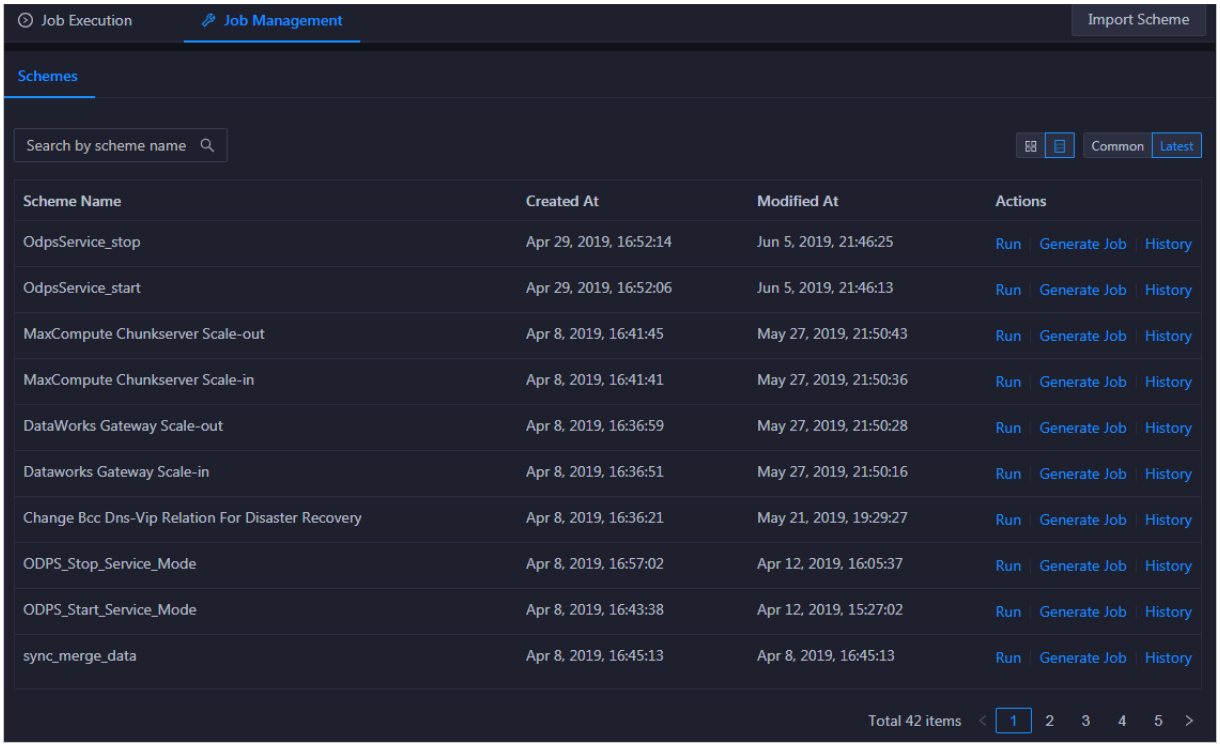

The **Job Management** page provides the following features:

- 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.

11.1.5.2.2. Jobs

11.1.5.2.2.1. 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 Job parameters.

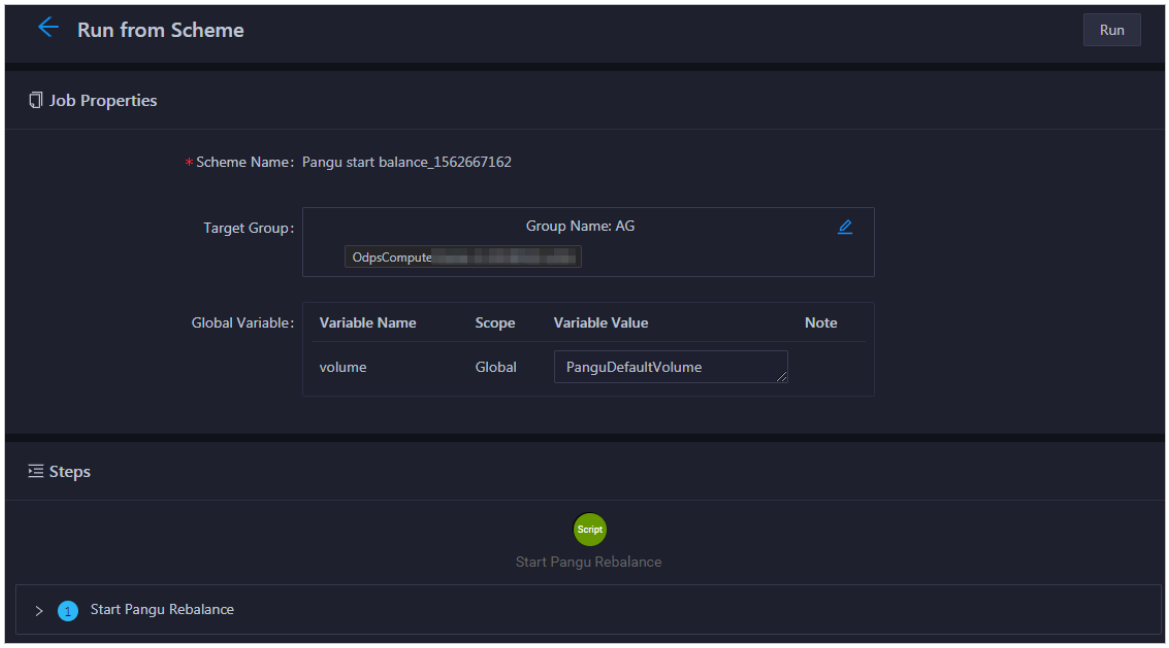

Job parameters

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**.

Parameter	Description
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**.

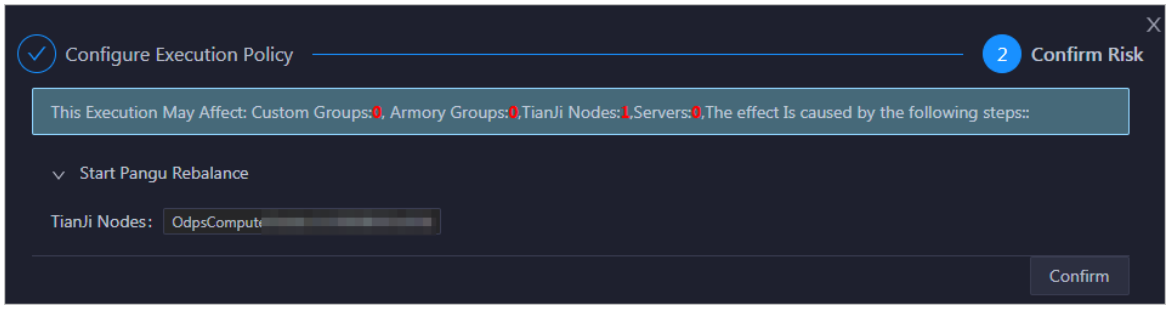

After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see View the execution history.

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

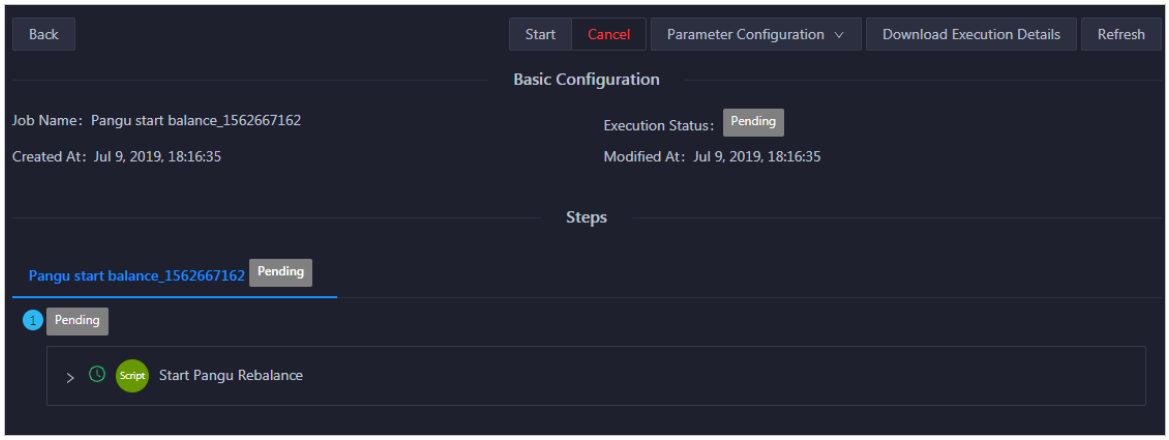

If you do not perform any operation and exit the job execution page, you can find a job record on the **Execution History** page. Click **View** to go to the job execution page again.

11.1.5.2.2.2. 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 Big Data 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 Job properties.

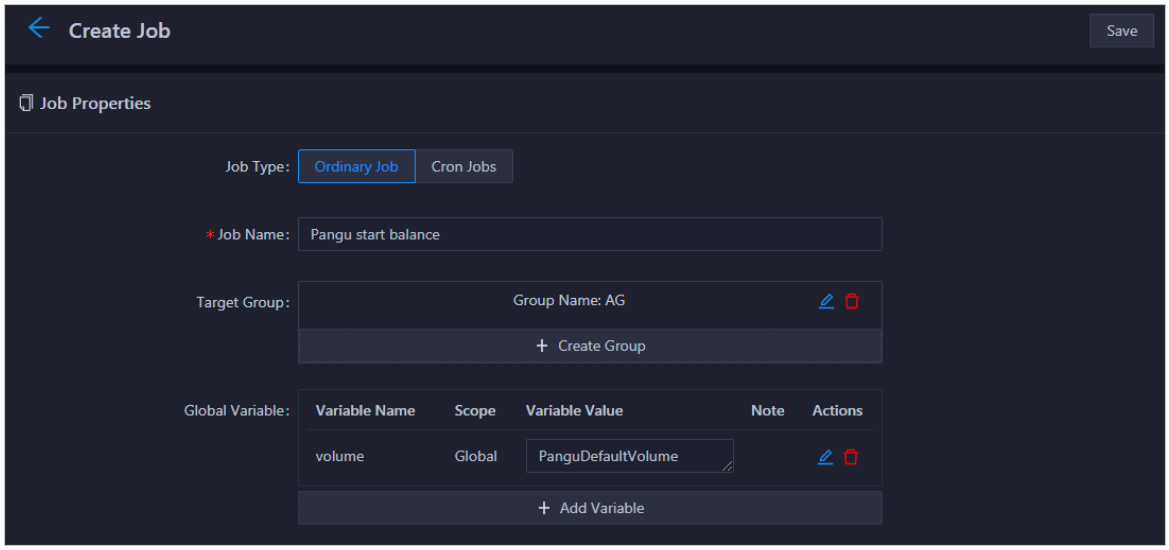

Job properties

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.
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**.

Parameter	Description
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.

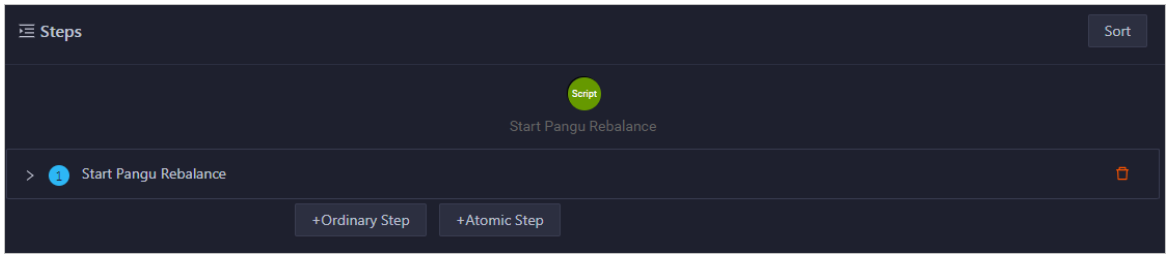

The steps include ordinary steps and atomic steps.

- **+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. Parameters of command execution steps, Parameters of script execution steps, Parameters of file push steps, Parameters of API call steps, and Parameters of manual steps describe the parameters for different types of ordinary steps.

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.
	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.

Section	Parameter	Description
Advanced Configuration	**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.

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.
	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.

Section	Parameter	Description
Advanced Configuration	**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.

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.
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.

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.

Parameter	Description
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**.
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.

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.

11.1.5.2.2.3. 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**.

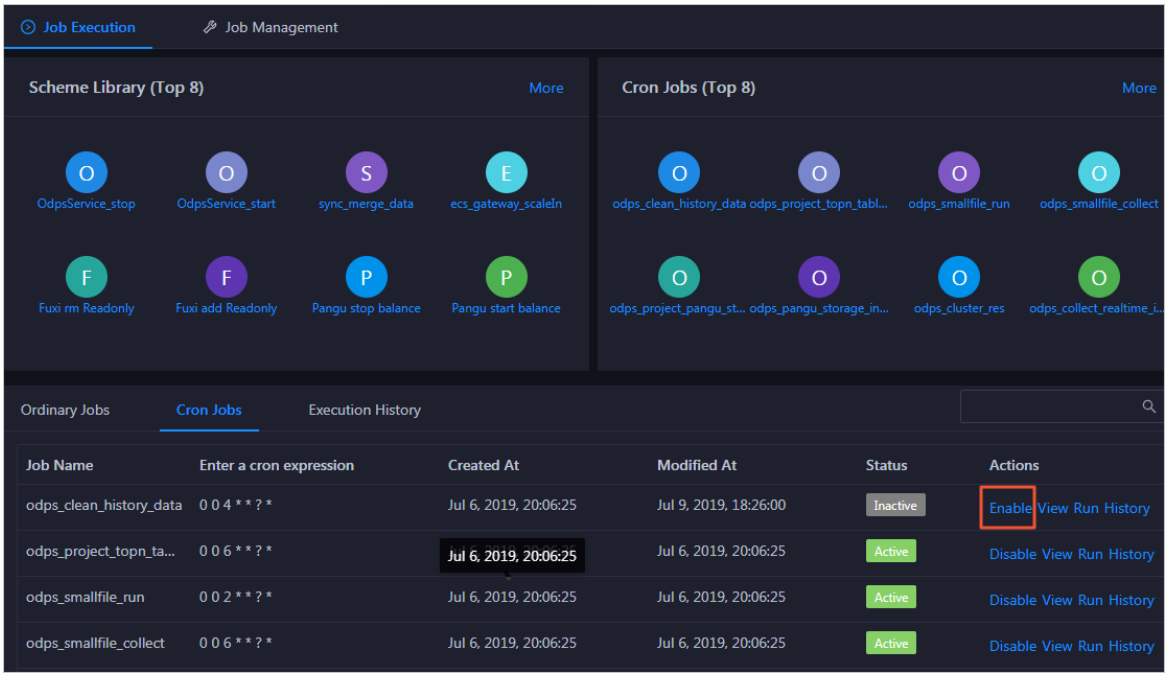

4. On the **Cron Jobs** page, you can enable or disable a cron job.
   - To enable a cron job in the inactive status, click **Enable** in the Actions column of the cron job.

     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**.

11.1.5.2.2.4. 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.

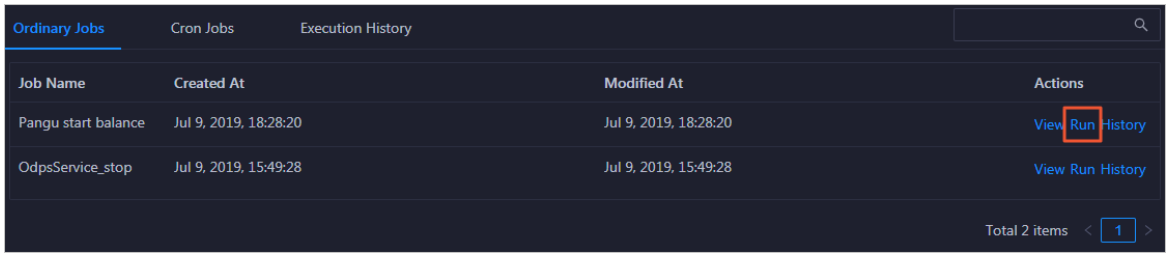

4. In the **Ordinary Jobs** list, click **Run** in the Actions column of a job.
5. Confirm the job risks in the dialog box that appears, and click **Confirm**.

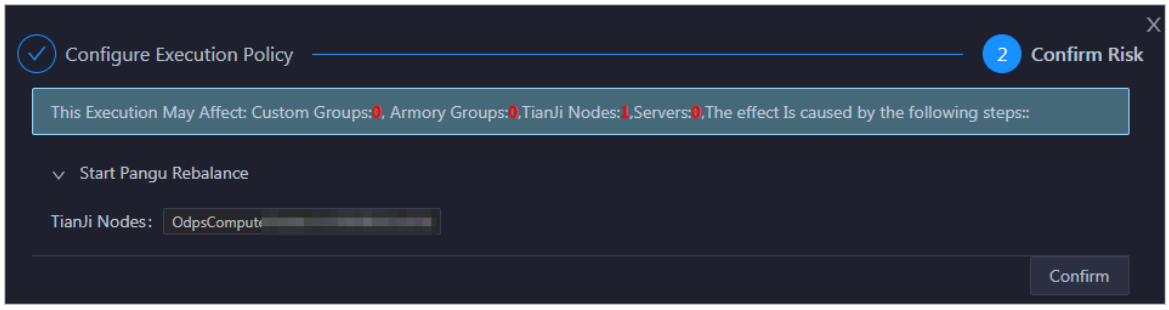

   After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see View the execution history.

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

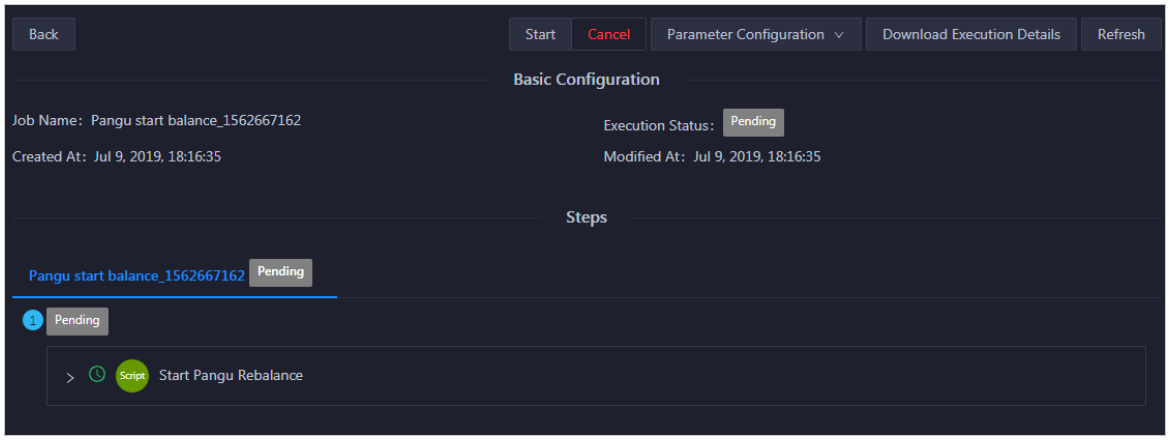

You can find the record about a job on the **Execution History** page, and click **View** to go to the detailed execution page.

11.1.5.2.2.5. 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.

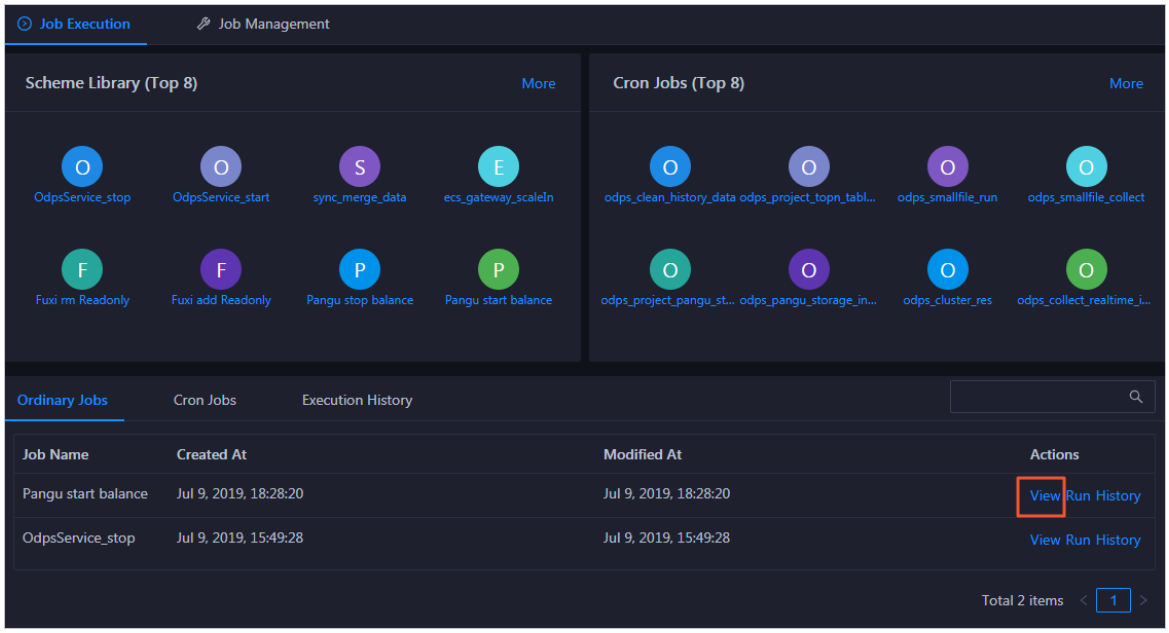

11.1.5.2.2.6. View the execution history of a job

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

Prerequisites

An ABM administrator account is obtained.

Context

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

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

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

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

Procedure

1. Log on to the ABM console.
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. Click the **Ordinary Jobs** tab on the **Job Execution** page.
4. On the **Ordinary Jobs** page, click **History** in the Actions column of an ordinary job. The **Execution History** page appears.You can view the execution history of this job on the **Execution History** page. For more information, see View the execution history.

11.1.5.2.3. Schemes

11.1.5.2.3.1. 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 Big Data 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.

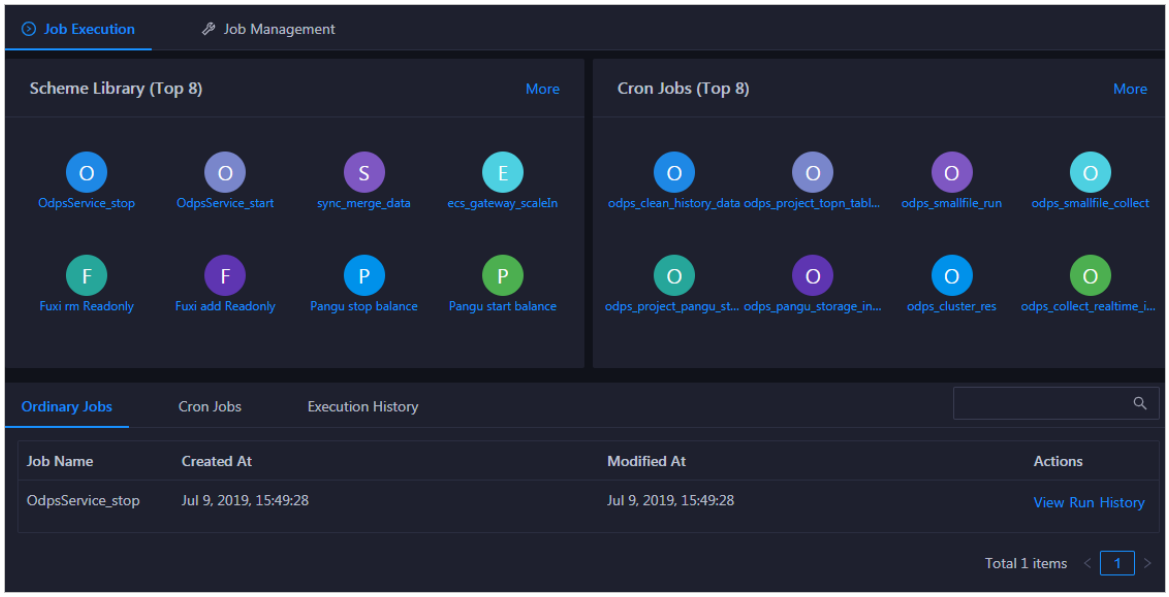

5. On the **Job Details** page, click **Save as Scheme** in the upper-right corner. The system prompts that you have saved the scheme.

Result

The new scheme has the same name as the job from which it was created and is listed on the **Schemes** page.

11.1.5.2.3.2. View schemes

A scheme is displayed in the scheme list after it is created. Apsara Big Data 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**.

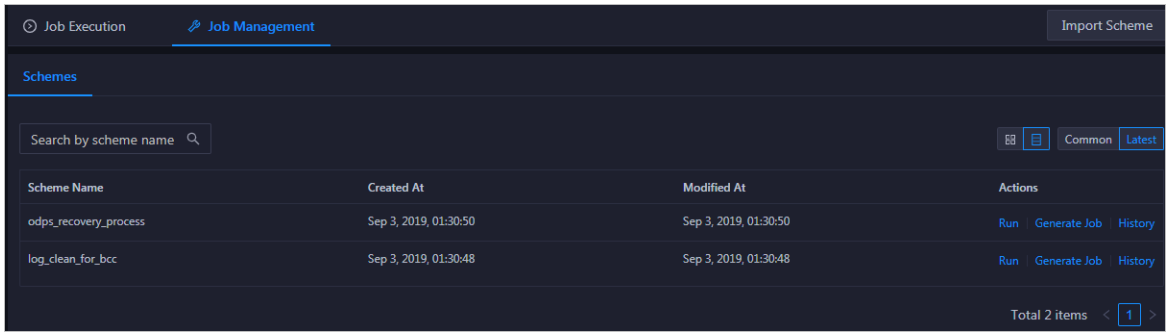

4. If there are too many schemes, you can enter the scheme name in the search bar to search for the required scheme.
5. Change the method for viewing schemes
   - View schemes in list (default): Click in the upper-right corner.
   - View schemes in cards: Click in the upper-right corner.

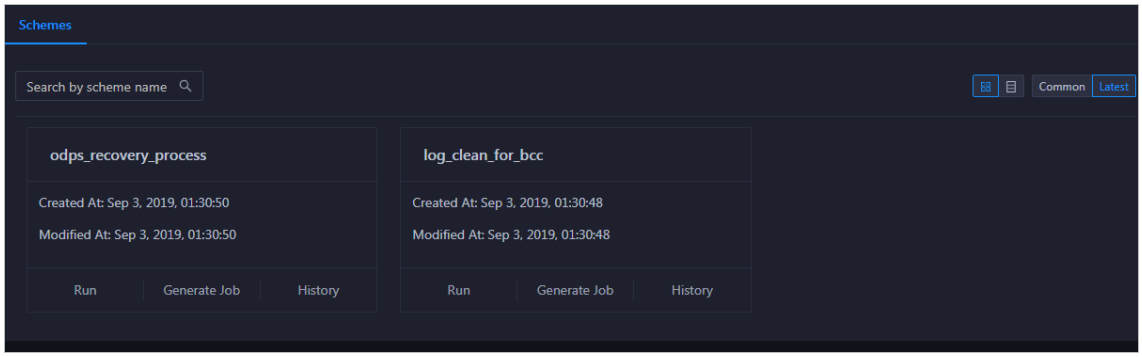

11.1.5.2.3.3. View the execution history of a scheme

Apsara Big Data 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.

11.1.5.2.4. View the execution history

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

Prerequisites

An ABM administrator account is obtained.

Context

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

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

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

Procedure

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

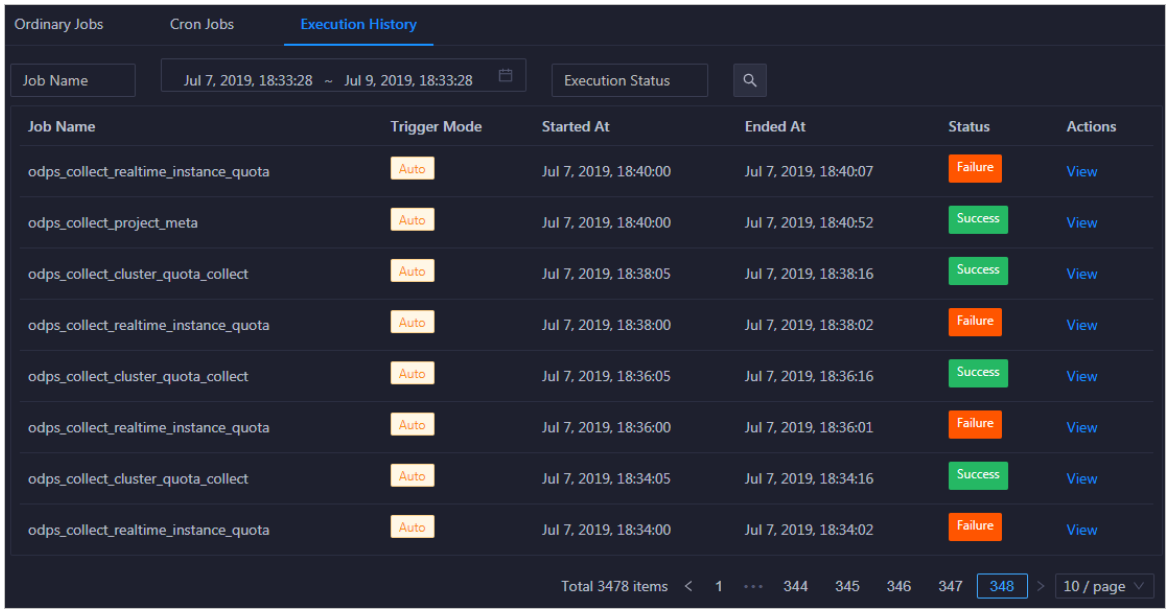

4. If there are too many execution records, filter them by a combination of one or more of the following filter conditions: job name, creator, execution status, and time range. Then, click to search for required records.
5. Click **View** in the Actions column of a record to view the execution details.

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

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

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

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

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

11.1.5.3. Account management

11.1.5.3.1. Terms

This topic describes the concepts of Apsara Stack accounts, AccessKey pairs, and roles.

Apsara Stack account

An Apsara Stack account is used to manage all your cloud resources. You can use an Apsara Stack account to manage all the resources under this account and access the Apsara Stack Operations (ASO) console and Apsara Big Data Manager (ABM).

role

In ABM, a role is a collection of operations and maintenance (O&M) permissions on cloud services. After you create an Apsara Stack account, you cannot use this account to perform O&M operations on cloud services. You must assign roles with the required permissions to the account. Permission granting and revoking take effect immediately. You cannot add, modify, or delete the built-in roles of ABM.

AccessKey

When you create an Apsara Stack account, the system generates a unique AccessKey pair for the account. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. An AccessKey pair is a username (AccessKey ID) and password (AccessKey secret) pair that is used to connect to cloud service instances. AccessKey pairs and Apsara Stack accounts are used in different scenarios. AccessKey pairs are used when you call the API operations of cloud services.

11.1.5.3.2. Log on to the ASO console

This topic describes how to log on to the Apsara Stack Operations (ASO) console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   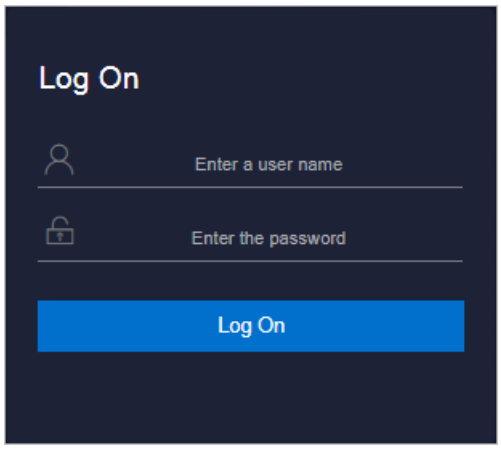

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.

11.1.5.3.3. Add a role

You can customize roles in the ASO console to achieve more flexible and efficient permission control.

Context

A role is a collection of access permissions. You can assign different roles to different users to meet requirements for system access control. Roles are classified into basic roles and custom roles. Basic roles, also known as atomic roles, are preset by the Operation Access Manager (OAM) system. Users cannot modify or delete these roles. Custom roles can be modified and deleted.

You can create different operations and maintenance (O&M) accounts for big data products such as MaxCompute, DataWorks, Realtime Compute, and DataHub. You can assign roles with the required permissions to the accounts. This eliminates the need to use the same account to perform O&M operations on all big data products.

ABM has the following built-in roles.

Role	Description
bcc_admin	The role of a super administrator. This role is used to grant the highest permissions of ABM to an account. The highest permissions include account management, permission management, and cloud service O&M.
bcc_admin_ads	The role of an AnalyticDB administrator. This role is used to grant the O&M permissions on AnalyticDB resources to an account.
bcc_admin_odps	The role of a MaxCompute administrator. This role is used to grant the O&M permissions on MaxCompute resources to an account.
bcc_admin_dataworks	The role of a DataWorks administrator. This role is used to grant the O&M permissions on DataWorks resources to an account.
bcc_admin_streamcompute	The role of a Realtime Compute administrator. This role is used to grant the O&M permissions on Realtime Compute resources to an account.
bcc_admin_dataapp	The role of an administrator for products other than Apsara, ABM, MiniLVS, MiniRDS, AnalyticDB, MaxCompute, DataWorks, and Realtime Compute. This role is used to grant the O&M permissions on the resources of other products such as Graph Analytics, Quick BI, DataHub, and Machine Learning to an account.
bcc_admin_biggraph	The role of a BigGraph administrator. This role is used to grant the O&M permissions on BigGraph resources to an account.
bcc_account_admin	The role of an account administrator. This role is used to grant the O&M permissions on account management and permission management to an account.

Procedure

1. Log on to the ASO console.

2. In the left-side navigation pane, click **System Management** and then **Roles** to go to the **Role Management** page.

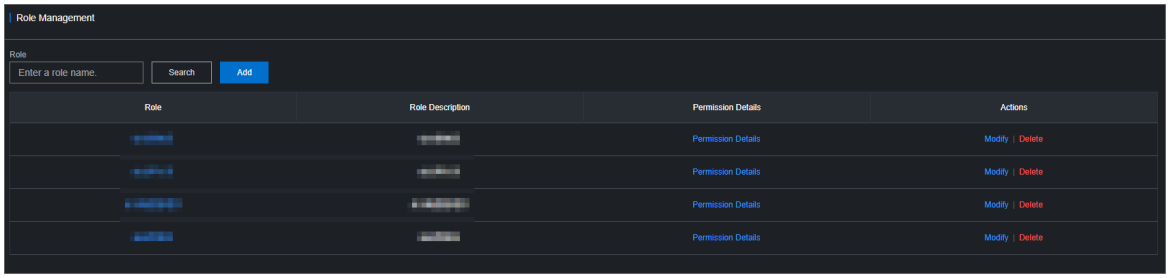

3. Add a role.

> **Note** To add a role in ASO, you must have the ASO security officer role.

   i. On the **Role Management** page, click **Add**.

   ii. In the **Add** dialog box, specify a **Role Name** and **Role Description**, select a **Base Role**, and configure menu permission settings.

> **Note** When you specify **Base Role**, you can enter the *bcc* keyword.

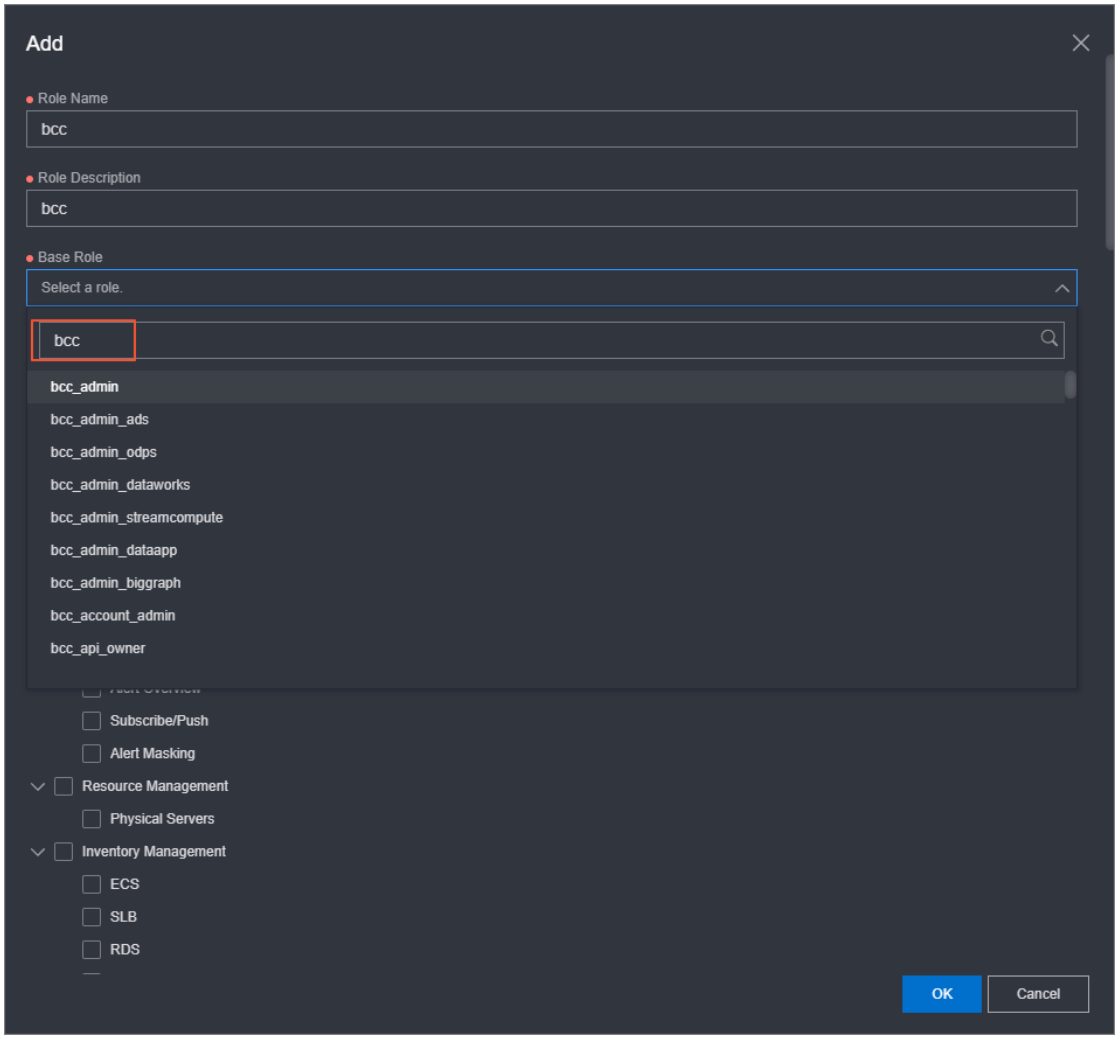

iii. After you configure the parameters, click **OK**.

11.1.5.3.4. Add an Apsara Stack account

An Apsara Stack account is used to manage all your cloud resources. You can use an Apsara Stack account to manage all the resources under this account and access the Apsara Stack Operations (ASO) console and Apsara Big Data Manager (ABM).

Prerequisites

- A department is created in the ASO console. For more information, see the **Department management** topic in *Alibaba Cloud Apsara Stack Enterprise Operations and Maintenance Guide.*
- O&M roles are created for big data products. For more information, see Add a role.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, click **System Management** and then **Users**. On the page that appears, click the **Users** tab.

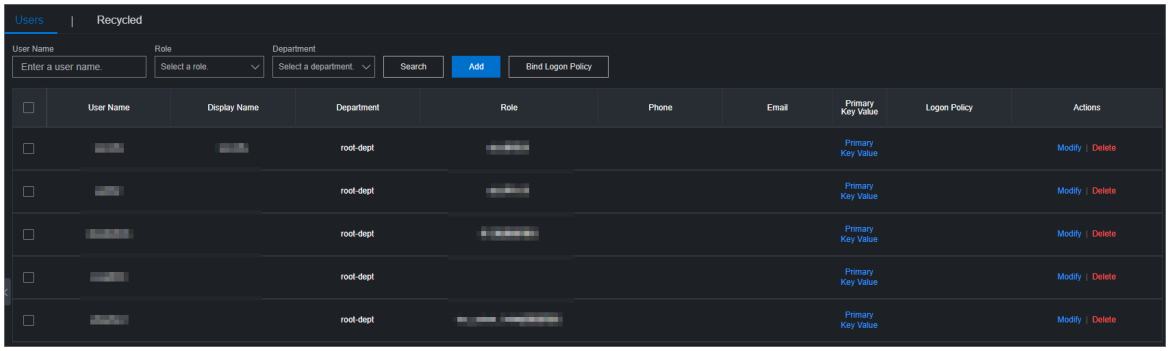

3. On the Users tab, click **Add**.

> **Note** To add a role in ASO, you must have the ASO security officer role.

4. In the **Add User** dialog box, enter a **User Name** and **Password**, and click **OK** to add a user.

11.1.5.3.5. Modify a role assigned to an account

You can modify a role that is assigned to an account.

Prerequisites

- You have created an Apsara Stack account that is used to manage big data products in the ASO console. For more information, see Add an Apsara Stack account.
- You have created roles in the ASO console that are used to grant permissions to the Apsara Stack account. For more information, see Add a role.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, click **System Management** and then **Users**. On the page that appears, click the **Users** tab.

3. On the **Users** tab, find the target account and click **Modify** in the **Actions** column.
4. In the **Modify User** dialog box, select a role from the **Role** drop-down list.
5. Click **OK**.

11.1.5.3.6. Grant permissions to an Apsara Stack account

After you create an Apsara Stack account, you cannot use this account to perform operations and maintenance (O&M) operations on cloud services. To obtain the O&M permissions, you must assign roles with the required permissions to the account. The granted permissions take effect immediately.

Prerequisites

You have obtained an account to which the super administrator permission or account management permission has been granted.

Context

> **Note** In ABM, all Apsara Stack accounts have permissions on the **Alert Configuration** and **API Management** modules.

ABM has the following built-in roles.

Role	Description
bcc_admin	The role of a super administrator. This role is used to grant the highest permissions of ABM to an account. The highest permissions include account management, permission management, and cloud service O&M.
bcc_admin_ads	The role of an AnalyticDB administrator. This role is used to grant the O&M permissions on AnalyticDB resources to an account.
bcc_admin_odps	The role of a MaxCompute administrator. This role is used to grant the O&M permissions on MaxCompute resources to an account.
bcc_admin_dataworks	The role of a DataWorks administrator. This role is used to grant the O&M permissions on DataWorks resources to an account.
bcc_admin_streamcompute	The role of a Realtime Compute administrator. This role is used to grant the O&M permissions on Realtime Compute resources to an account.
bcc_admin_dataapp	The role of an administrator for products other than Apsara, ABM, MiniLVS, MiniRDS, AnalyticDB, MaxCompute, DataWorks, and Realtime Compute. This role is used to grant the O&M permissions on the resources of other products such as Graph Analytics, Quick BI, DataHub, and Machine Learning to an account.
bcc_admin_biggraph	The role of a BigGraph administrator. This role is used to grant the O&M permissions on BigGraph resources to an account.

Role	Description
bcc_account_admin	The role of an account administrator. This role is used to grant the O&M permissions on account management and permission management to an account.

Procedure

1. Log on to the ABM console.
2. In the top navigation bar, click **Management**. On the page that appears, click **Accounts** in the left-side navigation pane.
3. Click the **Permissions** tab.
4. On the **Permissions** tab, select **Authorize** from the **Actions** drop-down list for the role that you want to assign to an account.

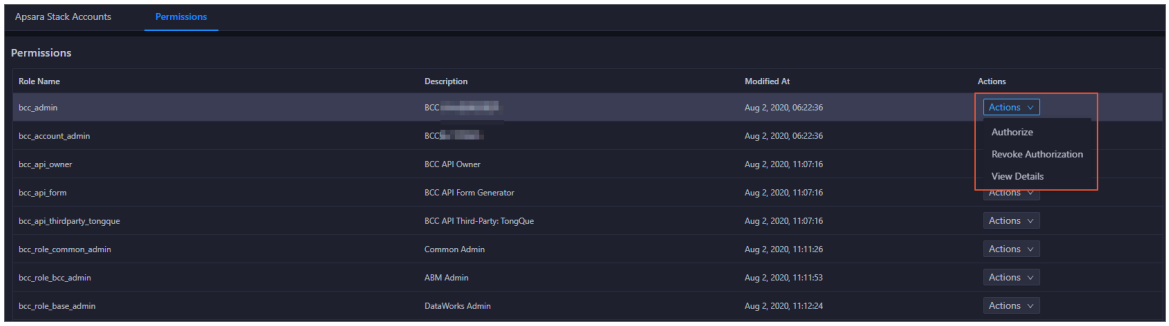

5. In the **Authorization Management** dialog box, select the accounts to which you want to assign the role. (Multiple choices are allowed.)

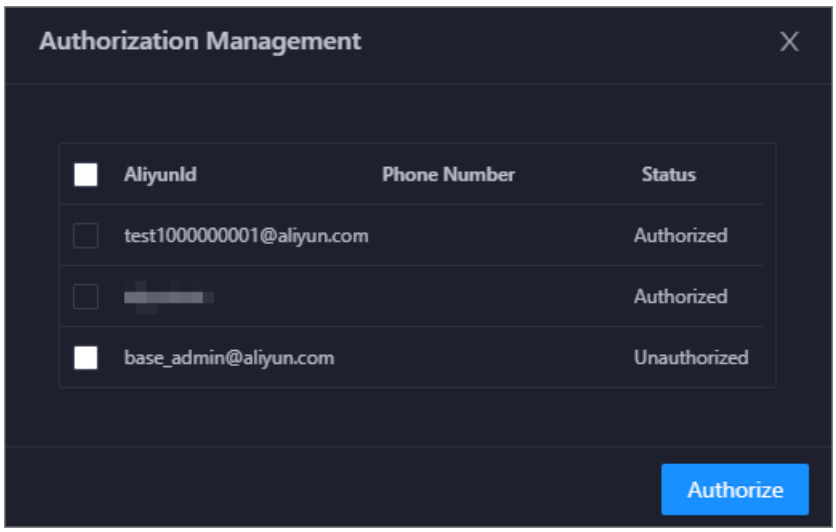

> **Note** If you have already assigned the role to an account, you cannot select this account.

6. Click **Authorize**. A message appears, which indicates that the role is assigned to the selected accounts.

11.1.5.4. Patch management

Apsara Big Data 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**.

11.1.5.5. Hot upgrade

Apsara Big Data 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**.

11.1.5.6. Health management

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

Prerequisites

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

Background

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

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

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

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

View details and mount points of scheduling items

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

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

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

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

Modify a scheduling item

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

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

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

> **Note** Set the **Execution Interval** parameter based on the **crontab** command.

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

View faulty hosts

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

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

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

Modify a monitoring item

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

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

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

Add a mount point for a monitoring item

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

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

The following table describes some key parameters.

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

11.1.5.7. Operation auditing

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

Prerequisites

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

Background

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

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

Procedure

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

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

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

State	Executable operation
Not Started	○ Click **Start** to start the task. ○ Click **Parameter Configuration** to view the parameter configuration of the task. ○ Click **Cancel** to cancel the task.
Pending	○ Follow the instructions and click **OK** to go to the next step. ○ Click **Rollback** to roll back to the complete status of the previous step. ○ Click **Parameter Configuration** to view the parameter configuration of the task. ○ Click **Cancel** to cancel the task.

State	Executable operation
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.
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.

11.1.6. Go to other platforms

Apsara Big Data Manager (ABM) provides the links to Apsara Stack Operation, Apsara Infrastructure Management Framework, and Alibaba Cloud Security to facilitate the O&M of big data products.

Prerequisites

You have obtained an ABM account that works properly and the corresponding password.

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.

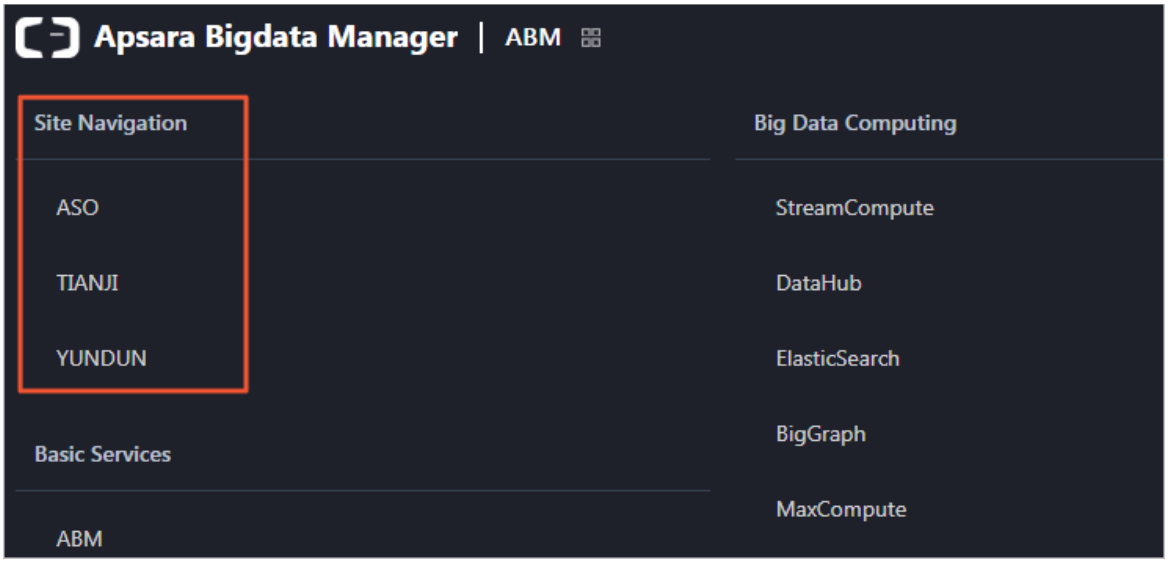

Result

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.

11.2. MaxCompute

11.2.1. Concepts and architecture

MaxCompute architecture shows the MaxCompute architecture.

MaxCompute architecture

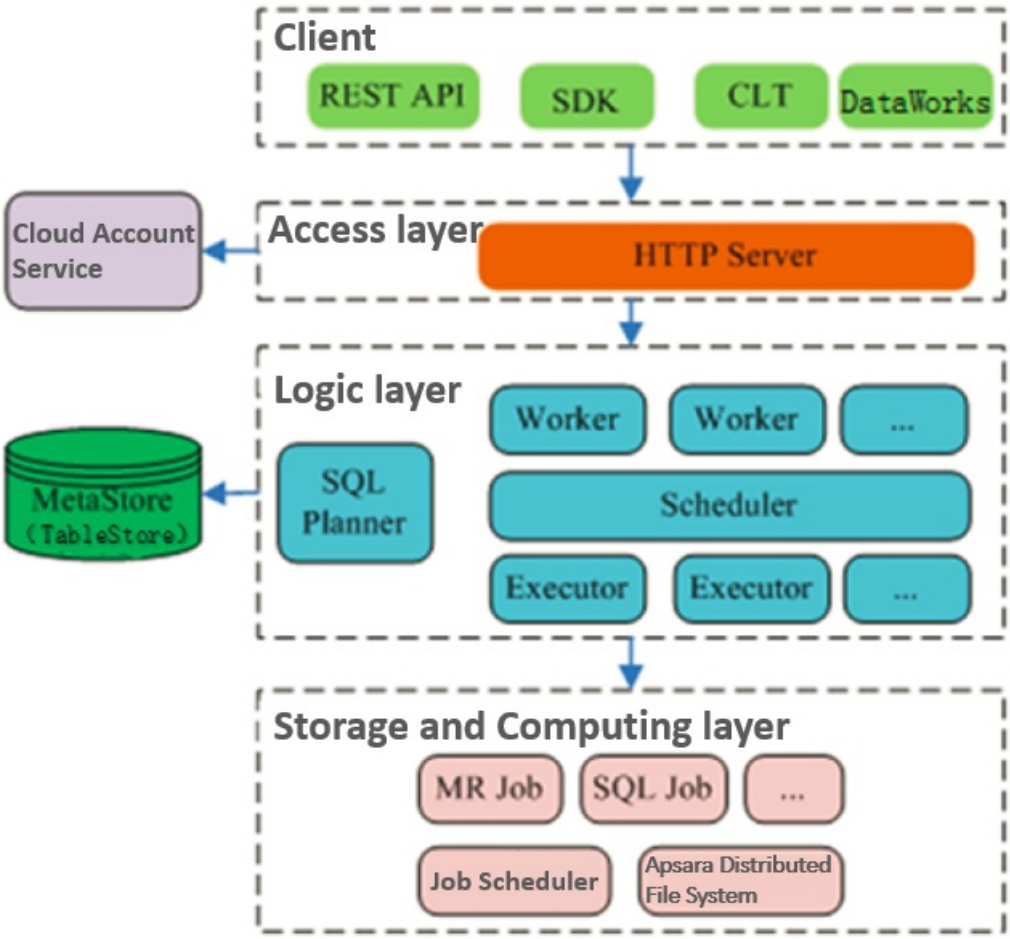

The MaxCompute service is divided into four parts: **client**, **access layer**, **logic layer**, and **storage and computing layer**. Each layer can be horizontally scaled.

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.

MaxCompute job execution process

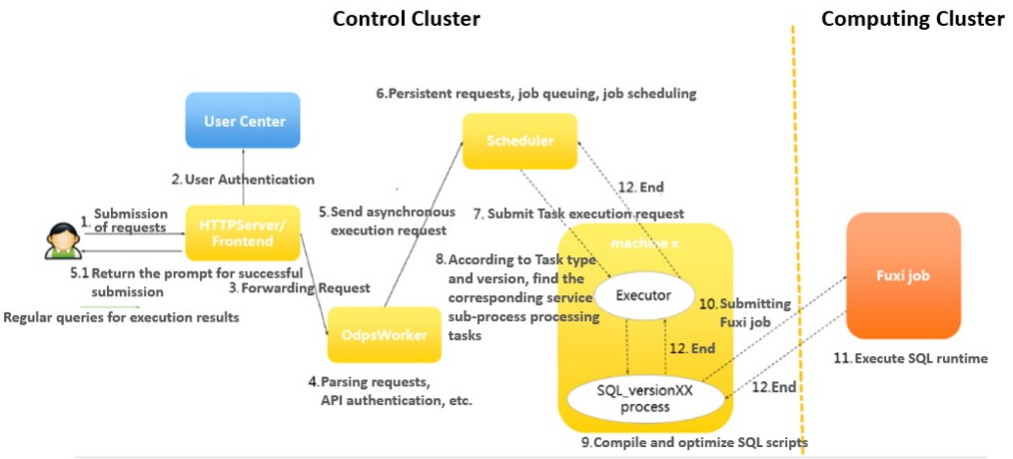

The following concepts are involved in the MaxCompute job execution process:

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.

11.2.2. O&M commands and tools

11.2.2.1. 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.

11.2.2.2. 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.

Client command options

```
$/apsara/odps_tools/clt/bin/odpscmd -h
Usage: odpscmd [OPTION]...
where options include:
    --help                                (-h)for help
    --config=<config_file>                specify another config file
    --project=<prj_name>                  use project
    --endpoint=<http://host:port>         set endpoint
    -u <user_name> -p <password>          user name and password
    --instance-priority=<priority>        priority scope[0-9]
    -M                                    read machine readable data
    -k <n>                                will skip begining queries and start from specified posi
ion
    -r <n>                                set retry times
    -f <"file_path;">                     execute command in file
    -e <"command;[command;]...">          execute command, include sql command
    -C                                    will display job counters
    -y                                    will not submit jobs to fuxi master
```

- **-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.

Commonly used commands

Command	Description
whoami;	Allows you to view your Apsara Stack tenant account and endpoint information.

Command	Description
show p;	Allows you to view information about all instances that have been run.
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.

11.2.2.3. 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.

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)
```

```
                    (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
```

11.2.2.4. LogView tool

11.2.2.4.1. 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.

11.2.2.4.2. 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.

A long string starting with logview

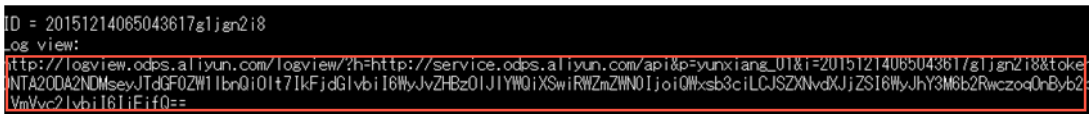

Enter the string with all carriage return and line feed characters removed in the address bar of the browser.

Composition of a LogView string

A LogView string consists of five parts, as shown in the following figure.

Composition of a LogView string

11.2.2.4.3. 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.

Relationships between MaxCompute tasks and Fuxi instances

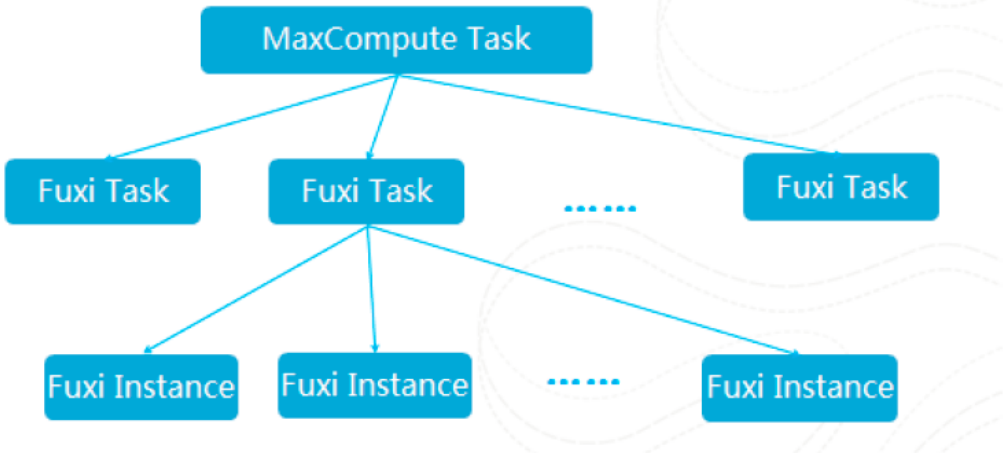

The following figures show the relevant information in LogView.

MaxCompute Instance

MaxCompute Instance

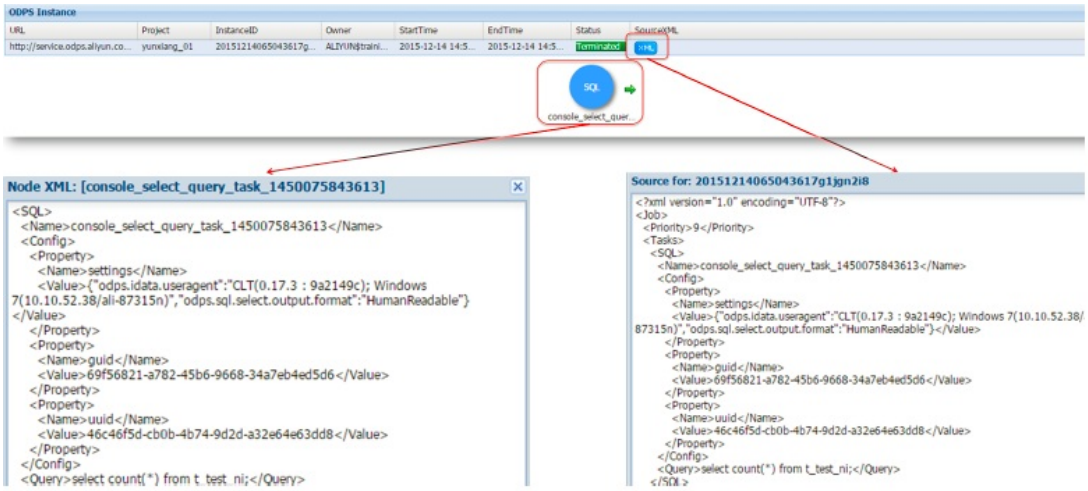

MaxCompute Task

MaxCompute Task

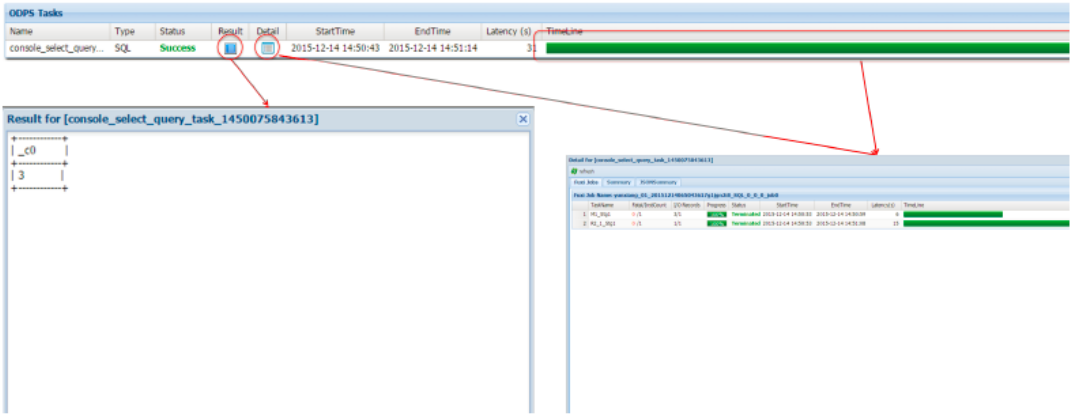

Task Detail - Fuxi Job

Task Detail - Fuxi Job(1)

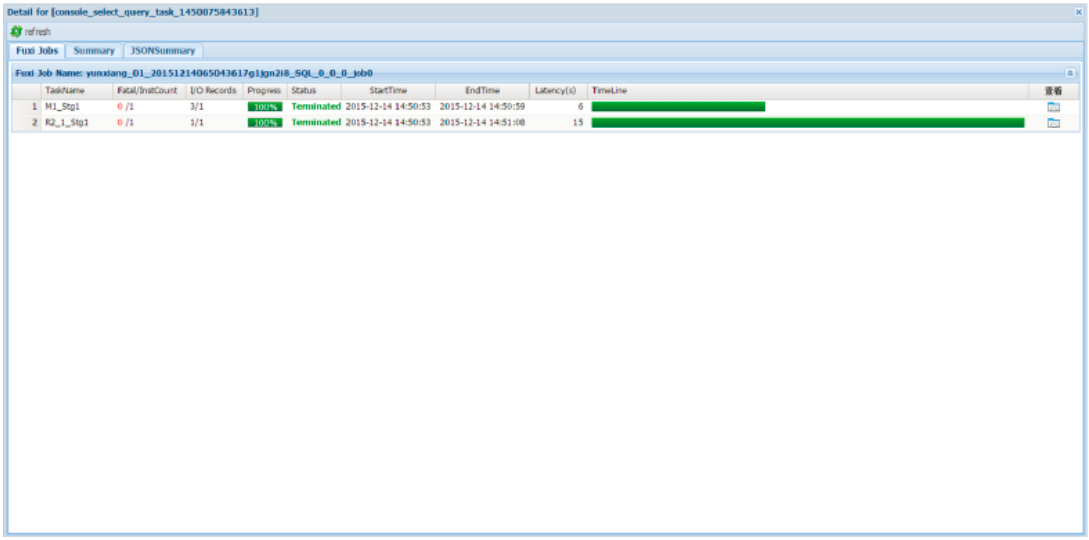

Task Detail - Fuxi Job(2)

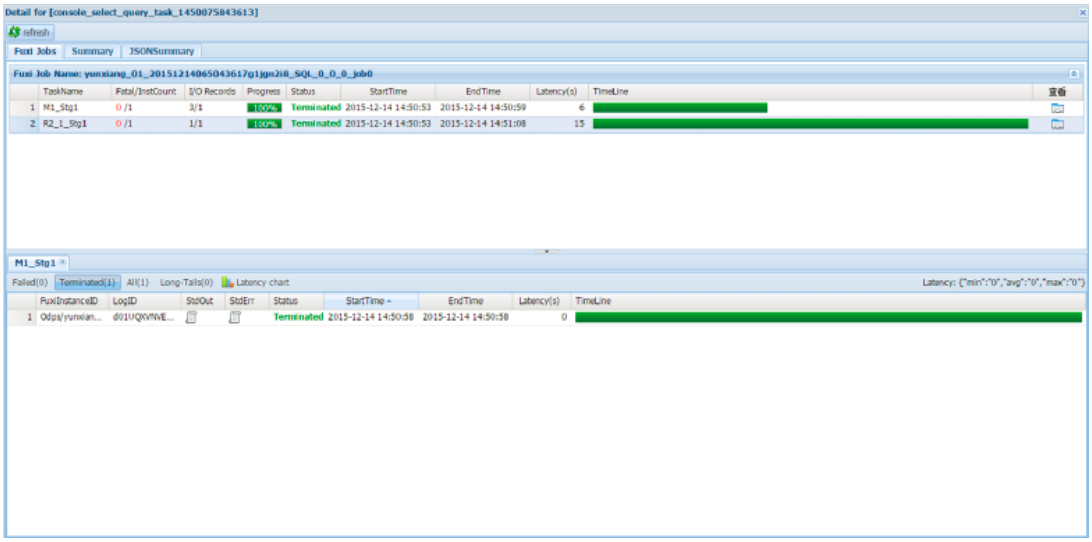

Task Detail - Summary

Task Detail - Summary

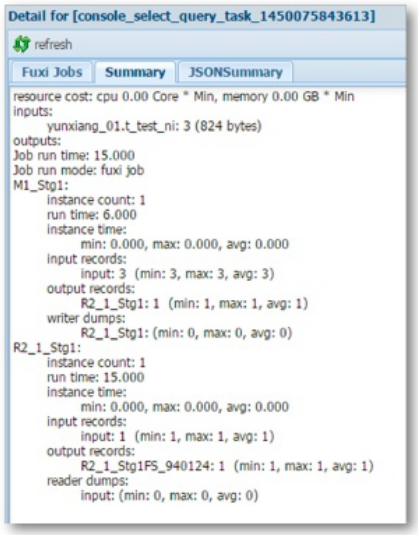

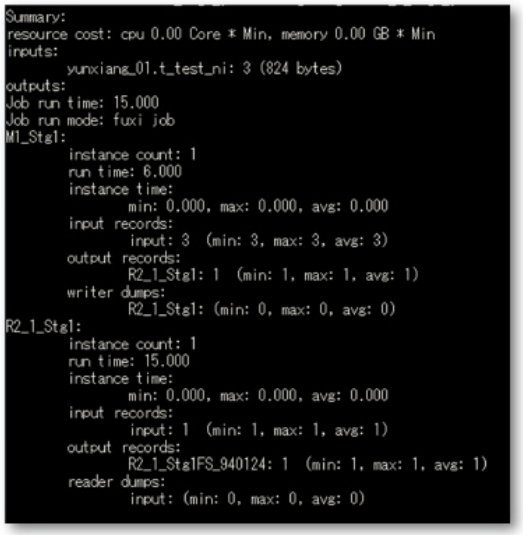

Task Detail - JSONSummary

Task Detail - JSONSummary

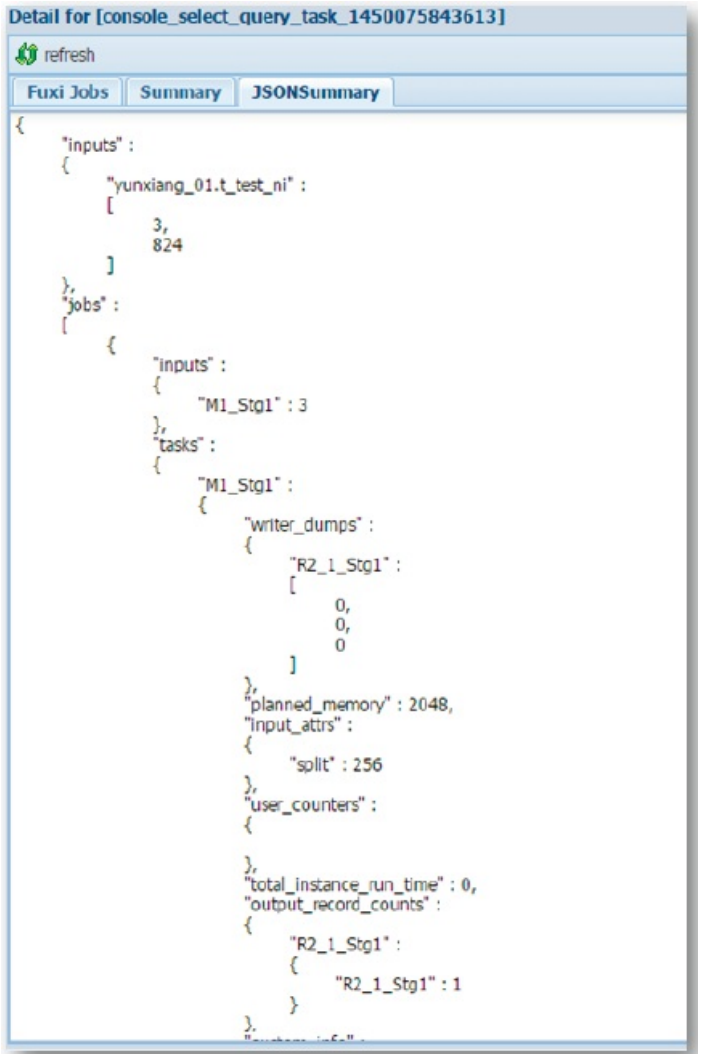

11.2.2.4.4. Basic operations and examples

View each point in time in the life cycle of a job.

View each point in time in the life cycle of a job

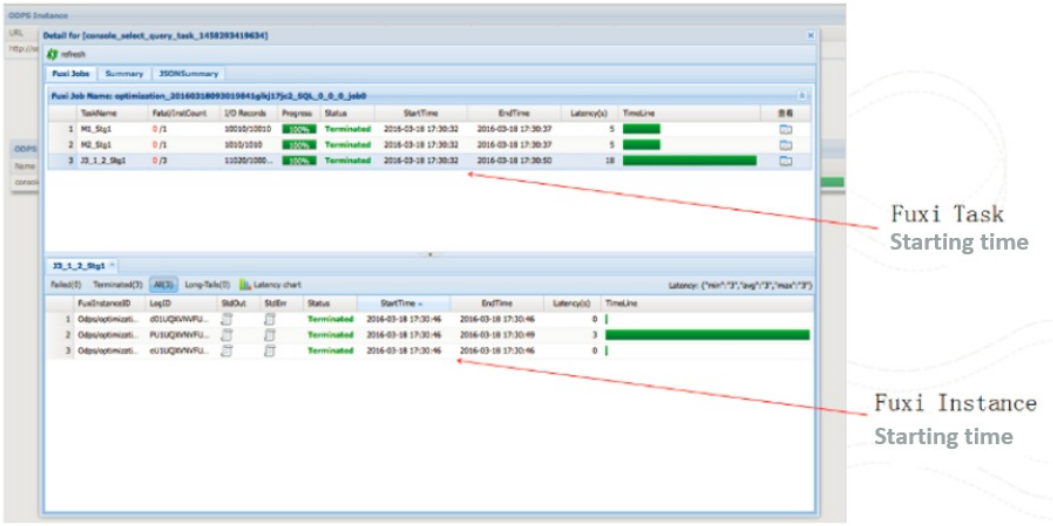

View the time it takes for Job Scheduler to schedule an instance.

View the time it takes for Job Scheduler to schedule an instance

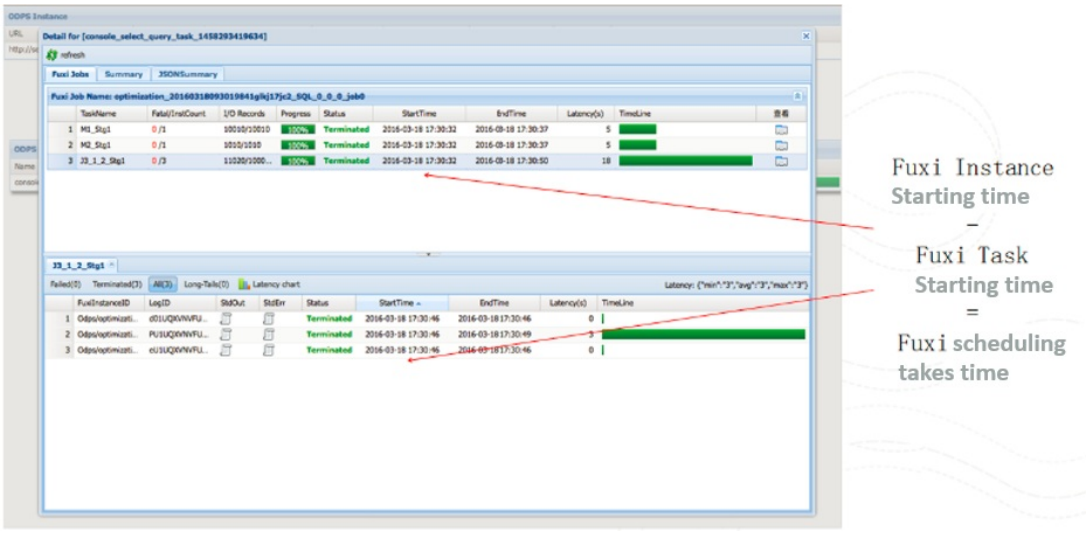

View the polling interval.

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
lbnQiOlt7IkFjdGlvbiI6WyJvZHBzOlJlYWQiXSwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2RwczoqOnBybzplY3RzL29wdGlt
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

Check for data skews

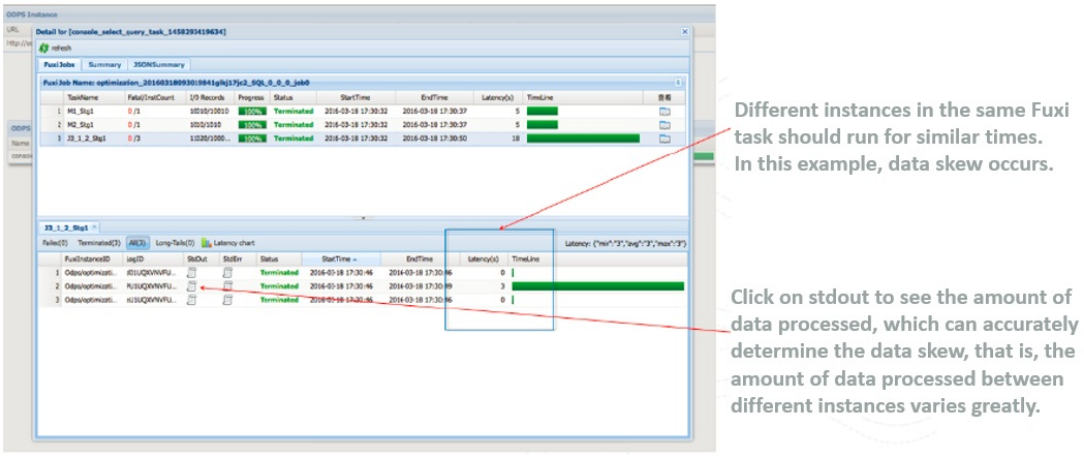

View the UDF and MR debugging information

View the UDF and MR debugging information

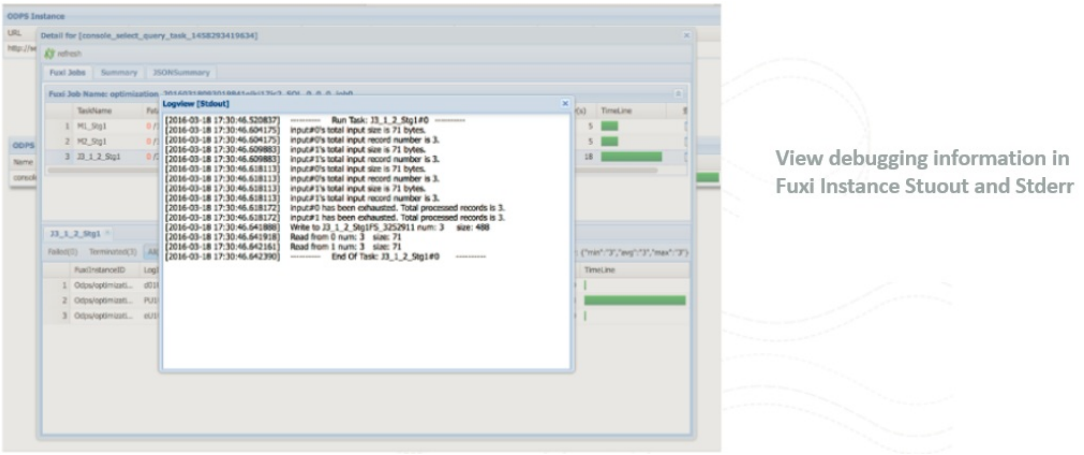

View the task status - Terminated

View the task status - Terminated

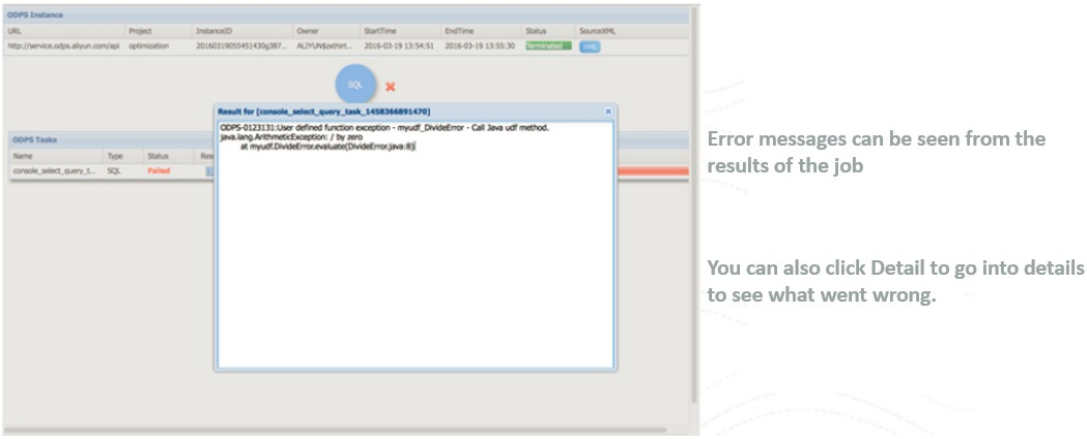

11.2.2.4.5. Best practices

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.

Locate LogView based on the instance ID

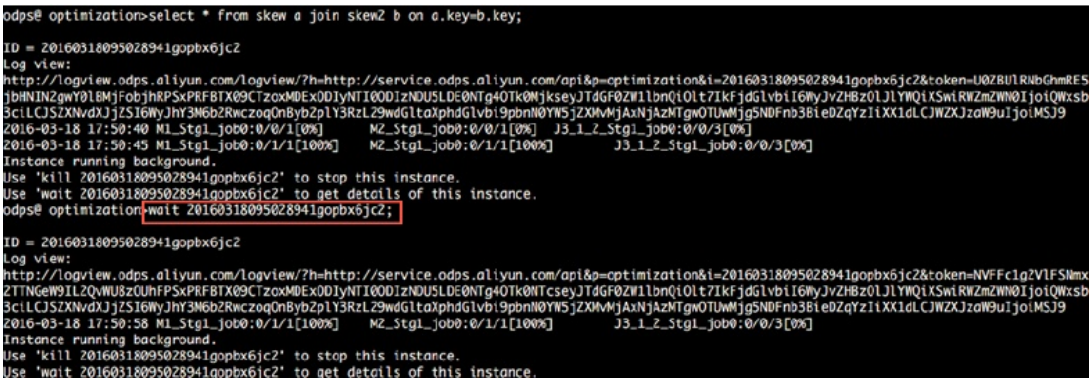

Locate running tasks

After you exit the control window, you can run the `show p;` command to locate currently running tasks and historical tasks.

Locate running tasks

11.2.2.5. Apsara Bigdata Manager

Apsara Bigdata Manager (ABM) supports O&M on big data services from the perspectives of businesses, services, clusters, and hosts. You can also update big data services, customize alert configurations, and view the O&M history in the ABM console.

On-site Apsara Stack engineers can use ABM to easily manage big data services by performing actions, such as viewing resource usage, checking and handling alerts, and modifying configurations.

For more information about how to log on to the ABM console and perform O&M operations in the console, see the *MaxCompute O&M* topic.

11.2.3. Routine O&M

11.2.3.1. 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.

11.2.3.2. Routine inspections

1. On the Cluster Operations page in Apsara Infrastructure Management Framework, check whether all machines have reached the desired state.
   i. 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.
   ii. 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.
   iii. 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;
   ```

```
odps@ odps_smoke_test>select count(*) from dual;

ID = 20180420061754827g78x7i
Log view:
http://logview.cn-hangzhou-env6-d01.odps.aliyun-inc.com:9000/logview/?h=http://s
180420061754827g78x7i&token=aEVmNTF1dm5GMnFOV1BSWjViZE0rOWRERnZFPSxPRFBTX09CTzox
SwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2RwczoqOnByb2plY3RzL29kcHNfc21va
J9
Job Queueing.
Summary:
resource cost: cpu 0.00 Core * Min, memory 0.00 GB * Min
inputs:
        odps_smoke_test.dual: 1 (1408 bytes)
outputs:
Job run time: 0.000
Job run mode: service job
Job run engine: execution engine
M1:
        instance count: 1
        run time: 0.000
        instance time:
                min: 0.000, max: 0.000, avg: 0.000
        input records:
                TableScan_REL5136522: 1  (min: 1, max: 1, avg: 1)
        output records:
                StreamLineWrite_REL5136523: 1  (min: 1, max: 1, avg: 1)
R2_1:
        instance count: 1
        run time: 0.000
        instance time:
                min: 0.000, max: 0.000, avg: 0.000
        input records:
                StreamLineRead_REL5136524: 1  (min: 1, max: 1, avg: 1)
        output records:
                ADHOC_SINK_5136527: 1  (min: 1, max: 1, avg: 1)

+------------+
| _c0        |
+------------+
| 1          |
+------------+
```

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=VS9hRzc4RjAzeXJ2bmRFOUtyYnNWSXFkNW0wPSxPRFBTX09C
iI6WyJvZHBzOlJlYWQiXSwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2RwczoqC
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:
   i. 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
```

   ii. 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
```

   iii. 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
```

iv. `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
```

v. `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       119872 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:

i. `puadmin fs -abnchunk -t none`

```
$puadmin  fs -abnchunk  -t  none
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type    FoundTime
```

ii. `puadmin fs -abnchunk -t onecopy`

```
$puadmin  fs -abnchunk  -t  onecopy
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type    FoundTime
```

iii. `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
```

```
$puadmin gems
ElectMasterStatus : ELECT_MASTER_OVER_ELECTION
PrimaryId         : tcp:/[redacted]
PreferedWorkerid  :
PrimaryLogId      : 617851602
TotalWokerNumber  : 3
ElectConsentNumber : 2
SyncConsentNumber  : 2
ElectSequence     : [935155f0-fb68-4cd9-bee9-08d23afe84eb,4,1328760004]
WorkerStatus    :
        tcp://[redacted] : ELECT_WORKER_STATUS_SECONDARY
        tcp://[redacted] : ELECT_WORKER_STATUS_SECONDARY
        tcp://[redacted] : ELECT_WORKER_STATUS_PRIMARY

[admin@vm010036032037 /home/admin]
$puadmin gss
PrimaryStatus : PRIMARY_STARTUP_SERVICE_STARTED
PrimaryCurrentLogId : 617852679
WorkerSyncStatus :
        tcp://[redacted][SyncedLogId:617852670, LastFailTime:2018-04-17 12:07:43, WorkerType: NORMAL]
        tcp://[redacted][SyncedLogId:617852638, LastFailTime:1970-01-01 08:00:00, WorkerType: NORMAL]
```

8. Perform daily inspections in Apsara BigData Manager to check disk usage.

11.2.3.3. 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 functioning 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 dis
played, 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.
   i. 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/S
   etGlobalFlag --Parameter={\"fuxi_Enable_BadNodeManager\":false}
   ```

   ii. Run the following command to check the hostnames in the existing blacklist:

   ```
   /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
   ```

   iii. 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
   ```

   iv. 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.

i. 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"
    }
  ]
}
```

ii. 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.

i. 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".

ii. 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.

i. 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"
}
]
}
```

ii. 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:1f147fec4883e082646715cb79c3710f7b2
ae9c6e6851fa9a9452b92b4b3366a 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
```

11.2.3.4. 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 dis
   played, 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.
   i. 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/S
      etGlobalFlag --Parameter={\"fuxi_Enable_BadNodeManager\":false}
      ```

   ii. Run the following command to check the hostnames in the existing blacklist:

      ```
      /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
      ```

   iii. 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
      ```

   iv. 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.

11.2.3.5. 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.

   role.conf file

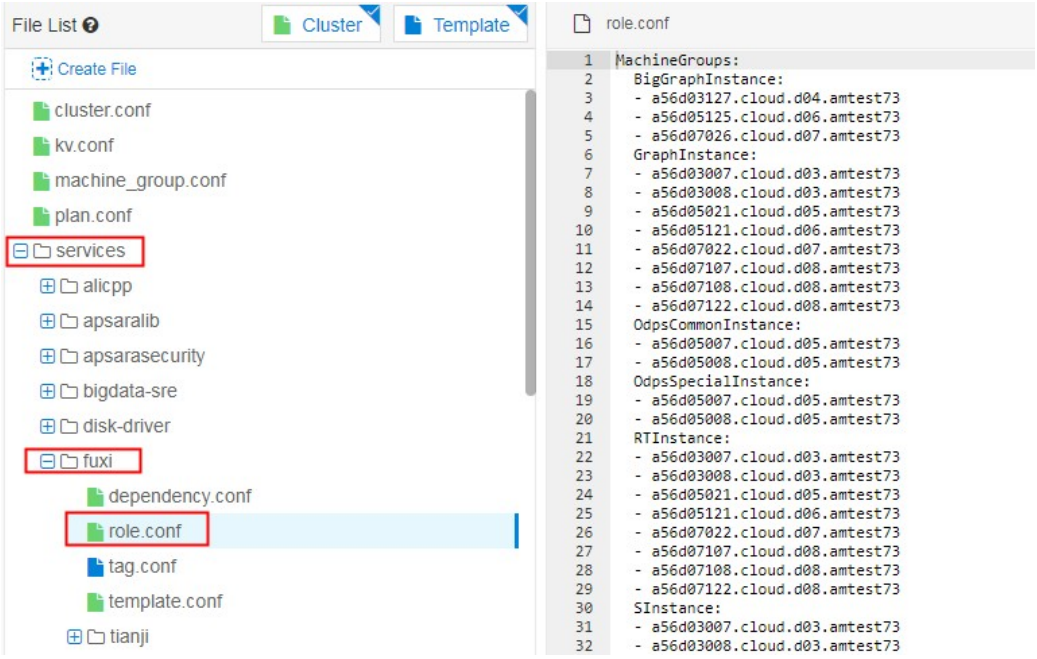

3. Adjust the machine tags on the right and click **Preview and Submit**.

   Adjust machine tags

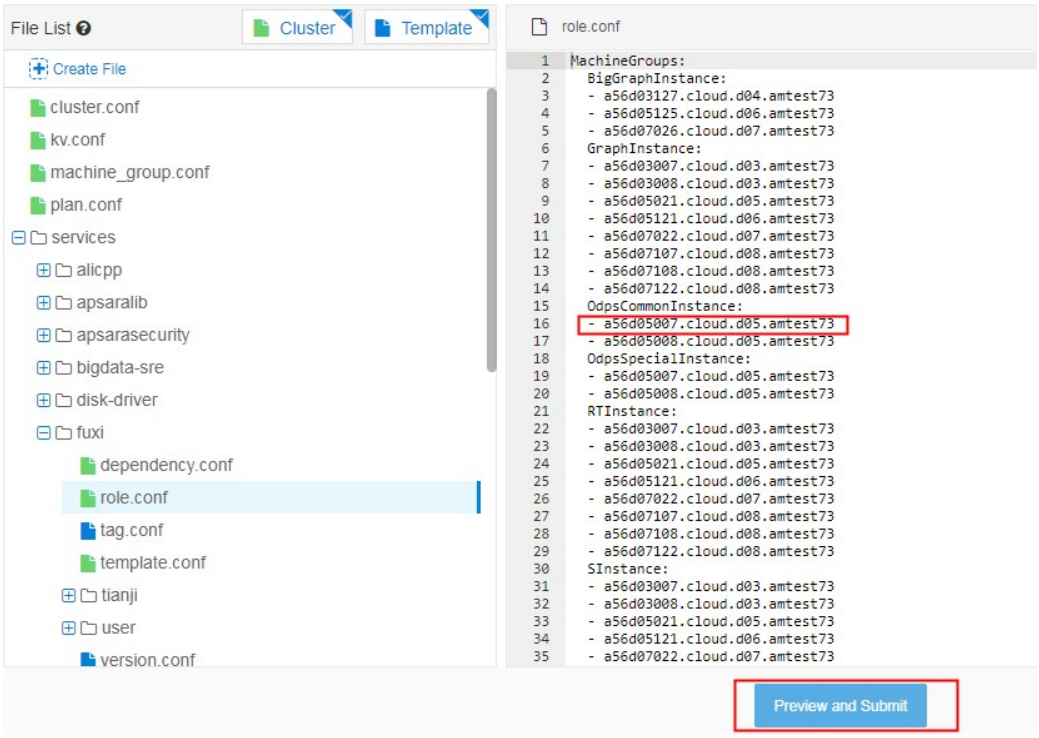

4. In the **Confirm and Submit** dialog box that appears, enter the change description and click **Submit**.

   Submit

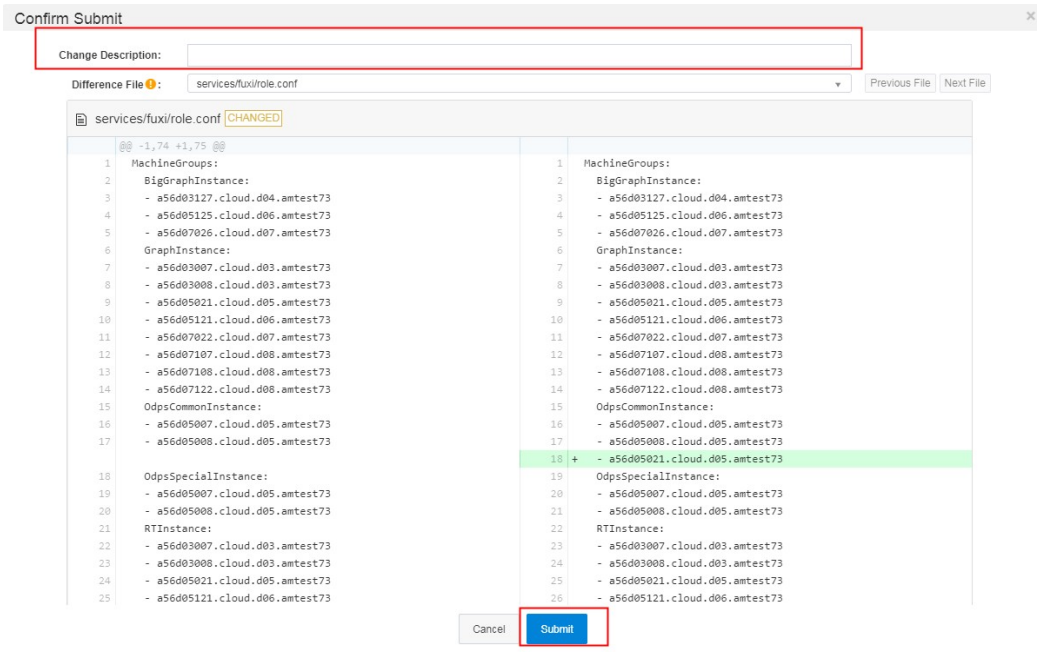

5. The cluster starts rolling and the changes start to take effect.

   > **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 `rttrl` command in the TerminalService window to confirm the

changes.

11.2.3.6. 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 pedeploymessagerservice
   -- 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 pedeployquotaservice
-- 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 pedeployreplicationservice
-- 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.
```

11.2.4. Common issues and solutions

11.2.4.1. 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.

Capacity information

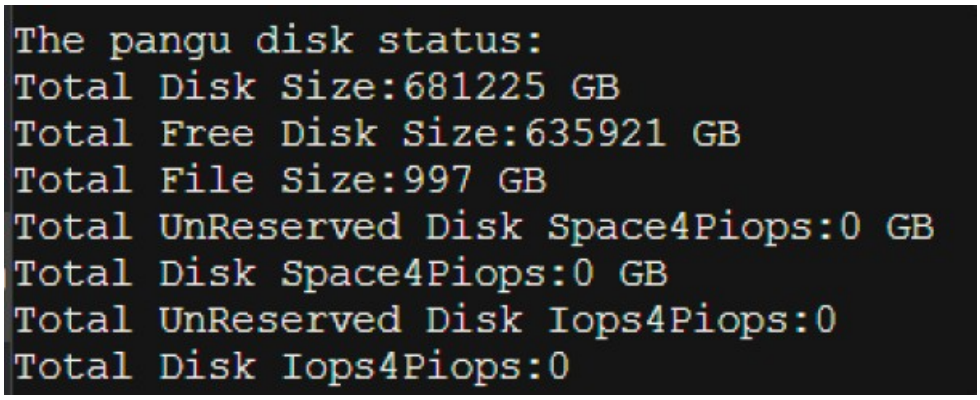

> **Note** Parameters:
>
> - Total Disk Size: the total amount of physical space. Each file is stored in three copies. The logical space is one third the size of the physical space.
> - Total Free Disk Size: the total size of available disks, excluding recycle bins on chunkservers.
> - Total File Size: 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.
  ```

  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.

Numbers of files that can be stored for different PanguMaster memory capacities

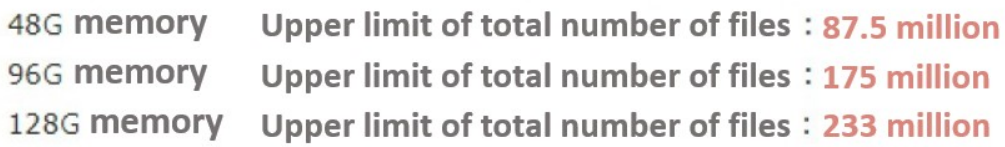

How to view the number of files stored in a MaxCompute cluster

- Run the `pu quota` command on the cluster AG to view the total number of files stored in a MaxCompute cluster.

  Total number of files

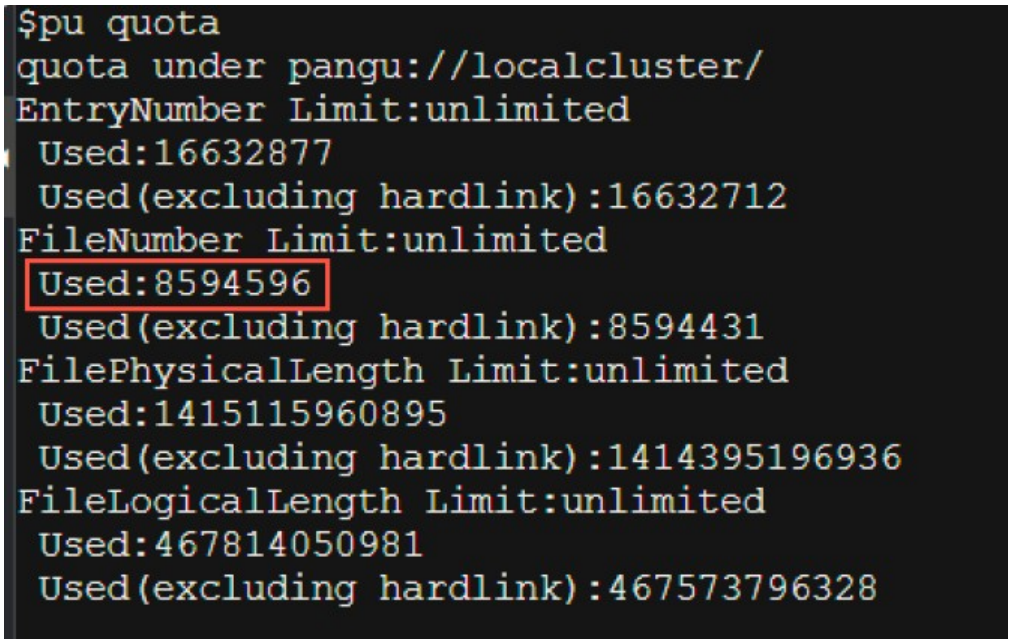

- This example uses the adsmr project to demonstrate how to view the number of files. Run the following command on the cluster AG to view the number of files for a single project in a MaxCompute cluster:

```
pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -A 4
```

Number of files for a single project

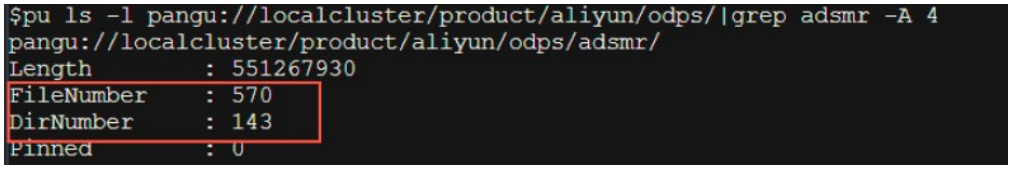

> **Note** Parameters:
>
> - FileNumber: the number of files used.
> - DirNumber: the number of directories used.
> - FileNumber + DirNumber = Number of files for the current project.

- 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 `rttrl` command on the cluster AG to view all computing resources.

All computing resources

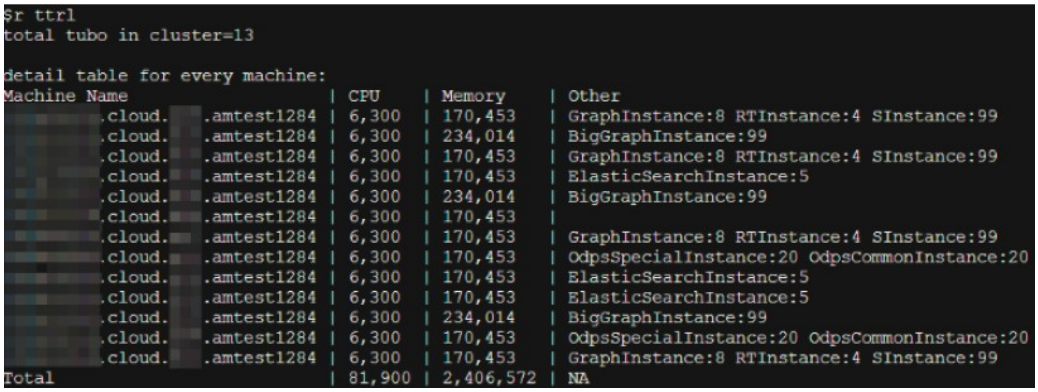

> **Note** In the command output, the domain name, total CPU capacity (Unit: U. 100 U = 1 core), and total memory (Unit: MB) of each Tubo machine, as well as the role of each Tubo machine in Job Scheduling System are listed in four columns.

- Run the `rtfrl` command on the cluster AG to view the remaining computing resources.

  ### Remaining computing resources

  ```
  $r tfrl
  Total tubo in cluster=13
  Detail table for every machine:
  Machine Name                    | CPU    | Memory    | Other
  .cloud.  .amtest1284 | 5,025  | 150,990   | GraphInstance:8 RTInstance:4 SInstance:81
  .cloud.  .amtest1284 | 6,090  | 226,874   | BigGraphInstance:98
  .cloud.  .amtest1284 | 5,285  | 153,634   | GraphInstance:8 RTInstance:4 SInstance:83
  .cloud.  .amtest1284 | 6,100  | 68,521    | ElasticSearchInstance:3
  .cloud.  .amtest1284 | 6,190  | 227,850   | BigGraphInstance:98
  .cloud.  .amtest1284 | 6,200  | 169,453   |
  .cloud.  .amtest1284 | 5,035  | 150,450   | GraphInstance:8 RTInstance:4 SInstance:83
  .cloud.  .amtest1284 | 4,600  | 131,565   | OdpsSpecialInstance:15 OdpsCommonInstance:12
  .cloud.  .amtest1284 | 6,200  | 104,921   | ElasticSearchInstance:4
  .cloud.  .amtest1284 | 6,000  | 67,521    | ElasticSearchInstance:3
  .cloud.  .amtest1284 | 5,790  | 218,634   | BigGraphInstance:97
  .cloud.  .amtest1284 | 5,400  | 133,089   | OdpsSpecialInstance:20 OdpsCommonInstance:13
  .cloud.  .amtest1284 | 5,485  | 157,634   | GraphInstance:8 RTInstance:4 SInstance:87
  Total                           | 73,400 | 1,961,136 | 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 `rcru` command on the cluster AG to view the resources used by all running jobs in MaxCompute.

  ### Resources used by all running jobs

  ```
  $r cru
  WorkItemName                                                                | CPU    | Memory  | VirturlResource
  Odps/DiskDriverService                                                      | 280    | 13,600  | {}
  Odps/odps_elasticsearch_elasticsearch_mdu_es_demo_20170509064623398g2q8q9d  | 200    | 1,024   | {}
  Odps/CGServiceControllerx                                                   | 1,980  | 66,660  | {'SInstance': 60}
  Odps/ReplicationServicex                                                    | 200    | 2,000   | {'OdpsSpecialInstance': 1}
  Odps/OdpsServicex                                                           | 1,400  | 45,128  | {'OdpsSpecialInstance': 4, 'OdpsCommonIns
  tance': 7}
  Odps/HiveServerx                                                            | 850    | 37,864  | {'OdpsCommonInstance': 4}
  Odps/XStreamServicex                                                        | 14,070 | 146,370 | {}
  Odps/QuotaServicex                                                          | 100    | 1,024   | {'OdpsSpecialInstance': 1}
  Odps/MessagerServicex                                                       | 300    | 3,092   | {}
  sm/sm used resource                                                         | 1,000  | 11,192  | {}
  Total Planned Resource                                                      | 20,380 | 327,954 | {'SInstance': 60, 'OdpsSpecialInstance':
  , 'OdpsCommonInstance': 11}
  ```

  > **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.

  Error messages

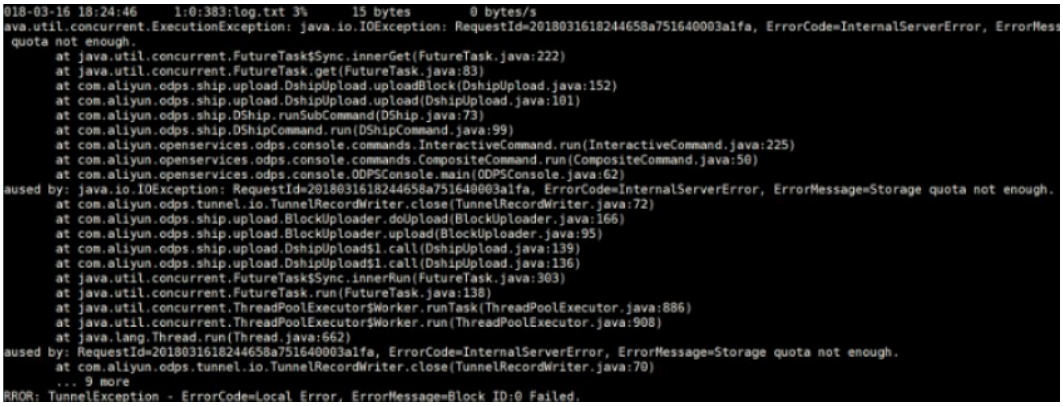

> **Notice** The sum of the storage capacity of all projects cannot exceed the total allowable storage capacity of a service. Similarly, the total file size of all projects cannot exceed the total allowable file size. Therefore, you must properly allocate the storage space and file size limit by project 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 `rquota` command on the cluster AG to view quota group settings.

    View quota group settings

    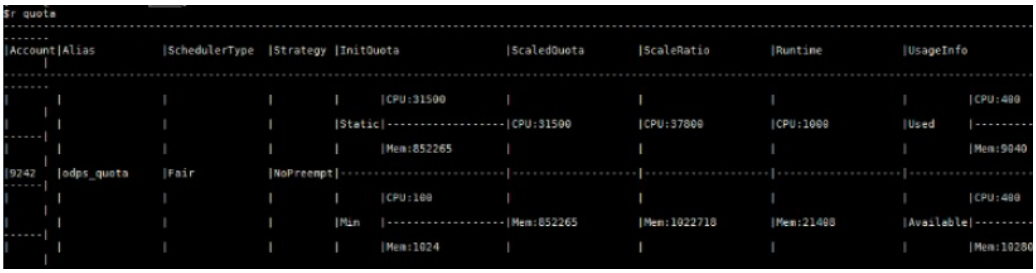

    You can run the following command on the cluster AG to create and modify a quota as needed:

```
sh /apsara/deploy/rpc_wrapper/rpc.sh setquota -i $QUOTAID -a $QUOTANAME -t fair -s $max_cpu_quot
a $max_mem_quota -m $min_cpu_quota $min_mem_quota
```

Note The command with $QUOTAID is used to modify a quota. The command without $QUOTAID is used to create a quota.

Create a quota

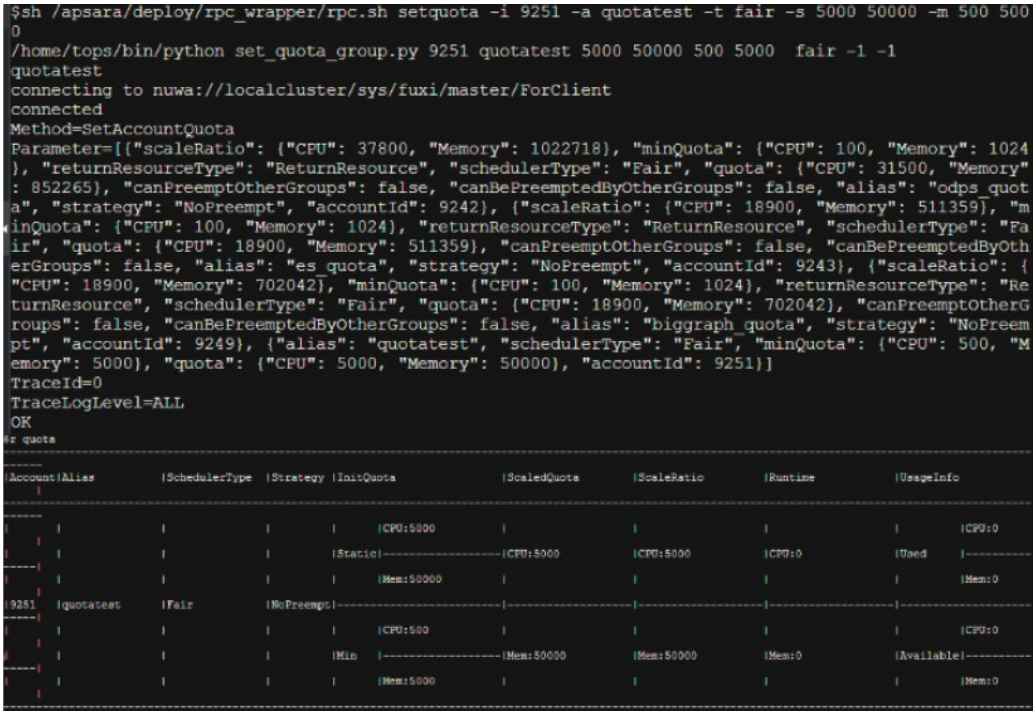

Modify a quota

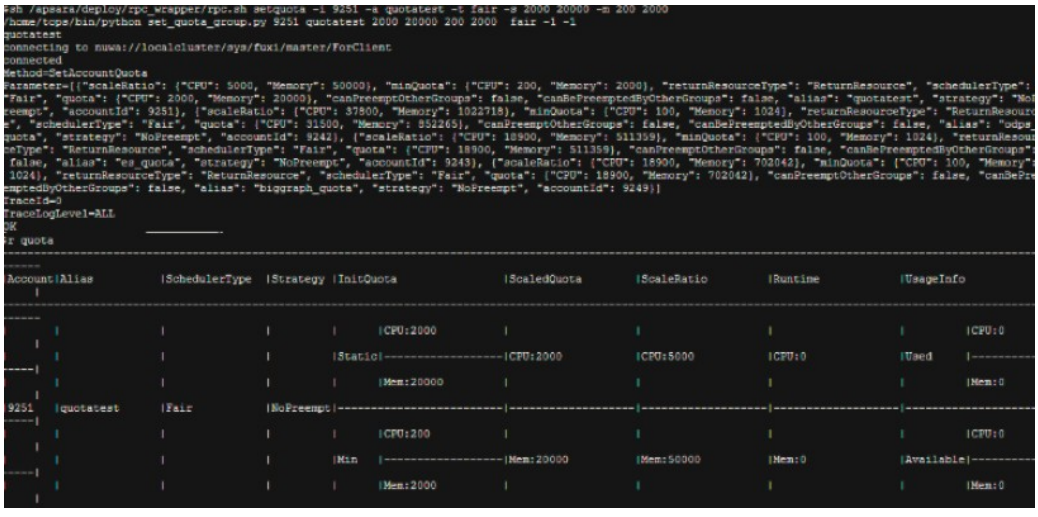

- How to divide quota groups

  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.

Jobs waiting for resources

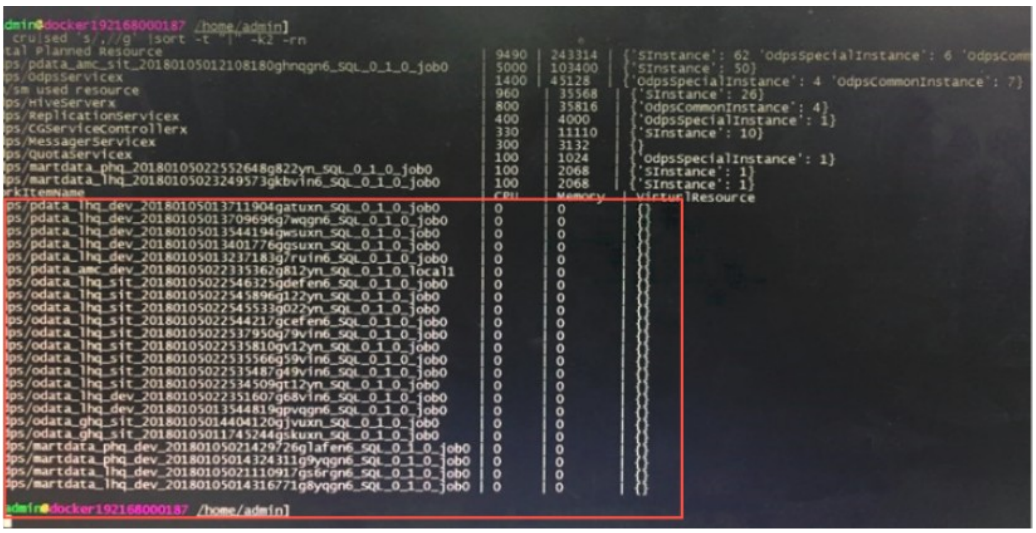

Quota used up

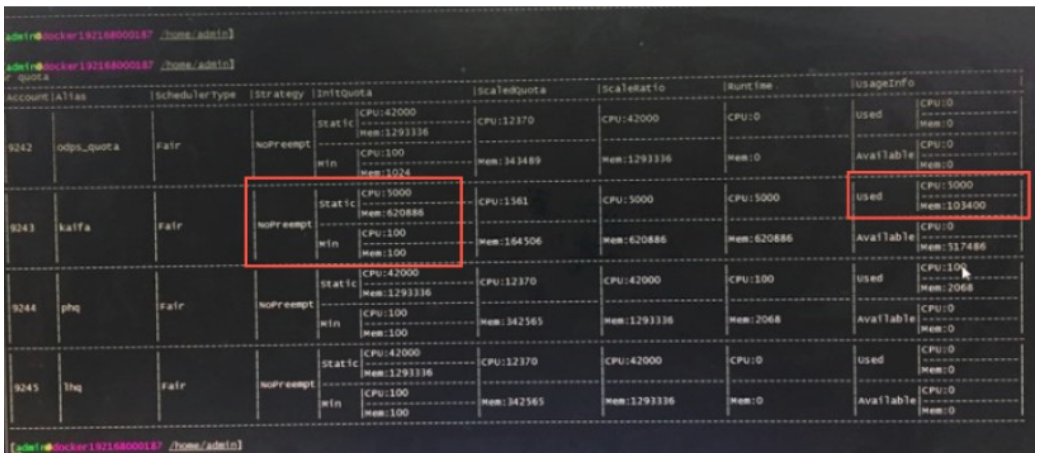

You must divide quota groups based on the following general principles:

- 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.

11.2.4.2. 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.

Symptom

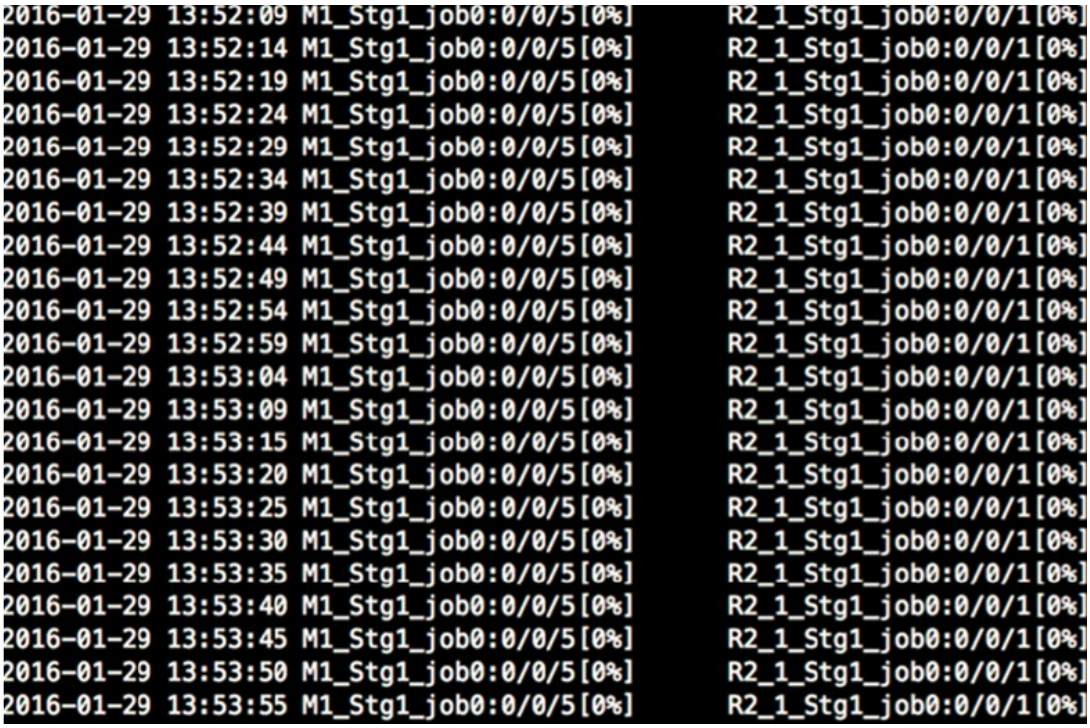

Cause: The issue is typically caused by insufficient resources. You can use LogView to determine the status of job resources (task instance status).

- 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.

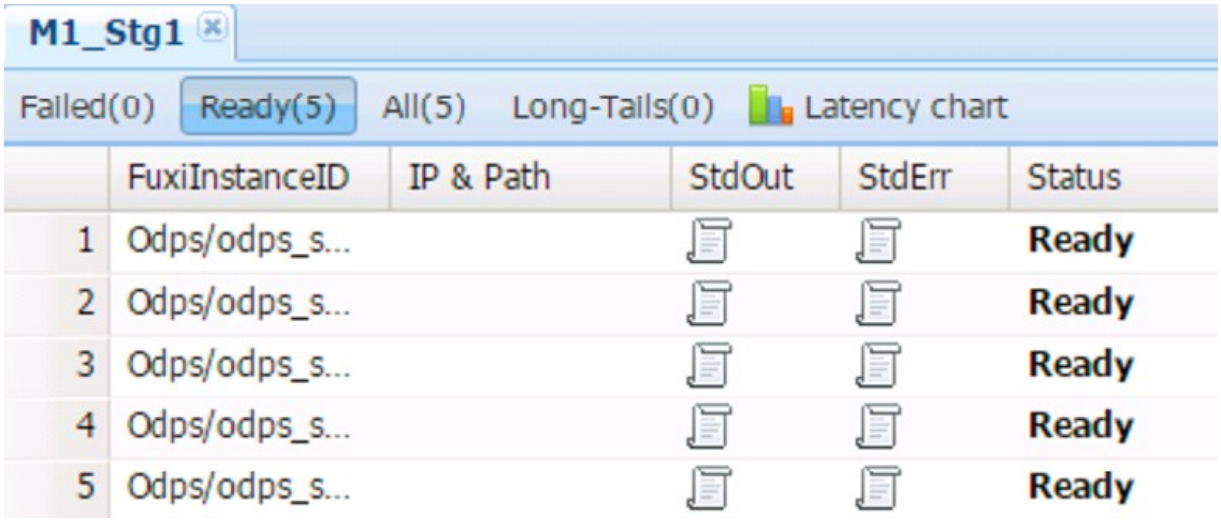

Solution:

- 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.

Cause analysis

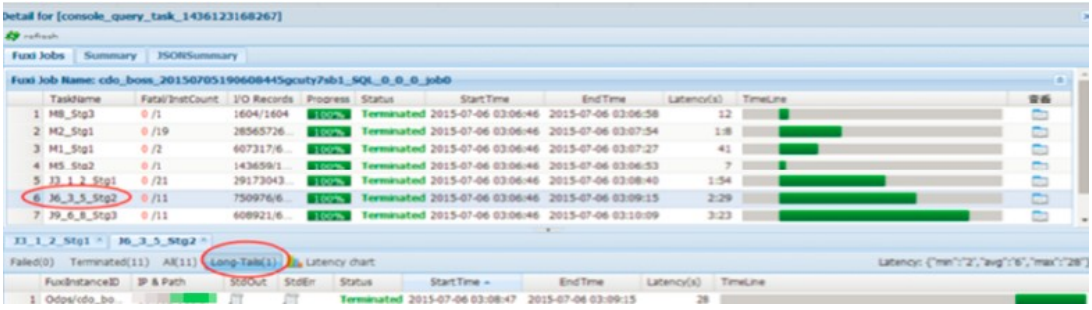

Further analysis: Analyze the job summary in LogView, and calculate the difference between the max and avg values of input and output records of a slow task. If the max and avg values differ by several orders of magnitude, it can be initially determined that the job data is skewed.

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 c
2, 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.

Storage information

```
odps@ odps_smoke_test>desc project odps_smoke_test -extended;
Name                                    odps_smoke_test
Description
Owner                                   ALIYUN$odpsadmin@aliyun.com
CreatedTime                             Fri Dec 25 00:43:06 CST 2015

Properties:
odps.table.lifecycle                    optional
odps.function.strictmode                false
odps.table.drop.ignorenonexistent       false
odps.instance.priority.level            3
odps.task.sql.write.str2null            false
odps.instance.priority.autoadjust       false
odps.table.lifecycle.value              37231
odps.task.sql.outerjoin.ppd             false
odps.optimizer.mode                     hbo
odps.instance.remain.days               30
READ_TABLE_MAX_ROW                      10000

Extended Properties:
tempDataLogicalSize                     3642
tempDataPhysicalSize                    10926
tableLogicalSize                        20530
usedQuotaPhysicalSize                   4162347
resourcePhysicalSize                    4043403
tempResourcePhysicalSize                0
tableBackupPhysicalSize                 38016
volumePhysicalSize                      0
volumeLogicalSize                       0
failoverPhysicalSize                    8412
tableBackupLogicalSize                  12672
failoverLogicalSize                     2804
tempResourceLogicalSize                 0
tablePhysicalSize                       61590
usedQuotaLogicalSize                    1387449
resourceLogicalSize                     1347801
```

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.

11.2.5. MaxCompute O&M

11.2.5.1. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   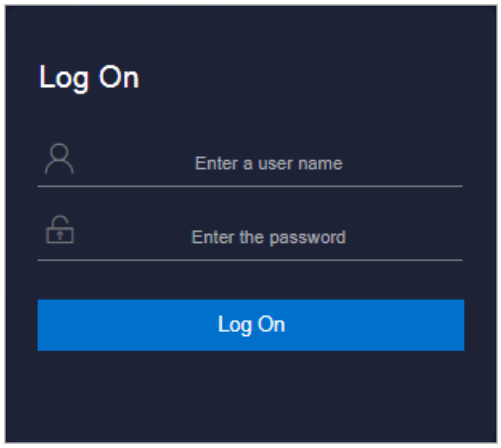

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.2.5.2. Business O&M

11.2.5.2.1. O&M overview and entry

This topic describes the business O&M features and how to go to the business O&M page.

Business O&M features

- Projects:
  - Project List: shows all projects and project details in a MaxCompute cluster. You can filter, query, and sort projects. You can also change the quota group of a project. If zone-disaster recovery is enabled, you can specify resource replication parameters and determine whether to enable resource replication for a project.
  - Authorize Package for Metadata Repository: 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 the switchover between the primary and secondary clusters. You can also determine whether to run scheduled tasks to synchronize resources between the primary and secondary clusters.
  - Project Migration: allows you to create, manage, and run project migration tasks, and view task details.
- Quota Groups: shows the quota groups of all projects in a MaxCompute cluster. It allows you to create quota groups, change quota groups, and view details about quota groups.
- Jobs: shows information about jobs in a MaxCompute cluster. You can filter and search for jobs. You can also view 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 the records of the tasks.
  - File Archiving: allows you to create file archive tasks for clusters and projects. You can also filter archive tasks and view the records of the tasks.
  - Resource Analysis: allows you to view the resource usage of the cluster from different dimensions.

Go to the business O&M page

1. Log on to the ABM console.
2. In the upper-left corner, click the  icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Business** tab. In the left-side navigation pane, choose **Projects > Project List**.

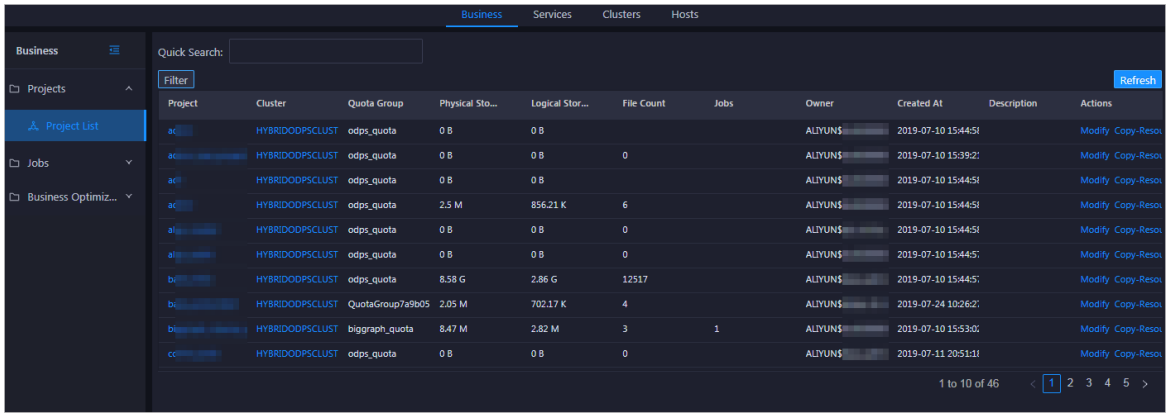

11.2.5.2.2. Project management

11.2.5.2.2.1. Project list

The Project List page shows all projects and project details in a MaxCompute cluster. You can filter, query, and sort projects. You can also change the quota group of a project. If zone-disaster recovery is enabled, you can specify resource replication parameters and determine whether to enable resource replication for a project.

Go to the Project List page

In the left-side navigation pane of the **Business** tab, choose **Projects > Project List** to view projects in a cluster.

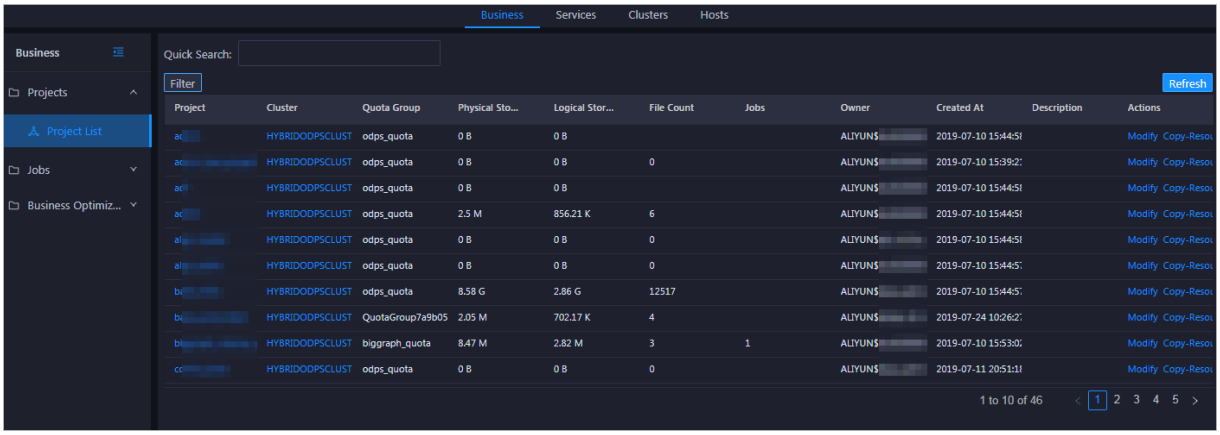

The **Project List** page shows the detailed information about all projects in a cluster. You can view the name, cluster, used storage, storage quota, storage usage, number of files, owner, and creation time of a project.

View project details

On the **Project List** page, click the name of a project to view its details. You can view the project overview, jobs, storage, configuration, quota group, and tunnel, as well as information about resource analysis and cross-cluster replication. For more information, see MaxCompute workbench. You can also grant access permissions on the metadata warehouse to project members and encrypt data of the project. For more information, see Grant access permissions on the metadata warehouse and Encrypt data.

Change a quota group

You can change the default quota group of a project.

1. On the **Project List** page, find the target project, click **Actions** in the Actions column, and select **Change Default Quota Group**. In the **Change Default Quota Group** pane, specify the required parameters.

   Parameters:

   - **Region**: the region of the project.
   - **Cluster**: the default cluster of the project. If the project belongs to multiple clusters, select a cluster from the drop-down list to serve as the default cluster.
   - **Quota Group**: the quota group to which the project belongs. To change the quota group, select a quota group from the drop-down list.
2. After you specify the parameters, click **Run**.

Modify the storage quota

You can modify the storage quota of a project.

1. On the **Project List** page, find the target project, click **Actions** in the Actions column, and select **Modify Storage Quota**. In the **Change Storage Quota** pane, specify the required parameters.

   Parameters:

   - **Region**: the region of the project
   - **Project**: the name of the project for which you want to modify the storage quota
   - **Cluster**: the default cluster of the project
   - **Target Storage Quota (TB)**: the new storage quota
   - **Reason**: the cause for the modification
2. After you specify the parameters, click **Run**.

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 configure data synchronization rules for the project to regularly synchronize data such as table data to a secondary cluster.

1. On the **Project List** page, find the target project, click **Actions** in the Actions column, and select **Resource Replication**. In the **Copy Resource** pane, specify the required parameters.

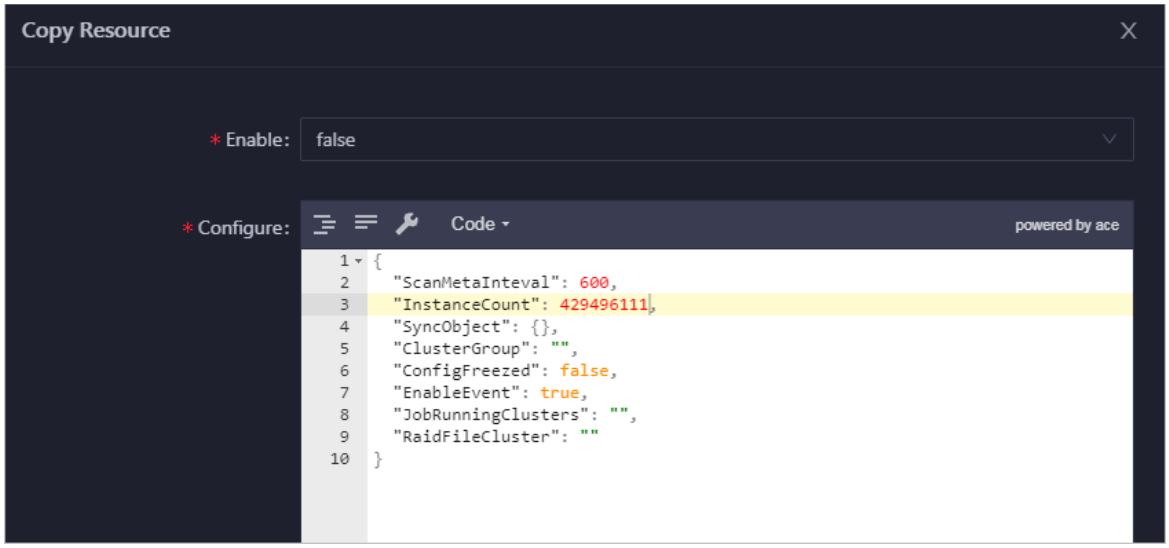

Parameters:

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

2. After you modify code in the **Configure** field, click **Compare Versions** to view the differences, which are highlighted.

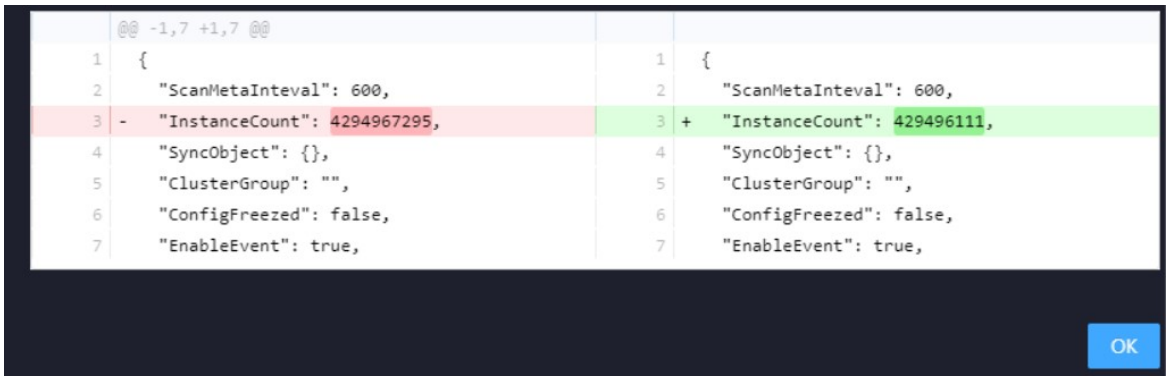

3. Click **Run**.

11.2.5.2.2.2. Project details

The Apsara Bigdata Manager (ABM) console shows your MaxCompute projects and project details. You can view the project overview, jobs, storage, configurations, quota groups, and tunnels, as well as information about resource analysis, storage encryption, and cross-cluster replication.

Go to the project details page

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Business** tab. The **Project List** page appears by default. Click the name of a project to view its details.

Overview

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

- Basic information, such as the default quota group, creator, creation time, service, and region
- Trend charts that show the trend lines of requested and used CPU and memory resources by minute in different colors
- Trend chart that shows the trend lines of CPU utilization and memory usage by day in different colors

Jobs

On the **Jobs** tab, 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 CPU utilization, maximum CPU utilization, minimum memory usage, maximum memory usage, DataWorks node, running status, start time, priority, and type. You can also view the operational logs of a job to locate its running faults.

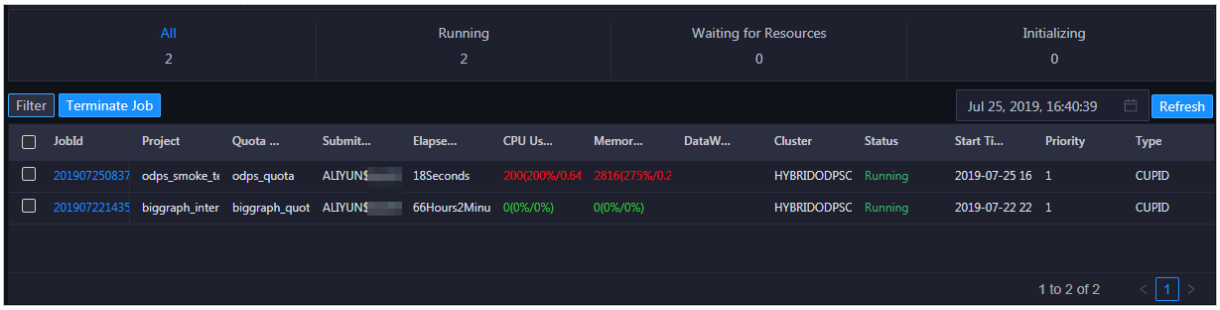

You can perform the following operations on the Jobs tab:

- Customize columns or sort job snapshots by column.
- View the operational logs of jobs or terminate jobs.

Storage

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

Note The Storage tab shows only information about storage resources. To query information about computing resources, go to the Quota Groups tab.

Configuration

On the **Configuration** tab, you can configure the general, sandbox, SQL, MapReduce, access control, and resource recycling properties of the project. You can configure package-based authorization to allow access to the metadata warehouse.

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

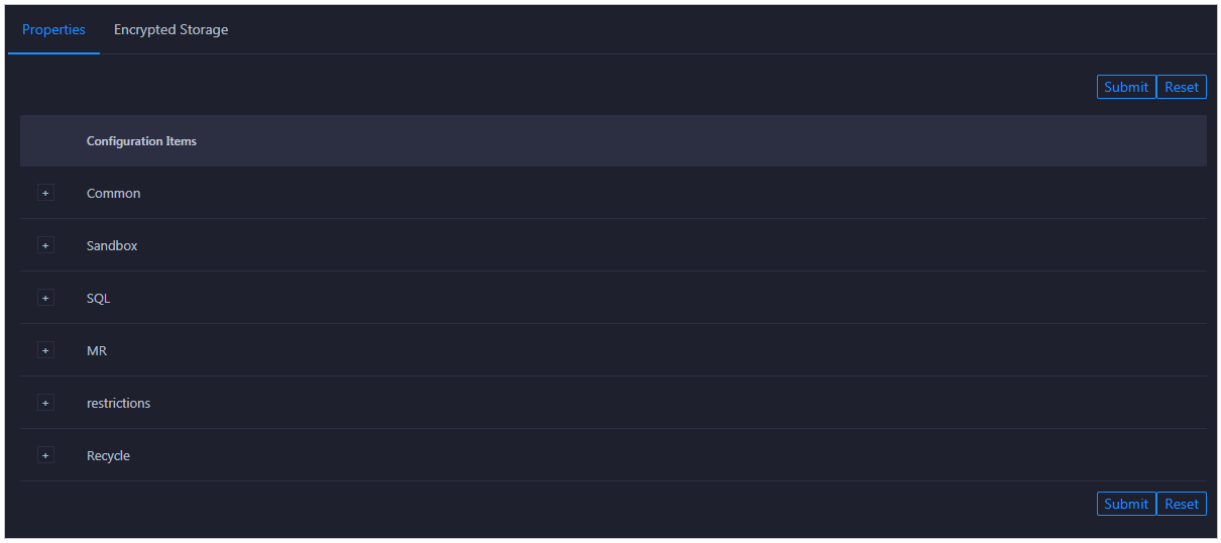

On the **Authorize Package for Metadata Repository** tab, you can install the package and perform package-based authorization.

Quota Groups

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

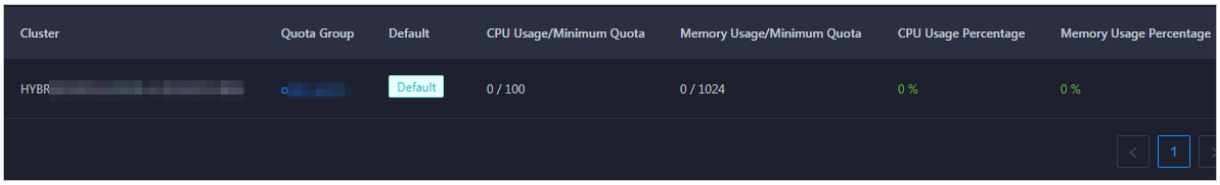

To view details about a quota group, click the quota group name in the **Quota** column. For more information, see Manage quota groups.

> **Note** The Quota Groups tab shows only information about computing resources. To query information about storage resources, go to the Storage tab.

Tunnel

On the **Tunnel** tab, you can view the tunnel throughput of the project in the unit of bytes per minute. The Tunnel Throughput (Bytes/Min) chart shows the trend lines of inbound and outbound traffic in different colors.

Resource Analysis

On the **Resource Analysis** tab, you can view the resource usage of the project from different dimensions, including tables, tasks, execution time, start time, and engines.

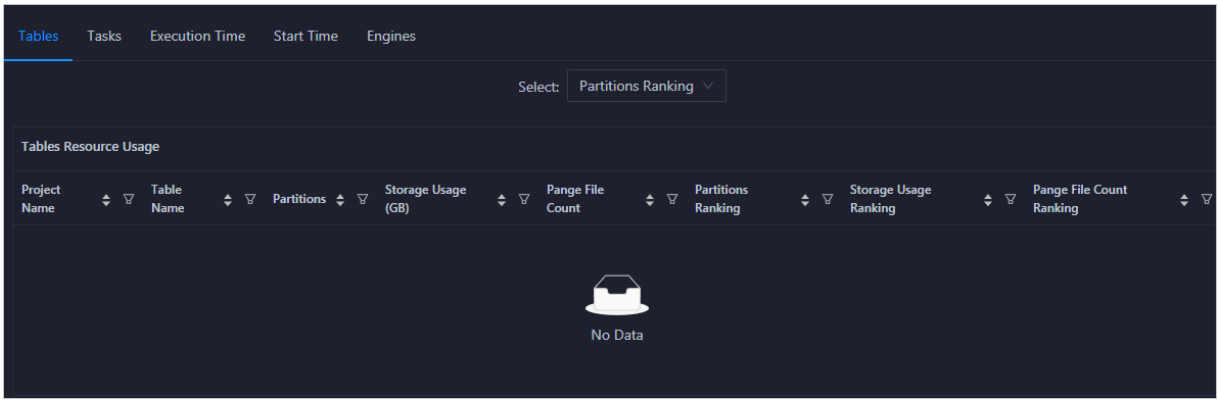

Encryption at Rest

On the **Encryption at Rest** tab, you can encrypt data by using the following encryption algorithms: AES-CTR, AES256, RC4, and SM4.

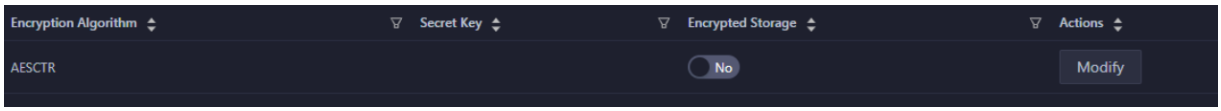

Cross-cluster Replication

On the **Cross-cluster Replication** tab, you can view the projects that have the cross-cluster replication feature enabled and the details and status of cross-cluster replication.

When you deploy multiple clusters to use MaxCompute, MaxCompute projects may be mutually dependent. In this case, data may be 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 threshold, MaxCompute adds the tables or partitions to the cross-cluster replication list.

Assume that Project1 in Cluster A depends on Table1 of Project2 in Custer B. In this case, Project1 directly reads data from Table1. If the duration of direct data reading reaches the specified threshold, MaxCompute adds Table1 to the cross-cluster replication list.

The **Cross-cluster Replication** tab consists of the **Replication Details** and **Replication Configuration** tabs.

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

11.2.5.2.2.3. Encrypt data

You can specify whether to encrypt the data stored in MaxCompute projects.

Prerequisites

If MaxCompute V3.8.0 or later is deployed, storage encryption is supported by default. If MaxCompute is upgraded to V3.8.0 or later, storage encryption is not supported by default. If you want to enable storage encryption, complete the configuration for your MaxCompute cluster.

Context

After storage encryption is enabled for a project, it cannot be disabled. After storage encryption is enabled, only the data that is newly written to the project is automatically encrypted. To encrypt historical data, you can create rules and configure tasks.

Before you encrypt historical data for a project, make sure that you understand the concepts of rules and tasks in Apsara Bigdata Manager (ABM). A rule is used to specify the time period of historical data that you want to encrypt in a specific project. After you create a rule, 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, only the latest rule 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. In the upper-left corner, click the  icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Business** tab. In the left-side navigation pane, choose **Projects > Project List**.

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 that is newly written to the project is automatically encrypted.

   i. On the **Encrypt** tab, click **Modify** in the Actions column. In the **Configure Encrypted Storage** pane, specify the **Encryption Algorithm**, **region**, and **project** parameters.

   > **Note** AES-CTR, AES256, and RC4 encryption algorithms are supported.

   ii. 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, perform the following steps:

   i. 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 message that appears, click **Run**. 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**.

   ii. Create a key rotate task.

   On the **Configure Task** tab, click **Add a key rotate task**. In the **Edit Key Rotate Task** pane, specify the required parameters and click **Run**.

Parameter	Description
Region	The region where the project whose data is to be encrypted resides. Select a region from the drop-down list.
Project Name	The name of the project whose 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 small value indicates a high priority.
Enabled	Specifies whether the task is enabled.
Bandwidth Limit	Specifies whether to limit the concurrency of merge tasks for the project. ■ **Yes**: indicates that merge tasks cannot be concurrently run. ■ **No**: indicates that merge tasks can be concurrently run.
Maximum Concurrent Tasks	The maximum number of merge tasks that can be run for the cluster of the selected project at the same time. 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 for the cluster of the selected project at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the cluster of the selected project, not only the merge tasks.

Merge Parameters	<pre>{
	"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"
	}</pre>

8. (Optional)View the history of data encryption in the project.

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

11.2.5.2.2.4. Grant access permissions on the metadata warehouse

You can grant access permissions on the metadata warehouse to projects and project members.

Prerequisites

- If MaxCompute V3.8.1 or later is deployed, the package of the metadata warehouse is installed by default. In this case, you can directly use Apsara Bigdata Manager (ABM) to grant access permissions on the metadata warehouse. If MaxCompute is upgraded to V3.8.1 or later, the package of the metadata warehouse is not installed by default. Before you grant access permissions on the metadata warehouse, you must manually install the package of the metadata warehouse.
- A project is created in DataWorks.

Context

To allow a project to access the metadata warehouse, grant the required permissions to the project and install the package to the project in the ABM console. When you install the package, ABM retrieves authentication information, such as the AccessKey pair, 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. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Business** tab. In the left-side navigation pane, choose **Projects > Project List**.
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. Then, click the **Authorize Package for Metadata Repository** tab.
6. Click **Authorize** in the Actions column. In the **Authorize Package** dialog box, click **Run**. A message appears, indicating that the permission is granted.
7. Click **Install** in the Actions column. In the **Install Package** dialog box, click **Run**. A message appears, indicating that the package is installed.After the package is installed, the switch in the **Authorized** column is turned on.

11.2.5.2.2.5. Perform disaster recovery

When a primary MaxCompute cluster fails, you can perform a primary/secondary switchover in the Apsara Bigdata Manager (ABM) console to restore services. This topic describes the prerequisites and procedure of disaster recovery. In this topic, disaster recovery indicates zone-disaster recovery.

Prerequisites

- The resource replication feature is disabled in the ABM console. To disable the feature, perform the following steps:
  i. Log on to the ABM console.
  ii. In the upper-left corner, click the icon and then **MaxCompute**.
  iii. In the left-side navigation pane of the **Business** tab, click **Projects** and then **Disaster Recovery**.
  iv. On the page that appears, turn off **Resource Synchronization Status**.
- The domain name of ABM is pointed to the IP address of the secondary ABM cluster. To point the domain name to the IP address, perform the following steps:
  i. Log on to the ABM console.
  ii. In the upper-left corner, click the icon and then **MaxCompute**.
  iii. On the MaxCompute page, click **Management** in the top navigation bar. In the left-side navigation pane of the page that appears, click **Jobs**. The **Jobs** tab appears by default.
  iv. Find the Change Bcc Dns-Vip Relation For Disaster Recovery job and click **Run** in the Actions column. The **Job Properties** section appears.
  v. Click the icon next to Group Name to configure the IP address of the Docker container.

> **Note** NewBccAGIp indicates the IP address of the Docker container under AG# for the bcc-saas service of the secondary ABM cluster. You must configure an IP address at the #Docker# level.

In the dialog box that appears, click the **Servers** tab. Enter the IP address of a server in the field and click **Add Server**. Then, click **OK**. The IP address is configured.

vi. In the upper-right corner, click **Run**. In the message that appears, click **Confirm**.

vii. On the page that appears, click **Start**. The switchover starts.

> **Note** If a step fails, click **Retry**. After all the steps are complete, the domain name of ABM is pointed to the IP address of the secondary ABM cluster.

- The secondary ABM cluster page is accessible. If this page is inaccessible, go to the */usr/loca/bigdatak/controllers/bcc/tool/disaster_recovery* directory of the Docker container in bcc-saa.AG# of the secondary ABM cluster. Then, execute the `/home/tops/bin/python change_dns_vip.py` script in the directory. If job_success appears, the execution succeeds. Then, execute the `/home/tops/bin/python disaster_init.py` script in the current directory. If job_success appears, the execution succeeds. After the scripts are successfully executed, you can go to the secondary cluster page.

  > **Note** If an exception occurs when you execute the scripts, click **Retry**.

- The Business Continuity Management Center (BCMC) switchover of MaxCompute is complete. The services on which MaxCompute depends are running normally. The services include AAS, Tablestore, and MiniRDS.
- By default, the data synchronization feature is disabled for MaxCompute projects because the computing and storage resources of the primary and secondary data centers are limited. To enable the data synchronization feature, submit a ticket.

Context

Take note of the following points for a disaster recovery switchover:

- By default, the logon to Apsara Bigdata Manager depends on the Apsara Stack Operations (ASO) console. If the ASO console has not reached the desired state, single sign-on is not supported. In this case, go to the */usr/loca/bigdatak/controllers/bcc/tool/disaster_recovery* directory of the Docker container in bcc-saa.AG#. Then, execute `change_login_by_bcc.sh` to switch the logon mode to the mode that is independent of the ASO console. After the ASO console has reached the desired state, execute `change_login_by_aso.sh` to switch the logon mode back to the mode that depends on the ASO console.
- An exception may occur in each step of the switchover process. If an exception occurs, click **Retry**. If the retry succeeds, proceed to the next step. If the exception persists after multiple retries, contact O&M engineers to perform troubleshooting. Then, click **Retry** to complete the step.
- For each switchover, the Apsara distributed operating system of the original primary MaxCompute cluster must be restarted. Otherwise, the admintask service may be faulty after the switchover is complete.
- In the Collect Unsynchronized Data step, an exception shown in the following figure may occur. If this occurs, click **Recollect Unsynchronized Data**.

Procedure

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

3. In the left-side navigation pane of the **Business** tab, click **Projects** and then **Disaster Recovery**.
4. In the upper-right corner, click **Switchover Process** to start the disaster recovery process.
5. Wait for resource replication to automatically stop.

   Wait for resource replication to automatically stop. After **Next** becomes blue, click **Next**.

   > **Note** If an error occurs, click **Retry**. If the retry is invalid, contact O&M engineers to perform troubleshooting and try again.

6. Switch control clusters.
   i. Wait for the primary/secondary switchover to complete for control clusters.

      > **Note** After the original primary cluster becomes the secondary cluster, the switchover is complete.

   ii. Click **Restart Standby Cluster**.

      > **Note** The MaxCompute clusters become abnormal.

   iii. After the MaxCompute clusters become normal, click **Restart Frontend Server** and wait until the restart result is returned.
   iv. After the restart succeeds, click **Test adminTask**.

      > **Note** If an exception occurs, click **Retry** and then **Test adminTask**. Alternatively, repeat from Step 6.b.

   v. After **Next** becomes blue, click **Next**.

      > **Note** The Switching message remains displayed until the test succeeds.

7. Switch computing clusters.

   The computing cluster switchover automatically starts for the projects that have two computing clusters. The switchover cannot be performed for the projects that have only one computing cluster. After the switchovers are completed for all the projects, click **Next**.

   > **Note** If the computing clusters of a project fail to be switched, contact O&M engineers to identify the cause of the exception. If the exception can be fixed, fix it and click **Retry** to continue the switchover. If the project is damaged or does not need a cluster switchover, click **Next** after you confirm that computing clusters of other projects are switched.

8. Switch the replication service to the secondary clusters.

   The script is automatically executed at the background. When a success message appears, click **Next**.

9. Collect unsynchronized data.

i. Wait for the system to collect statistics on projects that contain unsynchronized data.

> **Note** This step requires a long time to complete. The specific time depends on the data volume.

ii. After the collection is complete, click **Download Unsynchronized Data of Selected Projects** to download the unsynchronized data to your computer.

> **Note** The unsynchronized data that is obtained from this step is required for the Manually Fill in Missing Data step. The projects that are obtained from this step must be the same as those for the Repair Metadata and Manually Fill in Missing Data steps.

iii. After the unsynchronized data is downloaded, verify the data and click **Next**. If all data is synchronized, click **Next**.

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

10. Repair metadata.

    Select all projects, click **Repair Metadata of Selected Projects**, and then wait for results. If the metadata of some projects fails to be repaired, click **Download Last Execution Log** and send the logs to O&M engineers. The logs can be used to identify and analyze the cause of the exception. After the exception is fixed, repair the metadata of the projects again. If you do not need to repair the metadata of all projects, click **Next** after the metadata of required projects is repaired.

11. Manually supplement missing data.

    Use DataWorks or the odpscmd client to manually supplement the missing data based on the unsynchronized data that you downloaded. After you supplement the missing data, select all projects and click **Confirm Data Repair Complete**. Then, click **Next**.

12. Repair unsynchronized resources.

    i. Wait for the system to collect statistics on projects that contain unsynchronized resources.

    > **Note** This step requires a long time to complete. The specific time depends on the data volume.

    ii. Use DataWorks or the odpscmd client to manually supplement the missing resources based on the unsynchronized resources that you collected. If an exception occurs, send exception information to O&M engineers to perform troubleshooting. After all the project resources are repaired, click **Complete and Next**.

13. Wait for resource replication to automatically start.

    Wait for resource replication to automatically start. After **Next** becomes blue, click **Next**.

14. Exit the configuration wizard.

    After the switchover is complete, click **Back** to exit the wizard.

11.2.5.2.2.6. Migrate projects

Apsara Bigdata Manager (ABM) allows you to migrate MaxCompute projects across regions from one cluster to another. This allows you to balance the computing and storage resources of each cluster.

> **Note** The project migration feature is supported only when the clusters are deployed in multi-region mode.

Create a project migration task

1. In the left-side navigation pane of the **Business** tab, click **Projects** and then **Project Migration**.
2. In the upper part of the **Migration Mission** page, select the region where the project resides.

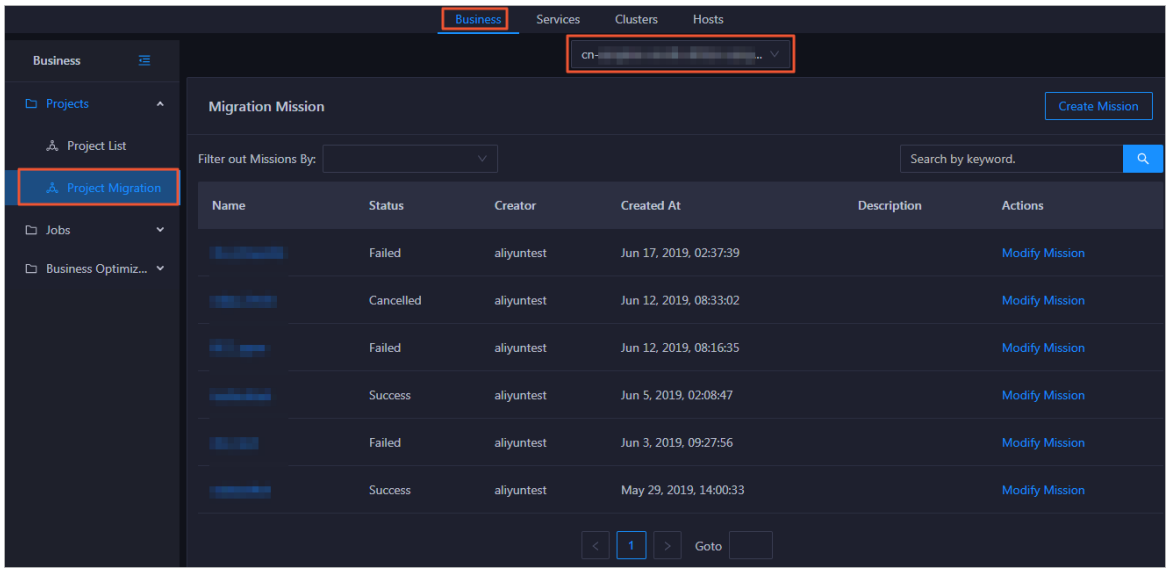

3. In the upper-right corner, click **Create Mission**. On the page that appears, specify the parameters in the **General**, **Source**, **Target Selection**, and **Cluster for Mission Execution** sections as prompted.

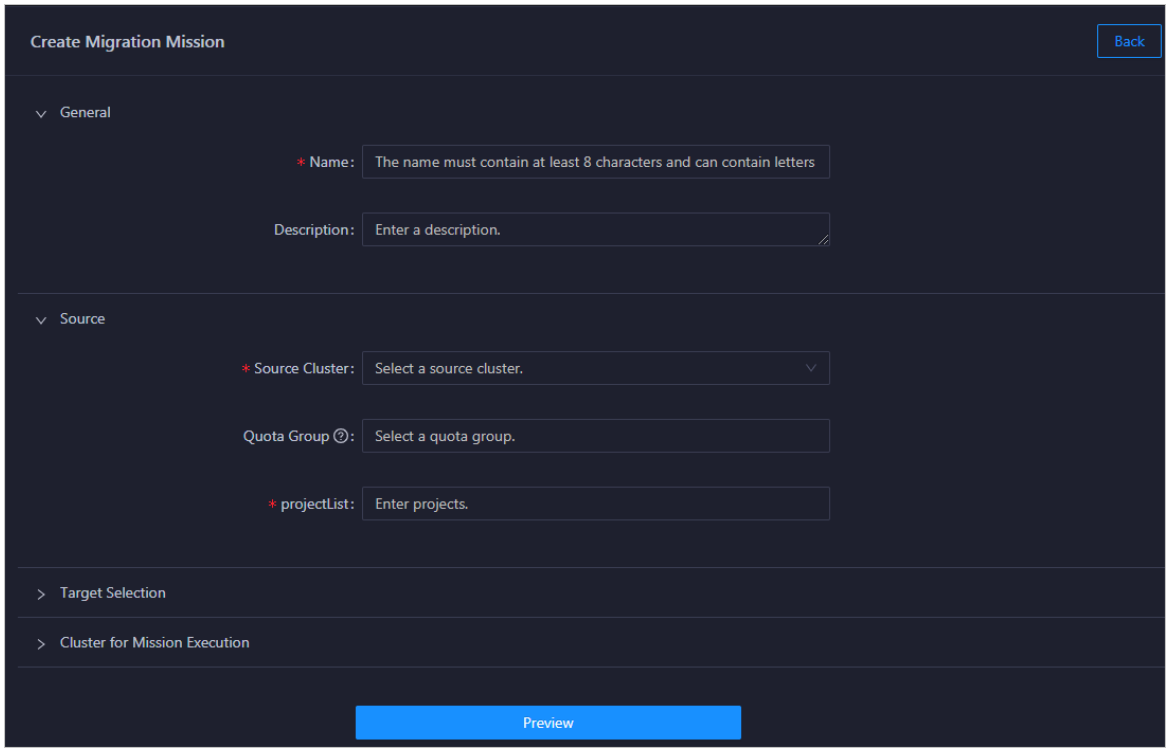

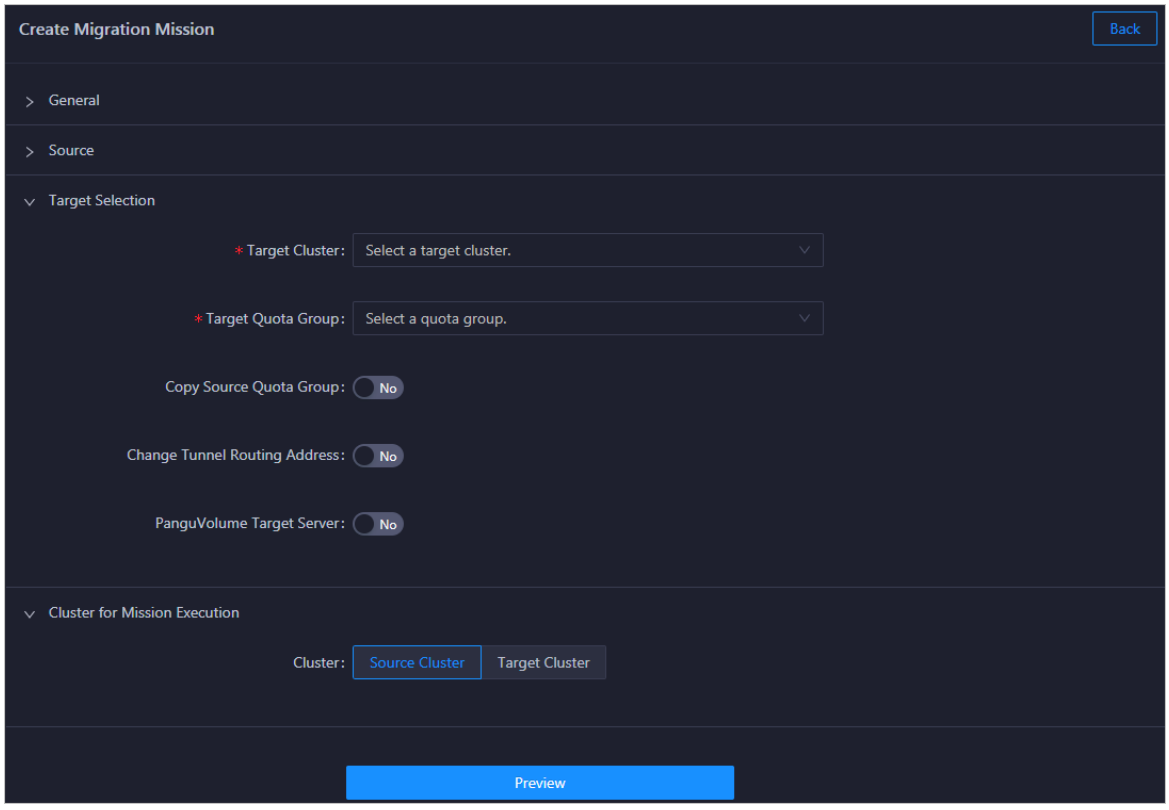

The following table describes the parameters.

Section	Parameter	Description
	Source Cluster	The name of the source cluster. Select a cluster from the drop-down list.

Section	Parameter	Description
Source	Quota Group	The quota group of the source cluster. Select a quota group from the drop-down list.
	projectList	The projects that you want to migrate. After **Quota Group** is specified, all the projects in the quota group are automatically loaded. You can migrate these projects at a time. If some projects in the quota group do not need to be migrated, you can remove the projects.
Target Selection	Copy Source Quota Group	Specifies whether the target cluster uses the same quota group as the source cluster. If you enable this feature, the **Target Quota Group** parameter cannot be specified.
	Change Tunnel Routing Address	Specifies whether to use a new Tunnel route. Tunnel provides highly concurrent upload and download services for offline data. Each project has a default Tunnel route. If you want to use a new Tunnel route after a project is migrated to a new cluster, enable **Change Tunnel Routing Address** and specify the new Tunnel route.
	PanguVolume Target Server	Specifies whether the target Apsara Distributed File System volume can be specified. Cross-volume project migration is not supported. Set this parameter to **No**.
Cluster for Mission Execution	Cluster	○ **Source Cluster**: indicates that the source cluster pushes the project to the destination cluster. ○ **Target Cluster**: indicates that the destination cluster pulls the project from the source cluster.

4. Click **Preview** to preview project migration details.

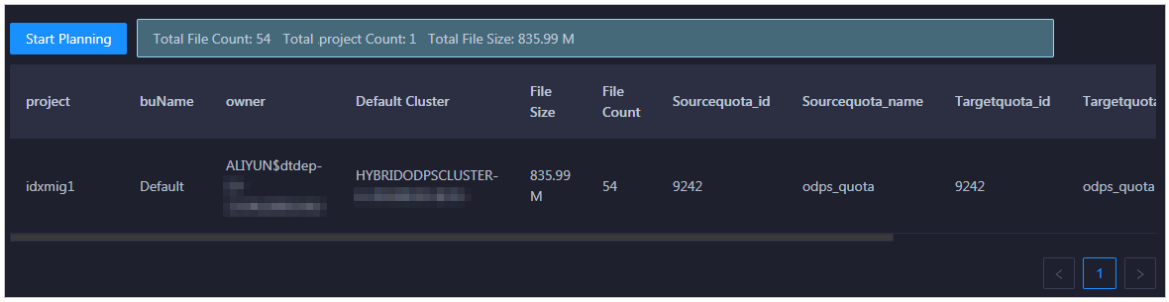

5. After you confirm the configuration, click **Start Planning** in the upper-left corner. A project migration task is generated. The migration details appear.

   It requires some time to generate the task.

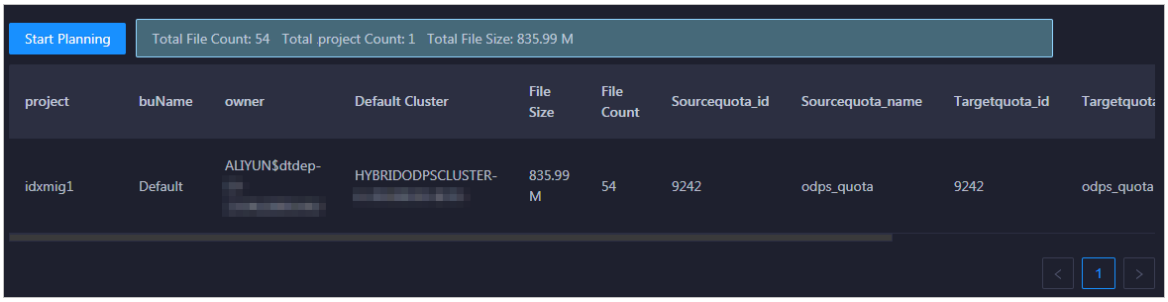

A standard project migration task generally includes five steps:

i. **Add Target Cluster**: Add the destination cluster to the cluster list of the project that you want to migrate.

ii. **Start to Replicate**: Replicate the project from the source cluster to the destination cluster.

iii. **Switch Default Cluster**: Change the default cluster of the project to the destination cluster. After the default cluster is changed, generated data is written to the destination cluster.

iv. **Clear Replication**: Clear the data replication list. During project migration, the migrated project in the source cluster and the corresponding project in the destination cluster synchronize data based on the data replication list. This ensures data consistency between the two projects. Data is continuously synchronized until the data replication list is cleared.

v. **Remove Source Cluster**: Delete the migrated project from the source cluster.

For more information about how to modify a task after it is generated, see Modify a project migration task.

Run the project migration task

After the project migration task is created, you can run the task on the **Migration Details** page.

1. Click the task name in the task list to go to the **Migration Details** page.
2. On the **Migration Details** page, click **Submit for Execution**.

   After the project migration task starts, the system automatically runs the **Add Target Cluster** and **Start to Replicate** steps in sequence.

   If you migrate multiple projects at a time, the process requires many steps to complete. Therefore, we recommend that you sort the steps by project to view the migration steps for each project. If the status of a step is **Success**, the step is complete. If the status of a step is **Failed**, the step fails.

   In the migration process, some steps can be run only after you click **OK**. If you do not need to run a step, click **Skip**. To confirm or skip multiple steps at a time, select the steps and click **OK** or **Skip** in the upper-left corner.

   You can also click the status of a migration step for a project. In the dialog box that appears, click **Yes** to skip the remaining steps.

3. When the **Start to Replicate** step is complete, check the difference in data volumes between the migrated project in the source cluster and the corresponding project in the destination cluster.

   > **Notice** We recommend that you run the next step only when the difference in data volumes does not exceed 5%.

To check the data volume of a project, log on to the admin gateway host in the cluster where the project resides and run the **pu dirmeta /product/aliyun/odps/${project_name}/** command.

4. If the difference in data volumes does not exceed 5%, perform either of the following operations:
   - Change the default cluster: Click **OK** in the Actions column of the **Switch Default Cluster** step. After this operation, the destination cluster becomes the default cluster of the migrated project. The default cluster is changed in this example.
   - Do not change the default cluster: Click **Skip** in the Actions column of the **Switch Default Cluster** step. After this operation, the source cluster is still used as the default cluster of the project.

   After the default cluster is changed, generated data is written to the destination cluster.

   > **Warning** During project migration, the migrated project in the source cluster and the corresponding project in the destination cluster synchronize data based on the data replication list to ensure data consistency. It requires some time for data synchronization to complete. Therefore, after the default cluster is changed, we recommend that you wait for about one week before you proceed to the next step.

5. Wait for about one week and check whether the data volume of the migrated project in the source cluster is the same as that of the corresponding project in the destination cluster.

   To check the data volume of a project, log on to the admin gateway host in the cluster where the project resides and run the **pu dirmeta /product/aliyun/odps/${project_name}/** command.

   > **Warning** Before you proceed to the next step, make sure that the data volume of the migrated project in the source cluster is the same as that of the corresponding project in the destination cluster. Otherwise, data may be lost.

6. To retain the migrated project in the source cluster, click **Skip** in the Actions column of the **Remove Source Cluster** step before you run the **Clear Replication** step.
7. After the data volume of the migrated project in the source cluster becomes the same as that of the project in the destination cluster, click **OK** in the Actions column of the **Clear Replication** step to clear the data replication list.

   After the data replication list is cleared, data is no longer synchronized between the migrated project in the source cluster and the corresponding project in the destination cluster.

   The system automatically runs the **Remove Source Cluster** step to delete all migrated projects from the source cluster. This releases storage and computing resources.

View migration details

You can view the details of a project migration task, including the steps, results, and debugging information.

1. If multiple migration tasks exist, search for a task or filter tasks on the **Migration Mission** page.
   - Filter tasks: Select a task state from the **Filter out Mission By** drop-down list. All tasks in this state are automatically filtered from the migration task list.
   - Search for a task: Enter the name of a migration task in the search box in the upper-right corner and click the search icon to search for the task.

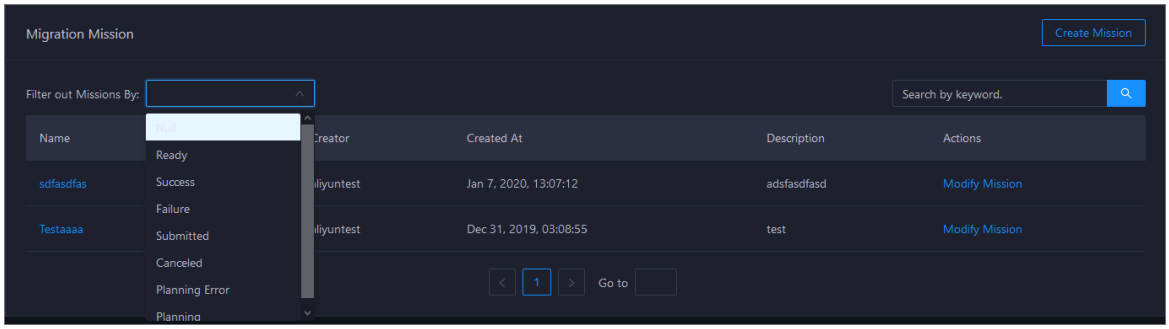

2. Click the name of a task. On the **Migration Details** page, view the details of the task.

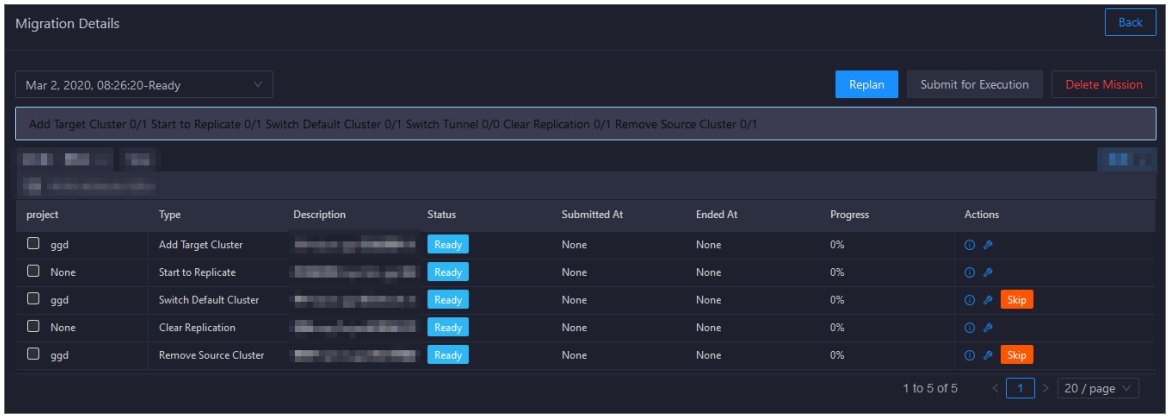

3. If a step fails, click the **Details** or **Debugging** icon in the Actions column to view the details or debugging information of the step. This allows you to identify the cause of the failure.
4. Perform other required operations.

   Click **Menu** in the upper-right corner. You can export the step list, change the column width to automatically fit the contents, or customize whether to show or hide a column.

   You can also right-click a cell in the step list and copy the cell content.

View step details and debugging information

If a step fails, you can view the step details and debugging information to identify the cause of the failure.

1. Find the step that fails to run during the migration of a project.

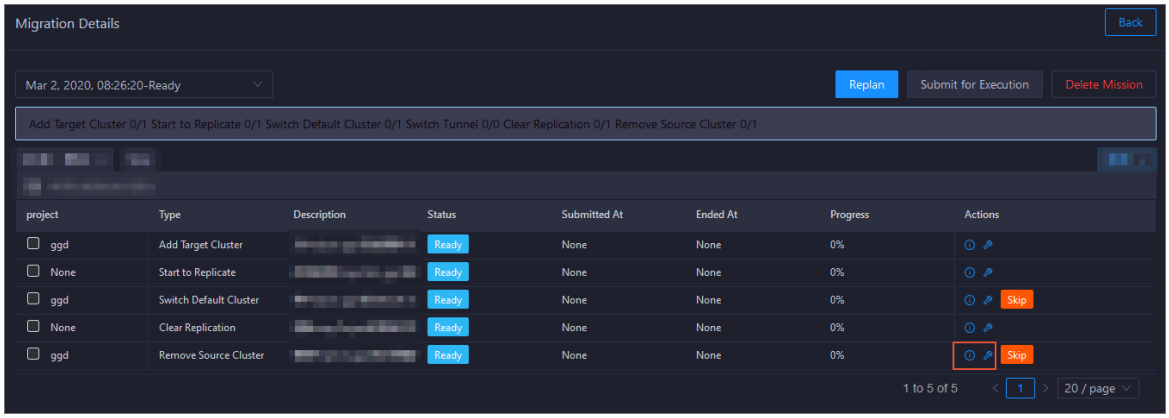

2. Click the **Details** icon in the Actions column to view the details of the step.

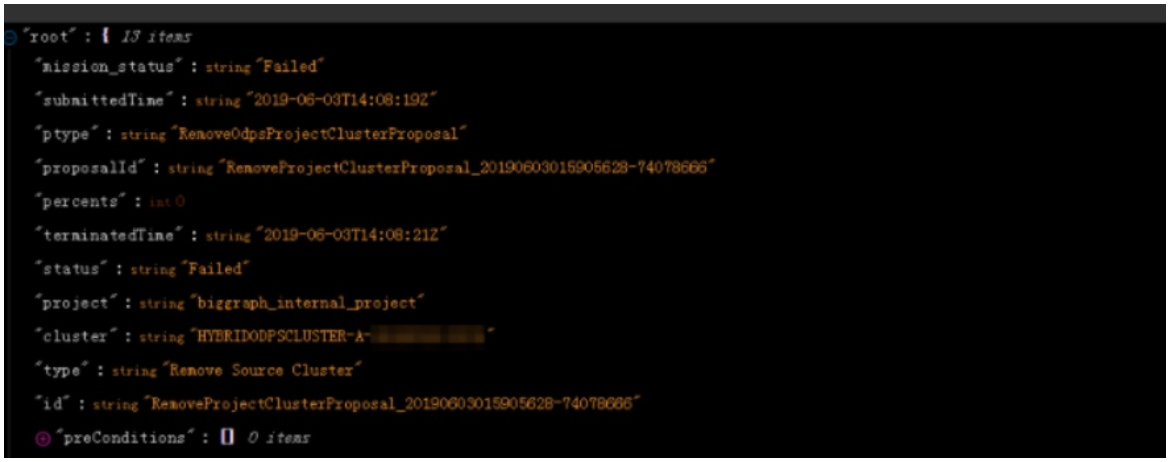

3. Click the **Debugging** icon in the Actions column to view the debugging information of the step.

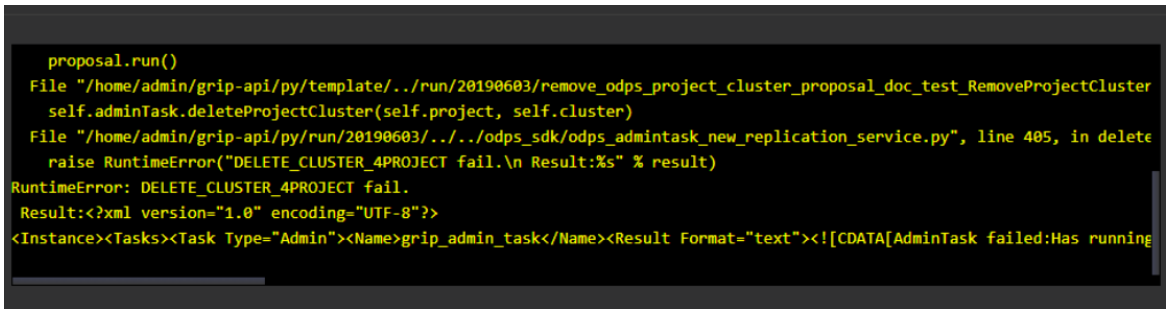

Modify a project migration task

After a project migration task is created, you can modify the task if the task does not meet your requirements.

To modify the task, find the target task, click **Modify Mission** in the Actions column, or click **Replan** on the **Migration Details** page.

11.2.5.2.3. Manage quota groups

Apsara Bigdata Manager (ABM) shows the quota groups of all projects in a MaxCompute cluster. It allows you to create quota groups, modify quota groups, and view details about quota groups.

Go to the Quota Groups page

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Business** tab. In the left-side navigation pane, click **Quota Groups** and then **Quota Groups**.

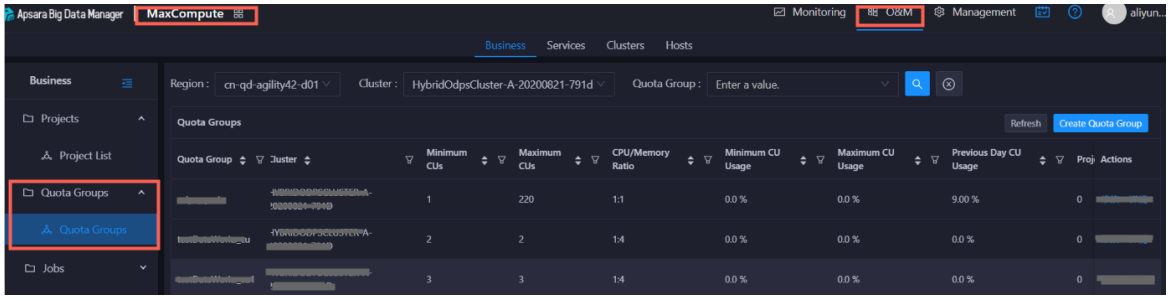

Create a quota group

In the upper-right corner of the **Quota Groups** page, click **Create Quota Group**. In the pane that appears, specify the required parameters and click **Run**.

Parameters

Parameter	Description
Quota Name	The name of the quota group.
Strategy	The policy of the quota group. Valid values: NoPreempt and Preempt.
Scheduler Type	The type of resource scheduling. Valid values: Fifo and Fair.
Minimum CUs	The minimum number of compute units that are provided by the quota group.
Maximum CUs	The maximum number of compute units that are provided by the quota group.
CPU/Memory Ratio	The ratio of CPUs to memory of servers in the quota group.

Modify a quota group

On the **Quota Groups** page, find the target quota group and click **Modify** in the Actions column. In the pane that appears, modify the settings and click **Run**.

View details about a quota group

On the **Quota Groups** page, find the target quota group and click **Details** in the Actions column. Then, you can view information about the resource usage and analysis of the quota group.

11.2.5.2.4. Job management

11.2.5.2.4.1. Job snapshots

The job snapshots feature allows you to manage the tasks that are created in MaxCompute and the merge tasks that are created in Apsara Bigdata Manager (ABM). You can also view job details by using Logview, terminate a job, and collect job logs.

View job snapshots

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

1. In the left-side navigation pane of the **Business** tab, click **Jobs** and then **Job Snapshots**. The **Job Snapshots** page appears.
2. In the upper-right corner, select the date and time to view job snapshots by day in the last week.
3. Click **All**, **Running**, **Waiting for Resources**, or **Initializing** to view job snapshots in the corresponding state on the specified date.
4. Find a target snapshot and click **Logview** in the Actions column. In the dialog box that appears, click **Run** to view the Logview of the job.

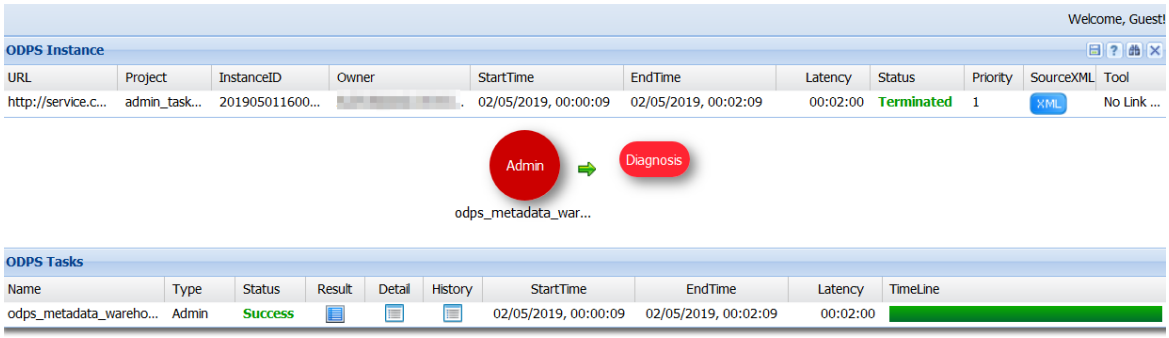

Terminate a job

1. In the left-side navigation pane of the **Business** tab, click **Jobs** and then **Job Snapshots**. The **Job Snapshots** page appears.
2. Select one or more jobs and click **Terminate Job** above the snapshot list. In the pane that appears, view information about the job or jobs that you want to terminate.

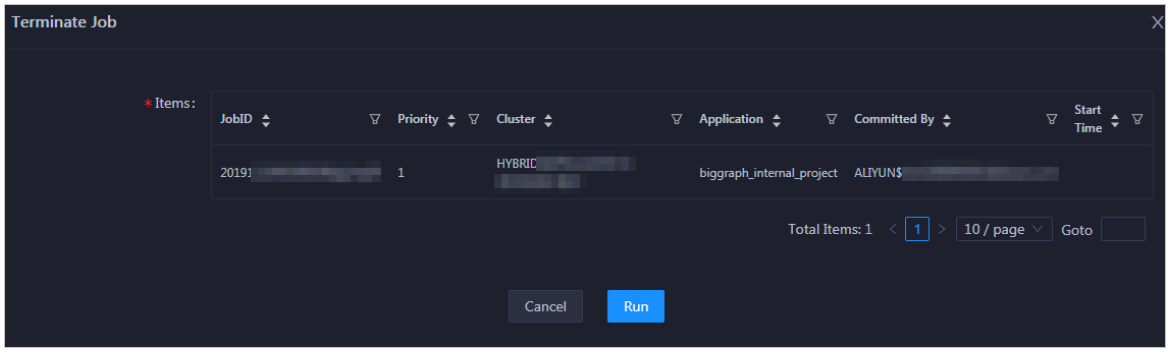

3. Click **Run**. A message appears, indicating the running result.

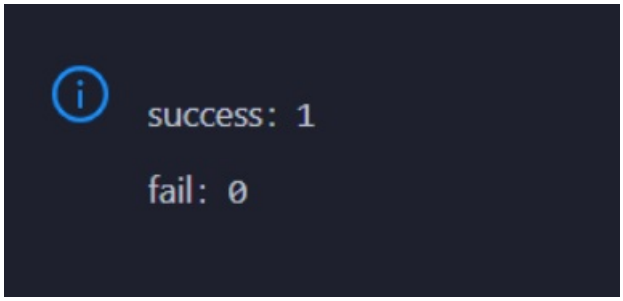

Collect job logs

When an exception occurs during job running, you can collect job logs to identify and analyze the exception.

1. In the left-side navigation pane of the **Business** tab, click **Jobs** and then **Job Snapshots**. The **Job Snapshots** page appears.
2. In the upper-right corner of the **Job Snapshots** page, click **Actions** and select **Collect Job Logs**.
3. In the **Collect Job Logs** pane, specify the required parameters.

   The following table describes the 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.
instanceid	Optional. The ID of the job instance.
requestid	Optional. The request ID returned when the job fails. If the value you specify is not a request ID, job logs that contain the specified value are collected.
Time Period	The time period to collect job logs.
Time Interval	Optional. The time interval to collect 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-right corner of the **Job Snapshots** page, click **Actions** and select **Execution History** next to **Collect Job Logs**. In the pane that appears, view the execution status and history of job log collection.

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails. If the status is **RUNNING**, click **Details** in the Actions column of a task to view the execution progress.
6. View the path to store job logs.

   In the **Execution History** pane, click **Details** in the Details column of an execution record to view the details. In the Steps section, view the path to store the job logs.

11.2.5.2.5. Business optimization

11.2.5.2.5.1. Merge small files

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

Create a file merge task for a cluster

If multiple small files exist in most projects of a MaxCompute cluster, you can create a task to merge these files in a centralized manner.

1. In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Merging**. The **Merge Tasks** tab appears.

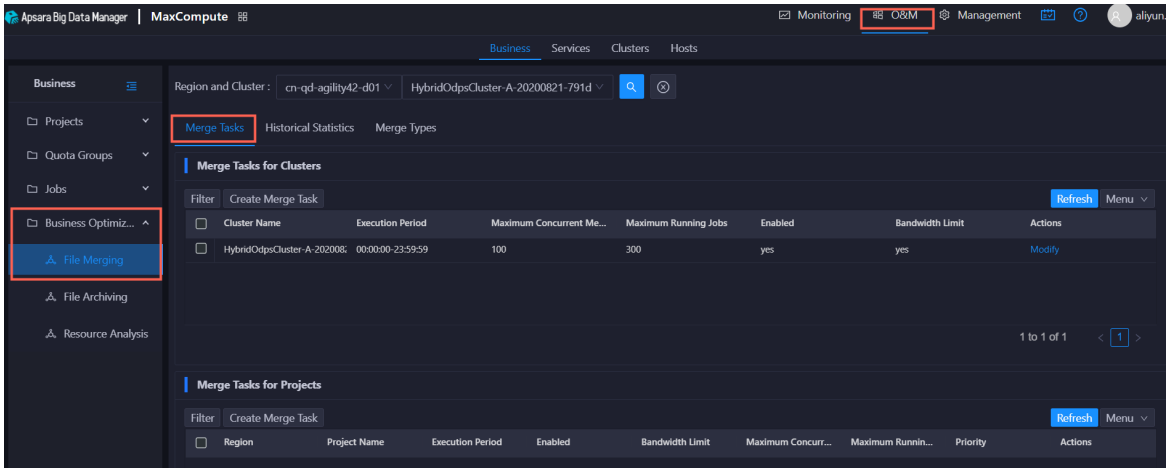

2. In the **Merge Tasks for Clusters** section, click **Create Merge Task**. In the pane that appears, specify the required parameters.

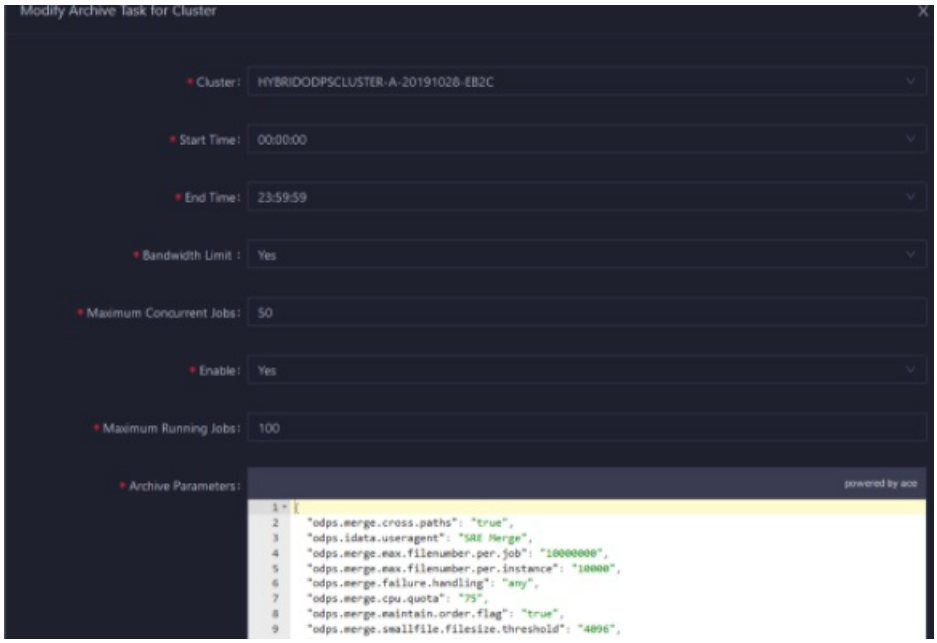

The following table describes the parameters.

Parameter	Description
Cluster	The cluster for which you want to run the merge task. Select a cluster from the drop-down list.
Start Time	The start time of the task.
End Time	The end time of the task.

Parameter	Description
Bandwidth Limit	Specifies whether to limit the concurrency of merge tasks for the cluster. ○ **Yes**: indicates that merge tasks cannot be concurrently run. ○ **No**: indicates that merge tasks can be concurrently run.
Maximum Concurrent Tasks	The maximum number of merge tasks that can be run for the selected cluster at the same time. This parameter is valid only when **Bandwidth Limit** is set to **No**.
Enabled	Specifies whether the task is enabled.

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 for the selected cluster at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the selected cluster, not only merge tasks.

3. Click **Compare Versions** below Merge Parameters to view the differences between the original and modified values.

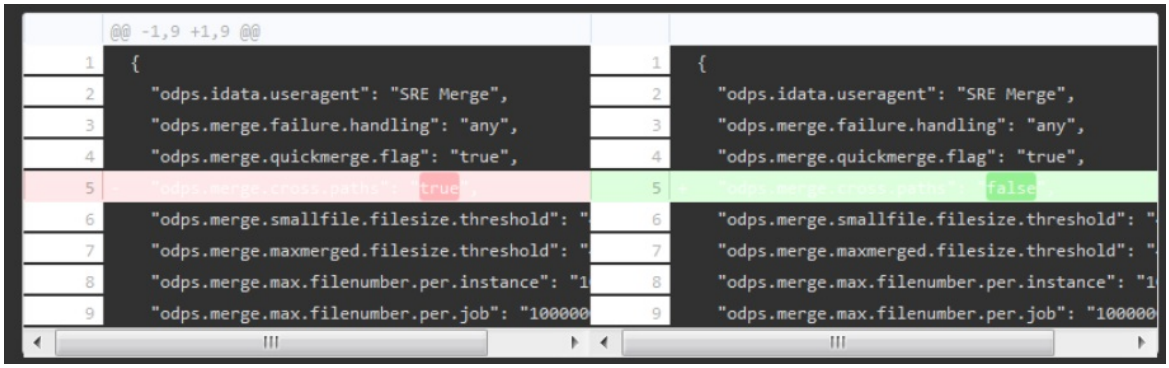

4. Click **Run**.

   The newly created merge task appears in the list of merge tasks for clusters.

Create a merge task for a project

If 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 a specific project.

1. In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Merging**. The **Merge Tasks** tab appears.

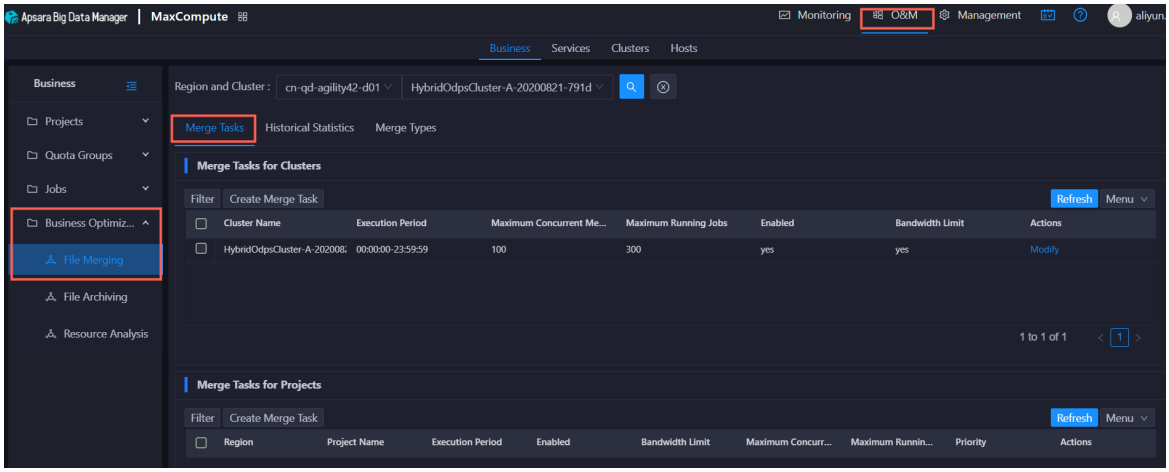

2. In the **Merge Tasks for Projects** section, click **Create Merge Task**. In the pane that appears, specify the required parameters.

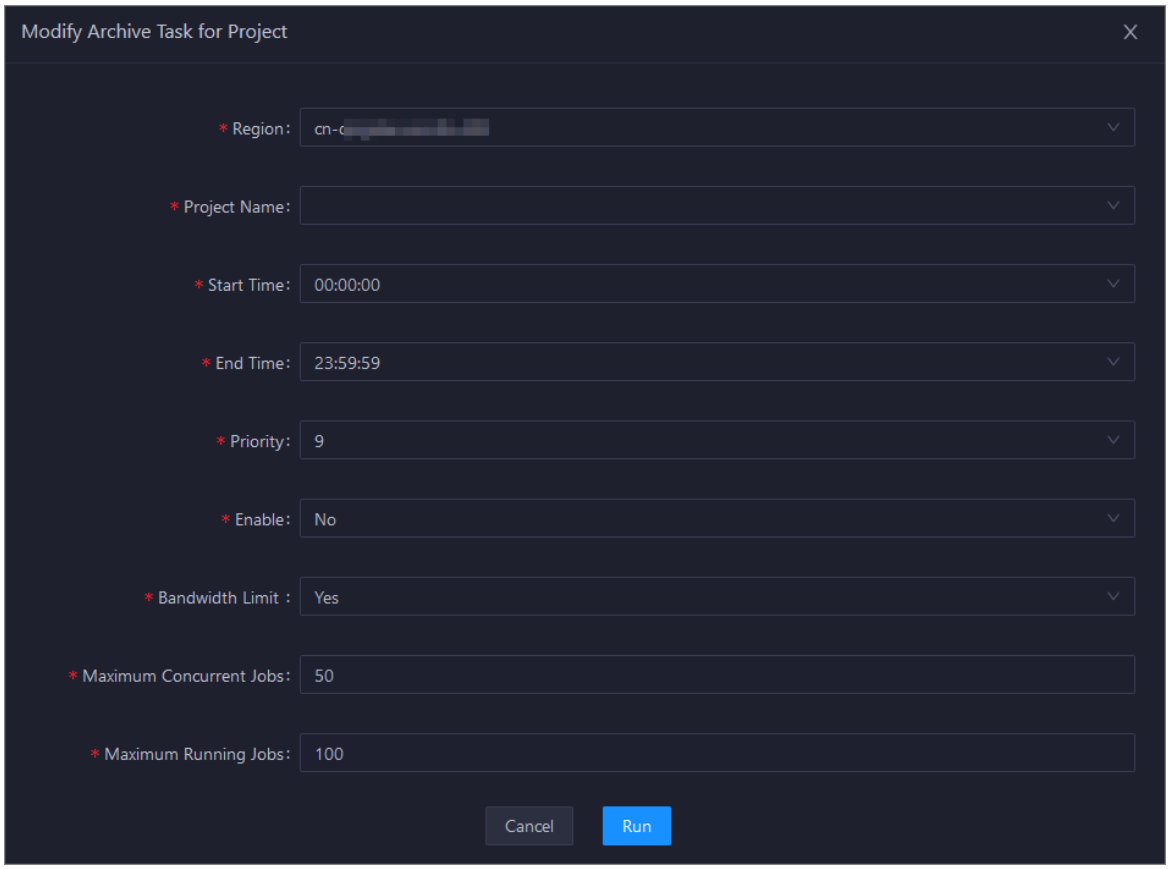

   The following table describes the parameters.

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

3. Click **Run**.

   The newly created merge task appears in the list of merge tasks for projects.

View merge task statistics

In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Merging**. Then, click the **Historical Statistics** tab to view the historical statistics of merge tasks for clusters and projects.

Merge Task Statistics

The trend chart for merge tasks shows statistics on the execution of all merge tasks for each day in the last month. It shows the numbers of running tasks, finished tasks, waiting tasks, timeout tasks, failed tasks, invalid tasks, merged partitions, and reduced files. It also shows the reduced data volume on physical storage, in bytes.

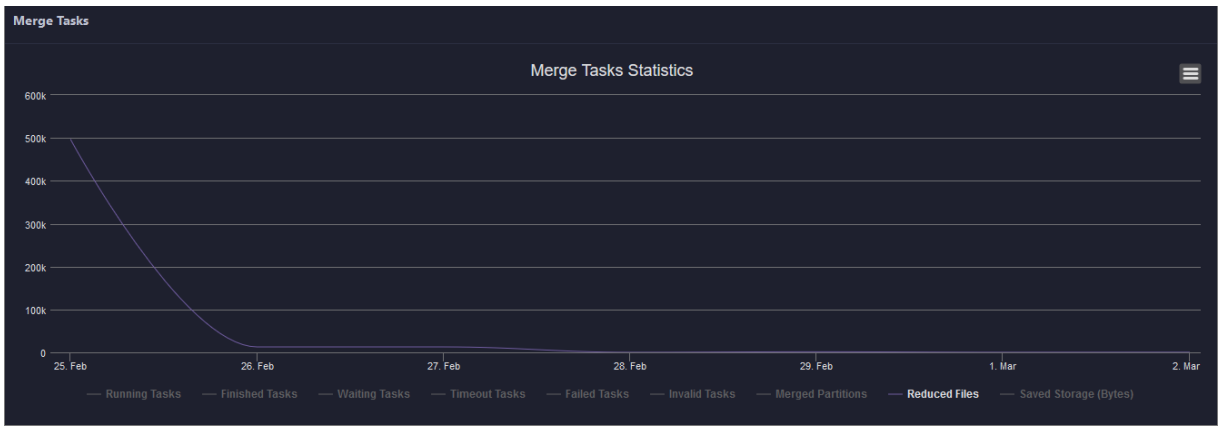

Merge Tasks for Clusters and Merge Tasks for Projects

The two tables show statistics on the execution of merge tasks for clusters and projects on a specific day in the last month. The tables show the numbers of running tasks, finished tasks, waiting tasks, timeout tasks, failed tasks, invalid tasks, merged partitions, and reduced files. The tables also show the reduced data volume on physical storage, in bytes.

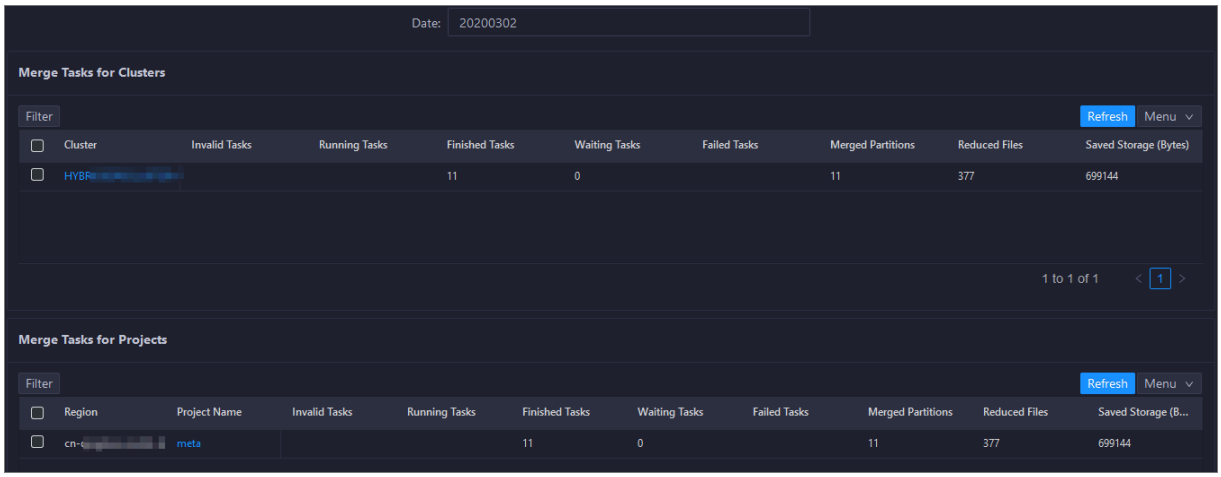

Manage merge types

In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Merging**. Then, click the **Merge Types** tab to view the existing merge types and merge parameters.

Create Merge Type

1. In the **Merge Tasks** section, click **Create Merge Type**. In the pane that appears, specify the required parameters.

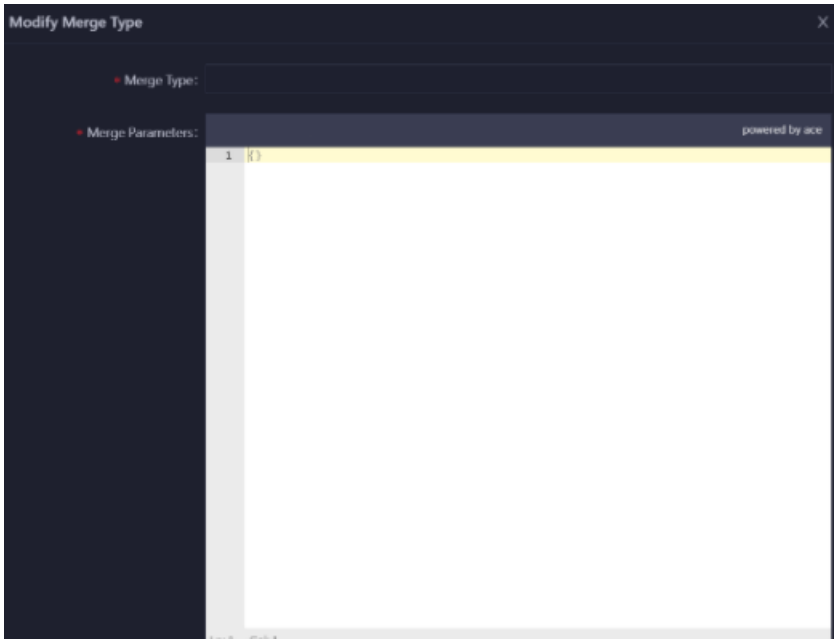

The following table describes the parameters.

Parameter	Description
Merge Type	The name of the merge type.
Merge Parameters	The merge parameters of the merge type.

2. Click **Compare Versions** below Merge Parameters to view the differences between the original and modified values.

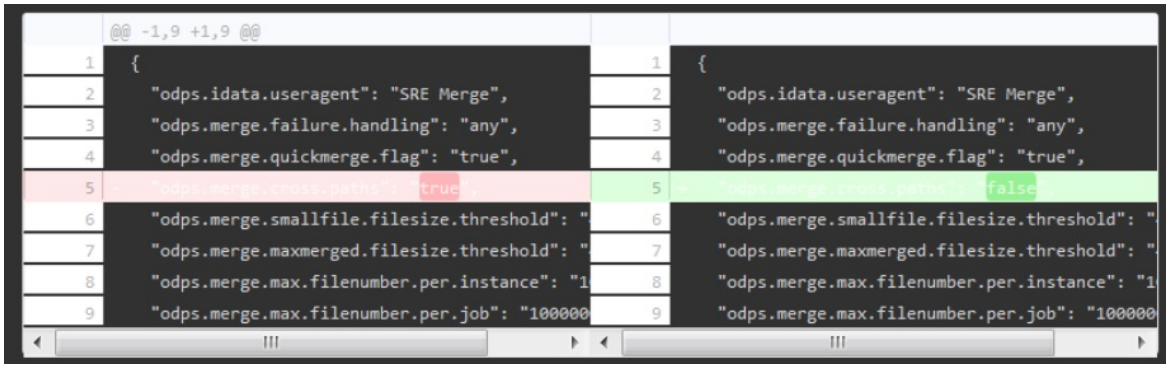

3. Click **Run**.

   The newly created merge type appears in the list of merge types.

11.2.5.2.5.2. Compress idle files

Apsara Bigdata Manager (ABM) allows you to create archive tasks to compress idle files in MaxCompute clusters and projects. This saves storage space for the clusters.

Definition

In a cluster, ABM sorts the tables or partitions created more than 90 days ago by storage space. Then, it compresses the first 100,000 tables or partitions.

Create an archive task for a cluster

If 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 centralized manner.

1. In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Archiving**. The **Archive Tasks** tab appears.

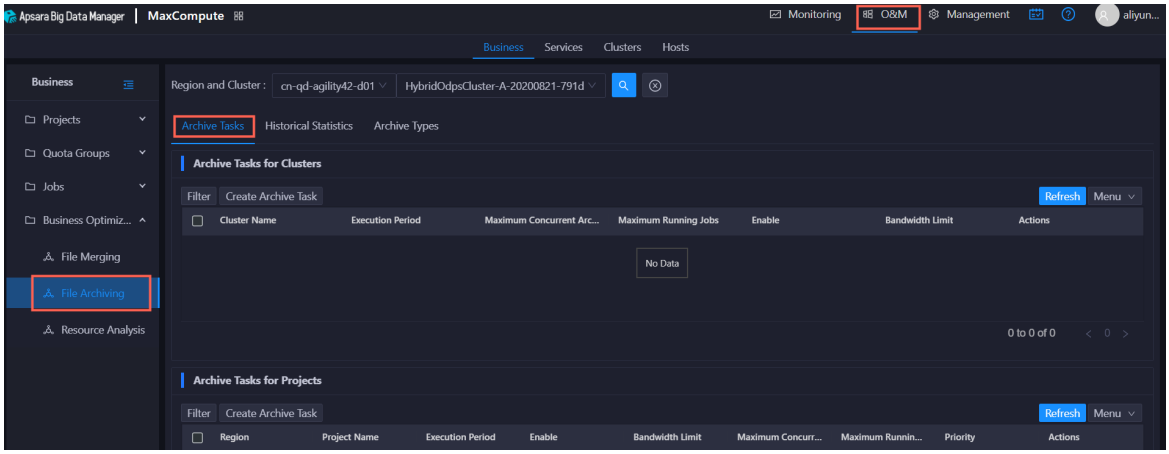

2. In the **Archive Tasks for Clusters** section, click **Create Archive Task**. In the pane that appears, specify the required parameters.

   The following table describes the parameters.

Parameter	Description
Cluster	The cluster for which you want to run the archive task. Select a cluster from the drop-down list.
Start Time	The start time of the task.
End Time	The end time of the task.
Bandwidth Limit	Specifies whether to limit the concurrency of archive tasks for the cluster. ○ **Yes**: indicates that archive tasks cannot be concurrently run. ○ **No**: indicates that archive tasks can be concurrently run.
Maximum Concurrent Jobs	The maximum number of archive tasks that can be run for the selected cluster at the same time. This parameter is valid only when **Bandwidth Limit** is set to **No**.
Enabled	Specifies whether the task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run for the selected cluster at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the selected cluster, not only 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** below Archive Parameters to view the differences between the original and modified values.
4. Click **Run**.

   The newly created archive task appears in the list of archive tasks for clusters.

Create an archive task for a project

If 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 a specific 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. In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Archiving**. The **Archive Tasks** tab appears.

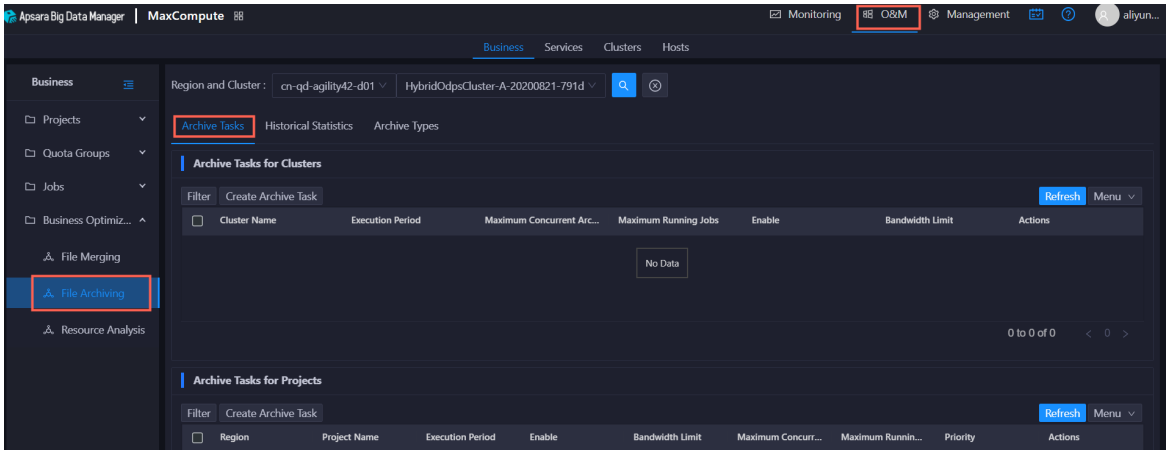

2. In the **Archive Tasks for Projects** section, click **Create Archive Task**. In the pane that appears, specify the required parameters.

   The following table describes the parameters.

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

3. Click **Run**.

   The newly created archive task appears in the list of archive tasks for projects.

View archive task statistics

In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Archiving**. Then, click the **Historical Statistics** tab to view the historical statistics of archive tasks for clusters and projects.

Archive Tasks

The trend chart for archive tasks shows statistics on the execution of all archive tasks for each day in the last month. It shows the numbers of running tasks, finished tasks, waiting tasks, timeout tasks, failed tasks, invalid tasks, merged partitions, and reduced files. It also shows the reduced data volume on physical storage, in bytes.

Statistics by Cluster and Statistics by Project

The two tables show statistics on the execution of archive tasks for clusters and projects on a specific day in the last month. The tables show the numbers of running tasks, finished tasks, waiting tasks, timeout tasks, failed tasks, invalid tasks, merged partitions, and reduced files. The tables also show the reduced data volume on physical storage, in bytes.

Manage archive types

In the left-side navigation pane of the **Business** tab, click **Business Optimization** and then **File Archiving**. Then, click the **Archive Types** tab to view the existing archive types and archive parameters.

Create Archive Type

1. In the **Archive Tasks** section, click **Create Archive Type**. In the pane that appears, specify the required parameters.

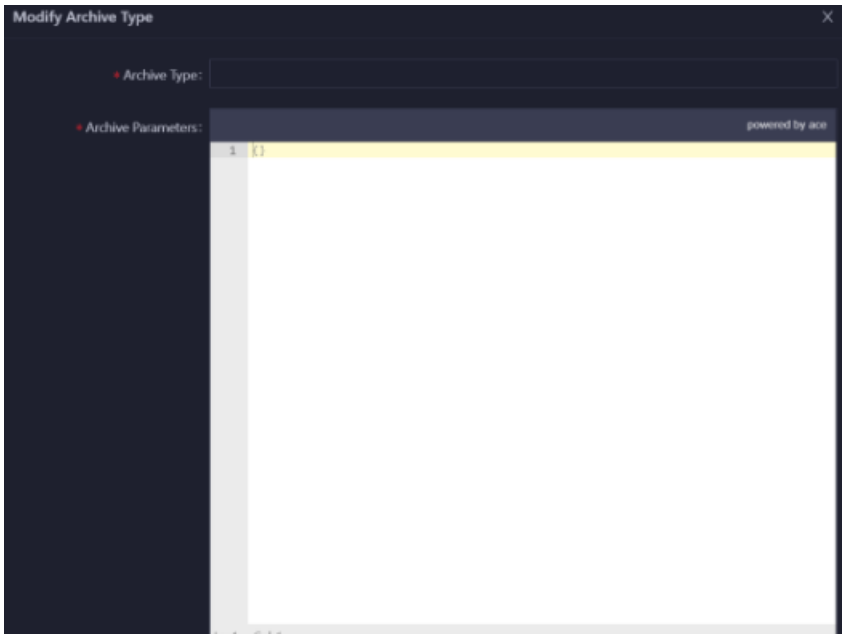

The following table describes the parameters.

Parameter	Description
Archive Type	The name of the archive type.
Archive Parameters	The archive parameters of the archive type.

2. Click **Compare Versions** below Archive Parameters to view the differences between the original and modified values.
3. Click **Run**.

   The newly created archive type appears in the list of archive types.

11.2.5.2.5.3. 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.

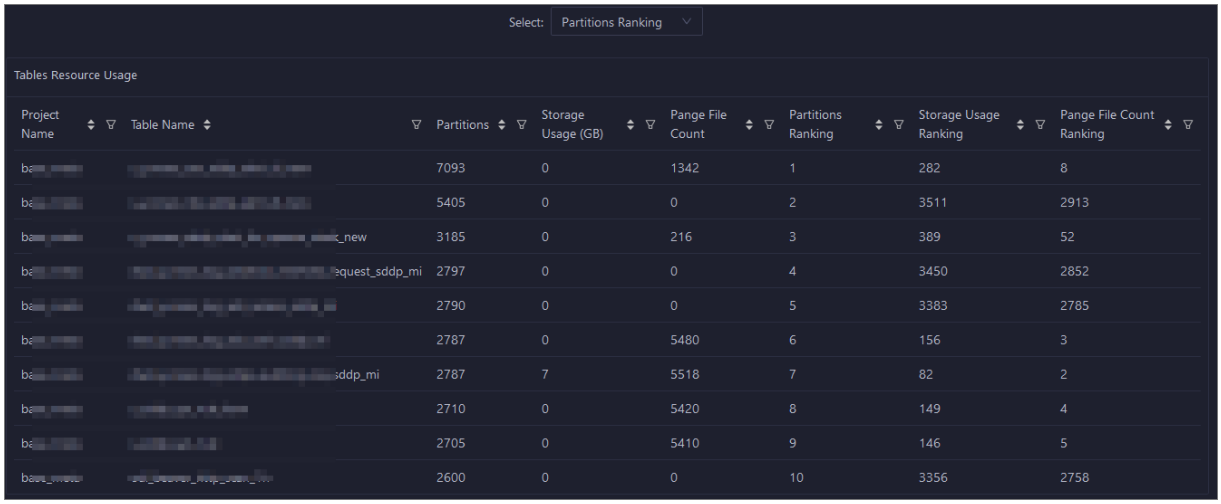

Projects

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.

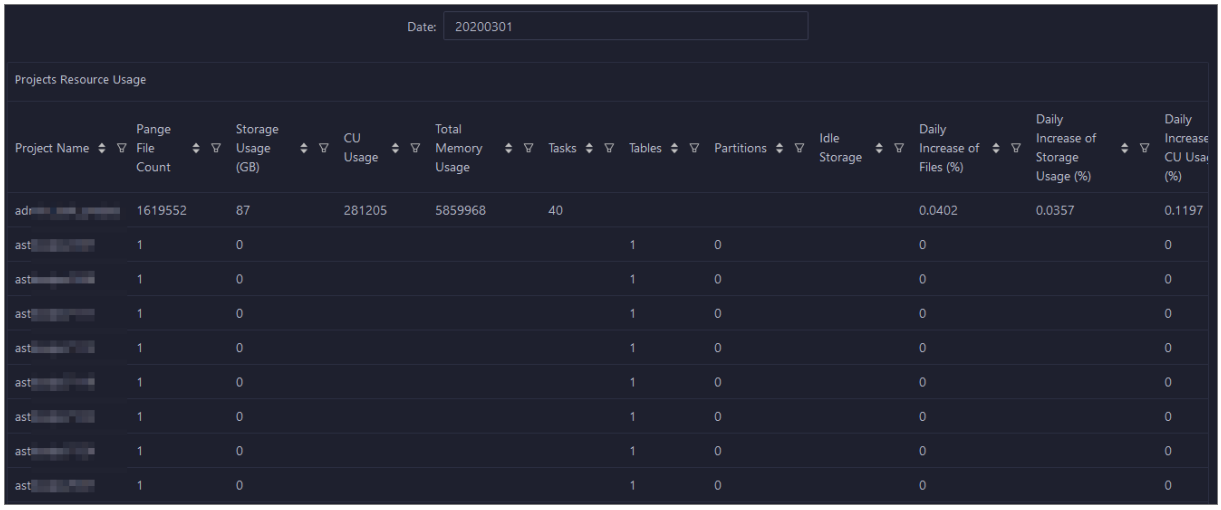

Tasks

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.

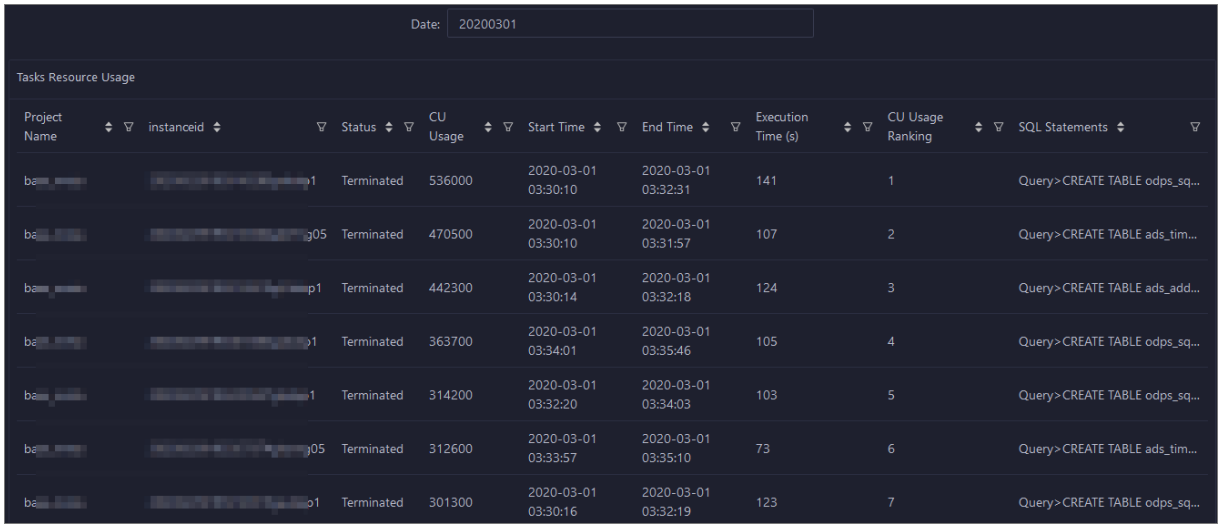

Execution time

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.

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_records), 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.

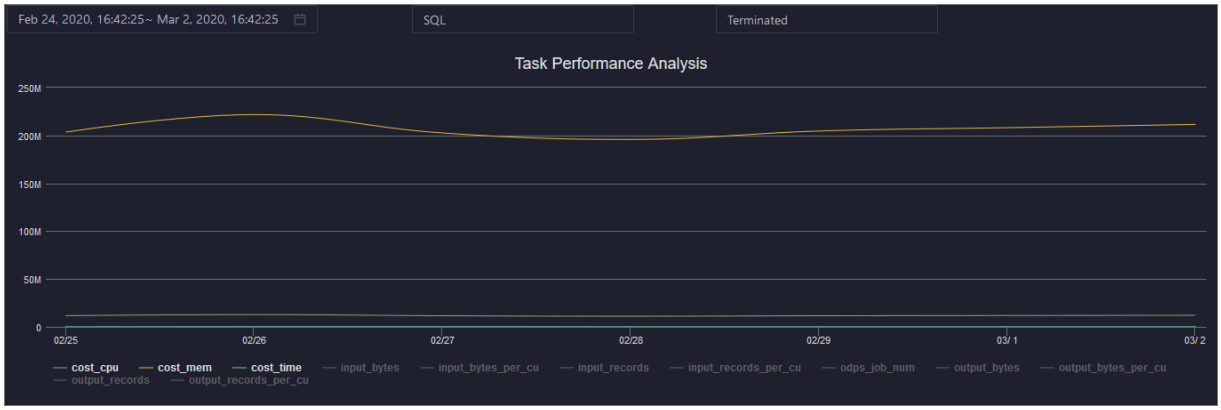

11.2.5.3. Service O&M

11.2.5.3.1. Control service O&M

11.2.5.3.1.1. O&M overview and entry

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

Control service O&M features

- Overview: shows the overall running information about the control service. You can view the service overview, service status, job running, executor pool size, and job status.
- Health Status: shows all checkers for the control service. You can query checker details, check results for hosts in a cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.
- Instances: shows information about the server roles of the control service. You can view the host, status, requested CPU resources, and requested memory of each server role.
- Configuration: provides the access entry to configure global computing, cluster-level computing, computing scheduling, and cluster endpoints.
- Metadata Repository: allows you to view the completion time and status of the output tasks of the metadata warehouse and the trend chart of the consumed time for running tasks in MaxCompute.
- Start Service Role or Stop Service Role: allows you to start or stop the server roles of the MaxCompute control service and view the execution history. If you fail to start or stop the server roles, you can identify the failure.
- Start Admin Console: allows you to start AdminConsole.
- Collect Service Logs: allows you to collect service logs for the specified time period. This enables you to identify the cause of a failure.

Go to the control service O&M page

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.

4. In the left-side navigation pane of the **Services** tab, click **Control**. The **Overview** tab for the control service appears.

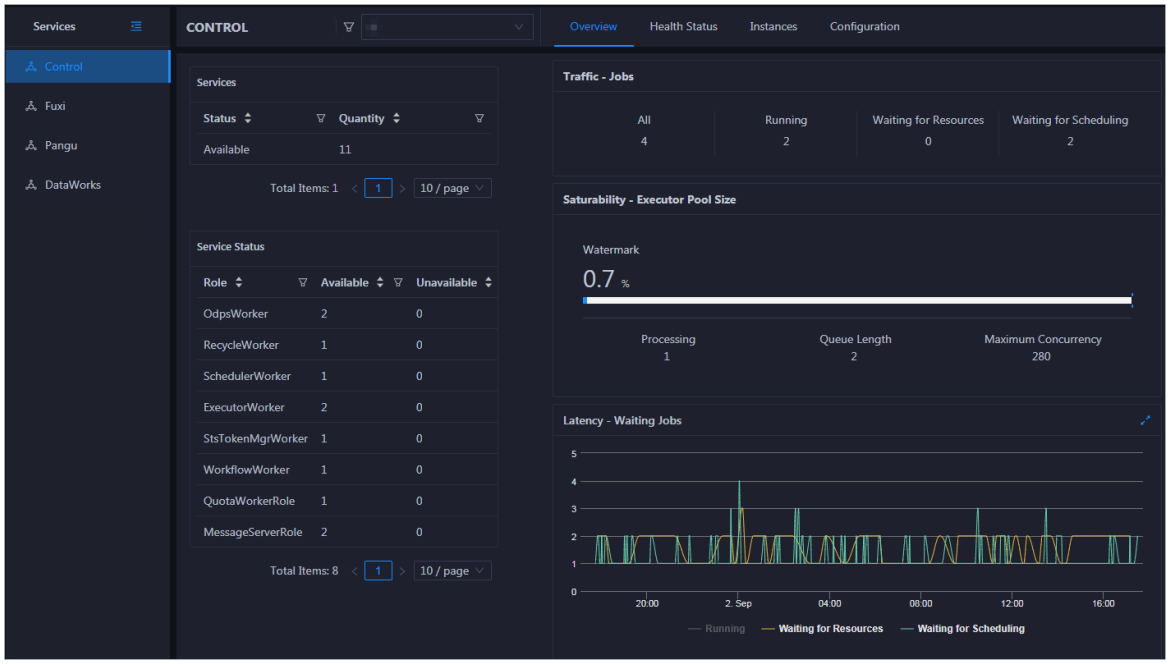

11.2.5.3.1.2. Control service overview

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.

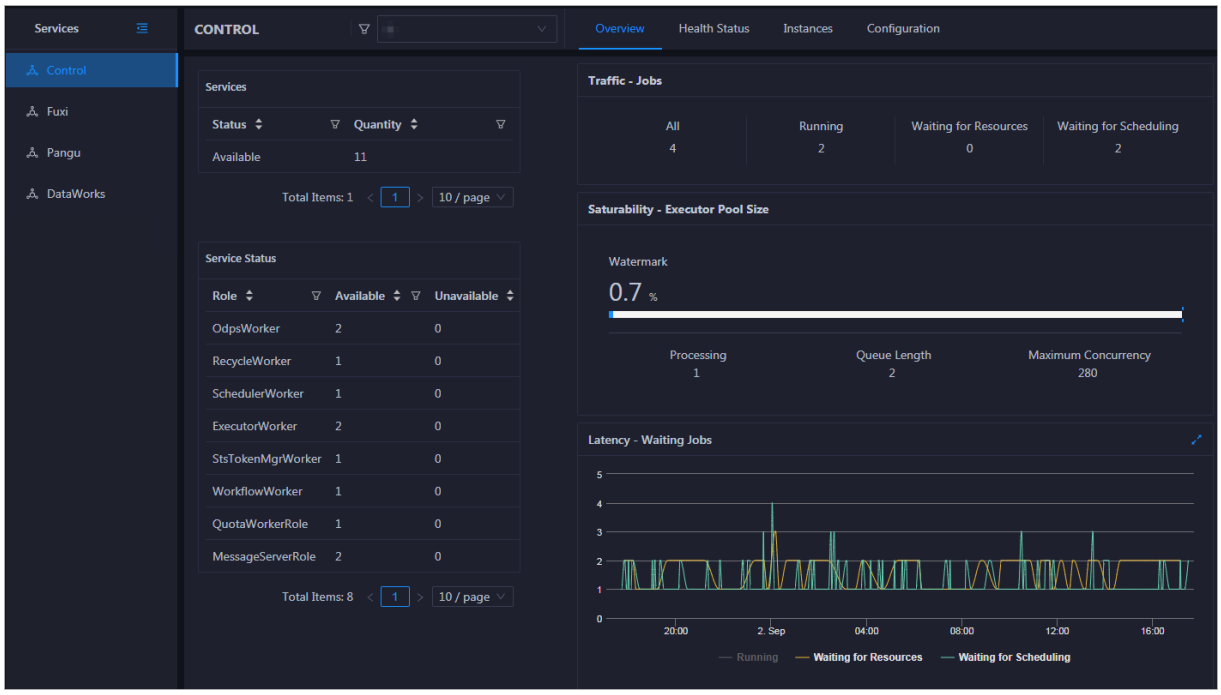

On the **Overview** page, you can view the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

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 respectively.

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.

11.2.5.3.1.3. 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.

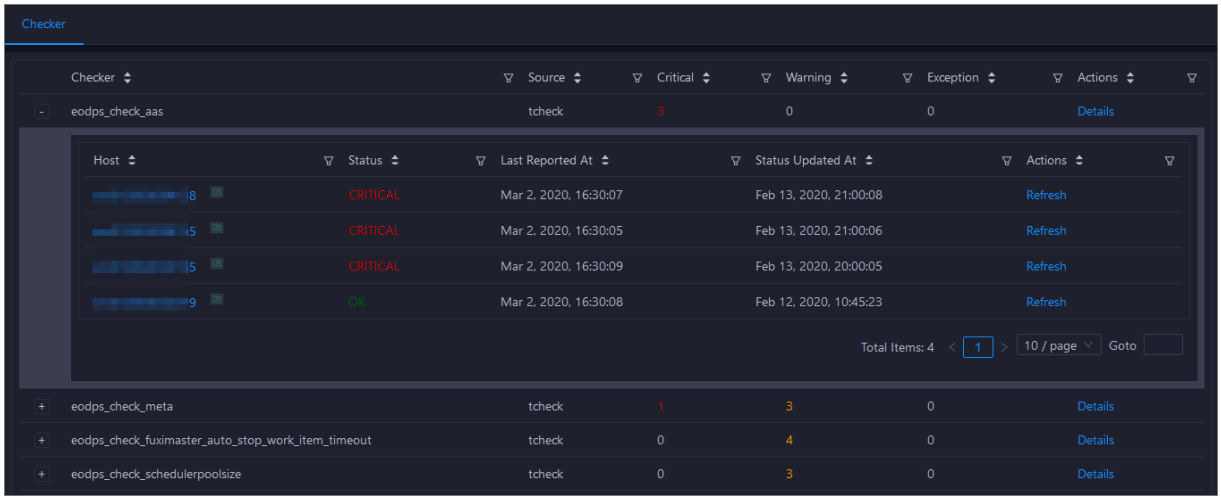

On the **Health Status** page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical**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

11.2.5.3.1.4. Instances

The Instances tab shows information about server roles, which includes the host, status, requested CPU resources, and requested memory of each server role.

Go to the Instances tab

In the left-side navigation pane of the **Services** tab, click **Control**. Then, click the **Instances** tab.

The Instances tab shows information about server roles, which includes the host, status, requested CPU resources, and requested memory of each server role.

11.2.5.3.1.5. 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.

11.2.5.3.1.6. 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.

11.2.5.3.1.7. Stop or start a server role

Apsara Bigdata Manager (ABM) allows you to start or stop the server roles of the MaxCompute control service and view the execution history. If you fail to start or stop the server roles, you can identify the failure.

Stop a server role

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Control**. In the upper-right corner of the tab that appears, click **Actions** and select **Stop Service Role**.
5. In the pane that appears, select a server role you want to stop and click **Run**.
6. In the upper-right corner, click **Actions** and select **Execution History** next to **Stop Service Role** to check whether the action is successful in the execution history.

   The Execution History pane shows the current status, submission time, start time, end time, and operator of each action.
7. Click **Details** in the Details column 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 computer.

Start a server role

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Control**. In the upper-right corner of the tab that appears, click **Actions** and select **Start Service Role**.
5. In the pane that appears, select a server role you want to start and click **Run**.
6. In the upper-right corner, click **Actions** and select **Execution History** next to **Start Service Role** to check whether the action is successful in the execution history.

   The Execution History pane shows the current status, submission time, start time, end time, and operator of each action.
7. Click **Details** in the Details column 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 computer.

Identify the cause of a failure

This section describes how to identify the cause of the failure to start a server role.

1. In the Execution History pane, click **Details** in the Details column of the task to view the details.
2. In the pane that appears, click **View Details** for a failed step to identify the cause of the failure.

   You can view the parameter settings, outputs, error messages, script, and runtime parameters to identify the cause of the failure.

11.2.5.3.1.8. 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. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Control**.
5. In the upper-right corner of the pane that appears, click **Actions** and select **Start Admin Console**.
6. In the **Start Admin Console** pane, click **Run**.

Step 2: View the execution status or progress

1. On any tab of the **Control** page, click **Actions** and select **Execution History** next to **Start Admin Console** to view the execution history.

RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

2. If the status is **RUNNING**, click **Details** in the Details column to view the execution progress.

Step 3: (Optional) Identify the cause of a failure

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

1. On any tab of the **Control** page, click **Actions** and select **Execution History** next to **Start Admin Console** in the upper-right corner to view the execution history.
2. In the pane that appears, 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 identify the cause of the failure.

11.2.5.3.1.9. Collect service logs

Apsara Bigdata Manager (ABM) allows you to collect service logs for the specified time period. This enables you to identify the cause of a failure.

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. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Control**.
5. In the upper-right corner of the tab that appears, click **Actions** and select **Collect Service Logs**.
6. In the **Collect Service Logs** pane that appears, specify the required parameters.

   The following table describes the 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. You can select multiple services.
Time Period	The time period in which the logs that you want to collect are generated.
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**.

Step 2: View the execution status or progress

1. On any tab of the **Control** page, click **Actions** and select **Execution History** next to **Collect Service Logs** in the upper-right corner to view the execution history.

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

2. If the status is **RUNNING**, click **Details** in the Details column to view the execution progress.

Step 3: (Optional) Identify the cause of a failure

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

1. On any tab of the **Control** page, click **Actions** and select **Execution History** next to **Collect Service Logs** in the upper-right corner to view the execution history.
2. In the Execution History pane, 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 identify the cause of the failure.

11.2.5.3.2. Job Scheduler O&M

11.2.5.3.2.1. O&M features and entry

This topic describes Job Scheduler O&M features. It also provides more information about how to go to the Job Scheduler O&M page.

Job Scheduler O&M features

- Overview: shows the key operating information of Job Scheduler. The information includes the service overview, service status, resource usage, compute node overview, and the trend charts of CPU utilization and memory usage.
- Health Status: shows all checkers for Job Scheduler. You can query checker details, check results for hosts in a cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.
- Quotas: allows you to view, create, or modify the quota groups in Job Scheduler.
- Instances: shows information about the master nodes and server roles of Job Scheduler and allows you to restart the master nodes.
- Compute Nodes: shows all compute nodes in Job Scheduler and allows you to add compute nodes to or remove compute nodes from a blacklist or read-only list.
- Enable SQL Acceleration or Disable SQL Acceleration: allows you to enable or disable SQL acceleration for Job Scheduler.
- Restart Fuxi Master Node: allows you to restart the primary and secondary master nodes for Job Scheduler.

Go to the Job Scheduler O&M page

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Fuxi**. The **Overview** tab appears.

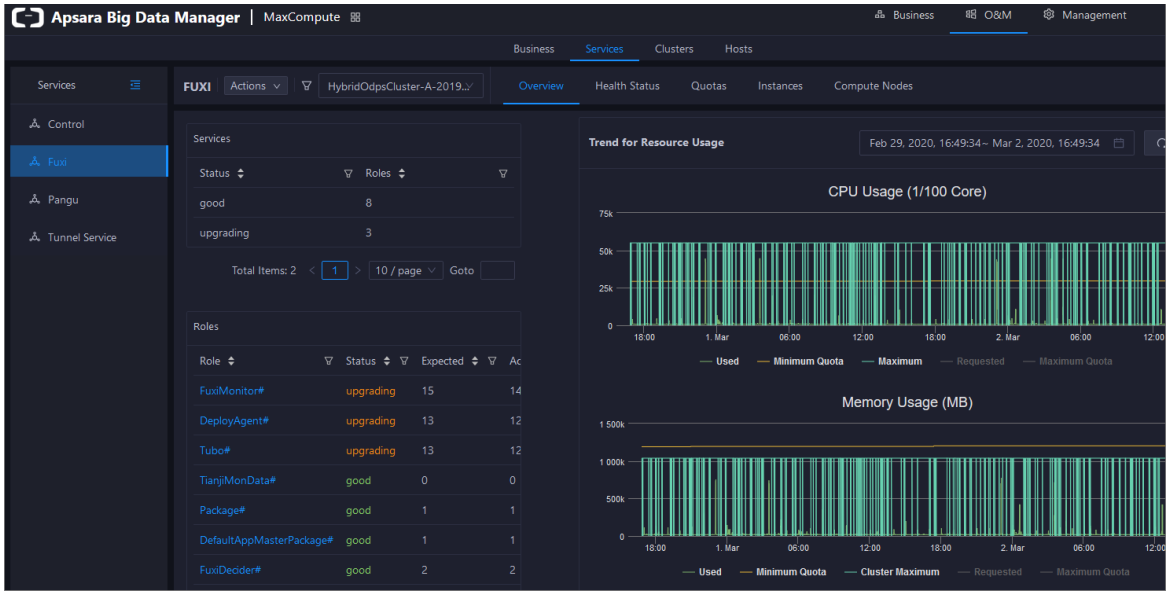

11.2.5.3.2.2. Overview

The Overview tab shows the key operating information of Job Scheduler. The information includes the service overview, service status, resource usage, compute node overview, and the trend charts of CPU utilization and memory usage.

Go to the Overview tab

1. In the left-side navigation pane of the **Services** tab, click **Fuxi**.
2. Select a cluster and click the **Overview** tab. The **Overview** tab for the selected cluster appears.

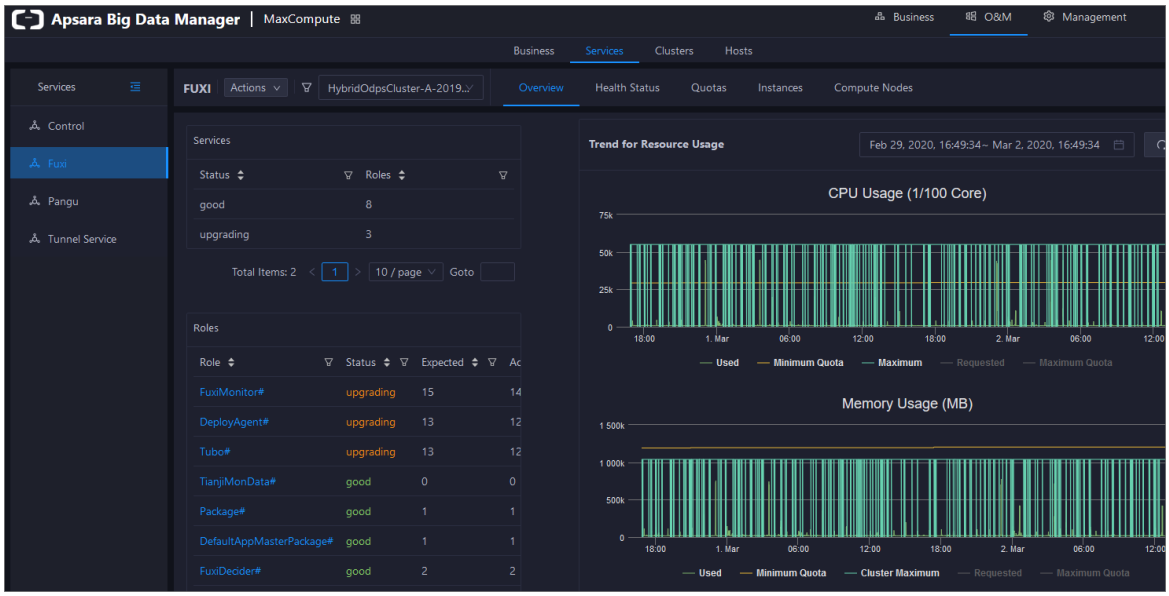

The **Overview** tab shows the key operating information of Job Scheduler. The information includes the service overview, service status, resource usage, compute node overview, and the trend charts of CPU utilization and memory usage.

Services

This section shows the numbers of available services, unavailable services, and services that are being updated.

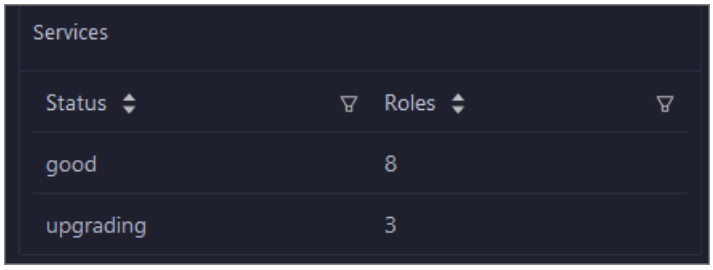

Roles

This section shows all Job Scheduler server roles and their states. You can also view the expected and actual numbers of machines for each server role.

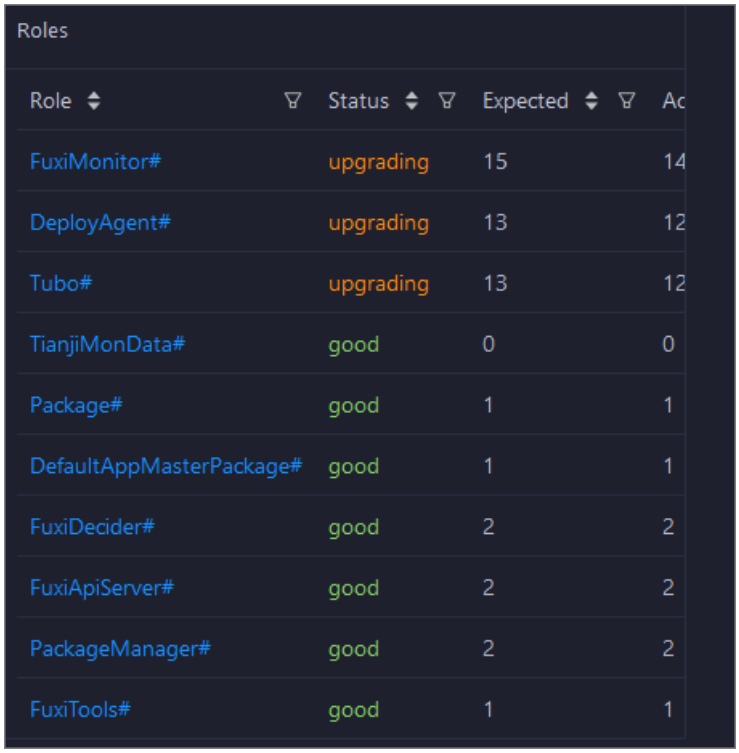

Click the name of a server role to go to the Apsara Infrastructure Management Framework console and view its details.

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

The Trend for Resource Usage section shows the trend charts of CPU utilization and memory usage for Job Scheduler. Each trend chart shows information about the used quota, minimum quota, maximum cluster quota, requested quota, and maximum quota in different colors. The trend charts are periodically refreshed. You can also manually refresh the trend charts. You can also view the trend charts of CPU utilization and memory usage for a specific period.

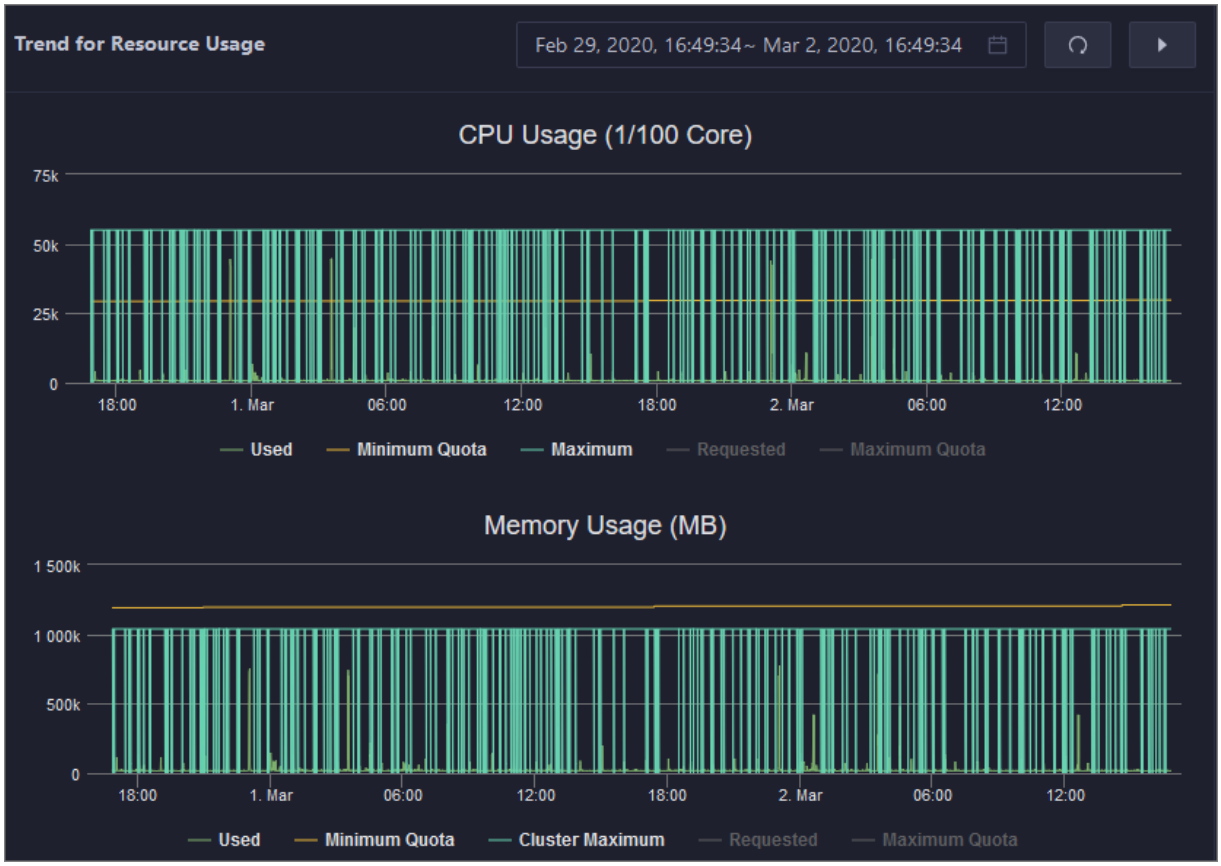

Saturability - Resource Usage

This section shows the allocation of CPU and memory resources.

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

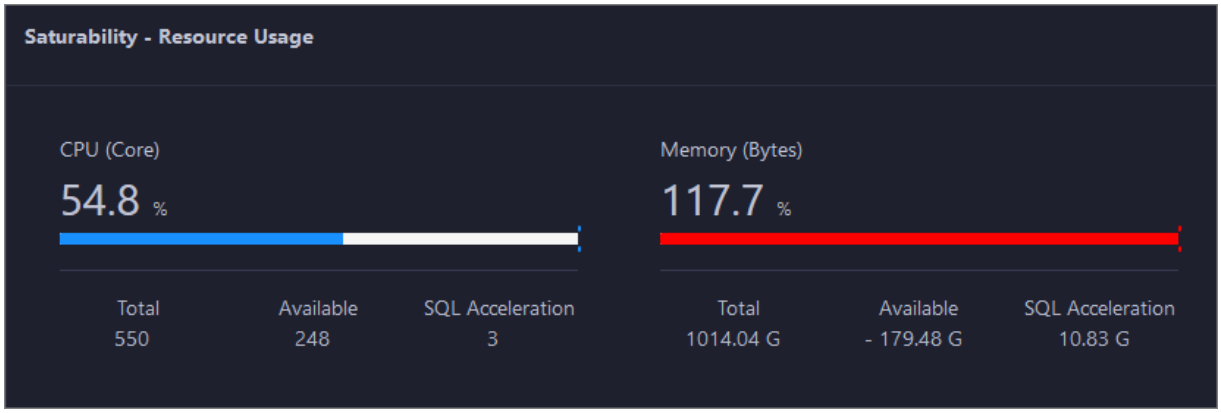

Compute Nodes

This section shows the details of compute nodes in Job Scheduler. The details include the percentage of online compute nodes, the total number of compute nodes, the number of online compute nodes, and the number of compute nodes in a blacklist.

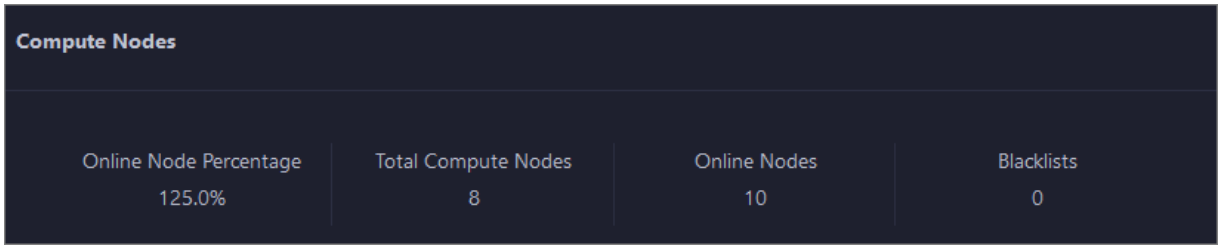

11.2.5.3.2.3. Job Scheduler health

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.

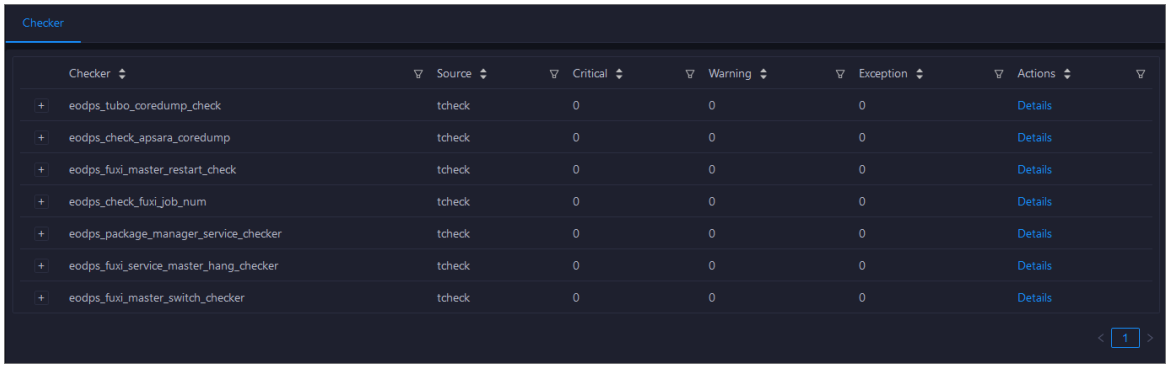

On the **Health Status** page, you can view all checkers of the Job Scheduler service and the check results for all hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

11.2.5.3.2.4. Quotas

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

Go to the Quotas tab

1. In the left-side navigation pane of the **Services** tab, click **Fuxi**.
2. Select a cluster and click the **Quotas** tab. The **Quotas** tab for the selected cluster appears.

The **Quotas** tab lists existing quota groups in Job Scheduler.

Create a quota group

1. In the upper-left corner of the **Quotas** tab, click **Create Quota Group**.
2. In the **Quota Group** pane, specify the required parameters.

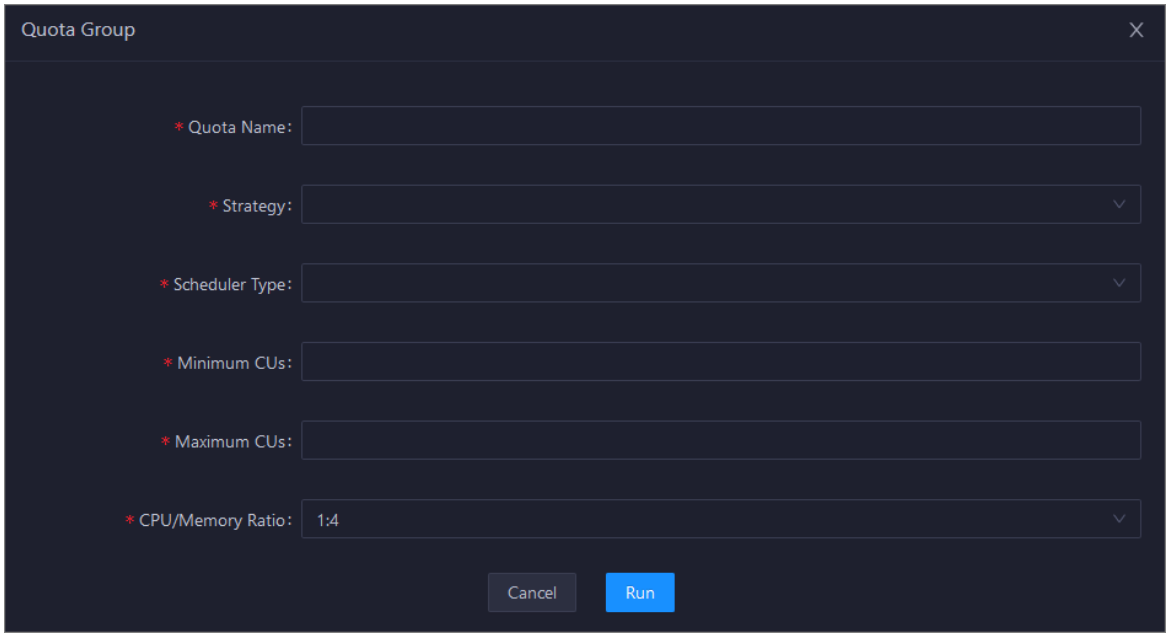

3. Click **Run**.

   The newly created quota group appears in the quota group list.

View quota group details

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

Resource usage

Applications

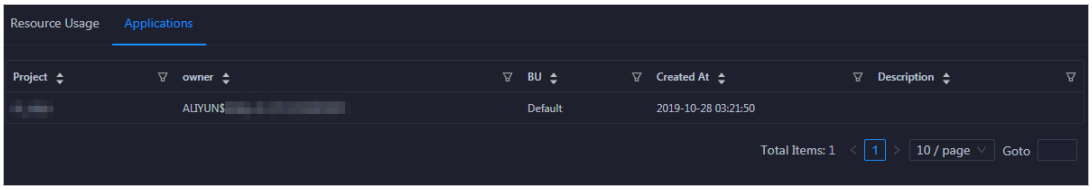

Modify a quota group

1. On the **Quotas** tab, find the quota group that you want to modify and click **Modify** in the Actions column. In the pane that appears, modify parameters as instructed.
2. Click **Run**.

   After the configuration is complete, you can check whether the quota group is modified in the quota group list.

11.2.5.3.2.5. Instances

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

Go to the Instances tab

1. In the left-side navigation pane of the **Services** tab, click **Fuxi**.
2. Select a cluster and click the **Instances** tab. The **Instances** tab for the selected cluster appears.

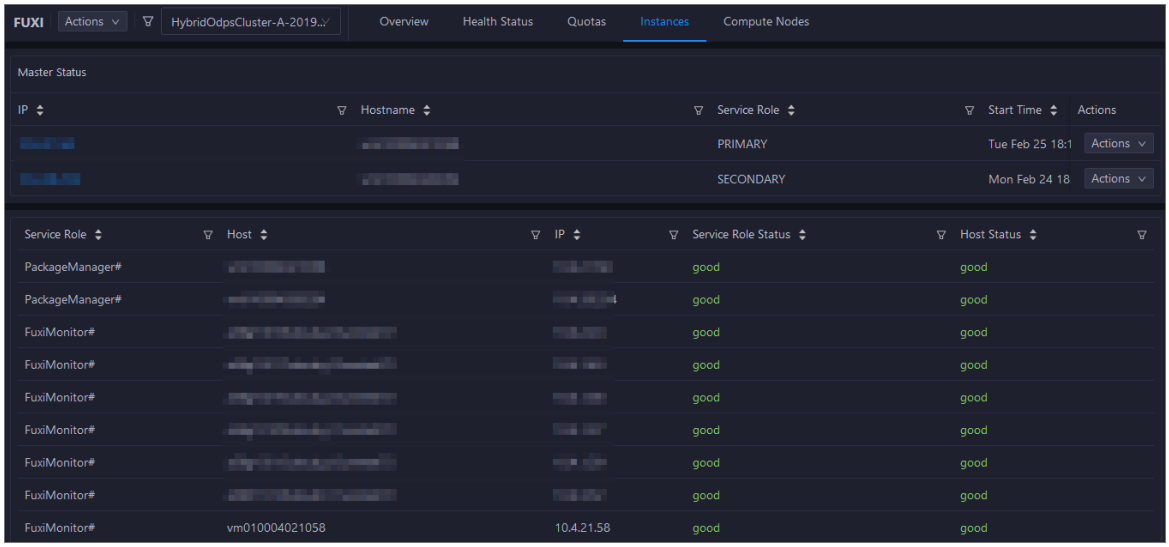

The **Instances** tab shows information about the master nodes and server roles of Job Scheduler. The information about the master nodes includes the IP address, hostname, server role, and start time. The information about a server role includes the role name, hostname, role status, and host status.

Supported operations

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

11.2.5.3.2.6. 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.

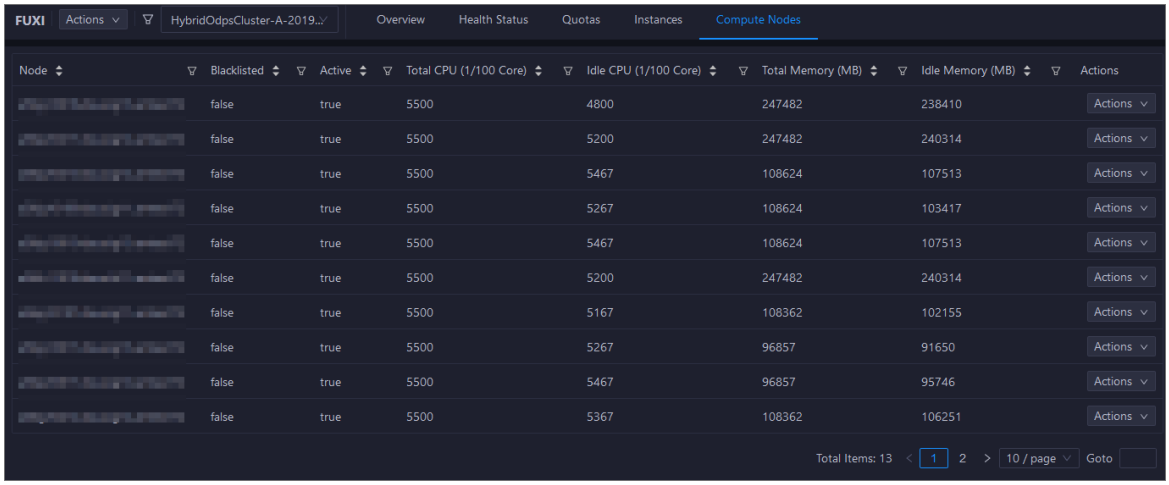

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including the total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active.

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.

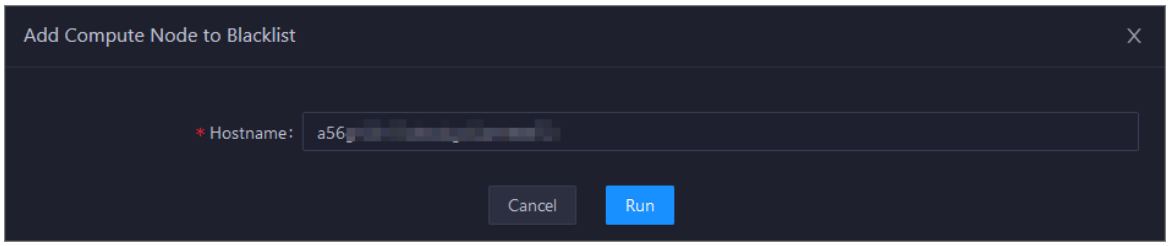

The value of the **Hostname** parameter is automatically filled. You do not need to specify a value for this parameter.

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

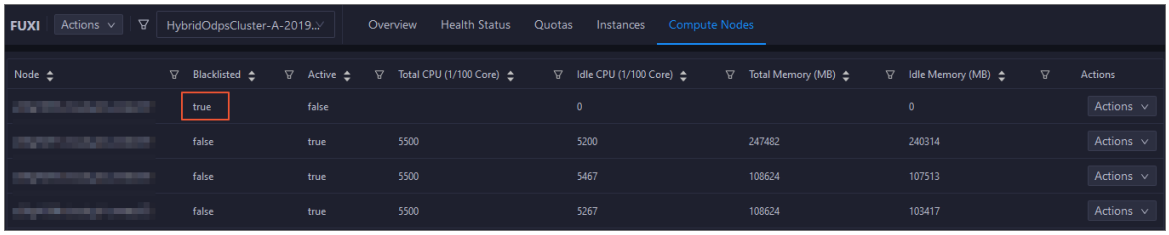

11.2.5.3.2.7. Enable and disable SQL acceleration

You can enable or disable SQL acceleration for Job Scheduler in the Apsara Bigdata Manager (ABM) console. The execution speed of SQL statements in Job Scheduler is greatly increased with SQL acceleration enabled, but more computing resources are consumed.

Enable SQL acceleration

1. In the left-side navigation pane of the **Services** tab, click **Fuxi**. Then, select a cluster.
2. In the upper-right corner of the tab that appears, click **Actions** and select **Enable SQL Acceleration**.
3. In the pane that appears, set the **WorkerSpans** parameter.

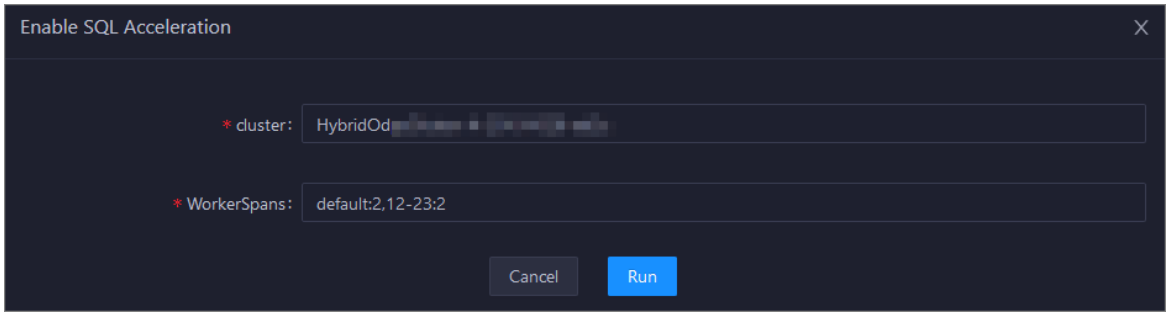

WorkerSpans: the default resource quota of the cluster and the resource quota for a specific period. Default value: **default:2,12-23:2**.

> **Note** 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, you can set this parameter to default:2,12-23:4 to increase the resource quota in peak hours.

4. Click **Run**.

Disable SQL acceleration

1. In the left-side navigation pane of the **Services** tab, click **Fuxi**. Then, select a cluster.
2. In the upper-right corner of the tab that appears, click **Actions** and select **Disable SQL Acceleration**.
3. In the pane that appears, click **Run**.

View the execution history of enabling or disabling SQL acceleration

After you submit the action of enabling or disabling SQL acceleration, you can view the execution history to check whether the action is complete. The system executes the action as a job. It provides execution records and logs for each execution so that you can identify faults encountered during its execution. This section describes how to view the execution history of enabling SQL acceleration.

1. In the left-side navigation pane of the **Services** tab, click **Fuxi**. Then, select a cluster.
2. In the upper-right corner of the tab that appears, click **Actions** and select **Execution History** next to **Enable SQL Acceleration**.
3. In the pane that appears, view the execution history of enabling SQL acceleration.

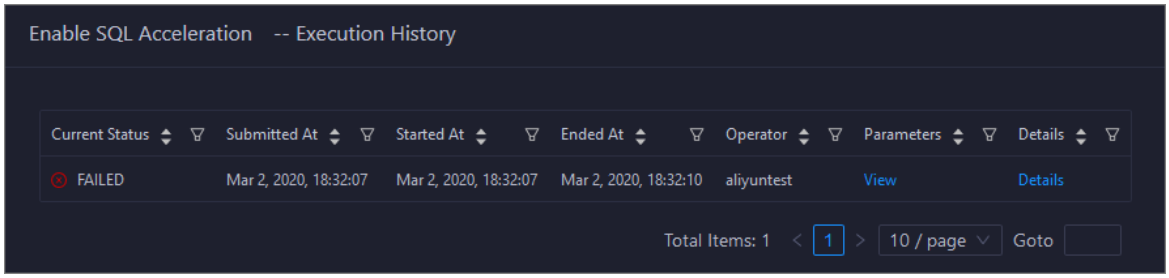

The execution history shows the current status, submission time, start time, end time, and operator of each execution.

4. If the execution fails, click **Details** to identify the cause of the failure.

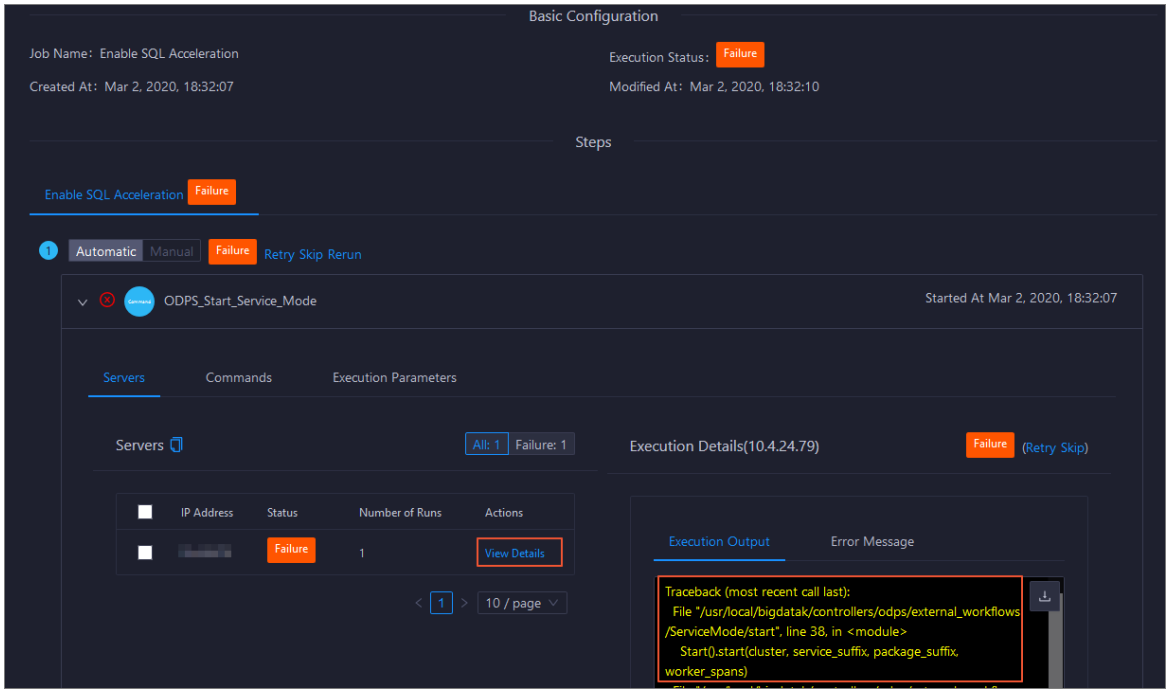

11.2.5.3.2.8. Restart a master node of Job Scheduler

Job Scheduler is the resource management and task scheduling system of the Apsara distributed operating 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. In the upper-right corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Fuxi**. Then, click the **Instances** tab.
5. On the **Instances** tab, 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** pane, click **Run**. The **Restart Fuxi Master Node** pane appears.

Step 2: View the execution status or progress

1. In the **Restart Fuxi Master Node** pane, check the execution history of restarting master nodes.

   The **Restart Fuxi Master Node** pane shows the restart history. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.
2. If the status is **RUNNING**, click **Details** in the Details column to view the execution progress.

Step 3: (Optional) Locate the cause of a failure

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

1. In the **Restart Fuxi Master Node** pane, check the execution history of restarting master nodes.
2. 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 identify the cause of the failure.

11.2.5.3.3. Apsara Distribute File System O&M

11.2.5.3.3.1. O&M features and entry

This topic describes the O&M features of Apsara Distributed File System. It also provides more information about how to go to the Apsara Distributed File System O&M page.

Apsara Distributed File System O&M features

- Overview: shows the key operating information of Apsara Distributed File System. The information includes the service overview, service status, storage usage, storage node overview, and the trend charts of storage usage and file count.
- Health Status: shows all checkers for Apsara Distributed File System. You can query checker details, check results for hosts in a cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.
- Instances: shows information about the master nodes and server roles of Apsara Distributed File System. You can change the primary master node or run a checkpoint on a master node of Apsara Distributed File System.
- Storage Nodes: shows information about the storage nodes of Apsara Distributed File System. You can set the status of a storage node to Disabled or Normal. You can also set the status of a disk on a storage node to Normal or Error.
- Change Primary Master Node: allows you to change the primary master node of Apsara Distributed File System in a cluster.
- Run Checkpoint on Master Node: allows you to run checkpoints on master nodes of Apsara Distributed File System to write memory data to disks.
- Empty Recycle Bin: allows you to clear the recycle bin of Apsara Distributed File System.
- Enable Data Rebalancing or Disable Data Rebalancing: allows you to enable or disable the data rebalancing feature of Apsara Distributed File System.

Go to the Pangu page

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Pangu**. Then, select a cluster. The **Overview** tab for the selected cluster appears.

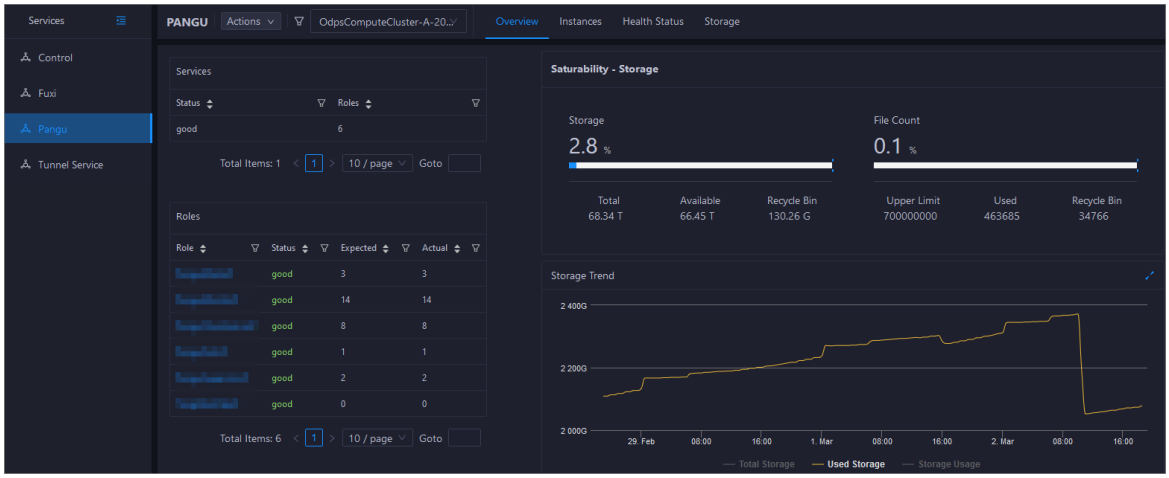

11.2.5.3.3.2. Overview

The Overview tab shows the key operating information about Apsara Distributed File System. The information includes the service overview, service status, storage usage, storage node overview, and the trend charts of storage usage and file count.

Go to the Overview tab

1. In the left-side navigation pane of the **Services** tab, click **Pangu**.
2. Select a cluster and click the **Overview** tab. The **Overview** tab for the selected cluster appears.

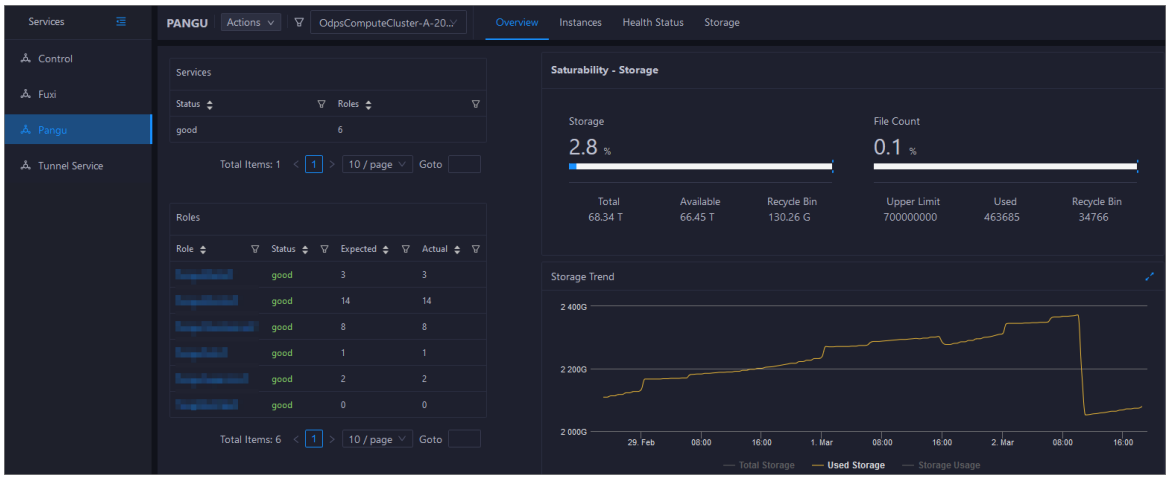

The **Overview** tab shows the key operating information about Apsara Distributed File System. The information includes the service overview, service status, health check result, health check history, storage usage, storage node overview, and the trend charts of storage usage and file count.

Services

This section shows the status of Apsara Distributed File System and the number of server roles.

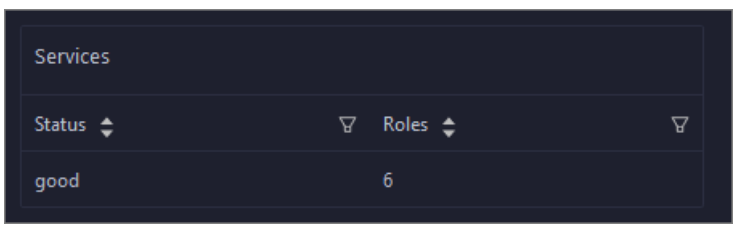

Roles

This section shows all server roles of Apsara Distributed File System and their states. You can also view the expected and actual numbers of hosts for each server role.

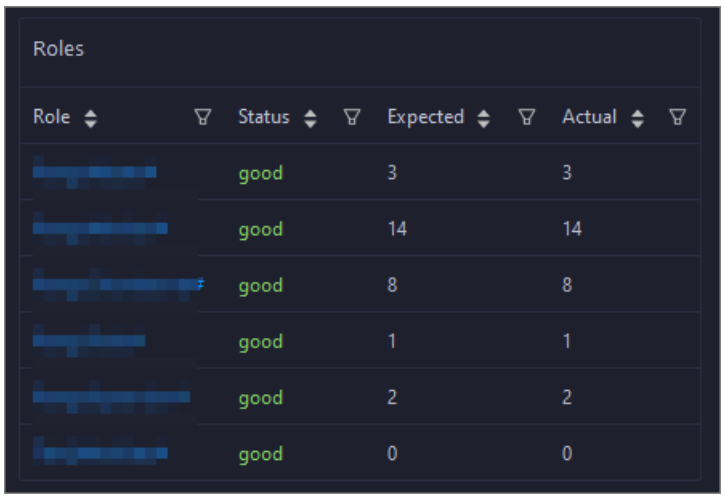

Saturability - Storage

This section shows the storage usage and file count.

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

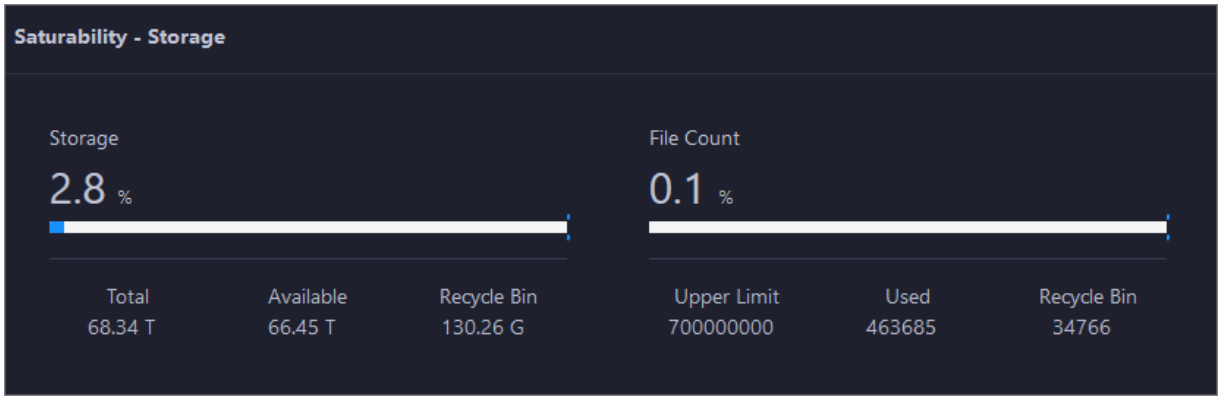

Storage Trend and File Count Trend

This section shows the trend charts of the storage usage and file count. The storage usage chart shows the trend lines of the total storage space, used storage space, and storage usage in different colors. The file count chart shows the trend line of the file count.

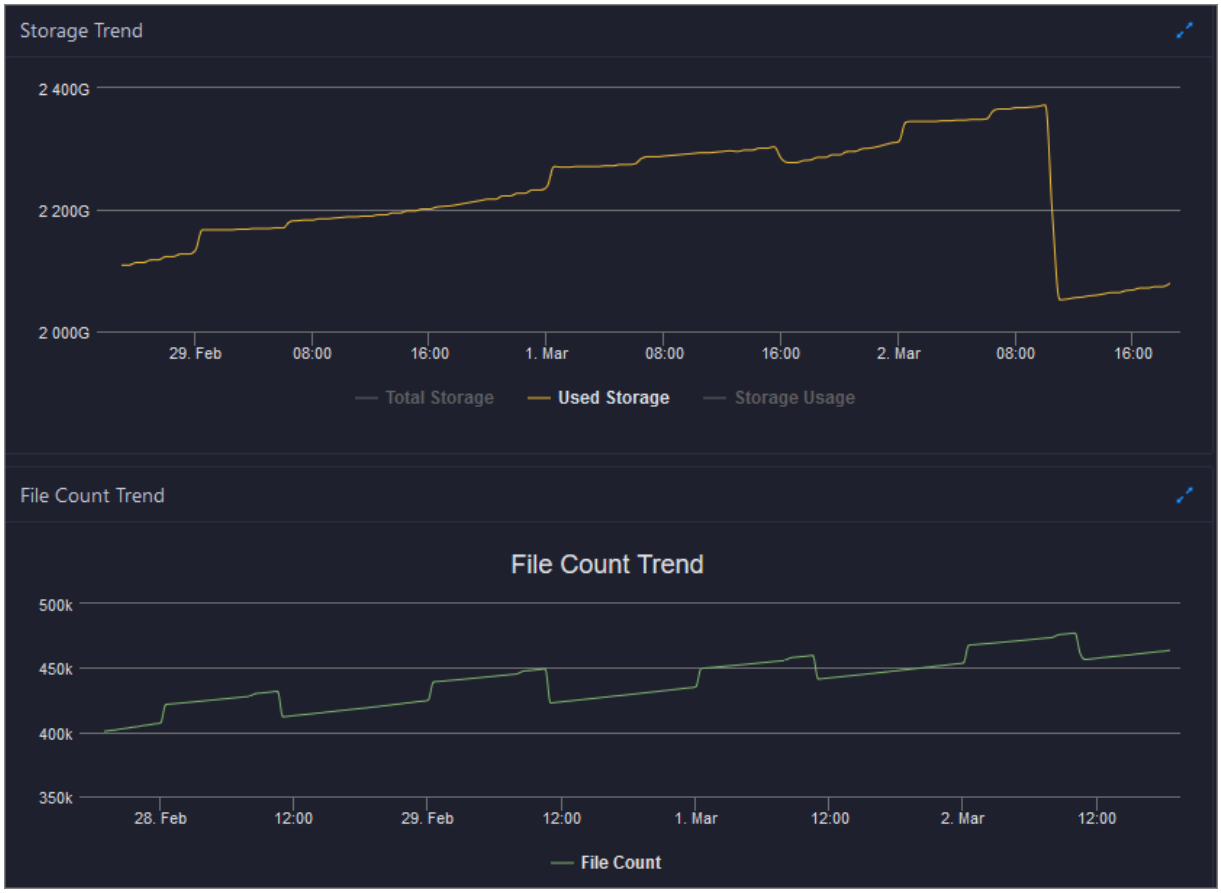

In the upper-right corner of the chart, click the icon to zoom in the chart. The following figure shows an enlarged chart of storage usage.

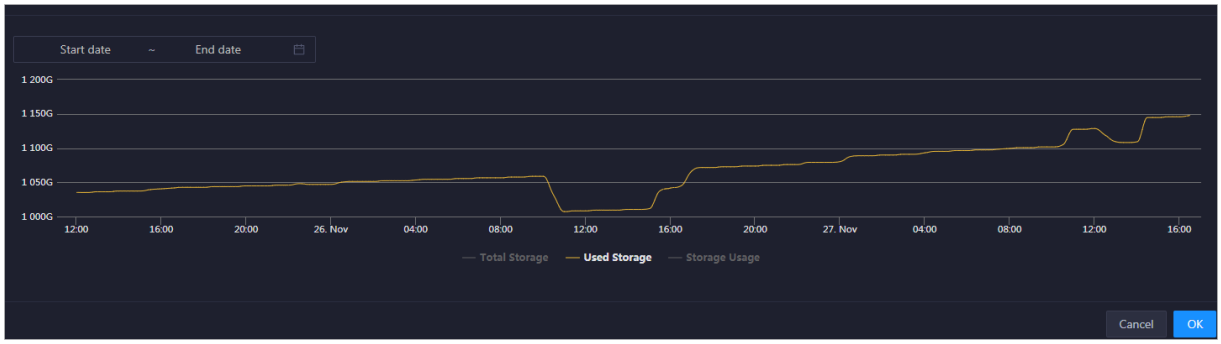

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

Storage Nodes

This section shows information about the storage nodes of Apsara Distributed File System. The information includes the numbers of data nodes, normal nodes, disks, and normal disks. You can also view the faulty node percentage and faulty disk percentage.

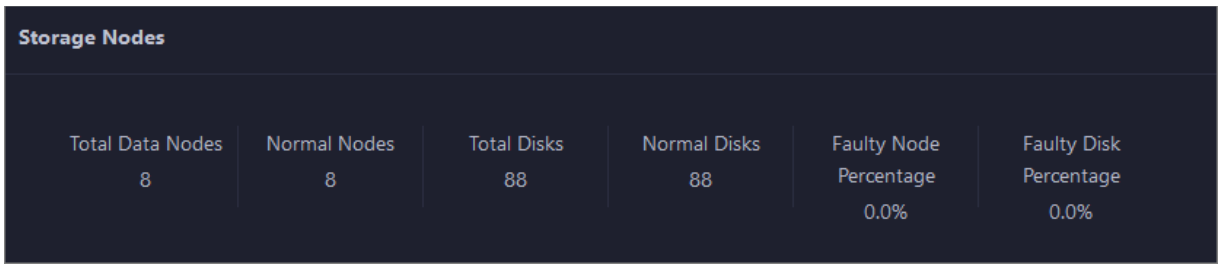

11.2.5.3.3.3. Instances

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

Go to the Instances tab

1. In the left-side navigation pane of the **Services** tab, click **Pangu**.
2. Select a cluster and click the **Instances** tab. The **Instances** tab for the selected cluster appears.

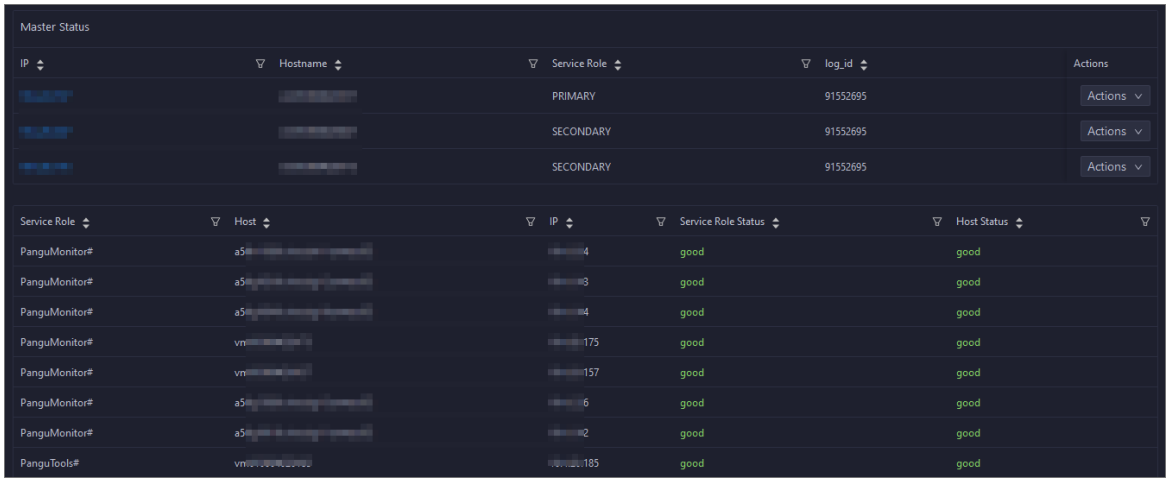

The **Instances** tab shows information about the master nodes and server roles of Apsara Distributed File System. The information about a master node includes the IP address, hostname, server role, and log ID. The information about a server role includes the role name, hostname, role status, and host status.

Supported 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 for Apsara Distributed File System and Run a checkpoint on the master nodes of Apsara Distributed File System.

11.2.5.3.3.4. 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.

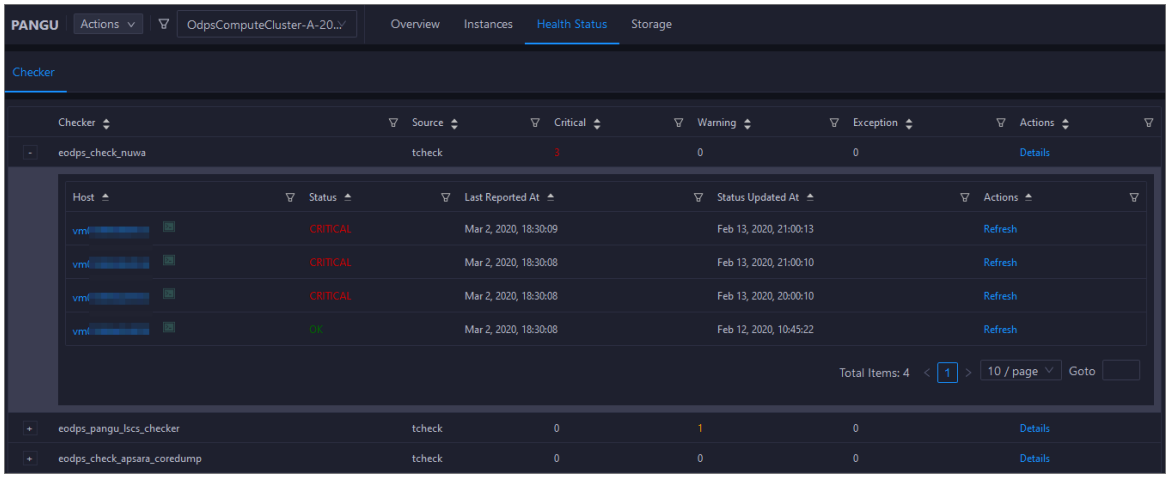

On the **Health Status** page, you can view all checkers of Apsara Distributed File System and the check results for all hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

11.2.5.3.3.5. 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.

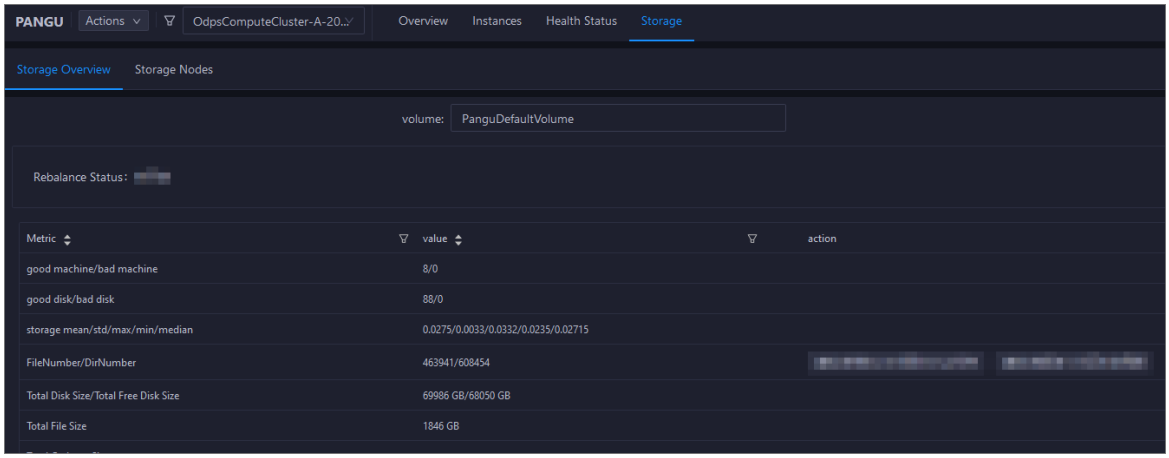

The **Storage Overview** page displays whether data rebalancing is enabled, key metrics and their values, suggestions to handle exceptions, and rack specifications of Apsara Distributed File System. The **Storage Nodes** page displays the information about all storage nodes of Apsara Distributed File 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.

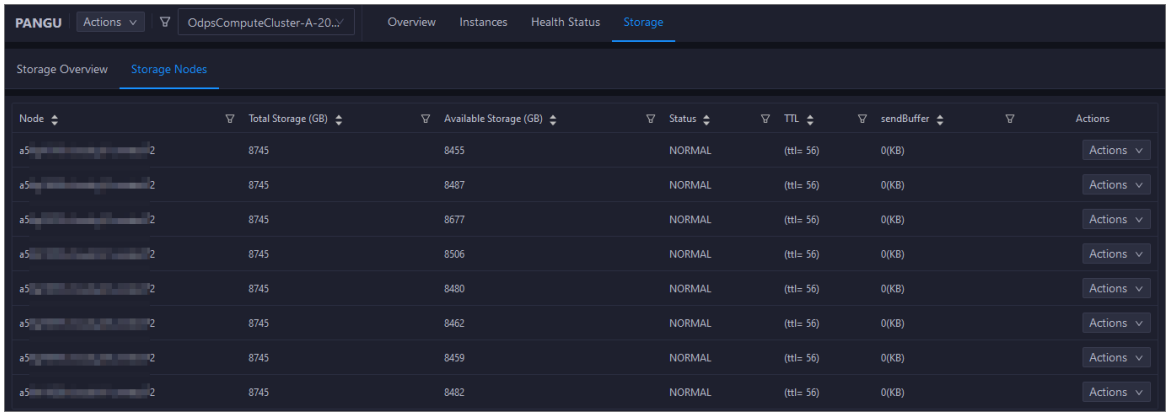

The **Storage Nodes** page displays the information about all storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, TTL, and send buffer size.

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.

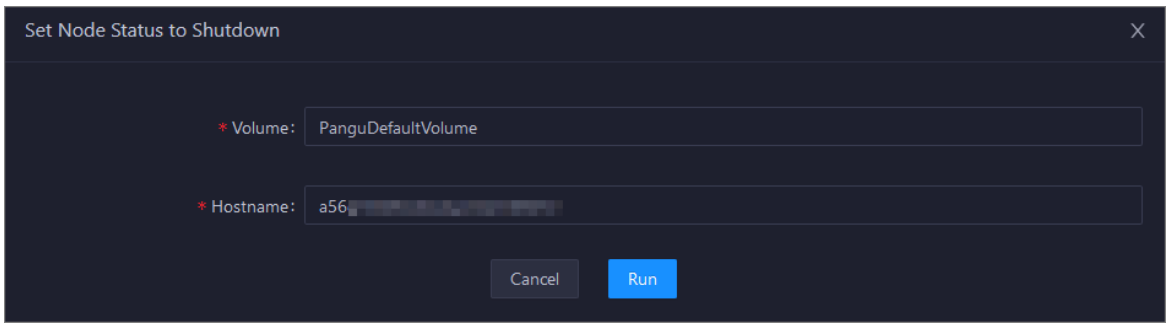

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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.

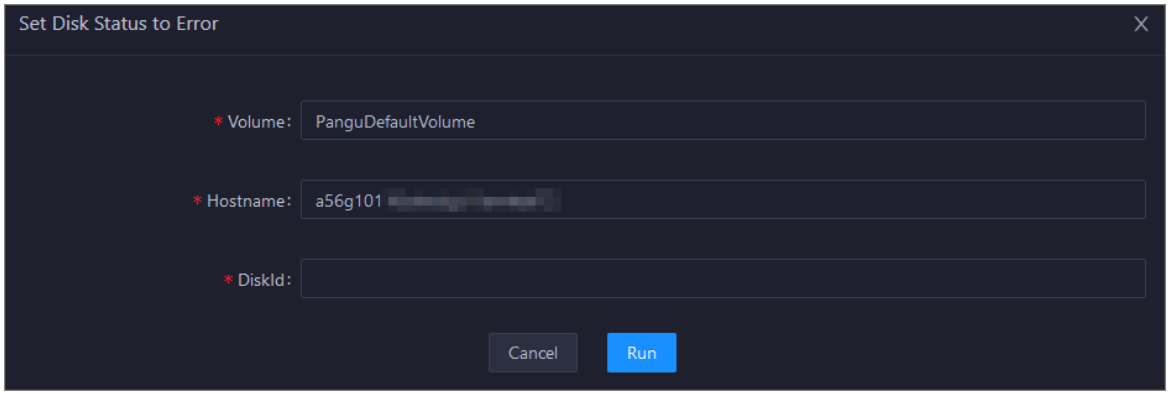

   The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.
3. Click **Run**. A message appears, indicating that the action has been submitted.

11.2.5.3.3.6. Change the primary master node of Apsara Distributed File System

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

Prerequisites

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

Background information

A volume in Apsara Distributed File System is similar to a namespace. The default volume is PanguDefaultVolume. If a cluster contains a large number of nodes, multiple volumes may exist. A volume has three master nodes. One of the nodes serves as the primary master node, and the other two nodes serve as secondary master nodes.

Procedure

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

5. In the **Master Status** section of the **Instances** tab, find the target master node, click **Actions** in the Actions column, and select **Change Primary Master Node**. In the pane that appears, specify the required parameters.

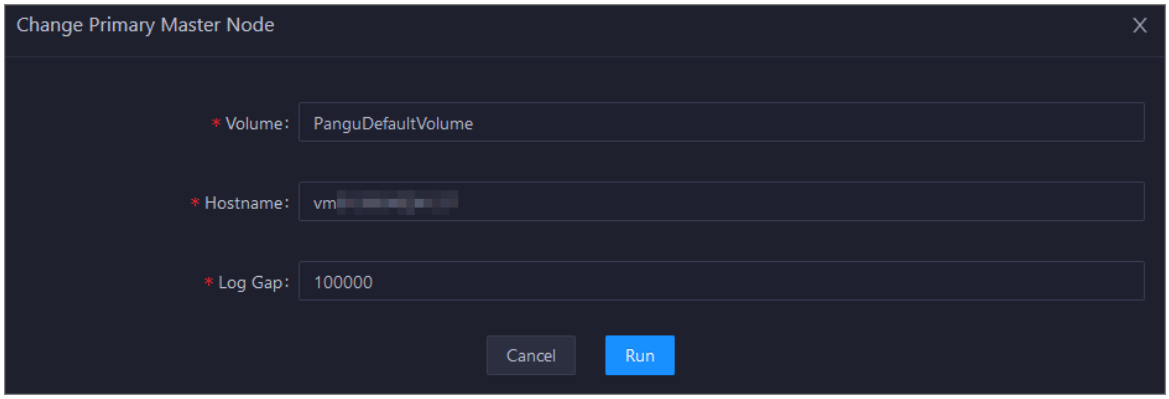

Parameter description:

- **Volume**: the volume whose primary master node needs to be changed. Default value: **PanguDefaultVolume**. If a cluster contains 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 you want to switch. During the switchover, the system checks the log number gap. If the gap is less than the specified value, the switchover is allowed. Otherwise, you cannot change the primary master node. Default value: **100000**.

6. Click **Run**. The **Change Primary Master Node** pane appears.

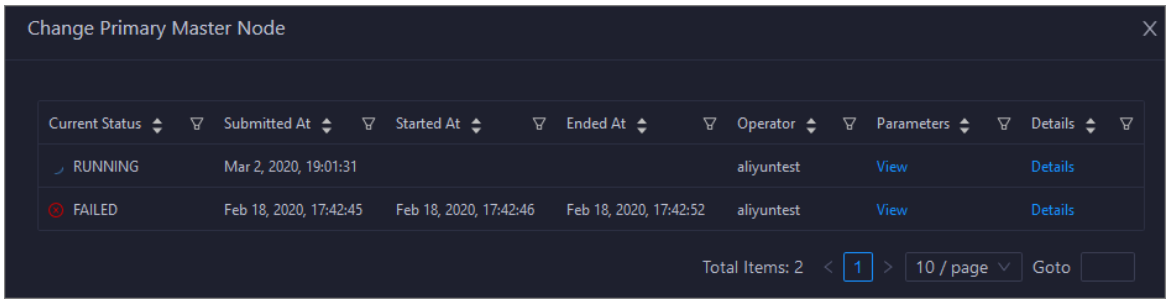

The **Change Primary Master Node** pane shows the switchover history. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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

11.2.5.3.3.7. Clear the recycle bin of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to clear the recycle bin of Apsara Distributed File System to release storage space.

Prerequisites

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

Procedure

1. In the left-side navigation pane of the **Services** tab, click **Pangu**. Then, select a cluster. The **Overview** tab for the selected cluster appears.
2. In the upper-right corner, click **Actions** and select **Empty Recycle Bin**.
3. In the pane that appears, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

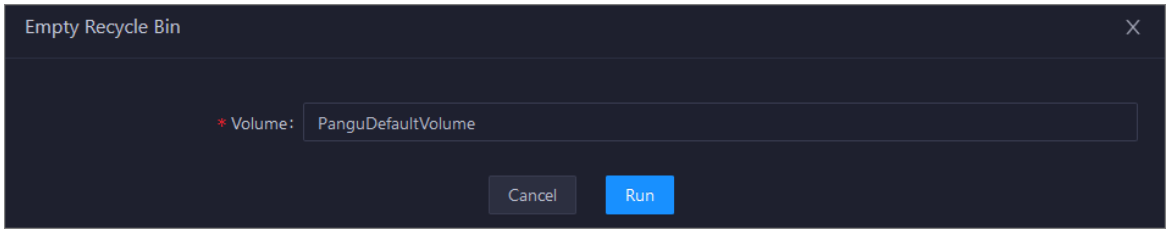

4. Click **Run**.
5. View the execution status.

   In the upper-right corner, click **Actions** and select **Execution History** next to **Empty Recycle Bin** to view the execution history.

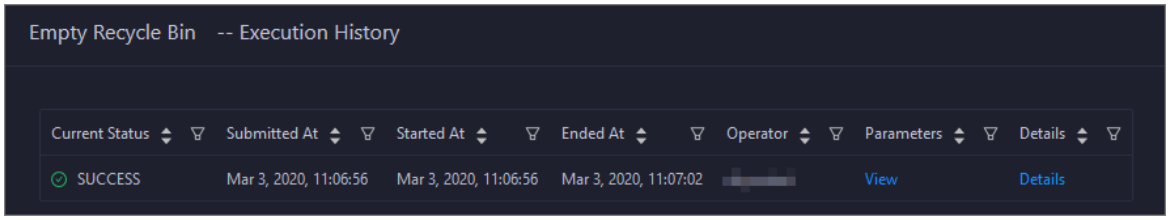

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

6. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

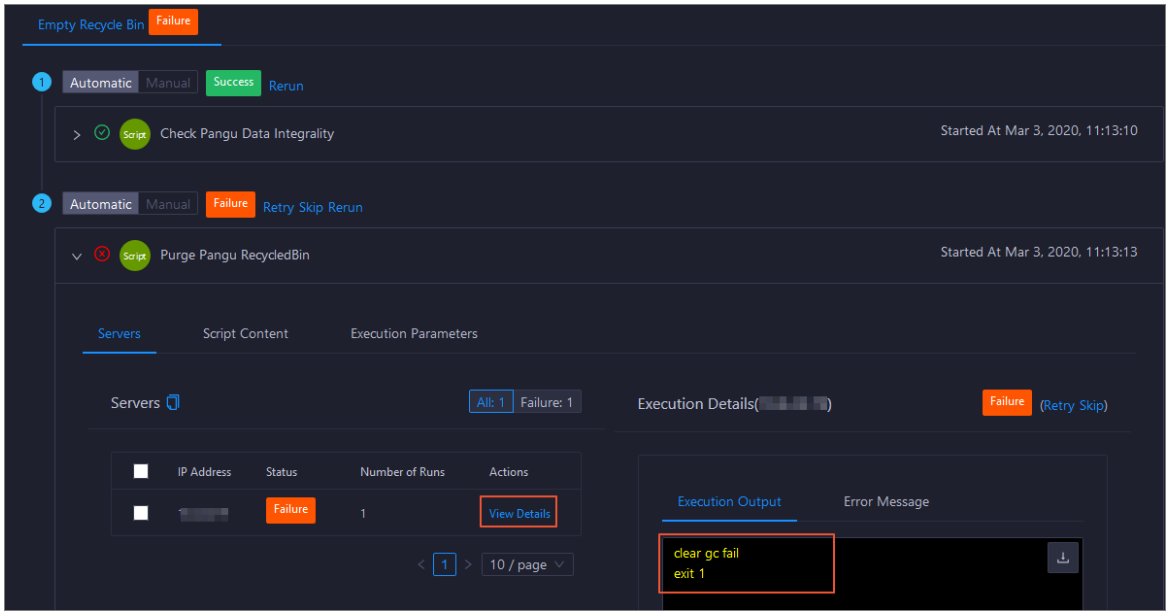

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

11.2.5.3.3.8. 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. In the left-side navigation pane of the **Services** tab, click **Pangu**. Then, select a cluster. The Overview tab for the selected cluster appears.
2. In the upper-right corner of the tab that appears, click **Actions** and select **Disable Data Rebalancing**.
3. In the pane that appears, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

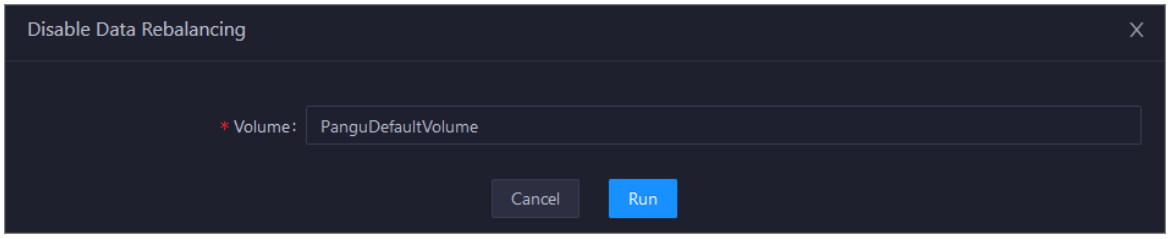

4. Click **Run**.
5. View the execution status.

   Click **Actions** and select **Execution History** next to **Disable Data Rebalancing** to view the execution history.

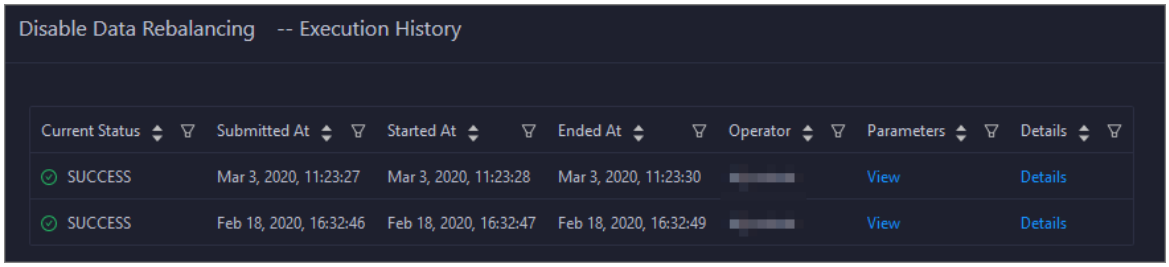

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

6. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of a failure.

Enable data rebalancing

1. In the left-side navigation pane of the **Services** tab, click **Pangu**. Then, select a cluster. The Overview tab for the selected cluster appears.

2. In the upper-right corner of the tab that appears, click **Actions** and select **Enable Data Rebalancing**.
3. In the pane that appears, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

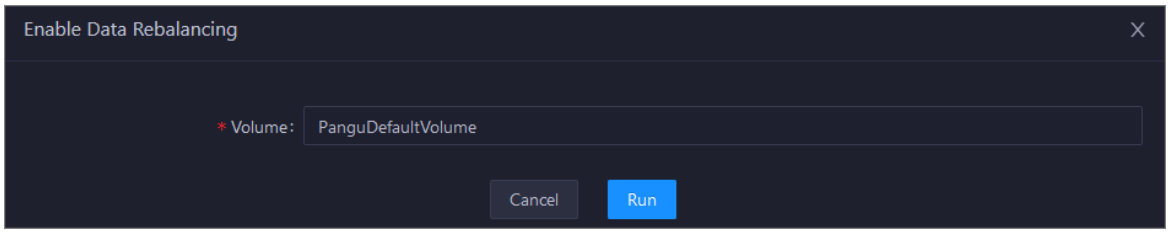

4. Click **Run**.
5. View the execution status.

   Click **Actions** and select **Execution History** next to **Enable Data Rebalancing** to view the execution history.

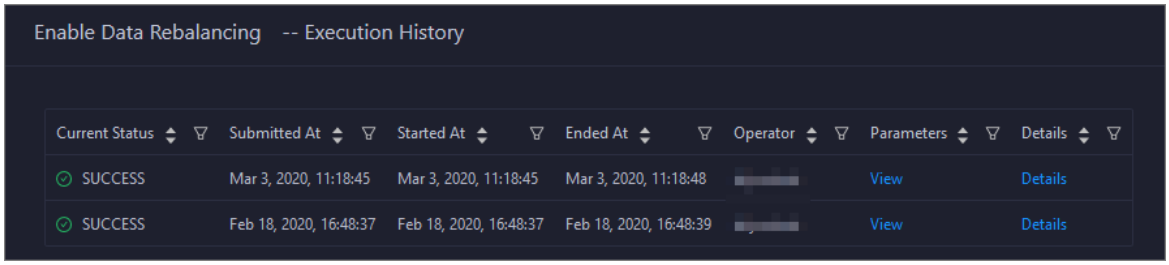

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.
6. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of a failure.

Identify the cause of a failure

This section uses the procedure of identifying the cause of the failure to enable data rebalancing as an example.

1. In the Execution History pane, click **Details** in the Details column for a failed execution.
2. In the pane that appears, click **View Details** for a failed step to identify the cause of the failure.

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

11.2.5.3.3.9. 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 to disks. If Apsara Distributed File System is faulty, you can use checkpoints to restore data to the status before the failure. This ensures 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. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Pangu**. Then, select a cluster and click the **Instances** tab.
5. In the **Master Status** section of the **Instances** tab, find the target master node, click **Actions** in the Actions column, and select **Run Checkpoint on Master Node**. In the pane that appears, set the **Volume** parameter.

   > **Note** The default value of **Volume** is **PanguDefaultVolume**.

6. Click **Run**. The **Run Checkpoint on Master Node** pane appears.

   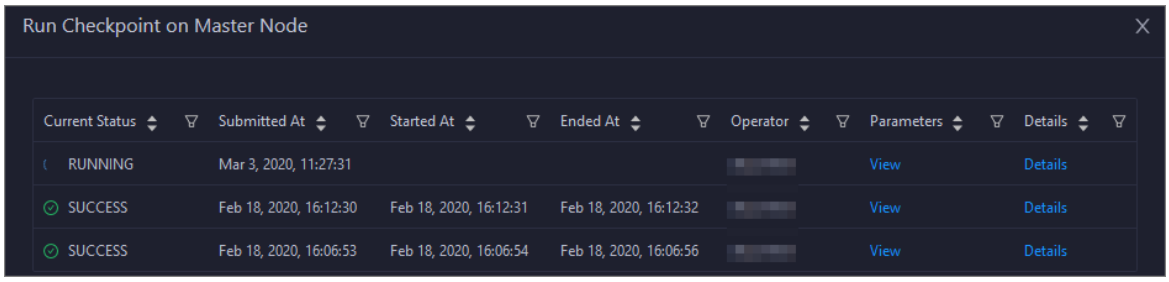

   The **Run Checkpoint on Master Node** pane shows the execution history of the checkpoint on the master node. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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

11.2.5.3.4. Tunnel service

11.2.5.3.4.1. O&M features and entry

This topic describes the definition and O&M features of the Tunnel service. It also provides more information about how to go to the Tunnel service O&M page.

Definition of the Tunnel service

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

Tunnel O&M features

- Overview: shows information about the Tunnel service. The information includes the service overview, service status, and throughput trend chart.
- Instances: shows information about the server roles of the Tunnel service.
- Restart Tunnel Server: allows you to restart one or more Tunnel servers.

Go to the Tunnel Service page

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Tunnel Service**. The **Overview** tab for the Tunnel service appears.

11.2.5.3.4.2. Overview

The Overview tab for the Tunnel service shows key operating information. The information includes the service overview, service status, and throughput.

Go to the Overview tab

In the left-side navigation pane of the **Services** tab, click **Tunnel Service**. The **Overview** tab for the Tunnel service appears.

The **Overview** tab shows key operating information about the Tunnel service. The information includes the service overview, service status, and throughput trend chart.

Services

This section shows the numbers of available services, unavailable services, and services that are being updated.

Roles

This section shows all Tunnel server roles and their states. You can also view the expected and actual numbers of hosts for each server role.

Tunnel Throughput

The Tunnel Throughput (Bytes/Min) chart shows the trend lines of the inbound and outbound traffic in different colors. The trend chart is periodically refreshed. You can also manually refresh the trend chart. You can also view the trend chart of Tunnel throughput in a specific period.

11.2.5.3.4.3. Instances

The Instances tab shows information about the Tunnel server roles. The information includes the role name, hostname, IP address, role status, and host status.

Go to the Instances tab

In the left-side navigation pane of the **Services** tab, click **Tunnel Service**. Then, click the **Instances** tab. The **Instances** tab for the Tunnel service appears.

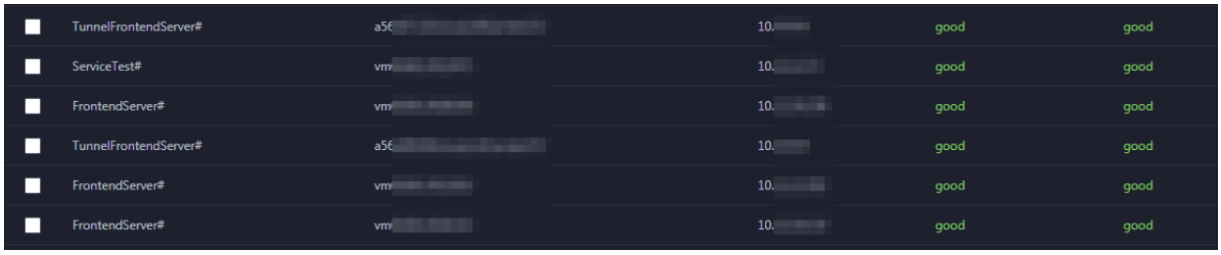

The **Instances** tab shows information about all Tunnel server roles, including the role name, hostname, IP address, role status, and host status. The status can be **good**, **error**, or **upgrading**.

11.2.5.3.4.4. Restart Tunnel servers

Apsara Bigdata Manager (ABM) allows you to restart Tunnel servers for the corresponding server roles.

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.

Step 1: Restart Tunnel servers

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Services** tab.
4. In the left-side navigation pane of the **Services** tab, click **Tunnel Service**. Then, click the **Instances** tab.

5. On the Instances tab, select one or more server roles for which you want to restart the Tunnel service. In the upper-right corner, click **Actions** and select **Restart Tunnel Server**.
6. In the **Restart Tunnel Server** pane, specify the required parameters.

   The following table describes the parameters.

Parameter	Description
Force Restart	Specifies whether to forcibly restart the Tunnel server for the selected server role. Valid values: ○ **no_force**: Do not forcibly restart the Tunnel server. If a server role is in the running state, the corresponding Tunnel server is not restarted. ○ **force**: Forcibly restart the Tunnel server. The Tunnel server is restarted regardless of the server role state.
Hostname	The hostname of the selected server role. The value is automatically provided. You do not need to specify a value for this parameter.

7. Click **Run**.

Step 2: View the execution status or progress

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

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.
2. If the status is **RUNNING**, click **Details** in the Details column to view the execution progress.

Step 3: (Optional) Identify the cause of a failure

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

1. On the **Overview** or **Instances** tab of the **Tunnel Service** page, click **Actions** in the upper-right corner. Then, select **Execution History** next to **Restart Tunnel Server** to view the execution history.
2. In the pane that appears, 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 identify the cause of the failure.

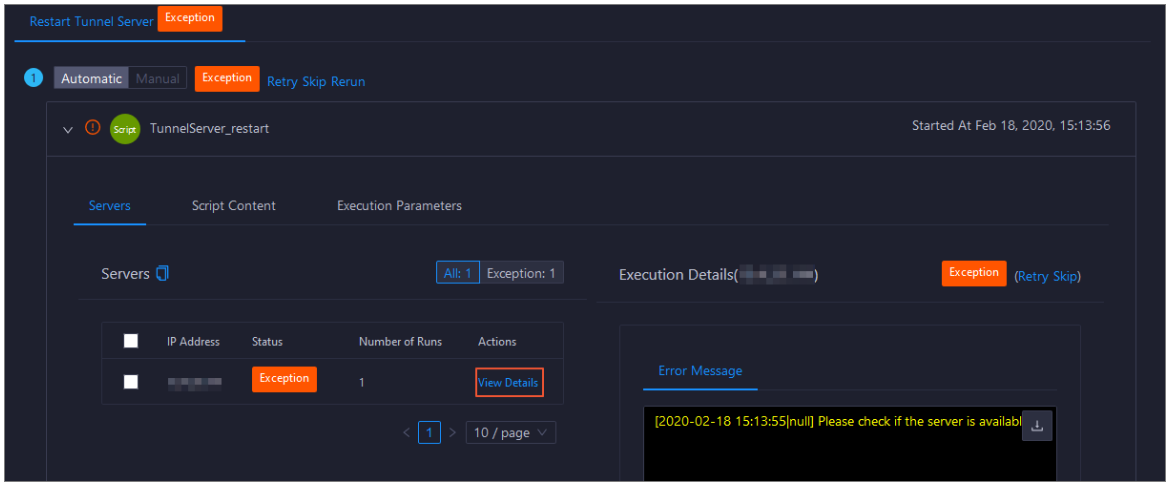

11.2.5.4. Cluster O&M

11.2.5.4.1. O&M features and entry

This topic describes MaxCompute cluster O&M features. It also provides more information about how to go to the MaxCompute cluster O&M page.

Cluster O&M features

MaxCompute cluster O&M features:

- Overview: shows the overall running information about a cluster. You can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU utilization, 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: shows all checkers for a cluster. You can query checker details, check results for hosts in the cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.
- Servers: shows information about hosts in a cluster. The information includes the hostname, IP address, role, type, CPU utilization, memory usage, root disk usage, packet loss rate, and packet error rate.
- Scale in Cluster or Scale out Cluster: allows you to scale in or scale out a MaxCompute cluster by adding or removing physical hosts.
- Enable Auto Repair: allows you to enable auto repair for MaxCompute clusters.
- Restore Environment Settings: allows you to restore environment settings for multiple hosts in the MaxCompute cluster at a time.

Go to the Clusters tab

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Clusters** tab.
4. In the left-side navigation pane of the **Clusters** tab, select a cluster. The **Overview** tab for the

cluster appears.

11.2.5.4.2. Overview

This topic describes how to go to the Overview tab of a MaxCompute cluster. It also shows the cluster overview and describes the operations that you can perform on this tab.

Go to the Overview tab

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Clusters** tab.
4. In the left-side navigation pane of the **Clusters** tab, select a cluster. The **Overview** tab for the selected cluster appears.

On the **Overview** tab, you can quickly log on to a host that is commonly used in MaxCompute cluster O&M. You can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, and packet transmission for the cluster.

Log on

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

1. In the **Log on** section, click the hostname in the **Hostname** column. The **Hosts** tab for the host appears.
2. In the left-side navigation pane, click the **Log On** icon of the host. The **TerminalService** page appears.

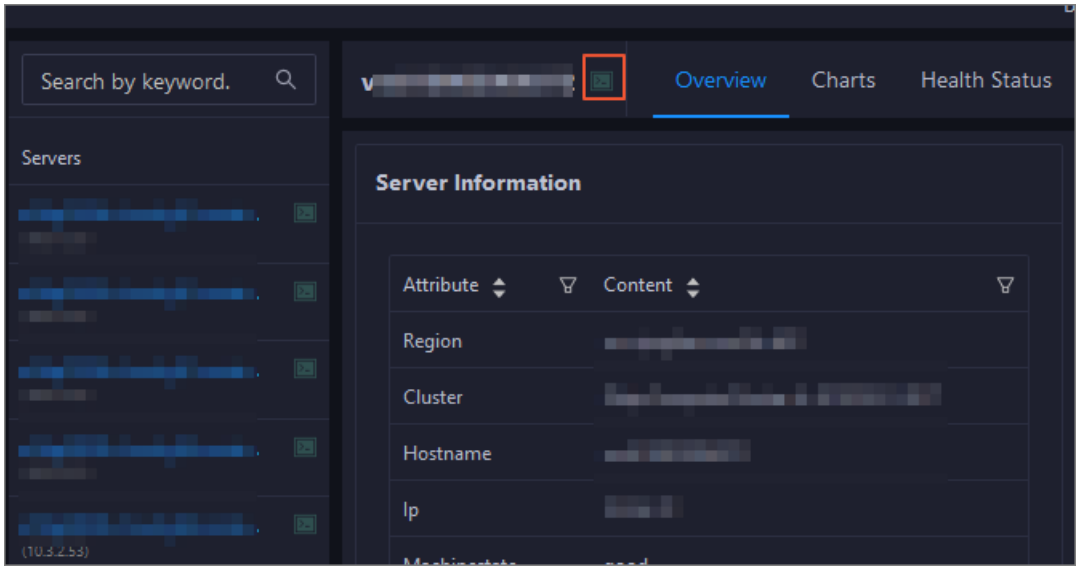

3. In the left-side navigation pane, click the hostname to log on to the host.

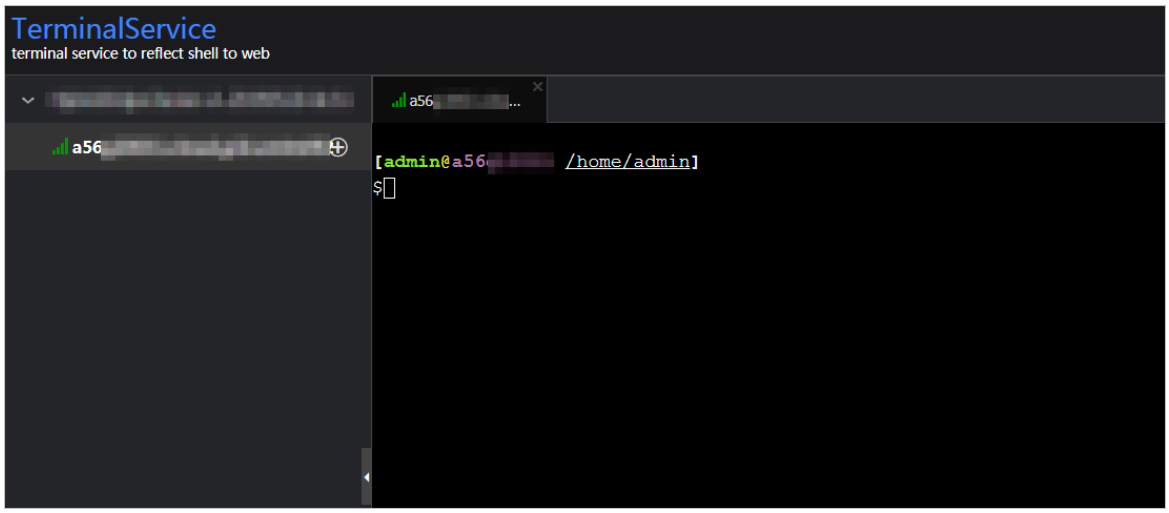

Servers

This section shows all host status and the number of hosts in each state. A host can be in the **good** or **error** state.

Services

This section shows all services deployed in the cluster and the numbers of services in the **good** and **bad** states.

CPU

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

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

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

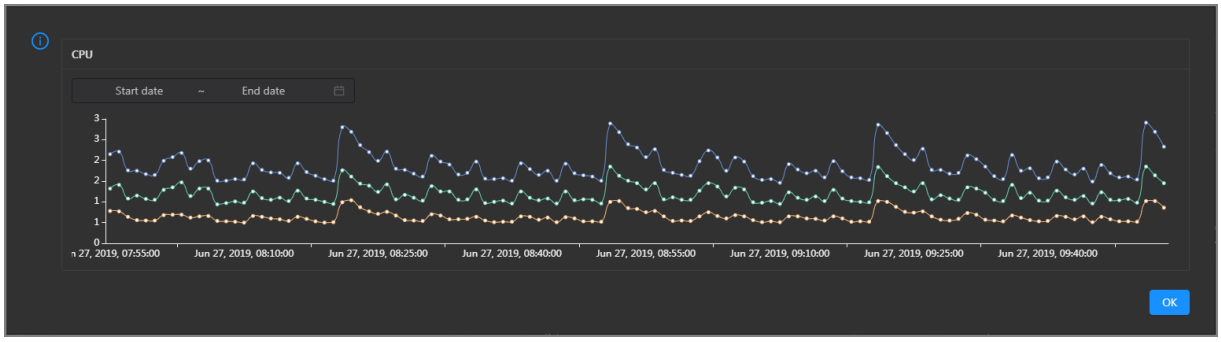

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

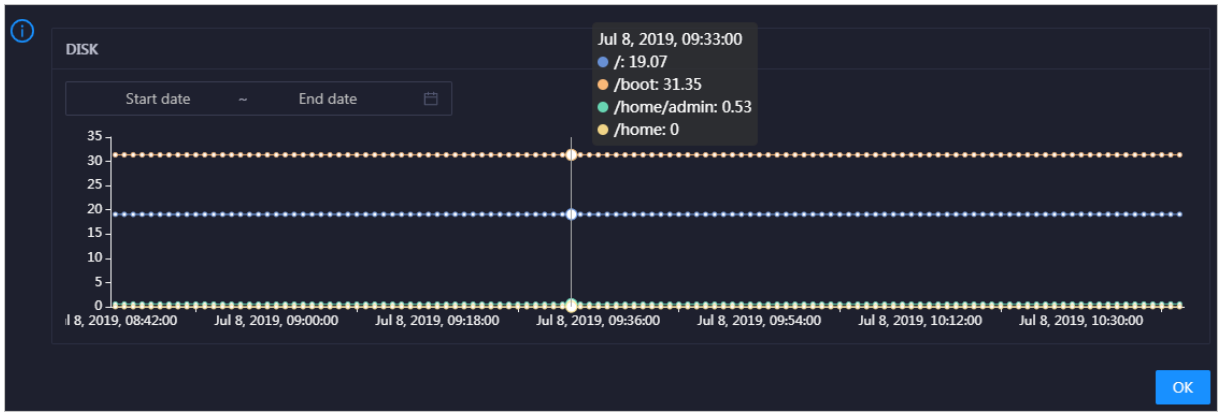

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

LOAD

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

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

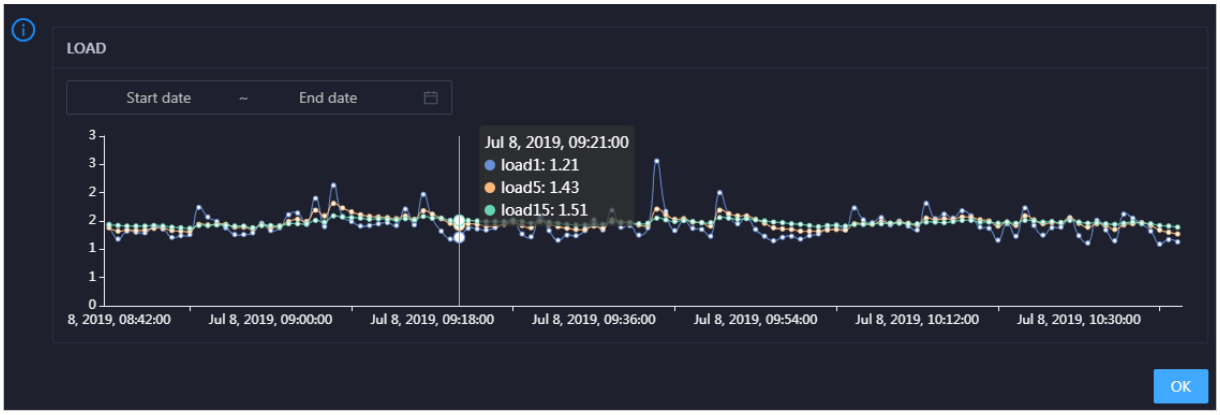

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

MEMORY

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

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

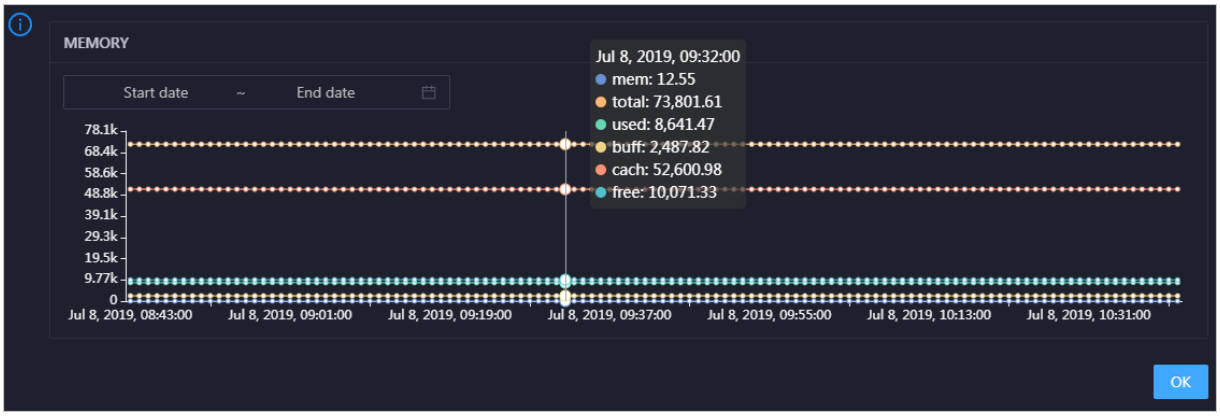

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

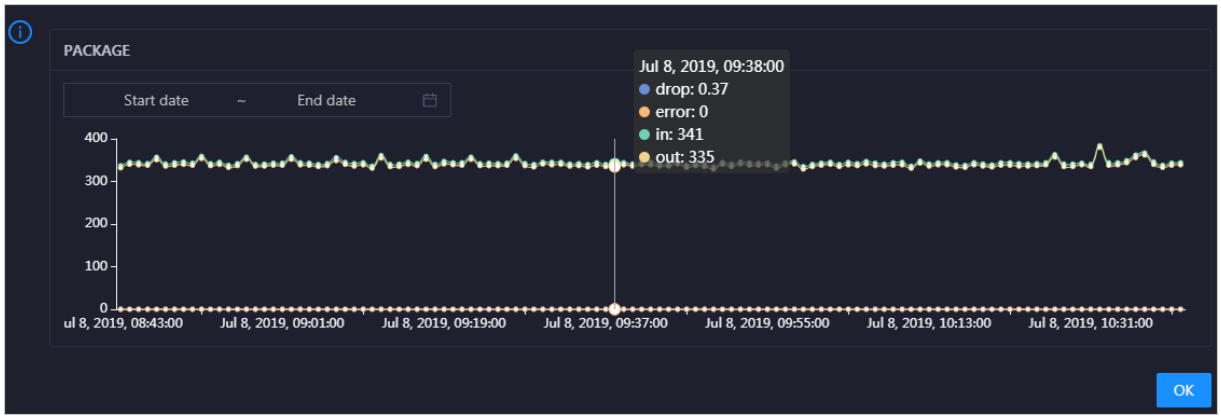

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

Health Check

This section shows the number of checkers for the cluster and the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

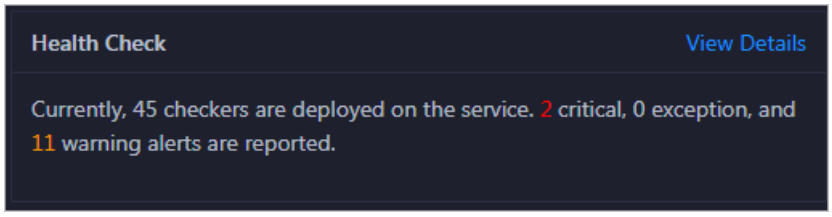

Click **View Details** to go to the Health Status tab. On this tab, you can view health check details. For more information, see Cluster health.

Health Check History

This section shows the records of the health checks performed on the cluster. You can view the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

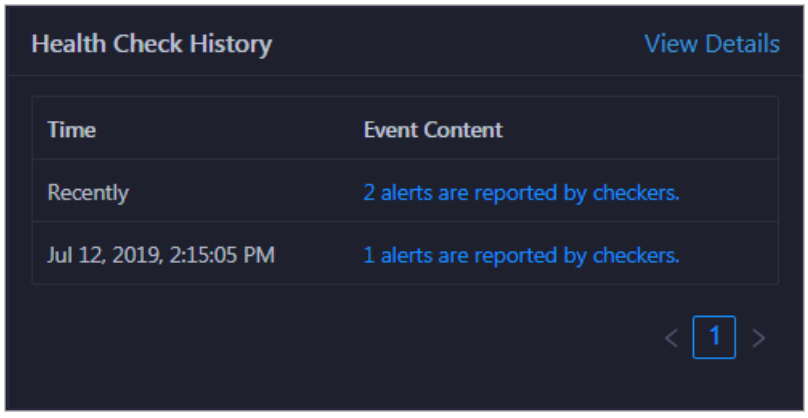

Click **View Details** to go to the Health Status tab. On this tab, you can view health check details. For more information, see Cluster health.

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

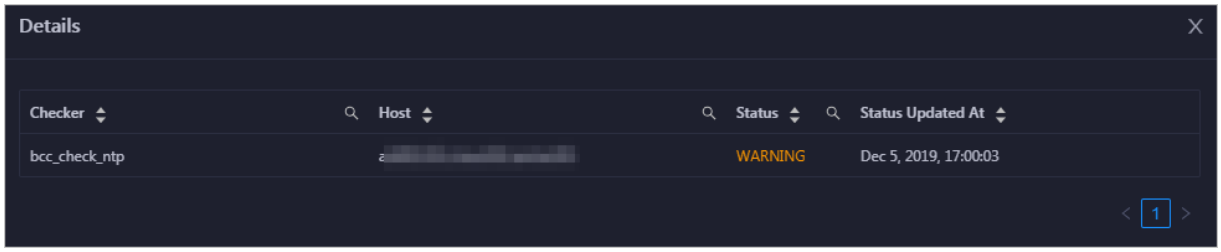

11.2.5.4.3. Cluster health

The Health Status tab shows all checkers for a cluster. You can query checker details, check results for hosts in the cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.

Go to the Health Status tab

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Clusters** tab.
4. In the left-side navigation pane of the **Clusters** tab, select a cluster. Then, click the **Health Status** tab. The **Health Status** tab for the selected cluster appears.

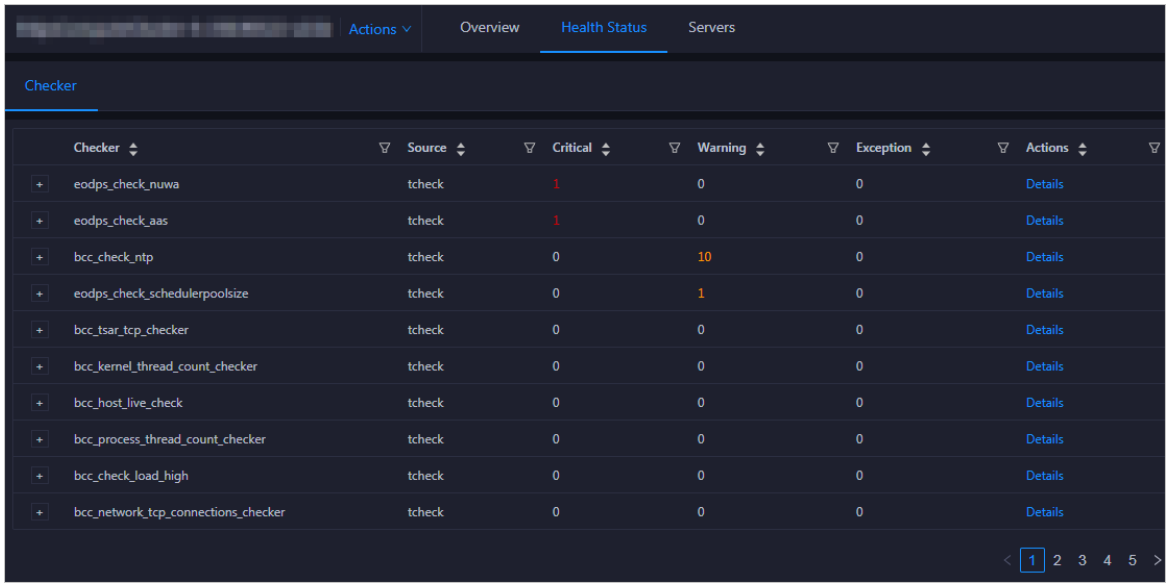

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are repesented in different colors. Handle the alerts in a timely manner,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

View checker details

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

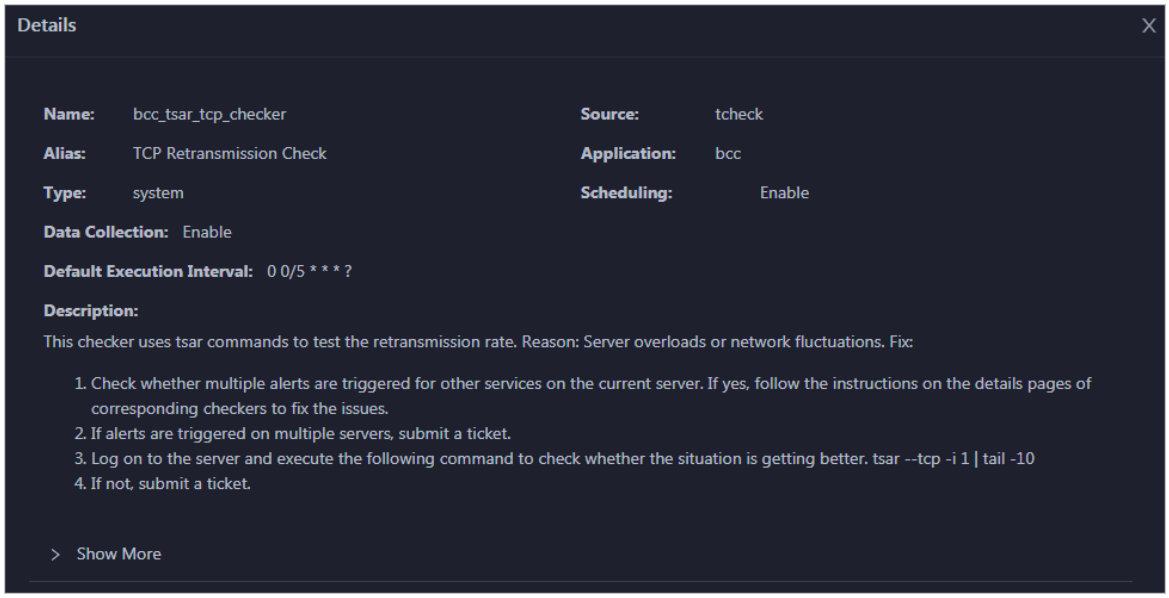

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

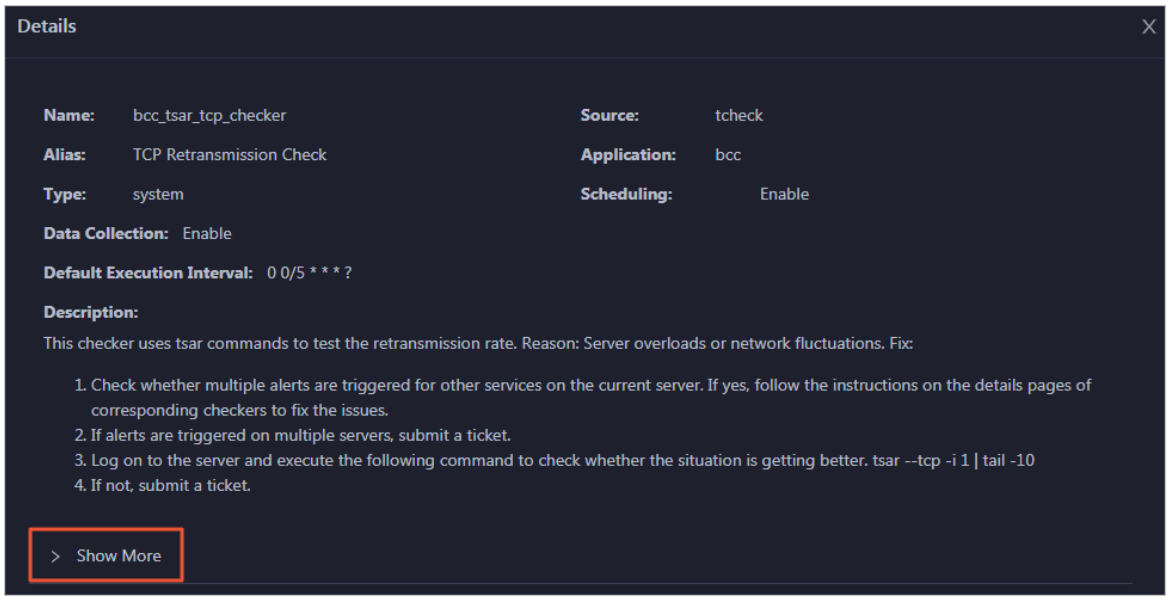

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

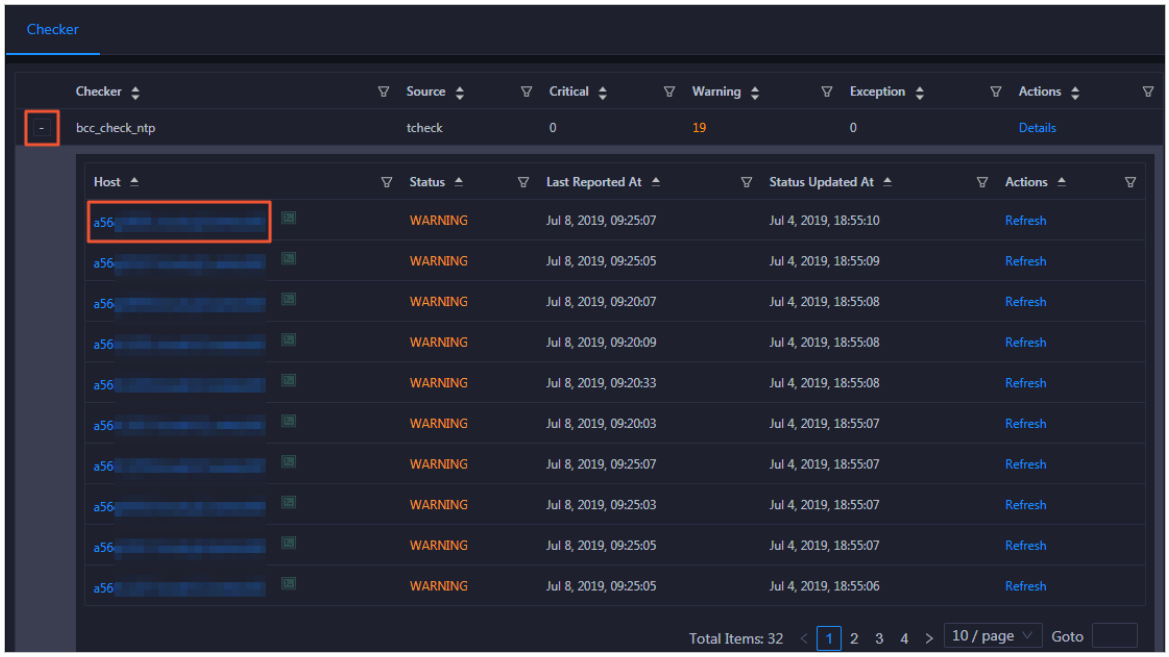

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

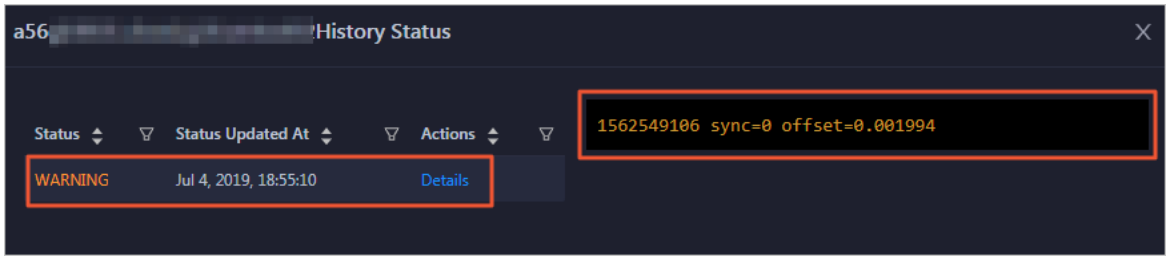

Clear alerts

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

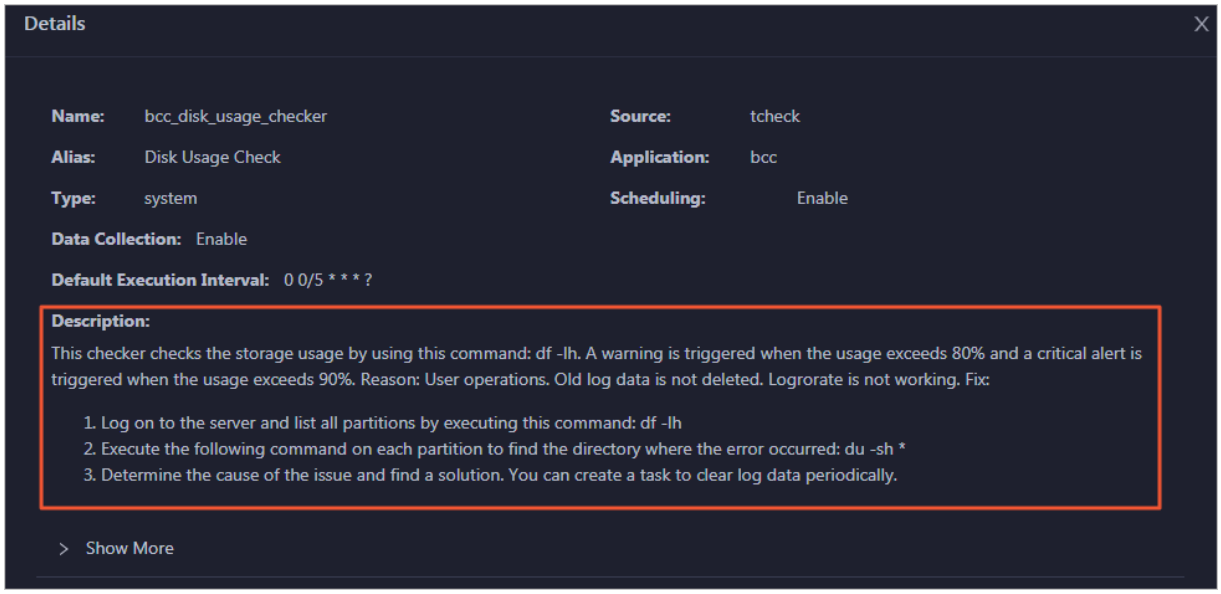

Log on to a host

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

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

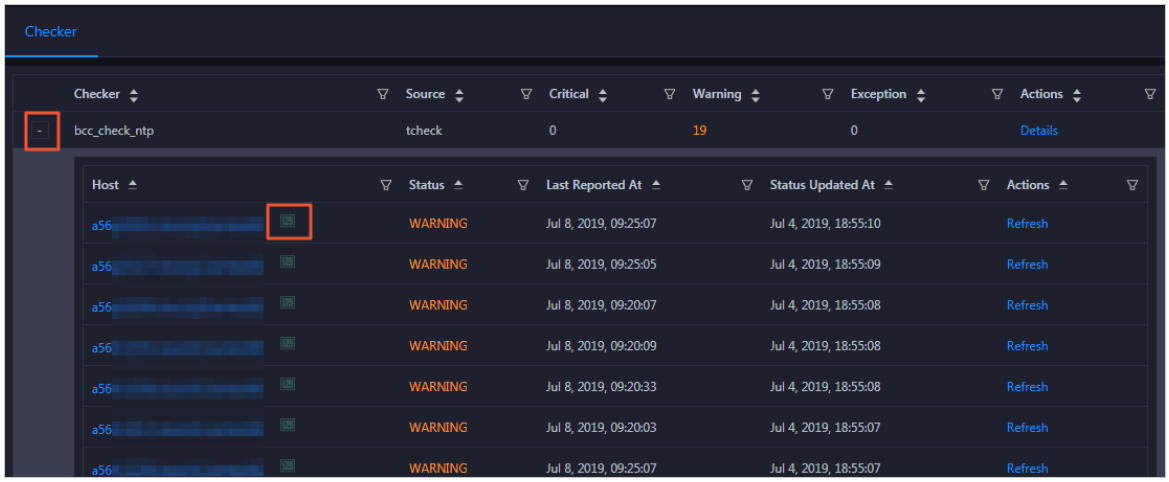

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

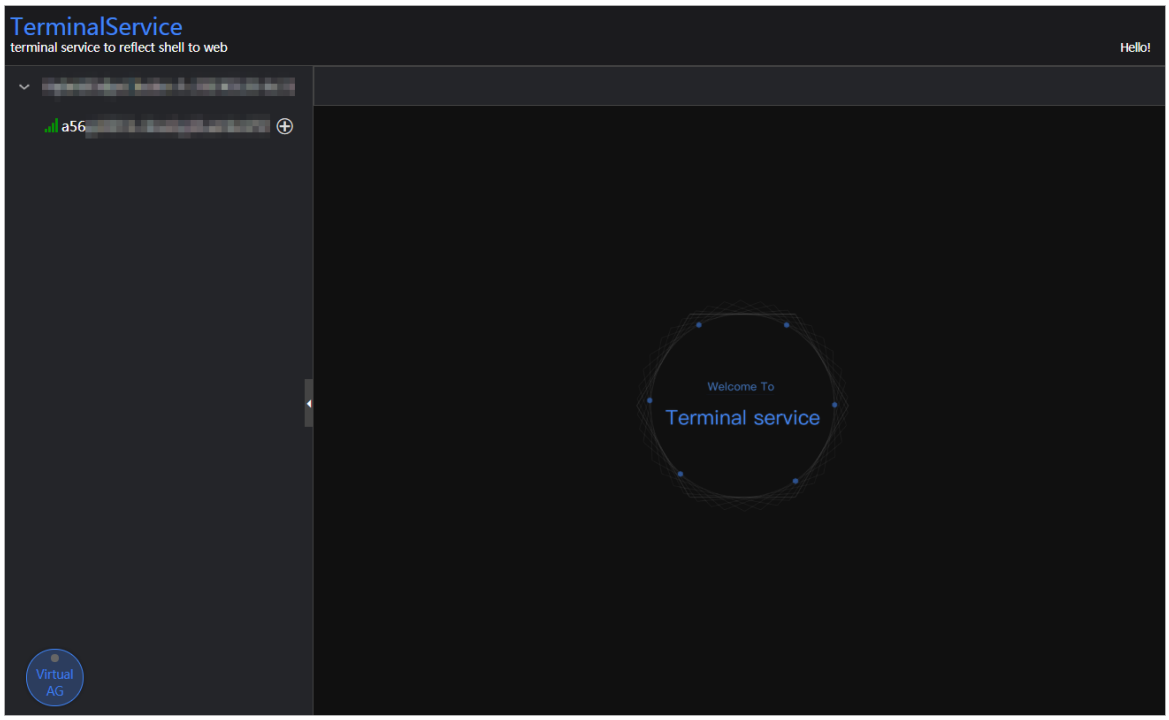

3. On the **TerminalService** page, click the hostname to log on to the host.

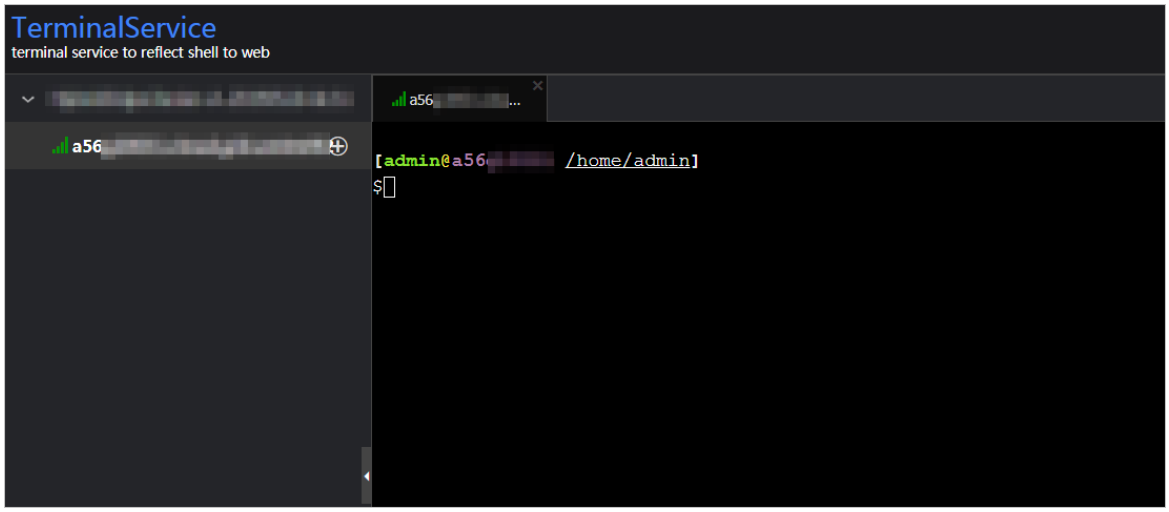

Run a checker again

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

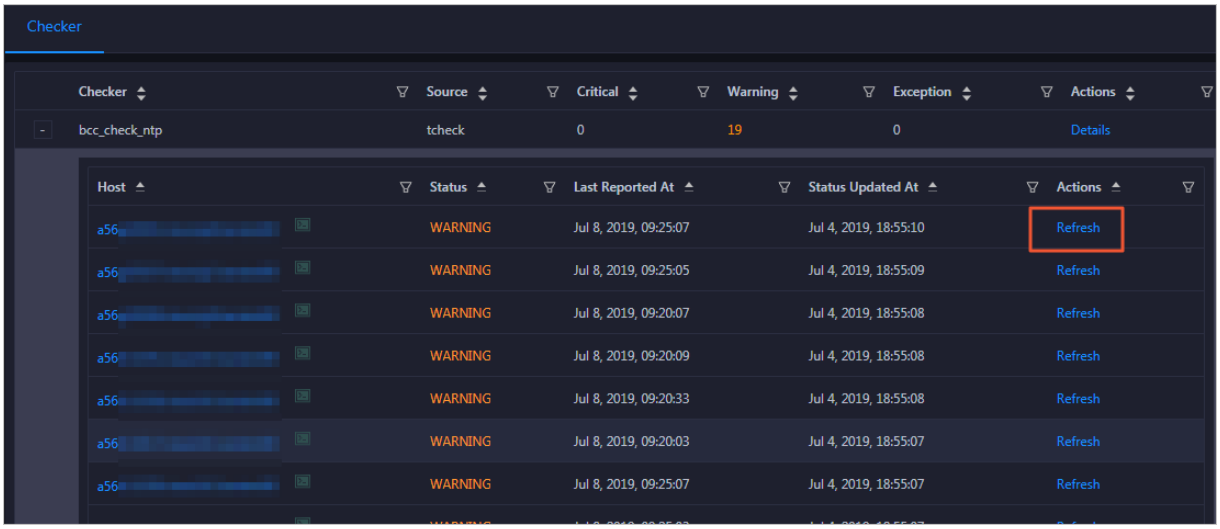

11.2.5.4.4. Cluster hosts

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.

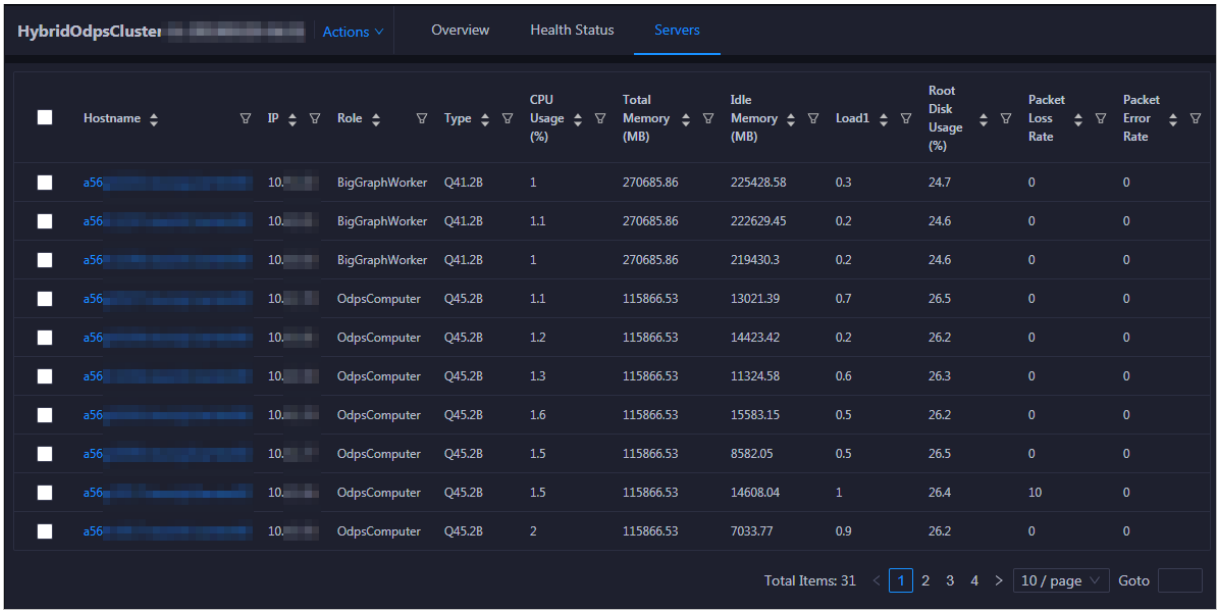

To view more information about a host, click the name of the host. The Host overview page appears.

11.2.5.4.5. Scale in and scale out a MaxCompute cluster

Apsara Bigdata Manager (ABM) supports MaxCompute cluster scaling. You can scale in a MaxCompute cluster by removing physical hosts from the MaxCompute cluster to the default cluster of Apsara Infrastructure Management Framework. You can scale out a MaxCompute cluster by adding physical hosts in the default cluster of Apsara Infrastructure Management Framework to the MaxCompute cluster.

Description

In Apsara Stack, scaling out a cluster involves complex operations. You must 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 is an idle resource pool that provides resources to scale out clusters. If you want to scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework. If you want to scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster.

You can use this method to scale out or in a MaxCompute cluster in the ABM console.

Prerequisites

- Scale-out: The physical host that you want to add is an **SInstance** host in the default cluster of Apsara Infrastructure Management Framework.
- Scale-out: The template host must be an **SInstance** host. You can log on to the admingateway host in a MaxCompute cluster to view SInstance hosts.
- Scale-in: The physical host that you want to remove is an **SInstance** host. You can log on to the admingateway host in a 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 add hosts to a MaxCompute cluster, you must specify an existing host as the template host. The hosts that you want to add copy configurations from the template host. This allows the hosts to 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.

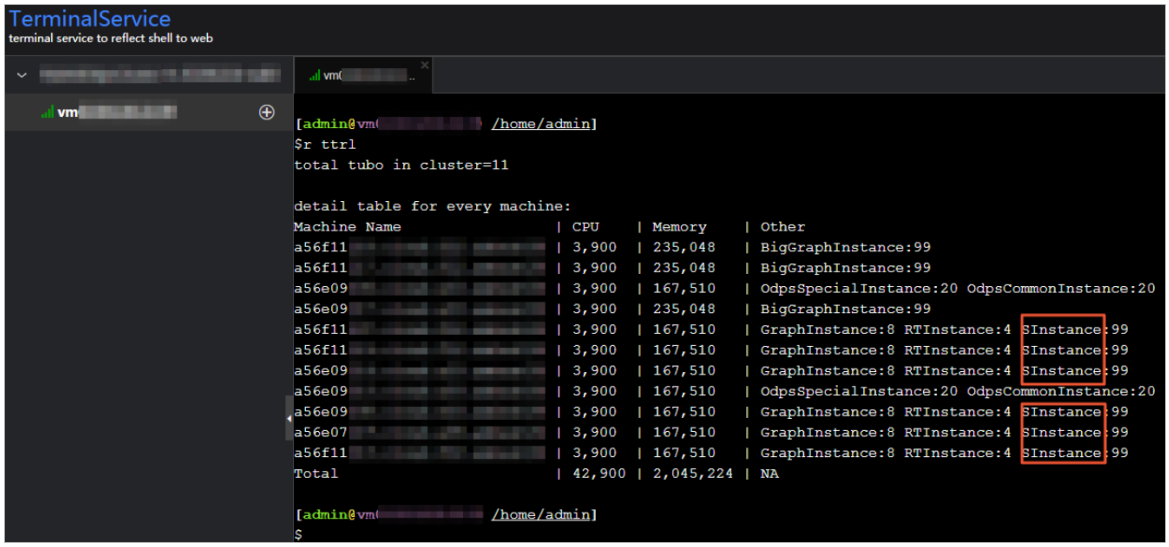

2. In the left-side navigation pane of the **Clusters** tab, select a cluster. Then, click the **Servers** tab. On the tab that appears, select an **SInstance** host and use it as the template host.

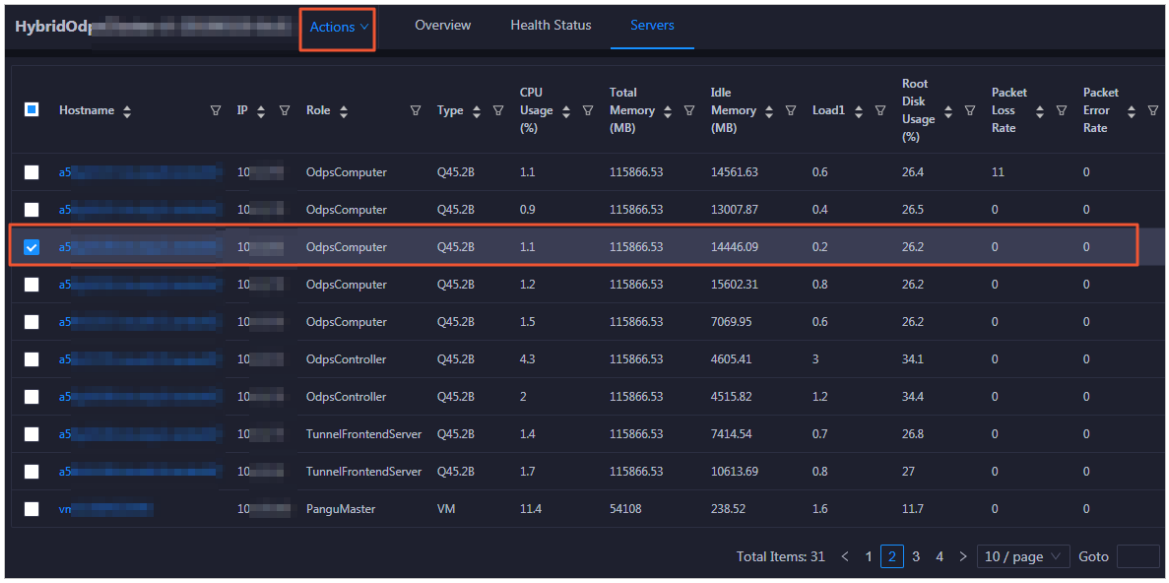

3. In the upper-right corner, click **Actions** and select **Scale out Cluster**. In the **Scale out Cluster** pane, specify the required parameters.

   Parameters:

   - **Region**: the region of the host that you want to add.
   - **Refer Hostname**: the name of the template host. By default, the name of the selected host is used.
   - **Hostname**: the name of the host that you want to add. The drop-down list displays all available hosts in the default cluster for scale-out operations. You can select one or more hosts

from the drop-down list.

4. Click **Run**.
5. View the scale-out status.

   In the upper-right corner, click **Actions** and select **Execution History** next to **Scale out Cluster** to view the scale-out history.

   It requires some time for the cluster to be scaled out. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **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 execution.

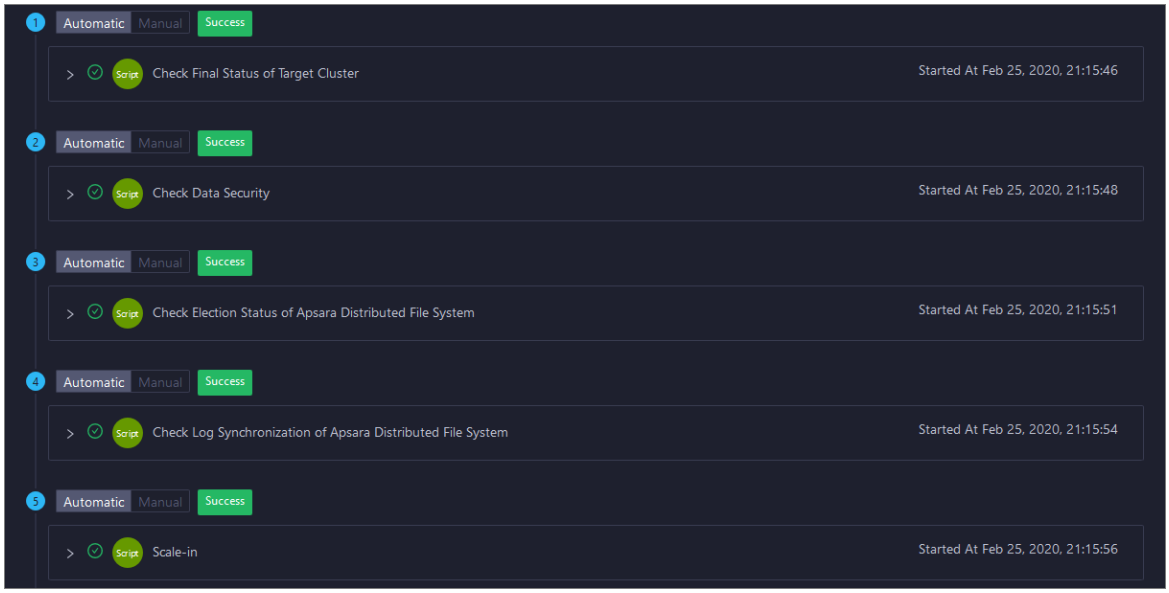

7. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of a scale-in or scale-out failure.

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.

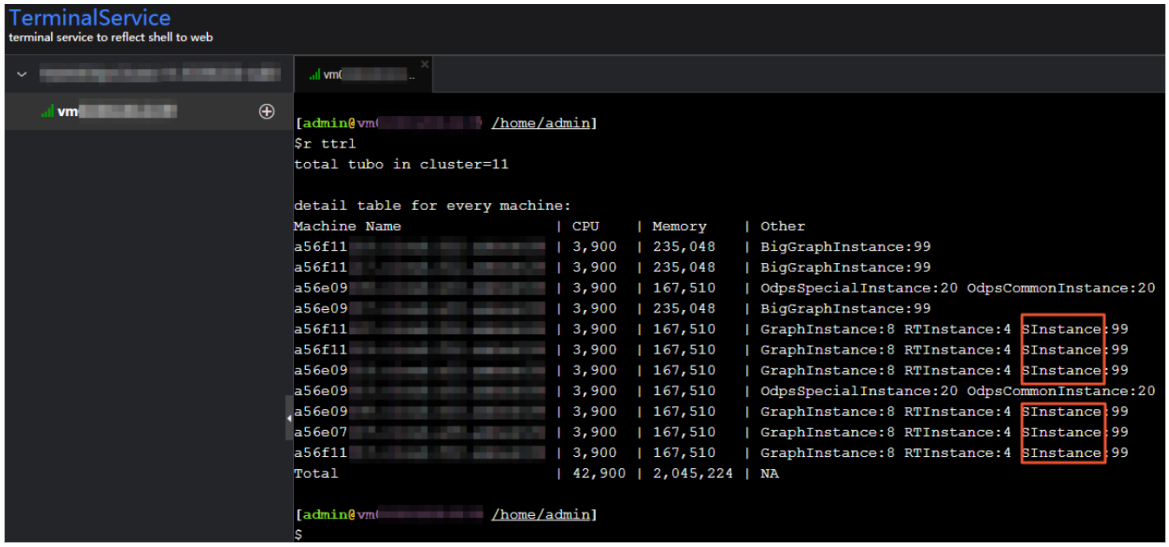

2. On the **Clusters** tab, select a cluster in the left-side navigation pane. Then, click the **Servers** tab. On the tab that appears, select one or more **SInstance** hosts that you want to remove.

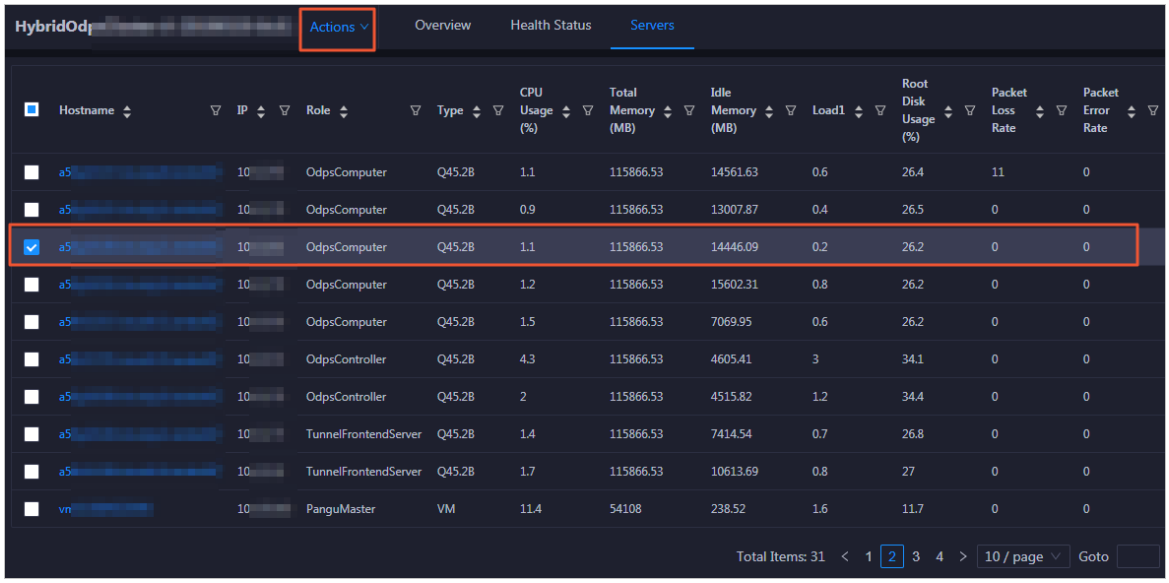

3. In the upper-right corner, click **Actions** and select **Scale in Cluster**. In the **Scale in Cluster** pane, set the **Hostname** parameter.

   Parameters:

   - **Region**: the region of the host that you want to remove.
   - **Hostname**: the name of the host that you want to remove. By default, the name of the selected host is used.

4. Click **Run**.

5. View the scale-in status.

   In the upper-right corner, click **Actions** and select **Execution History** next to **Scale in Cluster** to view the scale-in history.

   It requires some time for the cluster to be scaled in. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **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 execution.
7. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of a scale-in or scale-out failure.

Identify the cause of a scale-in or scale-out failure

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

1. In the upper-right corner of the **Clusters** tab, click **Actions** and select **Execution History** next to **Scale in Cluster** to view the scale-in history.
2. Click **Details** in the Details column of a failed operation to identify the cause of the failure.

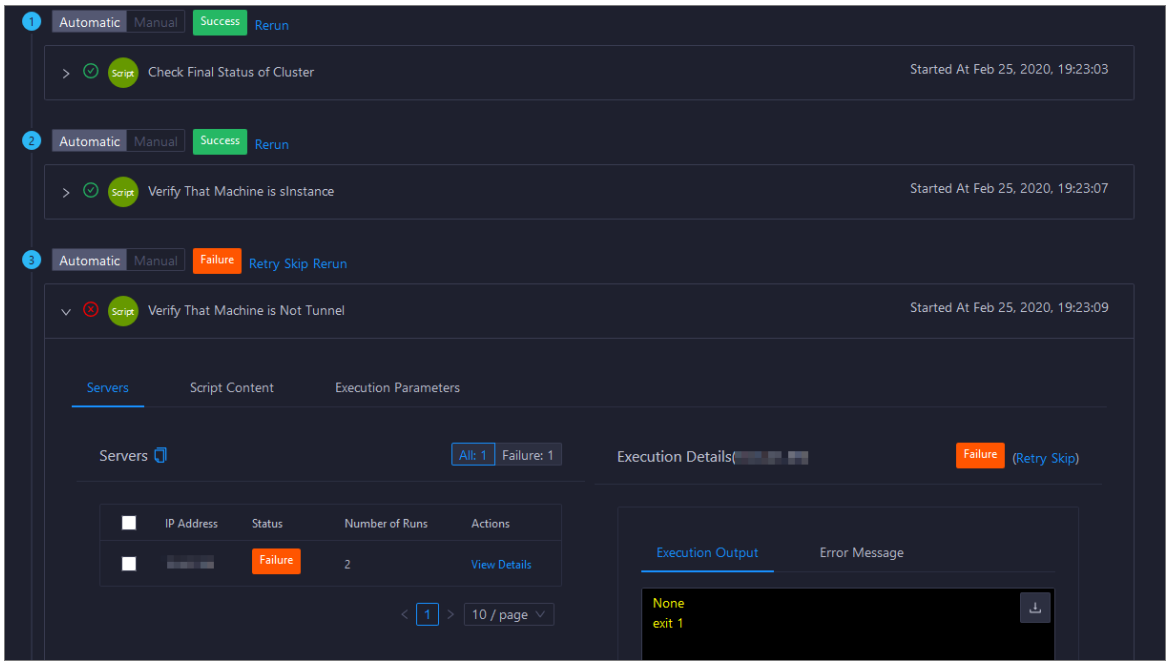

You can view information about parameter settings, host details, scripts, and runtime parameters to identify the cause of the failure.

11.2.5.4.6. Restore environment settings and enable auto repair

Apsara Bigdata Manager (ABM) allows you to restore the environment settings for multiple hosts in a MaxCompute cluster at a time. It also allows you to enable the auto repair feature for a MaxCompute cluster.

Restore environment settings

ABM allows you to restore the environment settings for multiple hosts in a MaxCompute cluster at a time.

1. In the upper-right corner of the **Clusters** tab, click **Actions** and select **Restore Environment Settings**. In the **Restore Environment Settings** pane, set the **Hosts** parameter.

> **Note** You can enter the names of multiple hosts and must separate the names with commas (,).

2. Click **Run**.
3. View the restoration status.

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

   It requires some time for the restoration to complete. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.
4. If the status is RUNNING, click **Details** in the Details column to view the steps and progress of the execution.
5. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure.

Enable auto repair

ABM allows you to enable the auto repair feature for a MaxCompute cluster. After this feature is enabled, repair tickets reported by Xunyangjian are automatically handled.

1. In the upper-right corner of the **Clusters** tab, click **Actions** and select **Enable Auto Repair**. In the **Enable Auto Repair** pane, set the **Cluster** parameter and set Auto Repair to Enable.

   Parameters:

   - **Cluster**: the name of the cluster for which you want to enable the auto repair feature.
   - **Auto Repair**: If you require the feature, set it to **Enable**. Otherwise, set it to **Disable**.
2. Click **Run**.
3. View the status of the feature.

   Click **Actions** and select **Execution History** next to **Enable Auto Repair** to view the feature-related operation history.

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.
4. If the status is RUNNING, click **Details** in the Details column to view the steps and progress of the execution.
5. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure.

11.2.5.5. Host O&M

11.2.5.5.1. O&M features and entry

This topic describes MaxCompute host O&M features. It also provides more information about how to go to the host O&M page.

Host O&M features

- Overview: shows brief information about hosts in a MaxCompute cluster. The information includes the server information, server role status, health check result, and health check history. You can also

view the trend charts of CPU utilization, disk usage, memory usage, load, and packet transmission for the host.

- Charts: shows the enlarged trend charts of CPU utilization, memory usage, disk usage, load, and packet transmission.
- Health Status: shows all checkers for a host. You can query checker details, check results for hosts in a cluster, and schemes to clear alerts (if any exists). You can also log on to a host and perform manual checks on the host.
- Services: shows the cluster, service instances, and service instance roles of a host.

Go to the Hosts tab

1. Log on to the Apsara Bigdata Manager (ABM) console.
2. In the upper-left corner, click the icon and then **MaxCompute**.
3. On the MaxCompute page, click **O&M** in the top navigation bar. Then, click the **Hosts** tab.
4. In the left-side navigation pane of the **Hosts** tab, select a host. The **Overview** tab for the host appears.

11.2.5.5.2. 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.

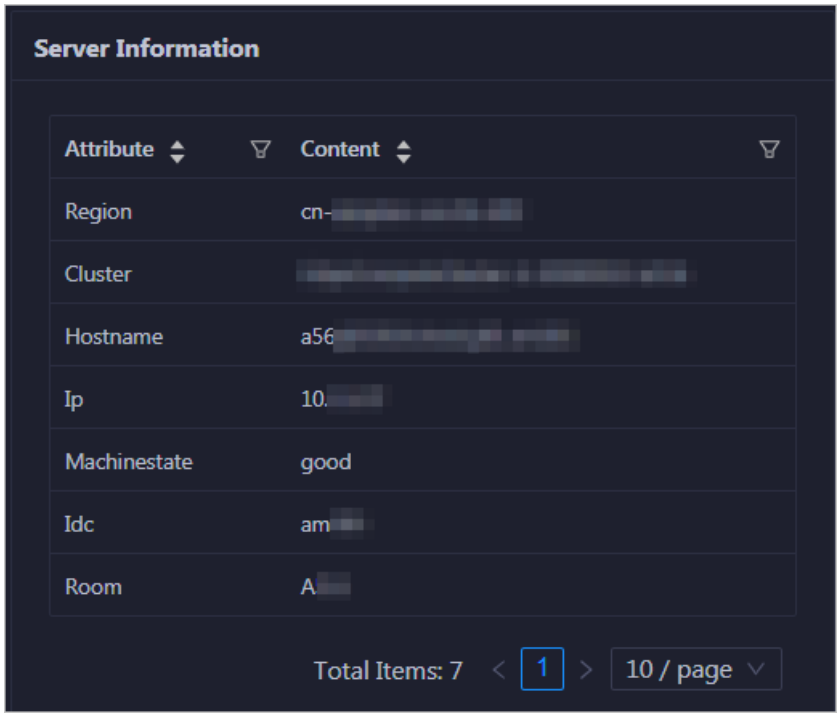

Service Role Status

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

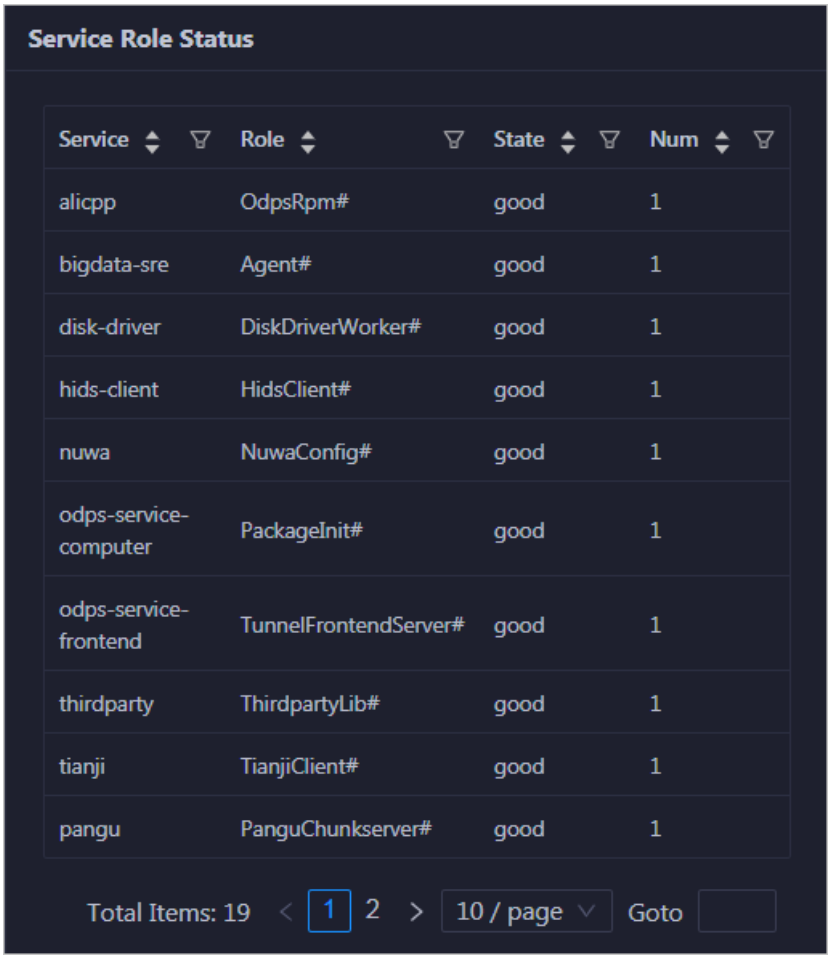

CPU

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.

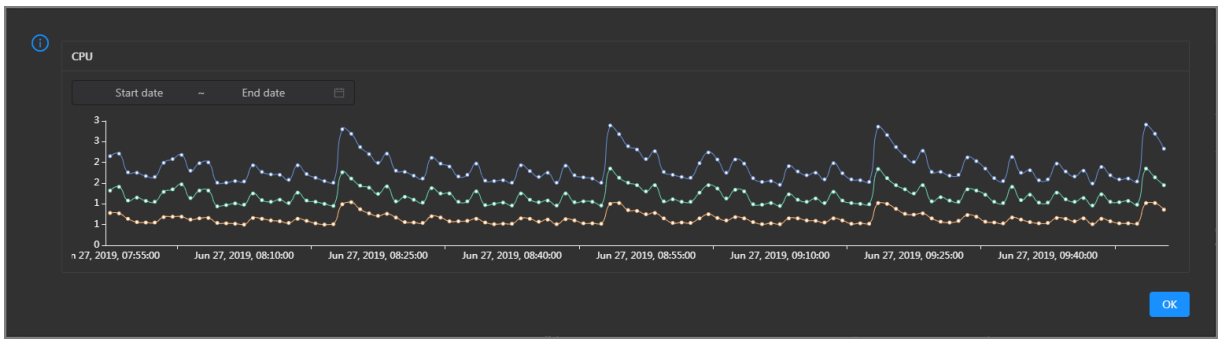

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

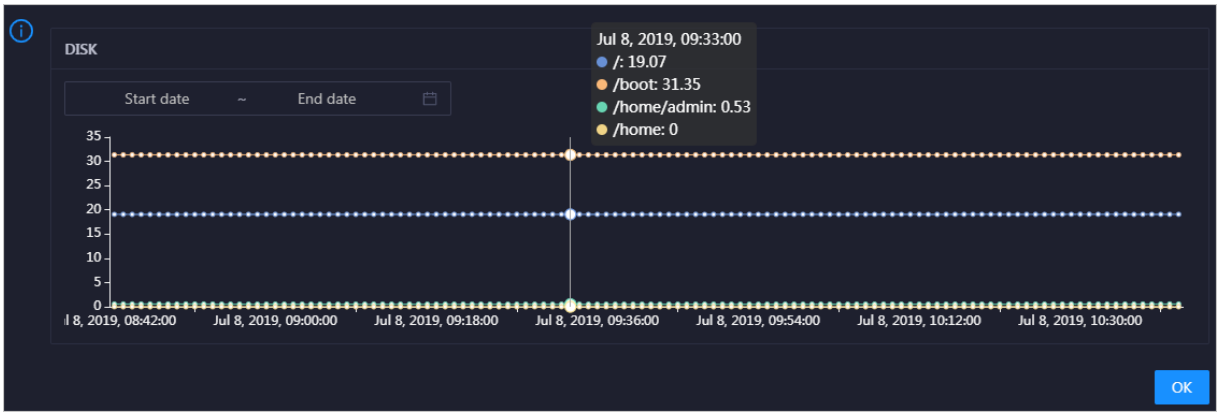

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

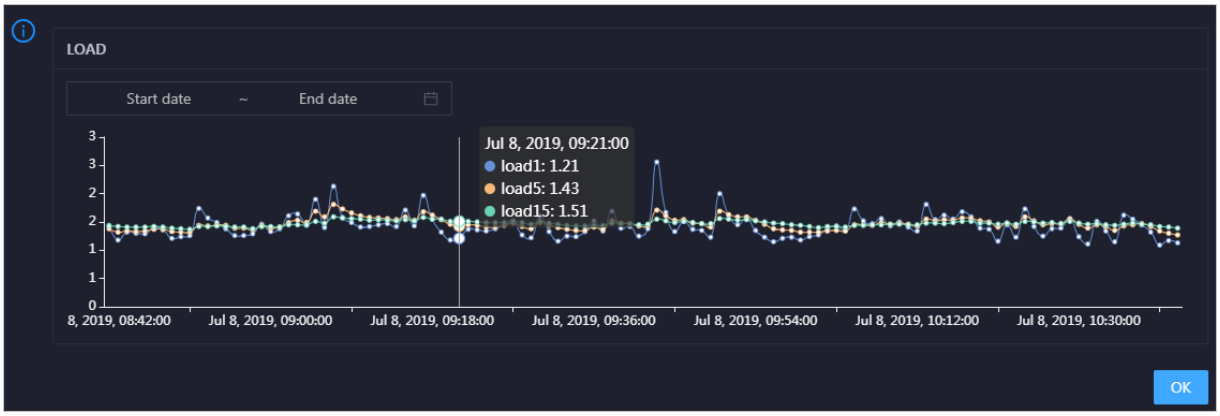

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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.

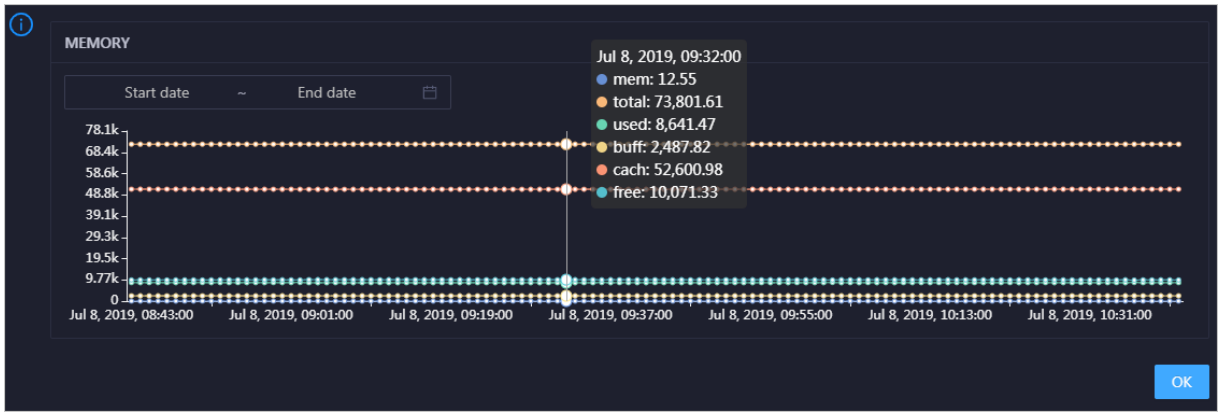

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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 transmission status of the host.

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

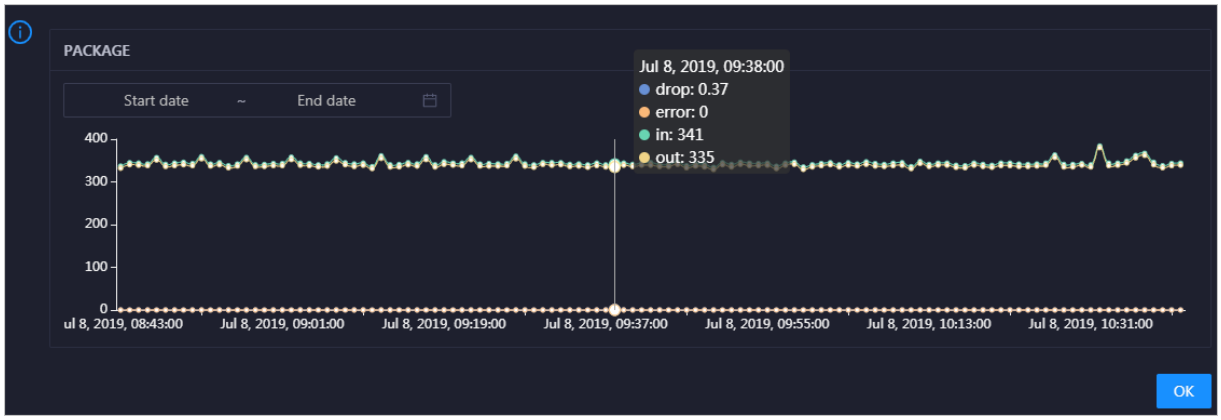

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

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

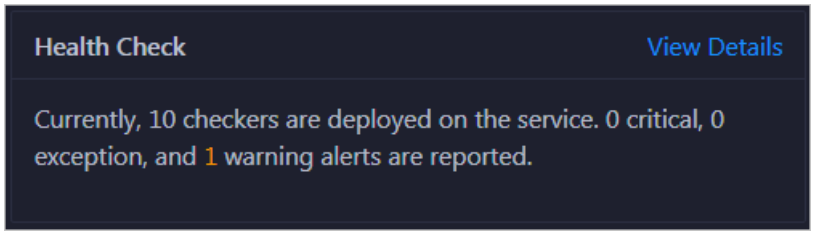

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

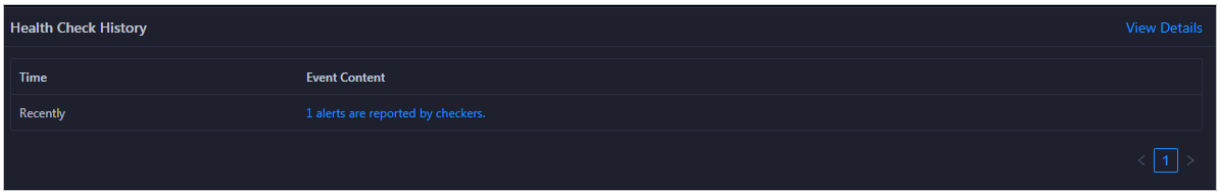

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

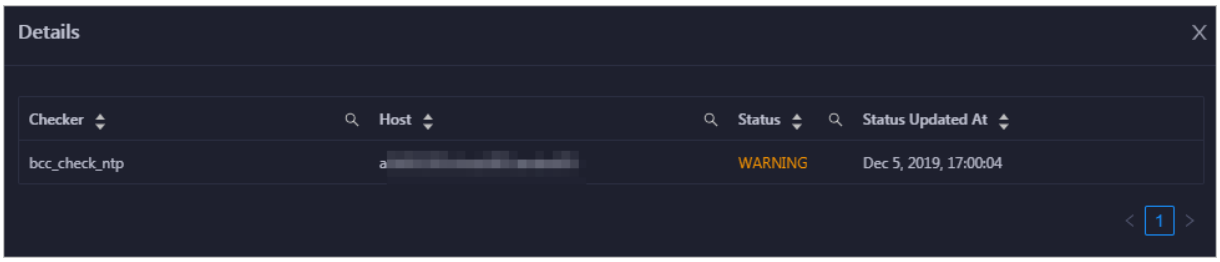

11.2.5.5.3. Host charts

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.

11.2.5.5.4. 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.

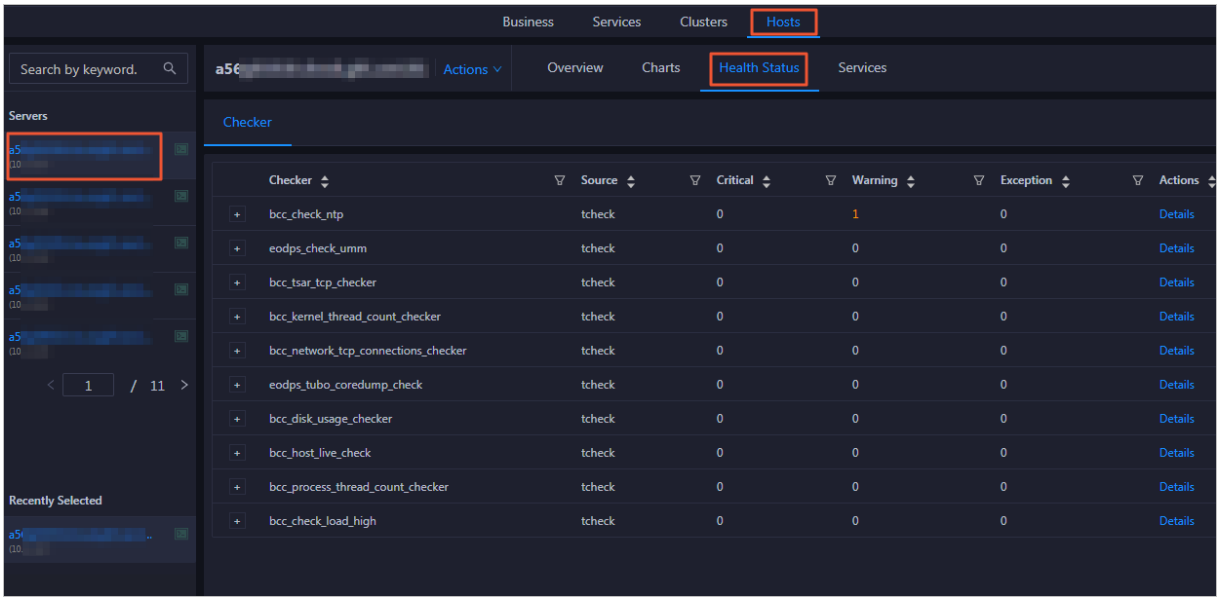

On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

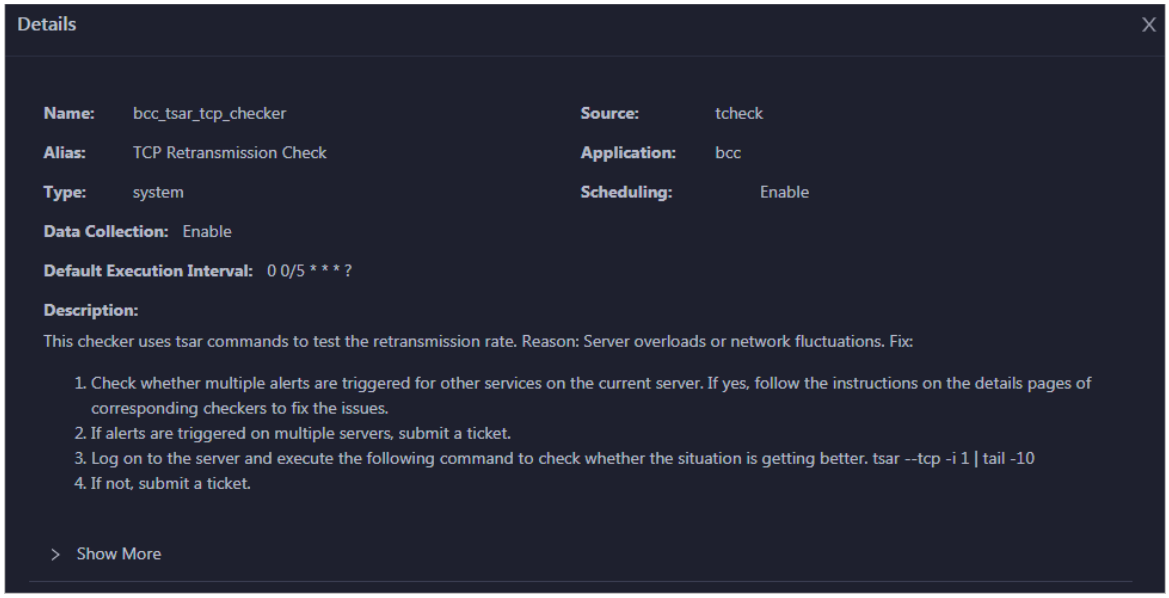

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

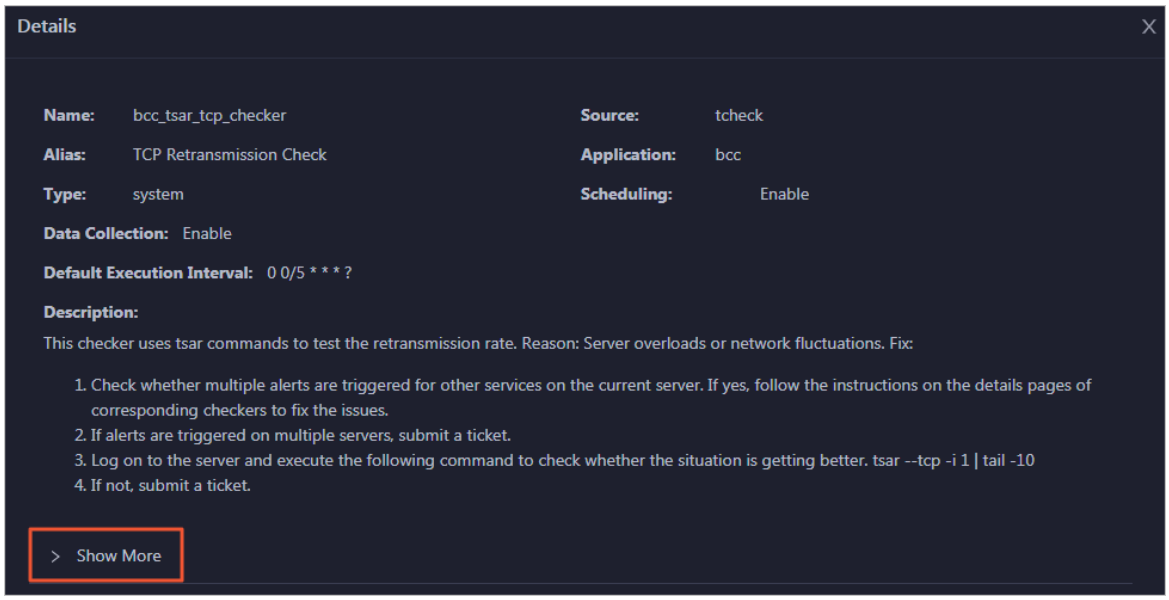

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

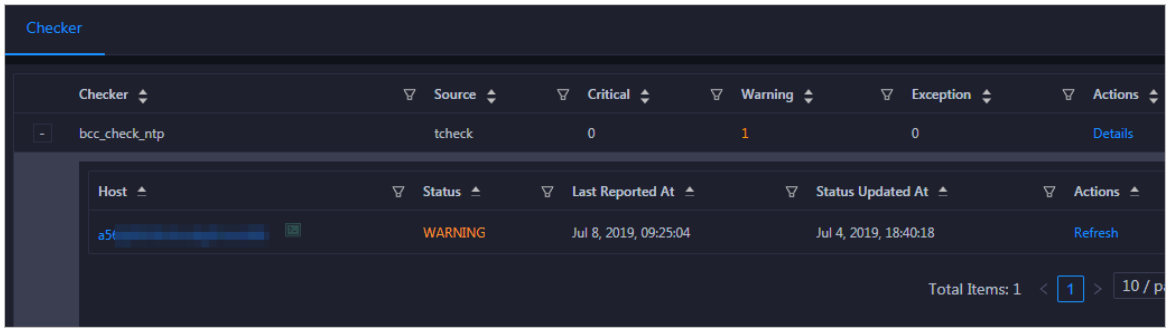

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

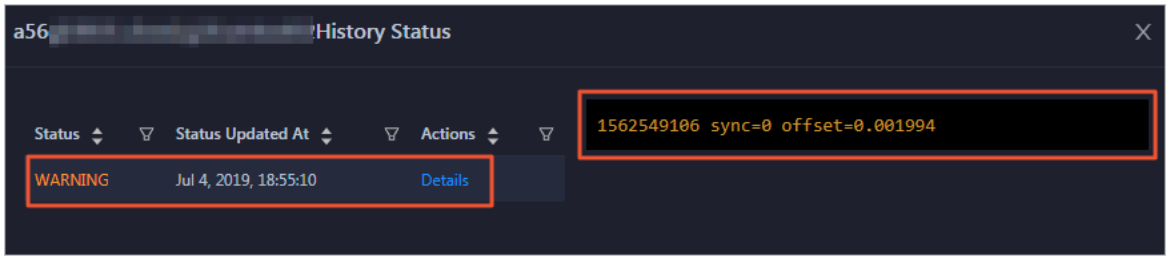

Clear alerts

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.

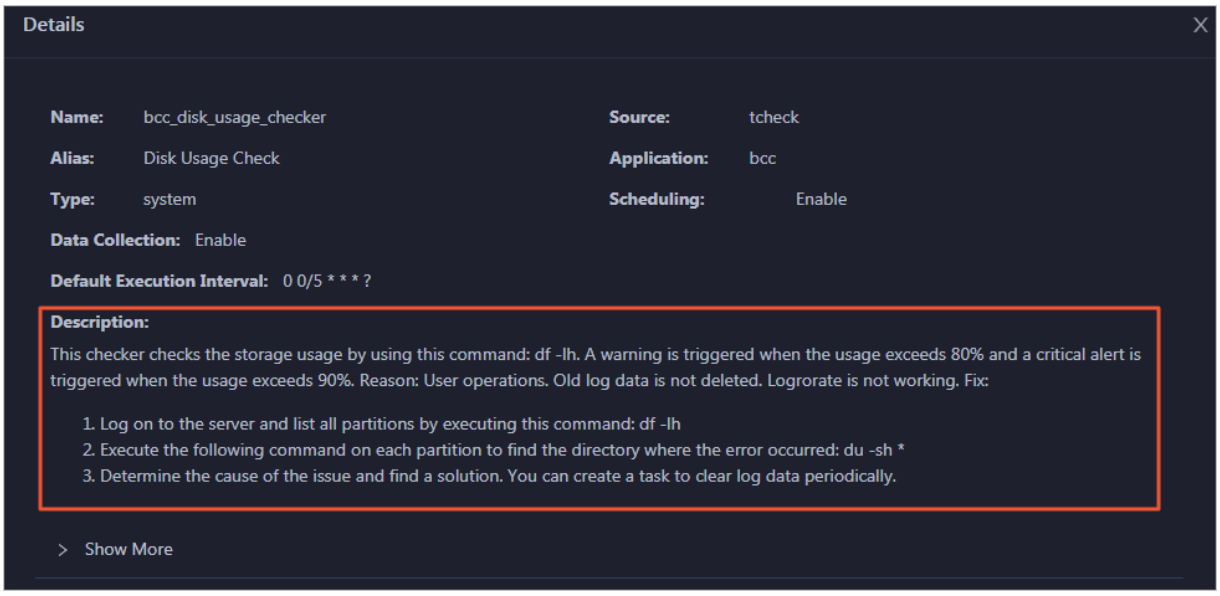

Log on to a host

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.

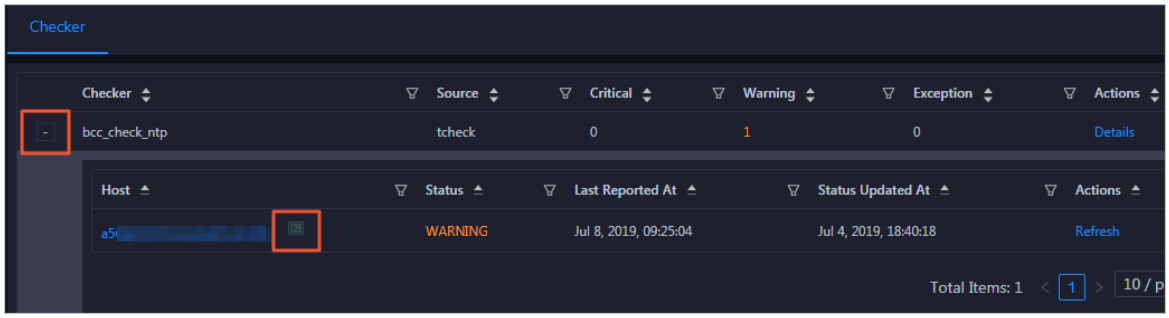

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

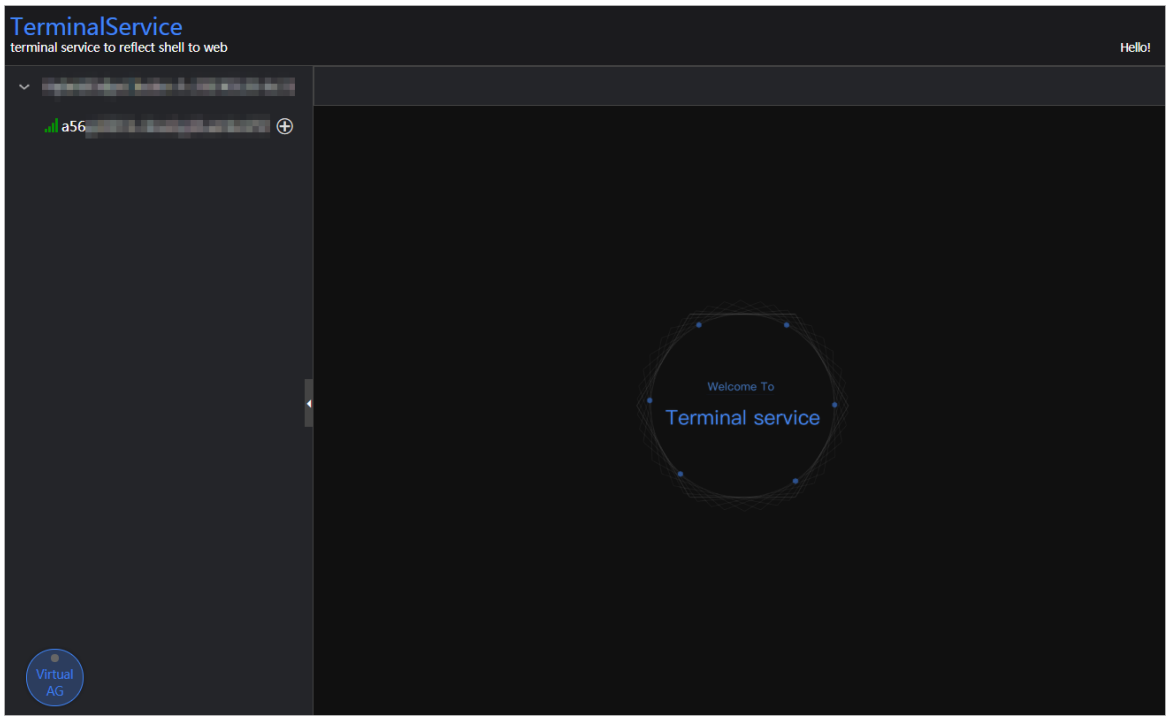

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

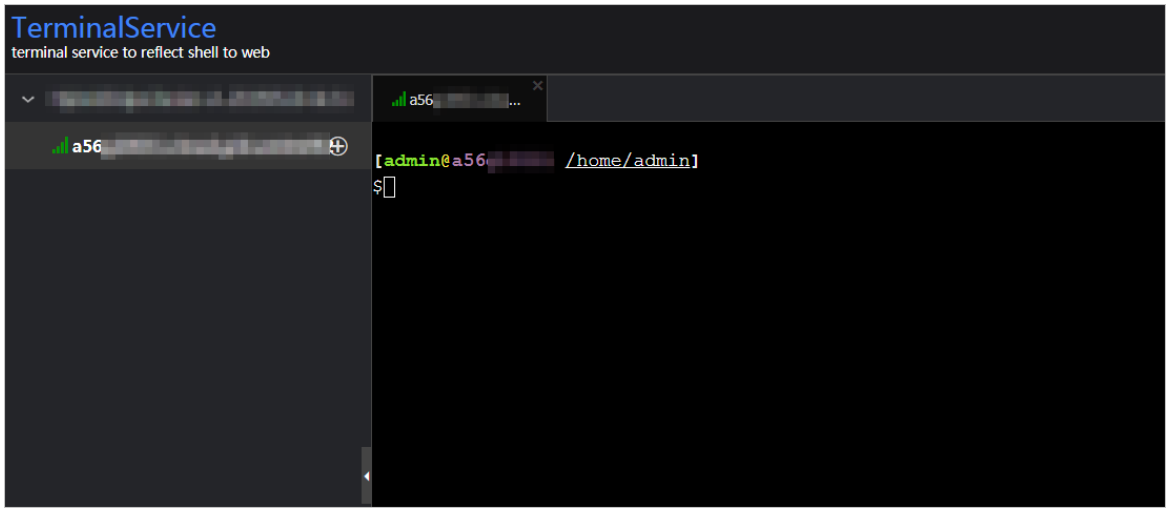

Run a checker again

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.

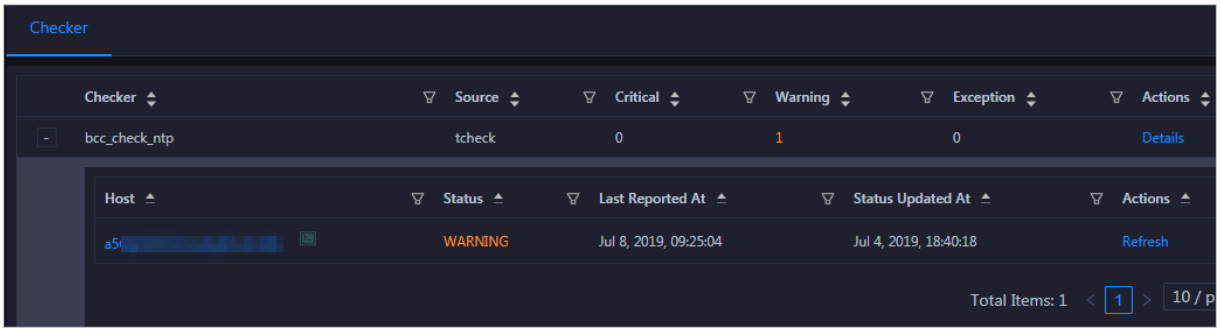

11.2.5.5.5. Host services

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.

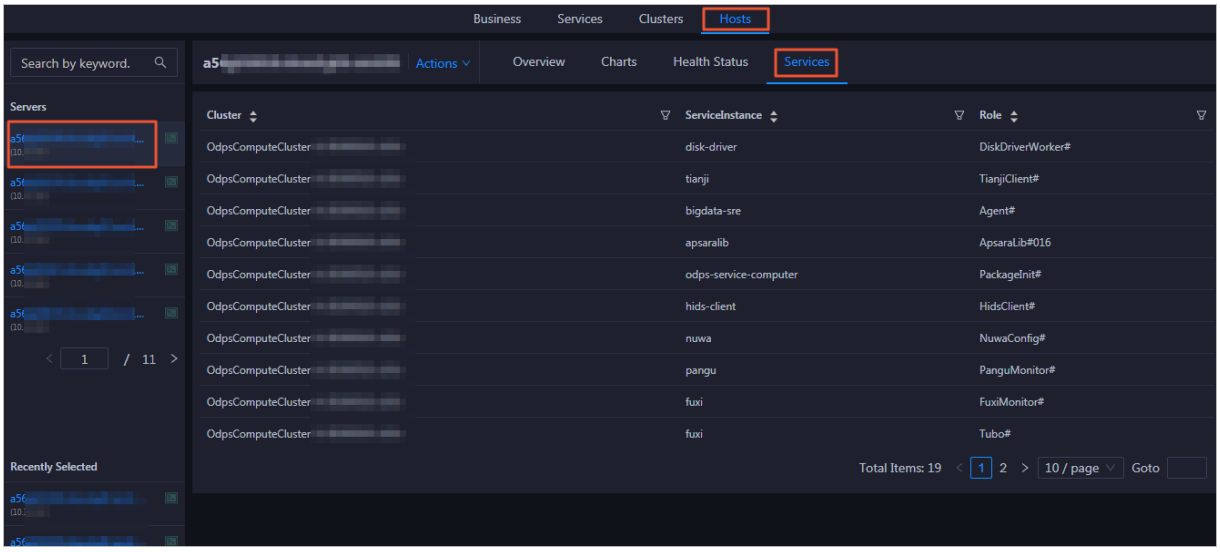

On the **Services** page, you can view the cluster, service instances, and service instance roles of the host.

11.3. DataWorks

11.3.1. Basic concepts and structure

11.3.1.1. What is DataWorks?

DataWorks is an end-to-end big data platform based on compute engines such as MaxCompute and E-MapReduce. It integrates all processes from data collection to data display and from data analysis to application running. DataWorks provides various features to help you complete the entire research and development (R&D) process in a quick and effective manner. The entire R&D process involves data integration, data development, data governance, data service provisioning, data quality control, and data security assurance.

DataWorks is an all-in-one solution for collecting, presenting, and analyzing data, and driving application development. It not only supports offline processing, analysis, and mining of large amounts of data, but also integrates core data-related technologies such as data development, data integration, production and operations and maintenance (O&M), real-time analysis, asset management, data quality control, data security assurance, and data sharing. In addition, it provides the DataService Studio and Machine Learning Platform for Artificial Intelligence (PAI) services.

In 2018, Forrester, a globally recognized market research company, named Alibaba Cloud DataWorks and MaxCompute as a world-leading cloud-based data warehouse solution. This solution is by far the only solution from a Chinese company to receive such an acknowledgment. Building on the success of the previous version, DataWorks V2.0 incorporates several new additions, such as workflows and script templates. DataWorks V2.0 supports dual workspaces for development, isolates the development environment from the production environment, adopts standard development processes, and uses a specific mechanism to reduce errors in code.

11.3.1.2. Benefits

This topic describes the benefits of DataWorks.

- Powerful computing capabilities

  DataWorks integrates with compute engines that can process large amounts of data.

  - DataWorks supports join operations for trillions of data records, millions of concurrent jobs, and petabytes (PB) of I/O throughput per day.
  - The offline scheduling system can run millions of concurrent jobs. You can configure rules and alerts to monitor the running statuses of nodes in real time.
  - DataWorks provides efficient and easy-to-use SQL and MapReduce engines, and supports most standard SQL syntax.
  - MaxCompute protects user data from loss, breach, or theft by using multi-layer data storage and access security mechanisms, including triplicate backups, read/write request authentication, application sandboxes, and system sandboxes.

- End-to-end platform

  DataWorks provides the graphical user interface (GUI) and allows multiple users to collaborate on a workspace.

  - DataWorks integrates all processes from data integration, processing, management, and monitoring to output.
  - You can create and edit workflows in a visual manner by using the workflow designer.
  - DataWorks provides a collaborative development environment. You can create and assign roles for varying nodes, such as development, online scheduling, maintenance, and data permission management, without locally processing data and nodes.

- Integration of heterogeneous data stores

  DataWorks supports batch synchronization of data among heterogeneous data stores at custom intervals in minutes, days, hours, weeks, or months. More than 400 pairs of heterogeneous data stores are supported.

- Web-based software

  DataWorks is an out-of-the-box service. You can use it on the Internet or an internal network without the need for installation and deployment.

- Multitenancy

  Data is isolated among different tenants. Each tenant controls permissions, processes data, allocates resources, and manages members in a unified and independent manner.

- Intelligent monitoring and alerting

  By setting monitoring thresholds, you can control the entire process of all nodes as well as monitor the running status of each node.

- Easy-to-use SQL editor

  The SQL editor supports automatic code and metadata completion, code formatting and folding, and pre-compilation. It offers two editor themes. These features ensure a good user experience.

- Comprehensive data quality monitoring

  DataWorks allows you to control the quality of data in heterogeneous data stores, offline data, and real-time data. You can check data quality, configure alert notifications, and manage connections.

- Convenient API development and management

  The DataService Studio service of DataWorks interacts with API Gateway. This makes it easy for you to develop and publish APIs for data sharing.

- Secure data sharing

  DataWorks enables you to de-identify sensitive data before you share it with other tenants, which ensures the security of your big data assets and maximizes their value.

11.3.1.3. Introduction to data analytics

This topic describes two typical scenarios of data analytics.

Scenario 1: data synchronization and analysis

Scenario 1 shows a typical scenario of data analytics.

1. Collect data from various databases to MaxCompute by using DataWorks.
2. Log on to DataWorks, create SQL, MapReduce, and shell nodes, and commit the nodes to MaxCompute for data analysis.
3. Use DataWorks to synchronize the analysis results from MaxCompute to the databases from which you collect data.

Scenario 1

> **Note** Base is the name of DataWorks from the technical perspective.

Scenario 2: data synchronization

DataWorks supports data synchronization between various databases. You can synchronize data by using DataWorks.

11.3.1.4. DataWorks architecture in Apsara Stack V3

This topic describes the framework and services of DataWorks.

DataWorks framework

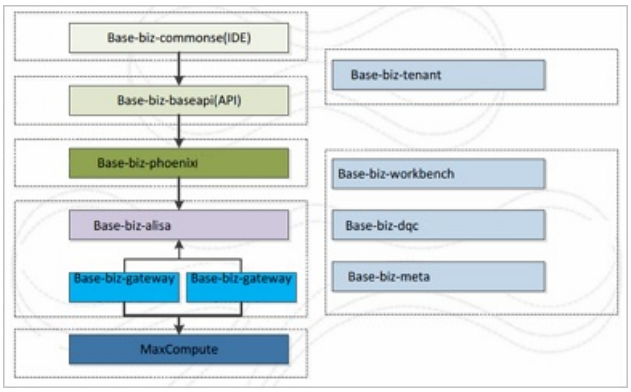

Services shown in the preceding figure play an important role for node scheduling and running. You can perform all O&M operations for DataWorks of Apsara Stack V3 in Apsara Infrastructure Management Framework. The following figure shows the services in DataWorks.

DataWorks services

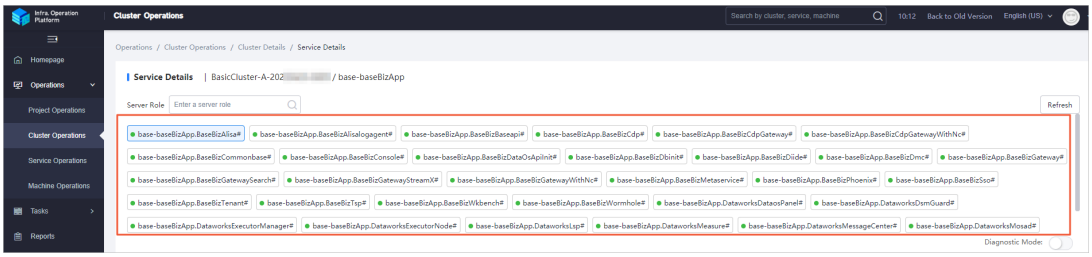

All services in DataWorks are deployed in Docker containers. You can log on to a host and run the docker ps command to view the containers in which the services are deployed.

Service architecture shows the architecture of each service except base-biz-gateway.

Service architecture

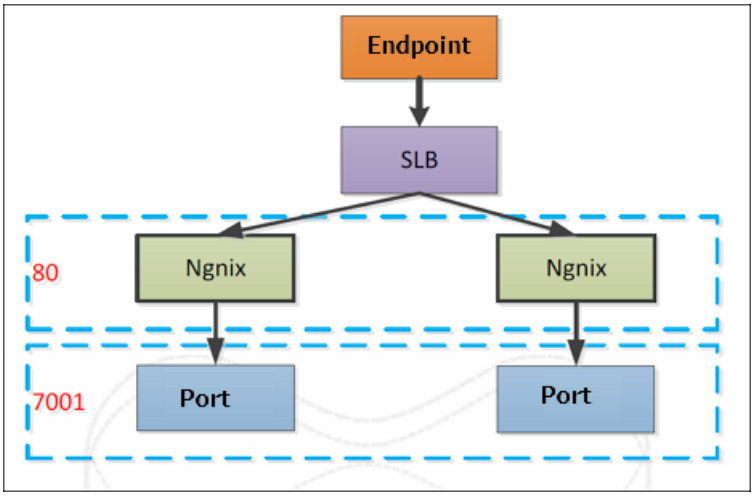

11.3.1.5. Service directories

This topic describes the directory structure of each service.

base-biz-gateway service

The base-biz-gateway service receives and runs nodes from the DataWorks integrated development environment (IDE) and the scheduling system.

- Logs directory: stores the operational logs of the base-biz-gateway service.
- taskinfo directory: stores the code run by user nodes and the execution logs.
- target directory: the main directory of the base-biz-gateway service. This directory stores the service code, start script, stop script, and configuration files.

base-biz-cdp service

The base-biz-cdp service is used to synchronize data.

- Logs directory: stores the operational logs of the base-biz-cdp service.
- Conf directory: stores the configuration files of the base-biz-cdp service.
- Bin directory: stores the start script.

Other services

The base-biz-alisa service directory is used as an example.

- Logs directory: stores the operational logs of the base-biz-alisa service.
- Conf directory: stores the configuration files of the base-biz-alisa service.
- Bin directory: stores the start script.

11.3.2. O&M by using Apsara Big Data Manager

11.3.2.1. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

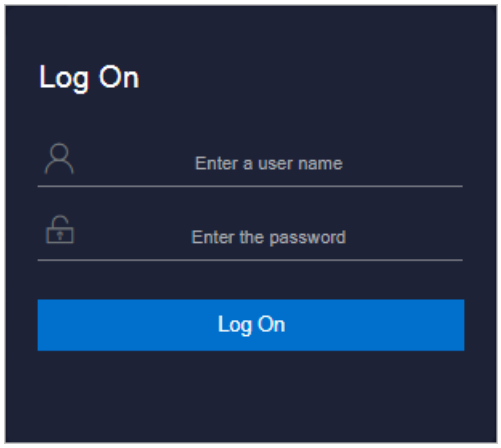

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.3.2.2. DataWorks O&M overview

This topic describes the features of DataWorks O&M supported by Apsara Big Data 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	Sub-module	Description

Module	Sub-module	Description
Data Warehouse under Services	Overview	Displays the key operation metrics, including service overview, service status, instance scheduling information, and slot usage. On this page, you can also view the trend chart of the total number of finished nodes.
	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 nodes.
	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.
Data Integration under Services	Overview	Displays overall information about Data Integration in the **Task Scheduling Overview**, **Today's Tasks**, **Third-party Dependencies - Response Time (milliseconds)**, **Third-party Dependencies - Total Requests**, and **Third-party Dependencies - Request Error Rate** sections.
	Task	Displays information about Data Integration nodes on the **Instances** and **Multi-dimensional Analysis** tabs.
	Historical Analysis	Displays historical analysis information about Data Integration on the **Multi-dimensional Analysis**, **Execution Time Analysis**, and **Task Rankings** tabs.
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.
	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	Sub-module	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 all checkers of a host, 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.

Go to the DataWorks O&M page

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

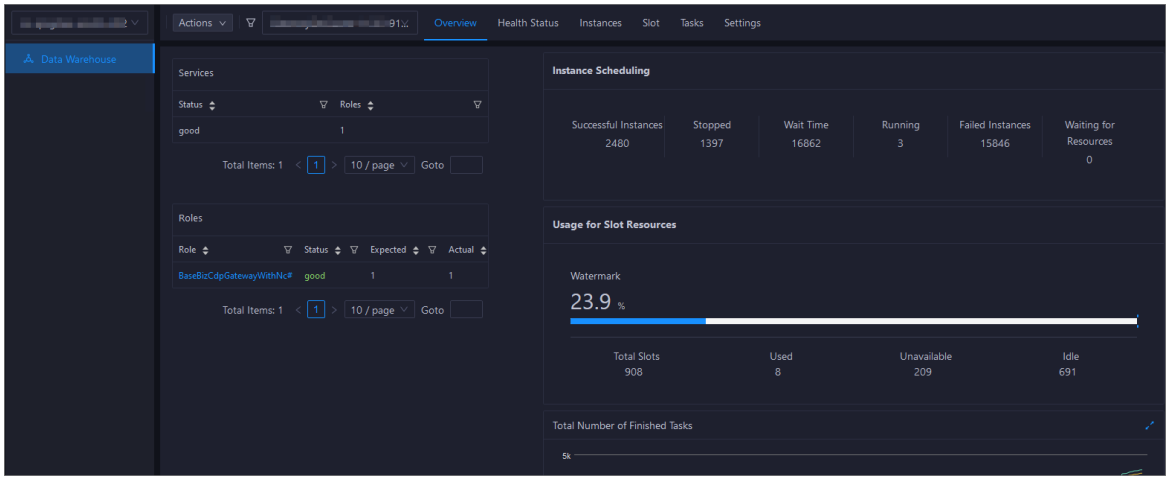

The **O&M** page includes three modules: **Services**, **Clusters**, and **Hosts**.

11.3.2.3. Service O&M

11.3.2.3.1. Data Warehouse

11.3.2.3.1.1. Service overview

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

Go to the Overview page under Services

1. Log on to the ABM console.

2. Click in the upper-left corner and select **DataWorks**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Overview** tab under the **Data Warehouse** page appears.

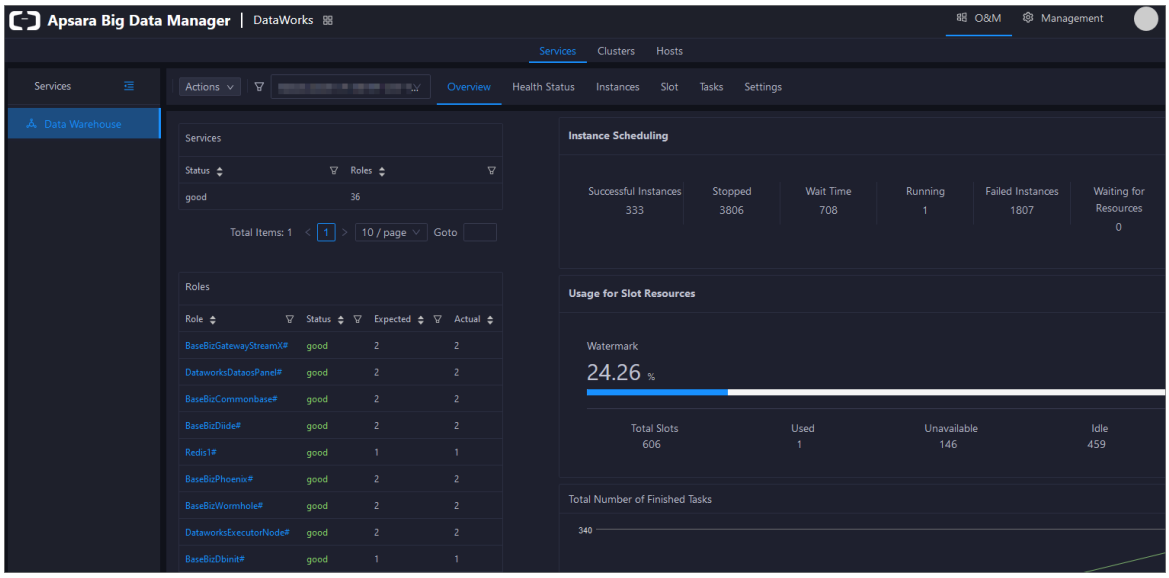

Services

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

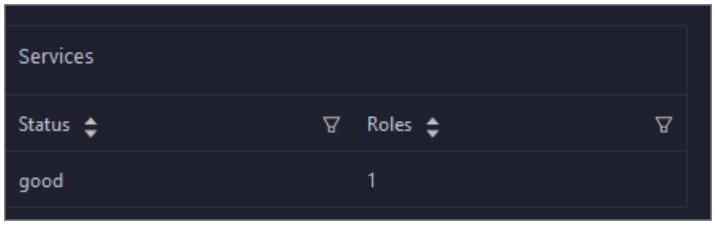

Roles

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

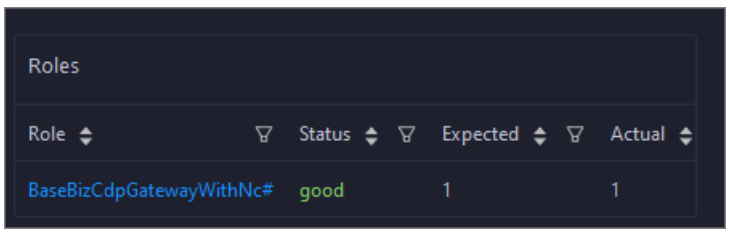

Instance Scheduling

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.

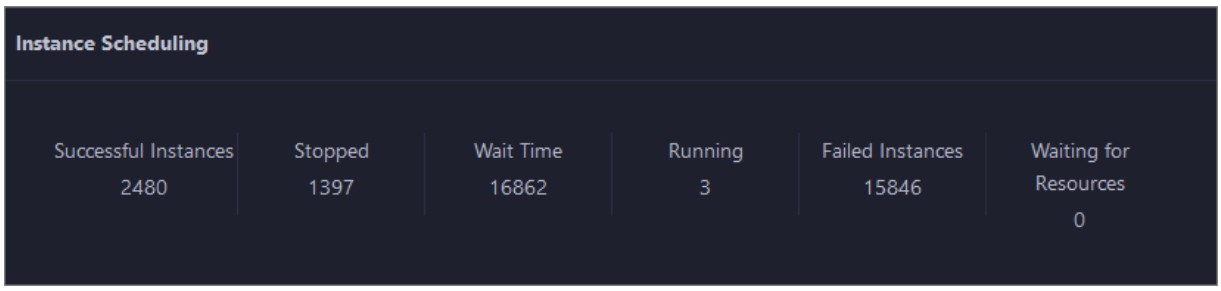

Usage for Slot Resources

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.

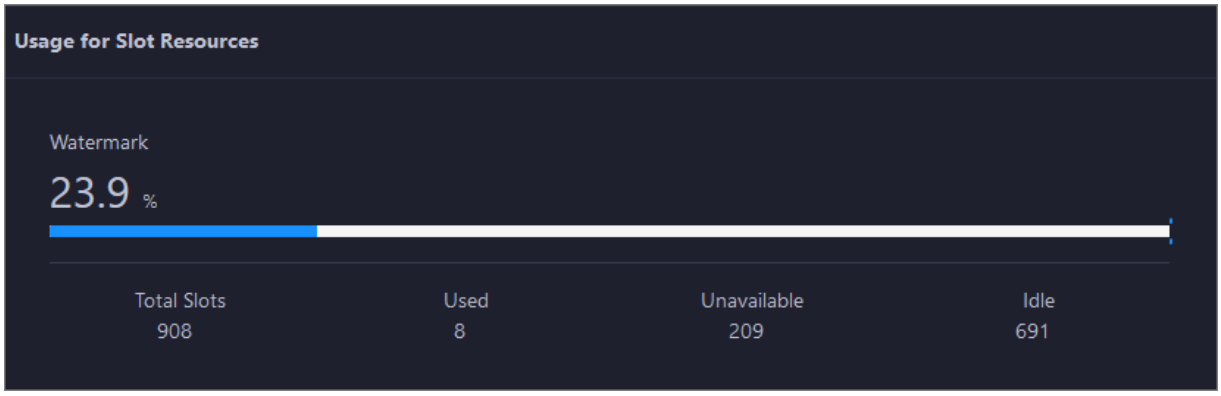

> **Note** Slots are resources that can be used by DataWorks for instance scheduling.

Total Number of Finished Tasks

This section displays the trend chart of the total number of finished nodes. The trend chart displays the trend lines of the number of nodes finished yesterday, the number of nodes finished today, and the average number of nodes finished each day over time in different colors.

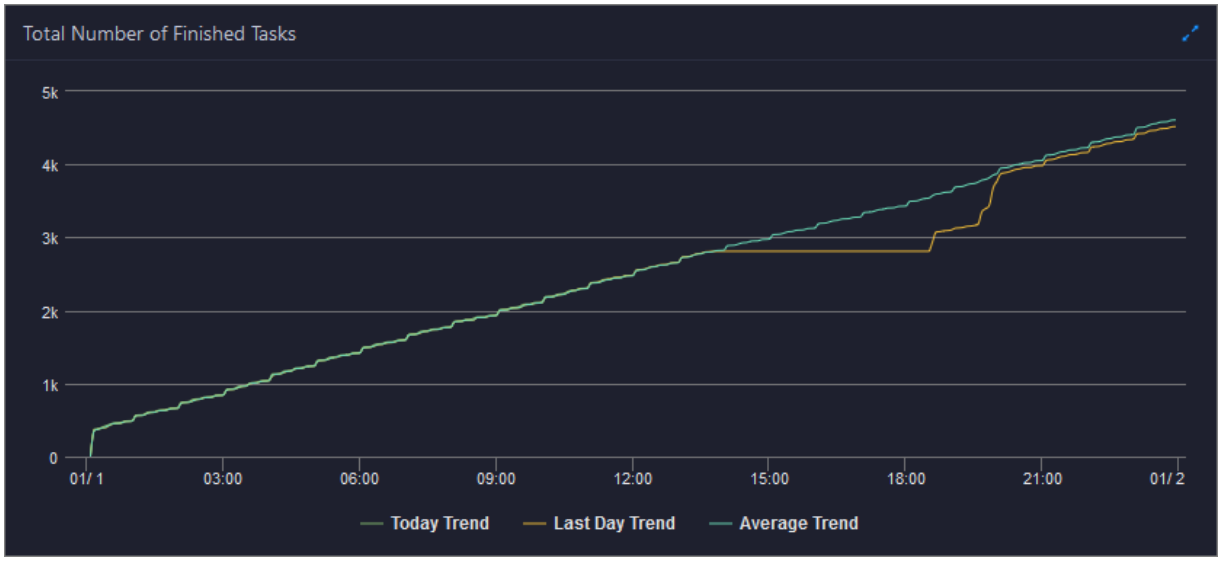

11.3.2.3.1.2. Service health

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.

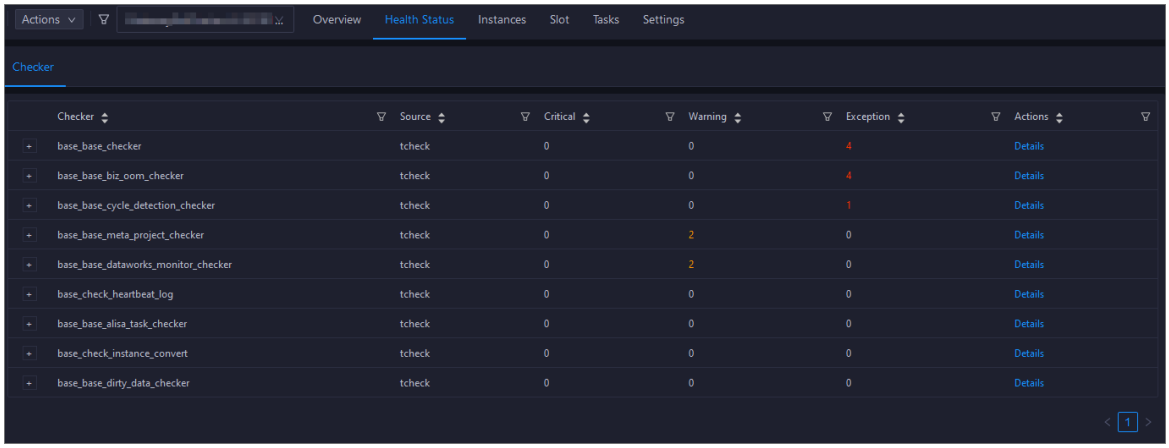

The **Health Status** page displays all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical**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

11.3.2.3.1.3. Service instances

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

Go to the Instances page under Services

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

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

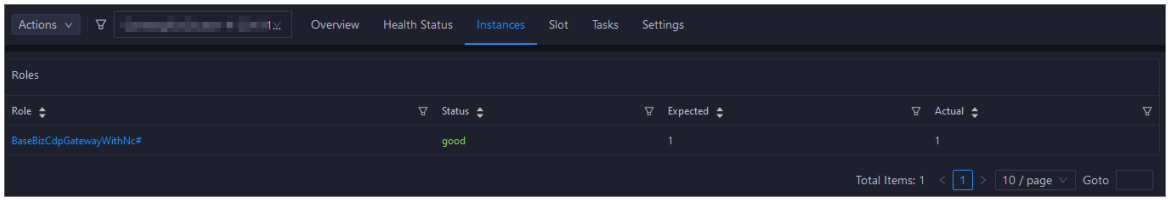

The **Instances** page displays information about all DataWorks service roles, including the status and the expected and actual numbers of hosts in the desired state. The statuses include **good**, **bad**, and **upgrading**.

Supported operations

You can filter or sort service roles based on a column to facilitate information retrieval on the **Instances** page.

11.3.2.3.1.4. 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.

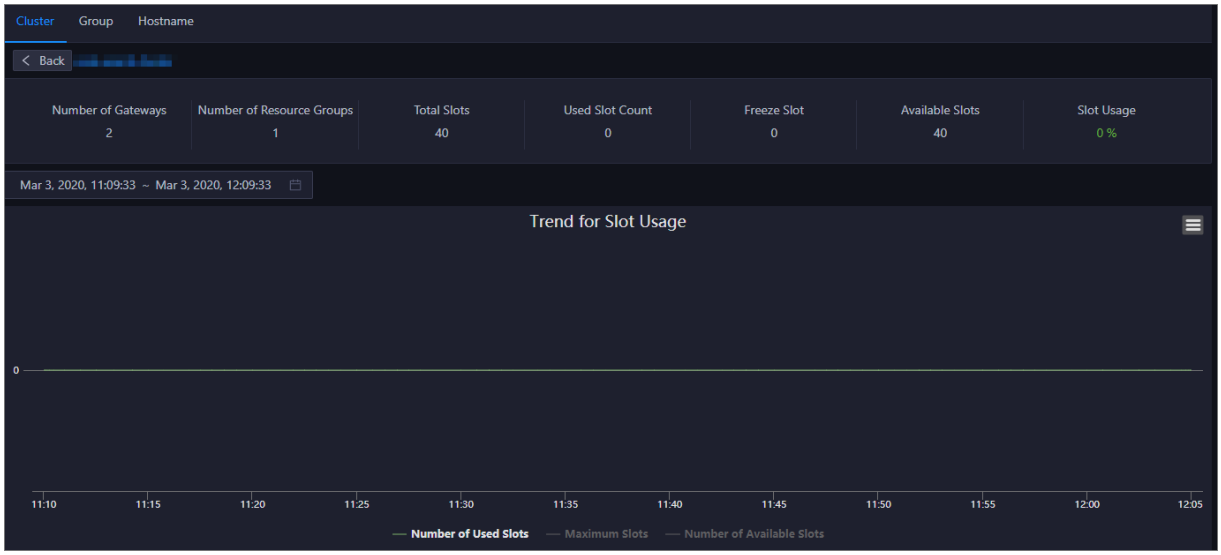

On the cluster details page, you can view the numbers of gateways, resource groups, slots, used slots, frozen slots, and available slots, and the slot usage of the cluster at the top. You can also view the trend chart of slot usage over time at the bottom. The trend chart displays the trend lines of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

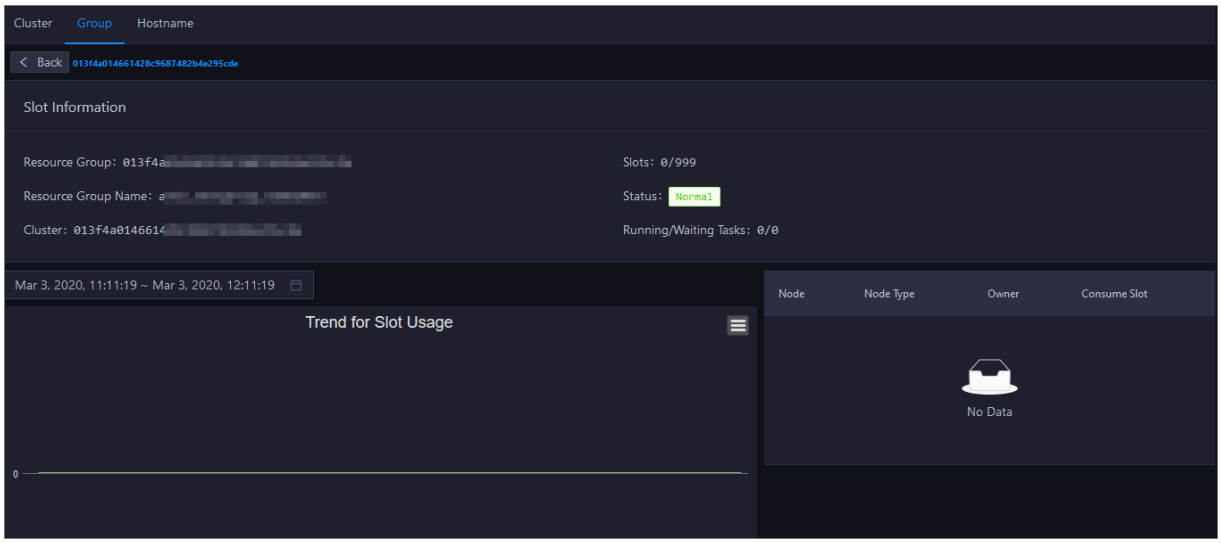

On the resource group details page, you can view the current slot information of the resource group, for example, the number of used slots and the maximum number of slots, at the top. You can also view the trend chart of slot usage over time, the nodes that occupy the slots, and the owners at the bottom.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

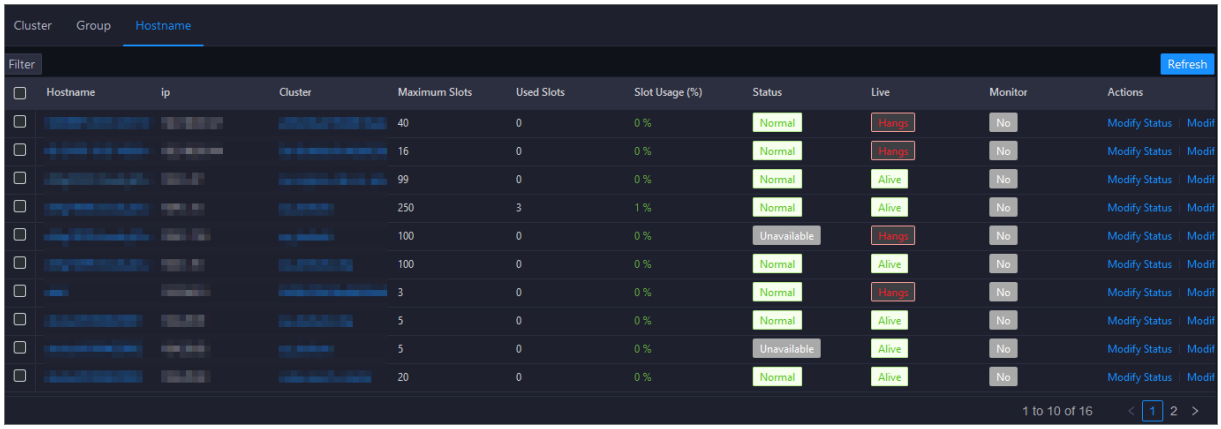

To view more information about slots of a specified host, click the name of the host.

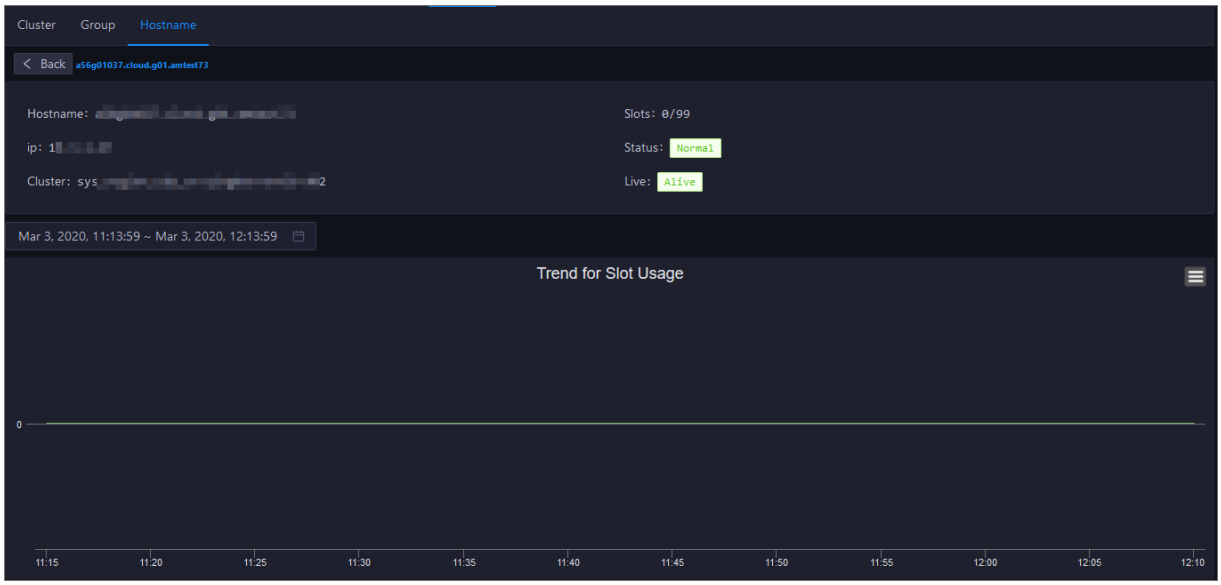

On the host details page, you can view the current slot information of the host, for example, the number of used slots and the maximum number of slots, at the top. You can also view the trend chart of slot usage over time at the bottom. The trend chart displays the trend lines of the number of used slots,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

11.3.2.3.1.5. Service nodes

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

Go to the Tasks page under Services

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

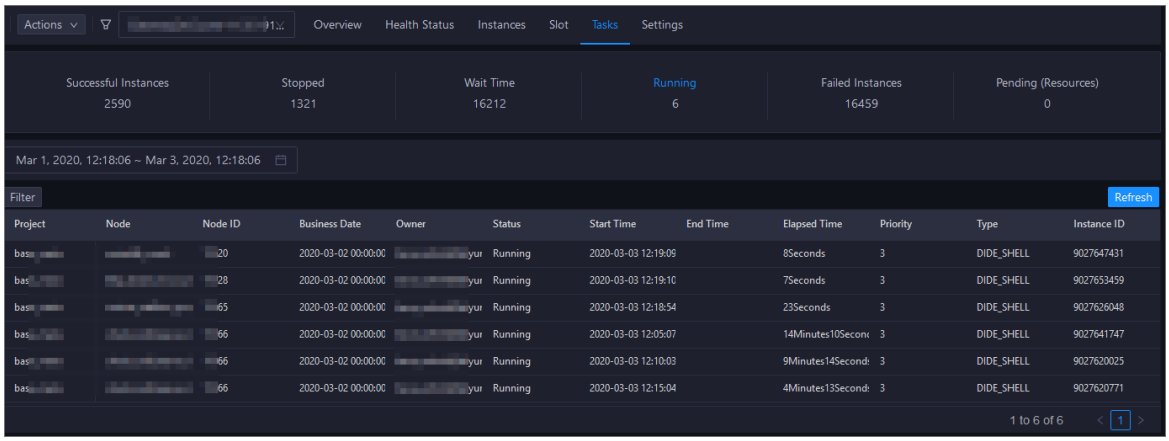

The **Tasks** page displays the node information of the current cluster, including the project name, node name, node ID, data timestamp, owner, running status, start time, end time, running duration, priority, type, and instance ID.

Filter nodes by status

On the **Tasks** page, the respective number of nodes in all statuses is displayed at the top. Click a node state to view corresponding nodes in the list. By default, nodes in the **Running** state appear.

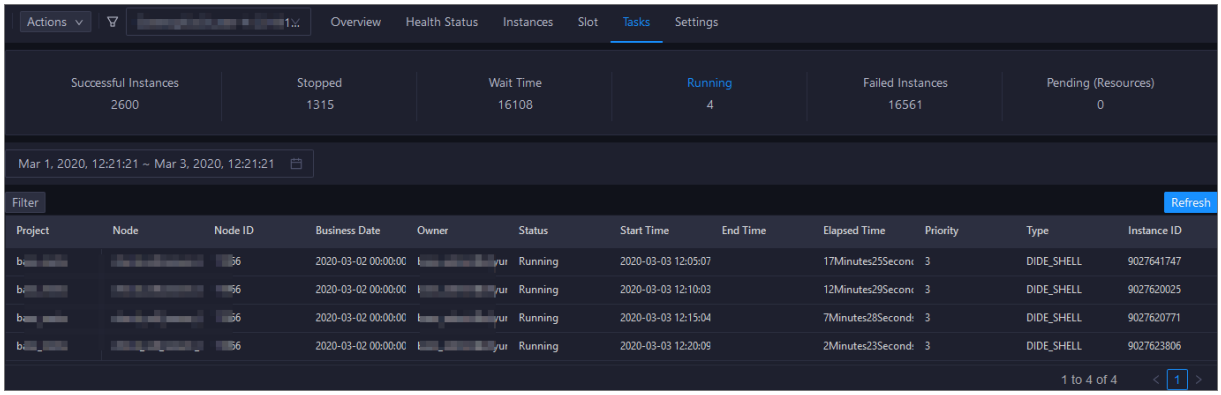

Filter nodes by time

Select a time period, including both the date and time, in the upper-left corner of the node list to view the nodes in the corresponding time period.

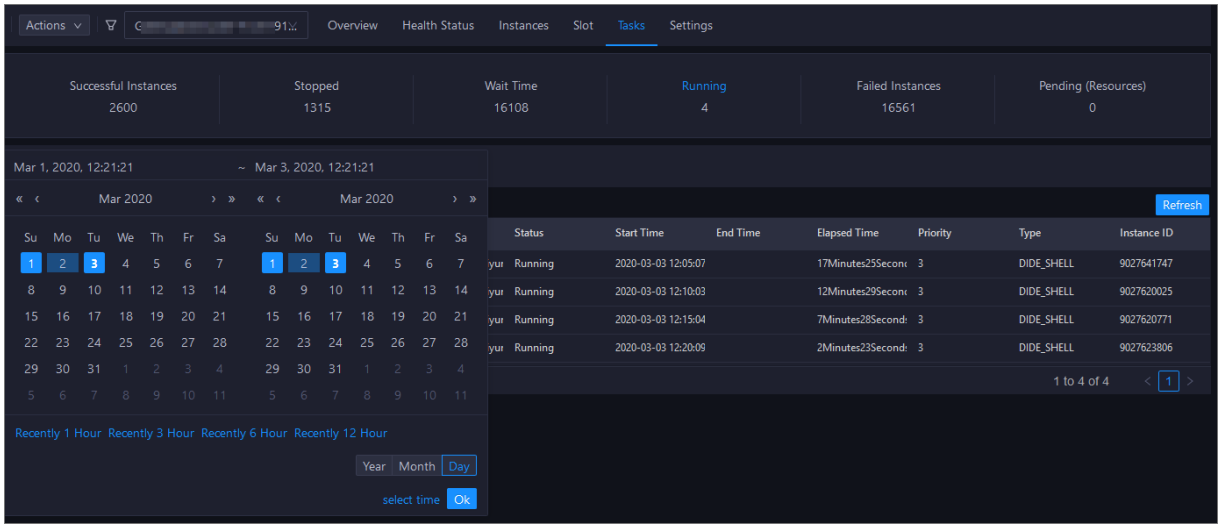

Other operations

You can filter nodes, sort nodes based on a column, and customize columns on the Tasks page.

11.3.2.3.1.6. 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.

11.3.2.3.2. Data Integration

11.3.2.3.2.1. Data integration overview

The Overview page of Data Integration displays information in the Task Scheduling Overview, Today's Tasks, Third-party Dependencies - Response Time (milliseconds), Third-party Dependencies - Total Requests, and Third-party Dependencies - Request Error Rate sections.

Procedure

1. Log on to the ABM console.
2. Click in the upper-left corner and select **DataWorks**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Overview** tab under the **Data Warehouse** page appears.
4. In the left-side navigation pane, click **Data Integration**. The **Overview** page appears.After you set the **Aggregation Period** parameter and select a time period, you can view the desired information in the following sections:
   - **Task Scheduling Overview**
   - **Today's Tasks**

   - **Third-party Dependencies - Response Time (milliseconds)**

   - **Third-party Dependencies - Total Requests**

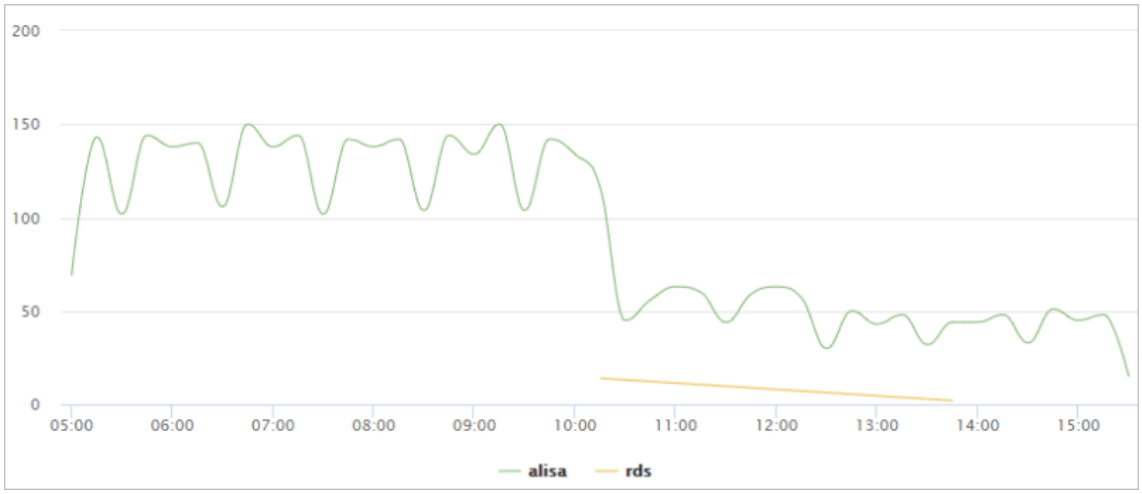

- Third-party Dependencies - Request Error Rate

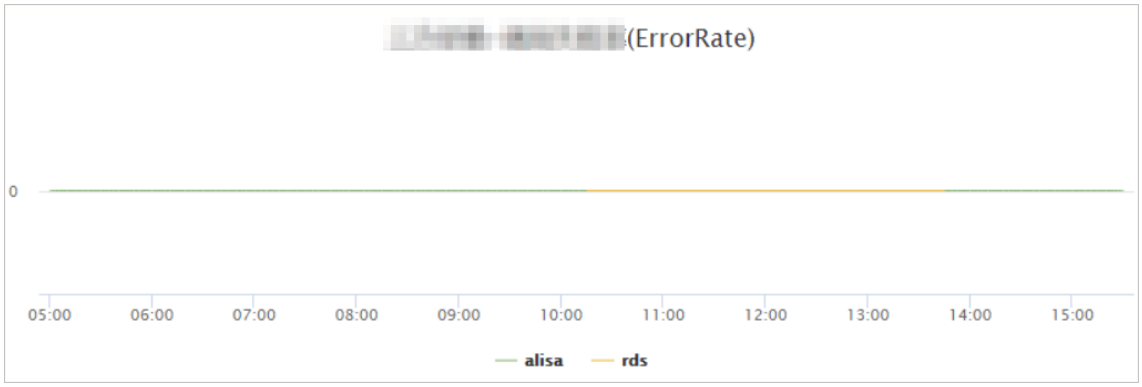

- Third-party Dependencies - Failed Requests

11.3.2.3.2.2. View Data Integration nodes

This topic describes how to view node information on the Tasks page of Data Integration, and obtain the required data such as the amount of synchronized data, synchronization speed, and node data volume.

Go to the Tasks page

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

4. In the left-side navigation pane, click **Data Integration**. The **Overview** page appears.
5. Click the **Tasks** tab. The **Instances** tab appears by default.

View instance information

On the **Instances** tab, you can filter instances by **Project Name**, **Resource Group**, **Host**, **Status**, **Read Plug-in Type**, **Write Plug-in Type**, and **Data Source**. If you need to customize more filter criteria, click **Advanced Search**.

You can also click the request ID, synchronization script number, or resource group name to view the corresponding details.

- After you click the request ID, you can view the event type, IP address from which the request is submitted, and start time of the instance.
- After you click the synchronization script number, you can view the following information:
  - On the **Job Statistics by Day** page, you can view the trends of the synchronized data volume, synchronization speed, and consumed time.
  - On the **Job Statistics by Run** page, you can view the trends of the synchronized data volume, synchronization speed, and consumed time.
  - On the **Jobs in the Final Status** page, you can view the trends of successful nodes, failed nodes, and killed nodes.
- After you click the resource group name, you can view the slot usage of the resource group.

After you view the corresponding details, you can click **Back** to return to the **Instances** page under **Task**.

On the Task page, you can also view the number of initialized nodes, submitted nodes, running nodes, failed, nodes, successful nodes, and nodes waiting to be scheduled.

View multi-dimensional analysis information

On the **Multi-dimensional Analysis** tab, you can filter historical analysis information by **Project Name**, **Resource Group**, **Host**, **Status**, **Read Plug-in Type**, **Write Plug-in Type**, or **Data Source** from the perspective of **Sync Data Size**, **Sync Speed**, or **Tasks**.

11.3.2.3.2.3. View historical analysis information

On the Historical Analysis page, you can view information about multi-dimensional analysis, execution time analysis, and nodes rankings.

Go to the Historical Analysis page

1. Log on to the ABM console.
2. Click  in the upper-left corner and select **DataWorks**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Overview** tab under the **Data Warehouse** page appears.
4. In the left-side navigation pane, click **Data Integration**. The **Overview** page appears.
5. Click the **Historical Analysis** tab. The **Multi-dimensional Analysis** page appears.

View multi-dimensional analysis information

On the **Multi-dimensional Analysis** tab, you can filter historical analysis information by **Project Name**, **Resource Group**, **Host**, **Status**, **Read Plug-in Type**, **Write Plug-in Type**, or **Data Source** from the perspective of **Sync Data Size**, **Sync Speed**, **Time**, or **Tasks**.

View execution time analysis information

On the **Execution Time Analysis** tab, you can filter required execution time information by **Project Name**, **Resource Group**, **Host**, **Status**, **Read Plug-in Type**, **Write Plug-in Type**, or **Data Source**.

View top 10 nodes

On the **Task Rankings** tab, you can filter top 10 nodes by **Project Name**, **Resource Group**, **Host**, **Status**, **Read Plug-in Type**, **Write Plug-in Type**, or **Data Source**.

11.3.2.3.3. 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.

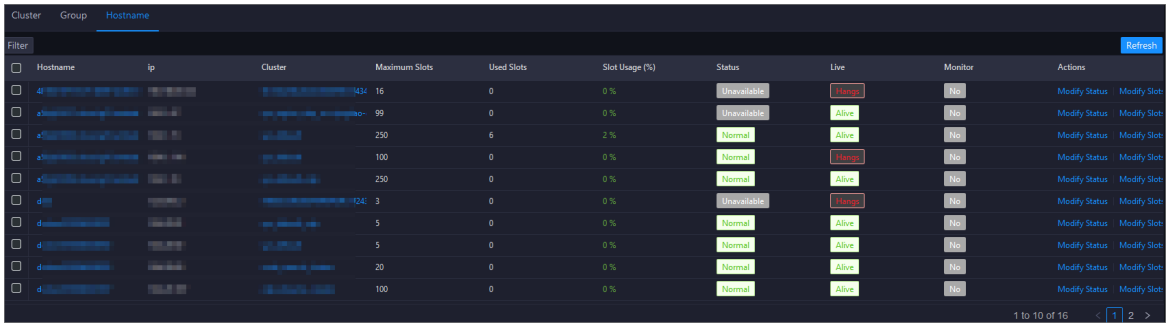

6. Choose **Actions > Scale-down for Normal Hosts** in the upper-left corner. In the **Scale-down for Normal Hosts** dialog box that appears, set relevant parameters.

   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.

11.3.2.4. Cluster O&M

11.3.2.4.1. 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, 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.

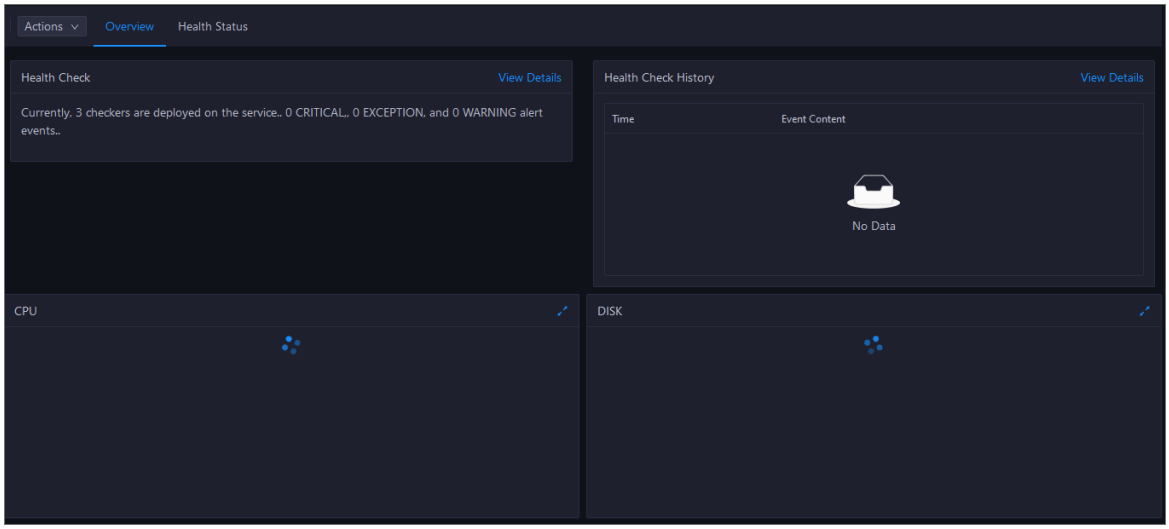

Health Check

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

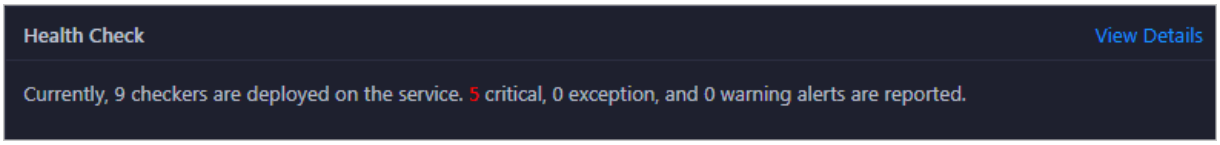

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Cluster health.

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.

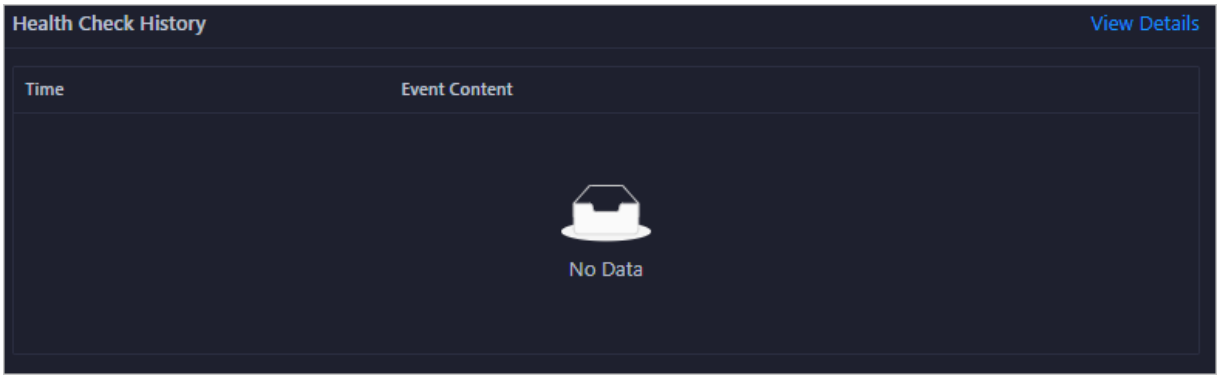

You can click the event content of a check to view the exception items.

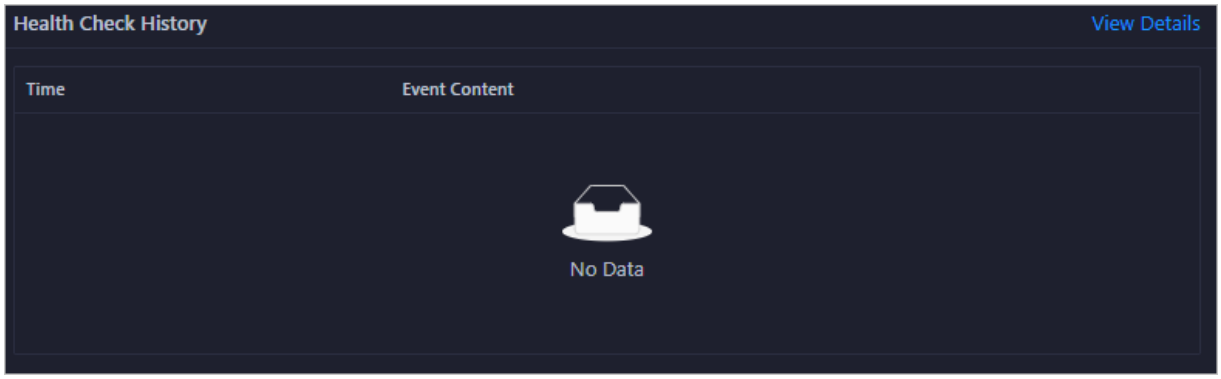

CPU

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

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

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

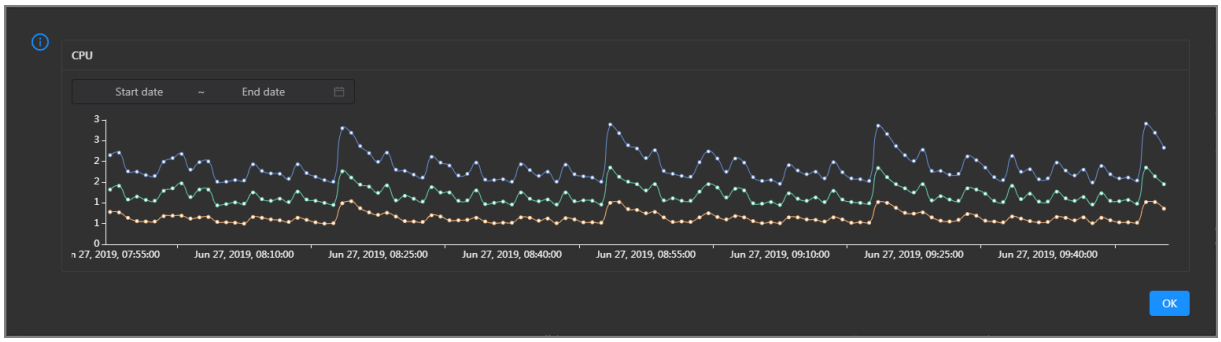

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

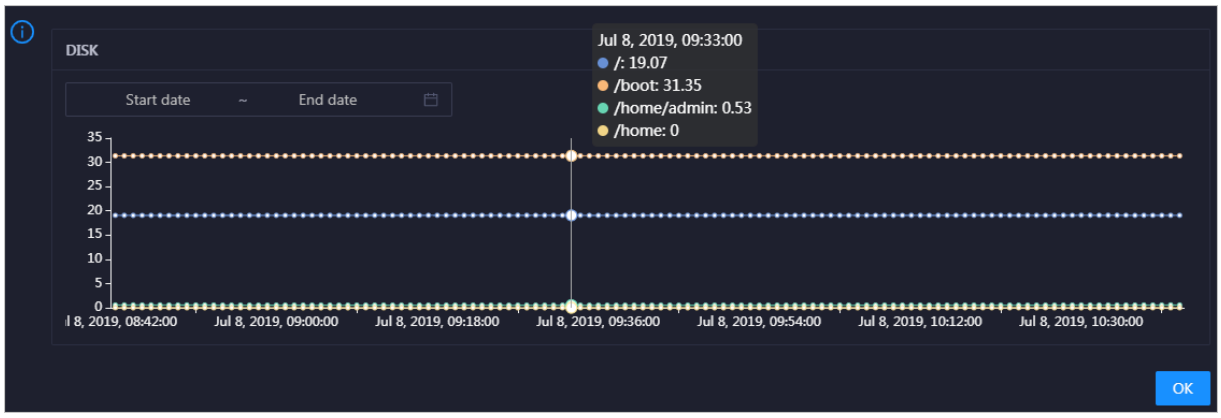

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

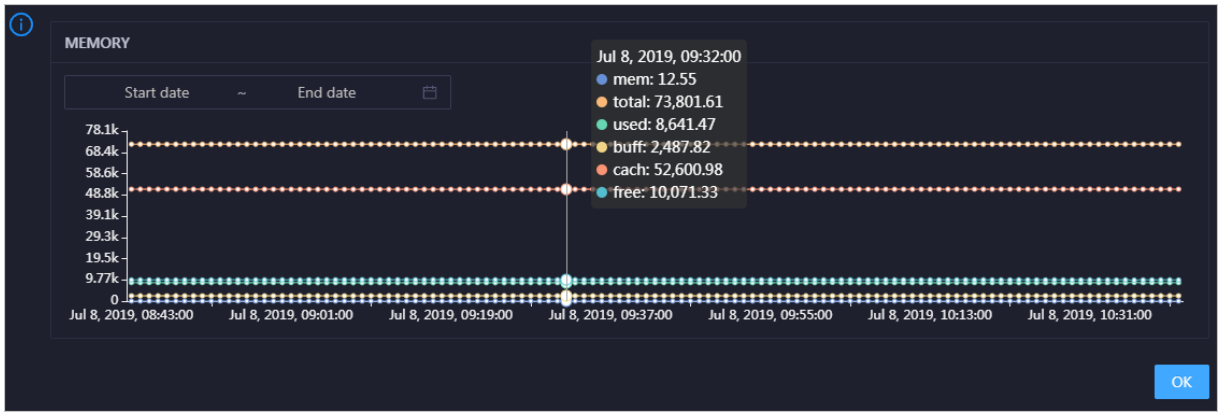

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

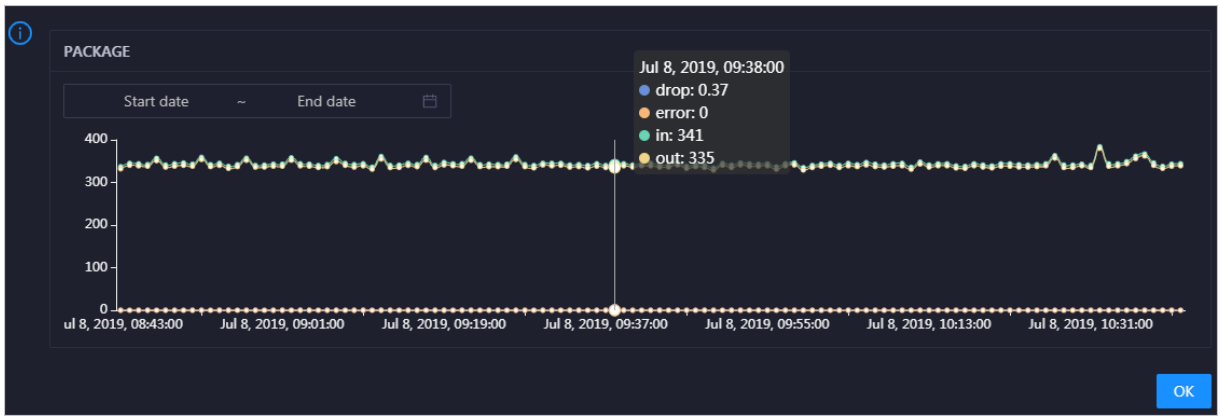

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

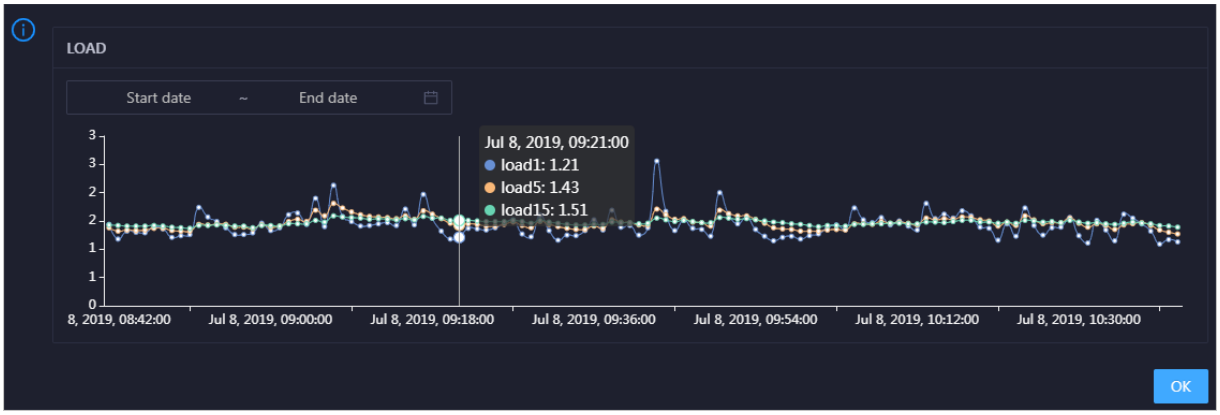

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.3.2.4.2. 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.

Go to the Health Status page under Clusters

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

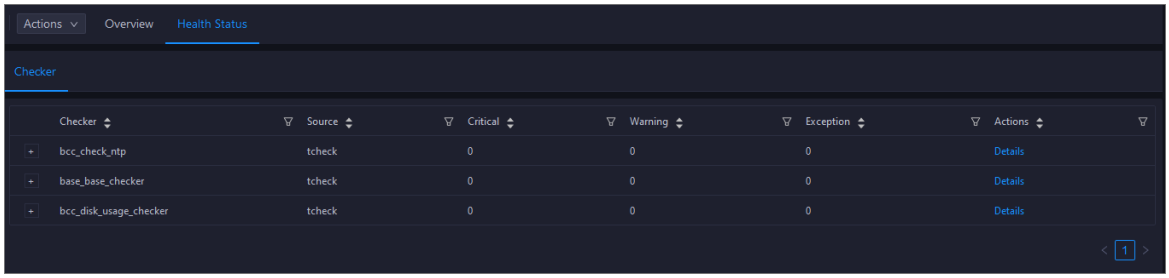

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are repesented in different colors. Handle the alerts in a timely manner,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

View checker details

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

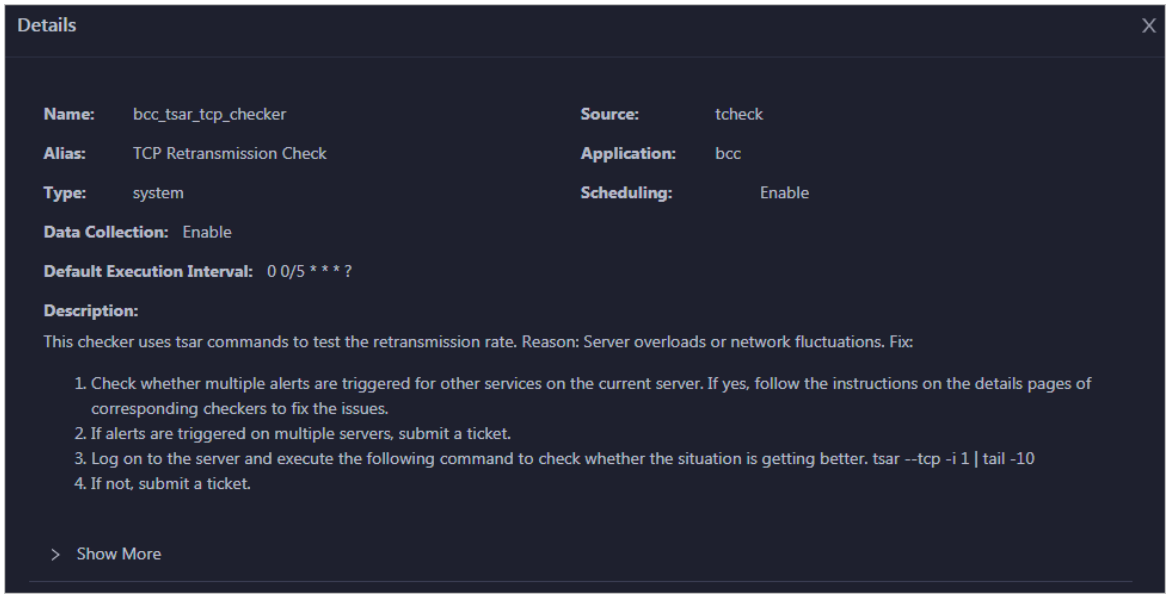

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

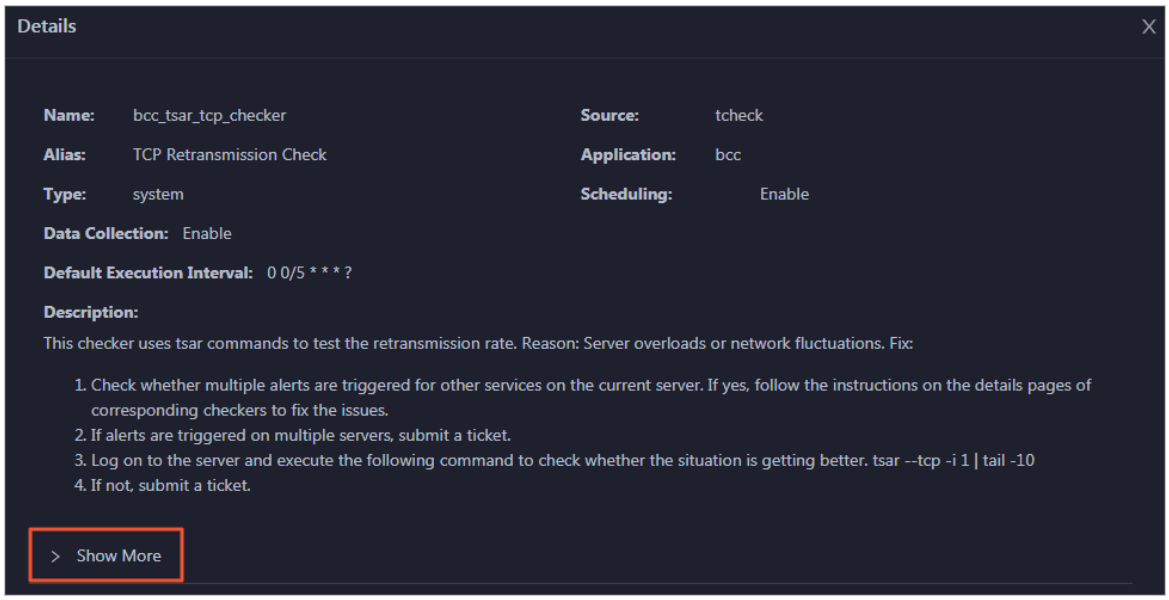

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

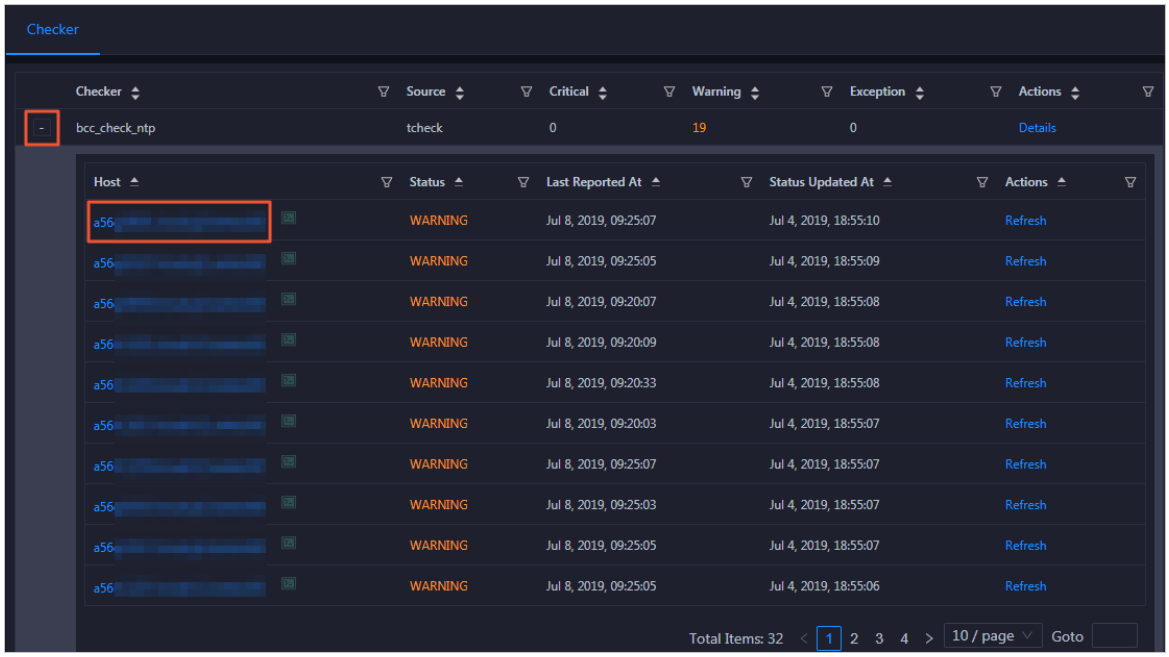

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

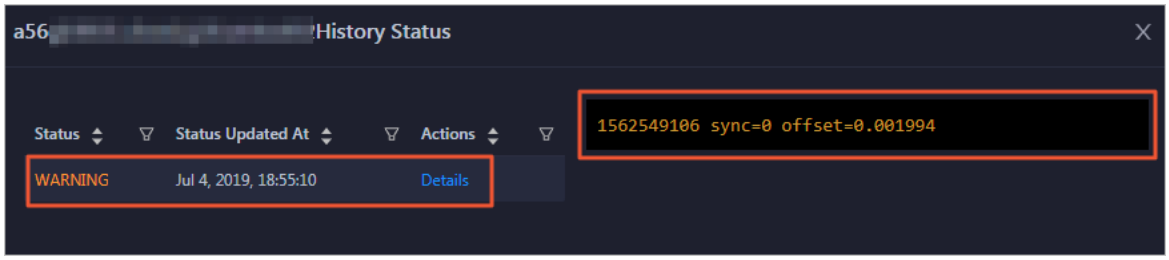

Clear alerts

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

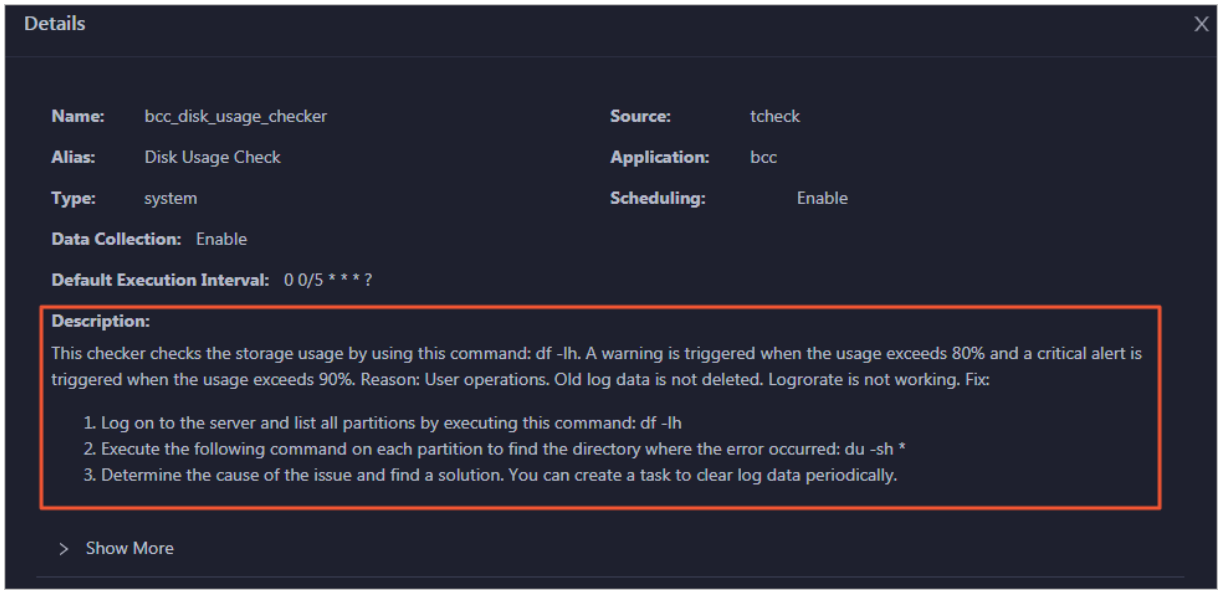

Log on to a host

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

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

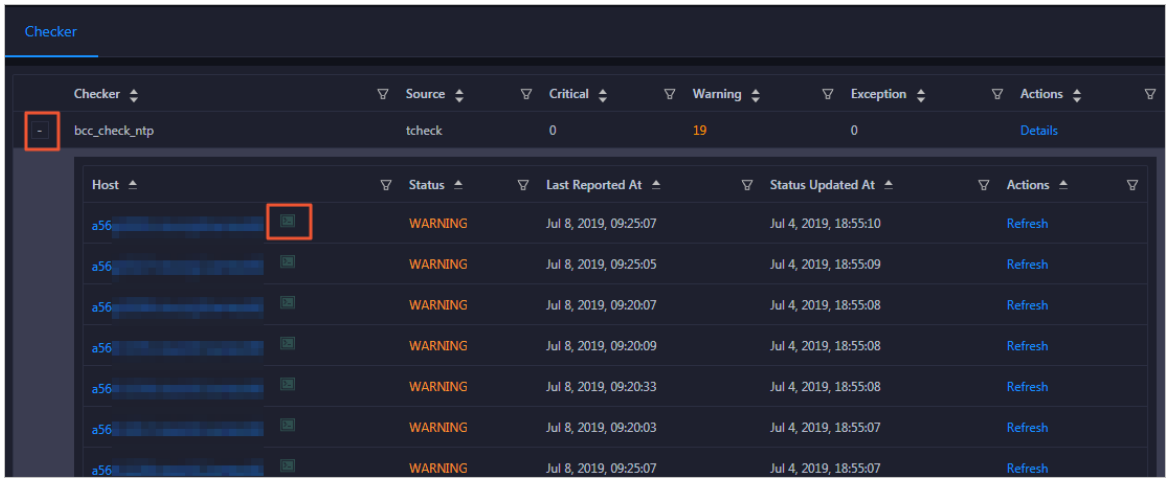

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

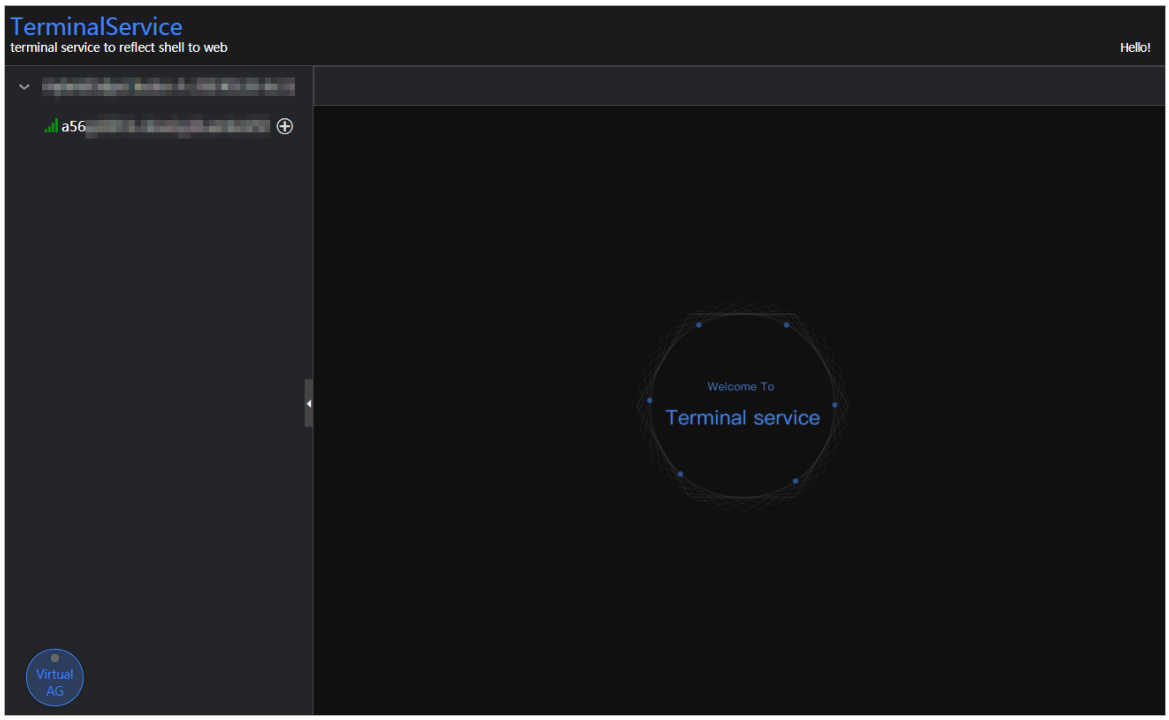

3. On the **TerminalService** page, click the hostname to log on to the host.

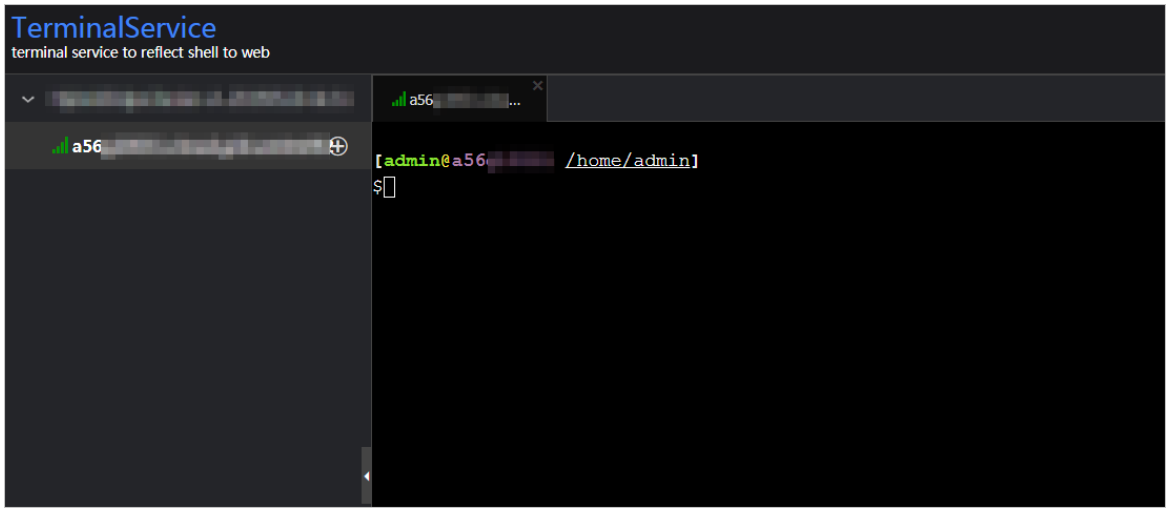

Run a checker again

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

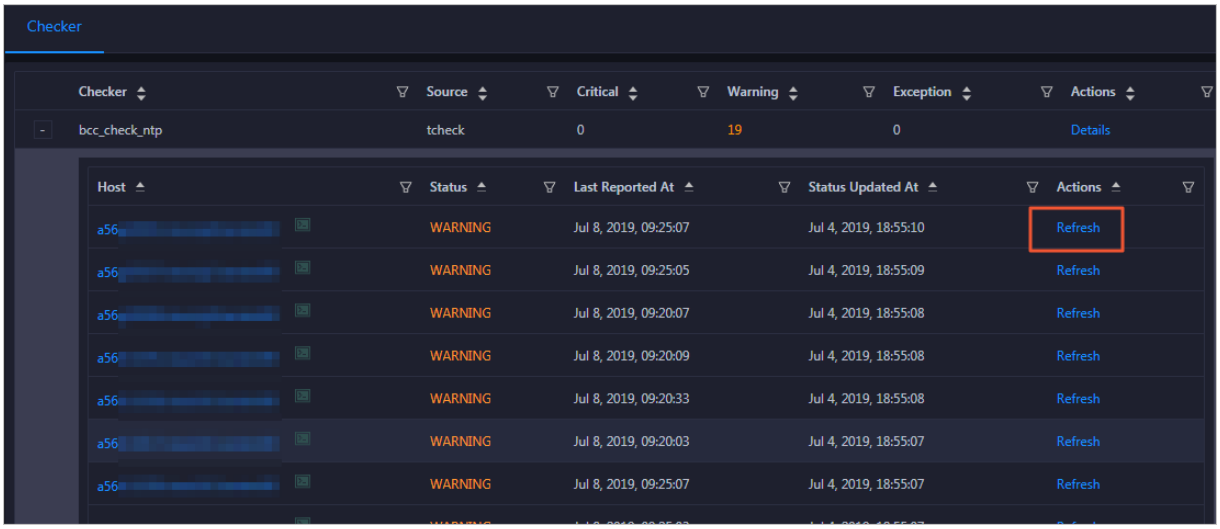

11.3.2.5. Host O&M

11.3.2.5.1. 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.

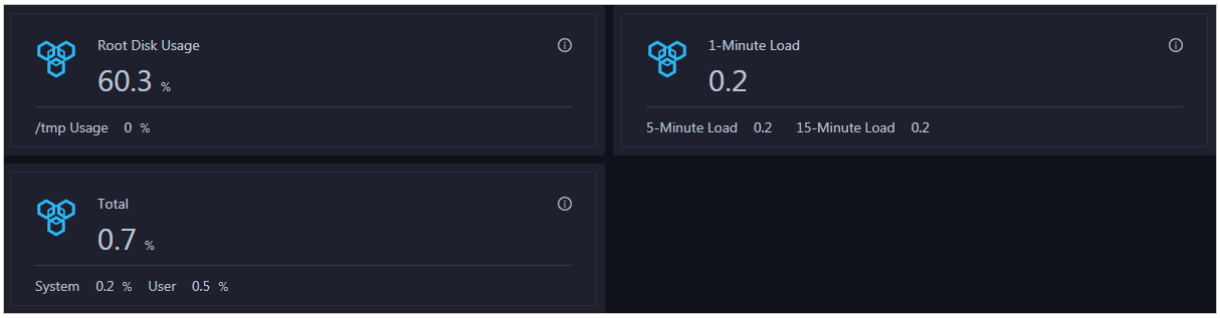

CPU

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.

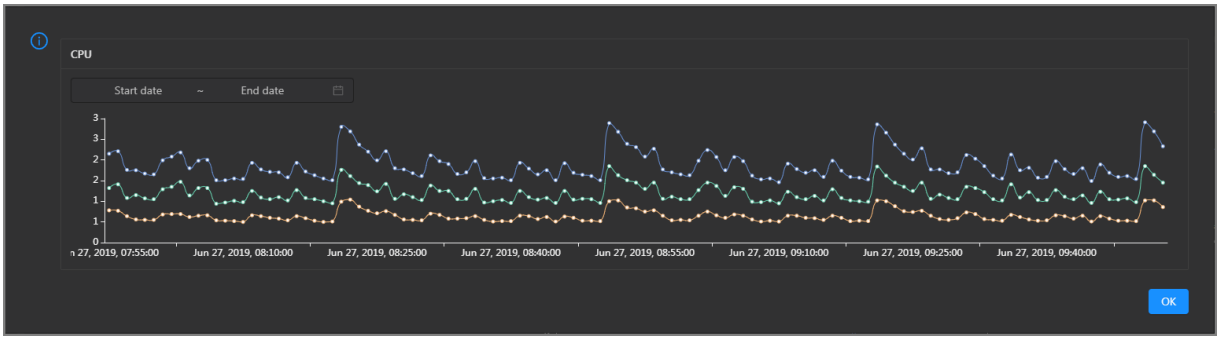

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

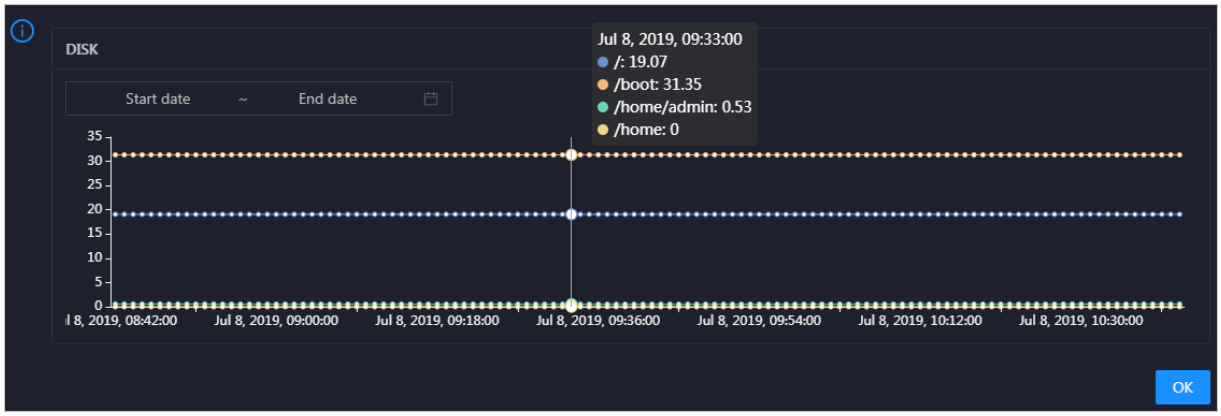

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

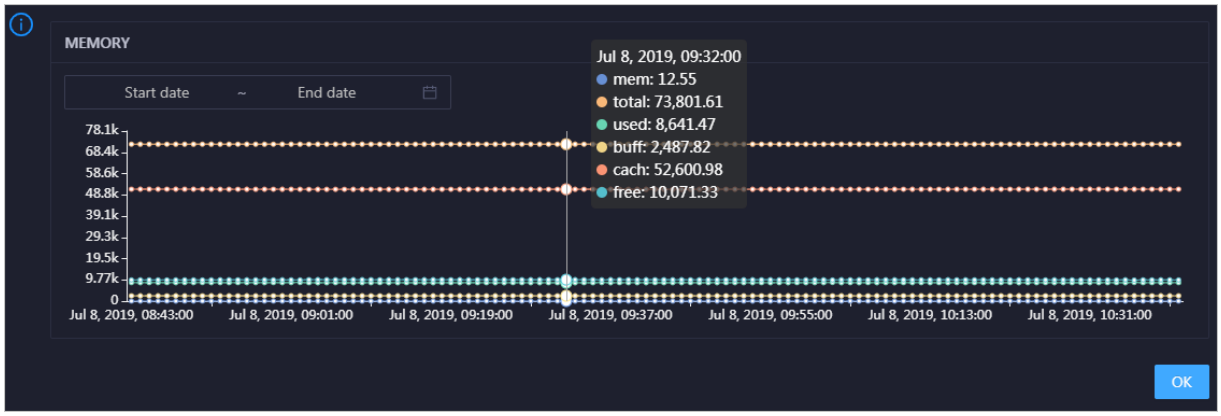

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

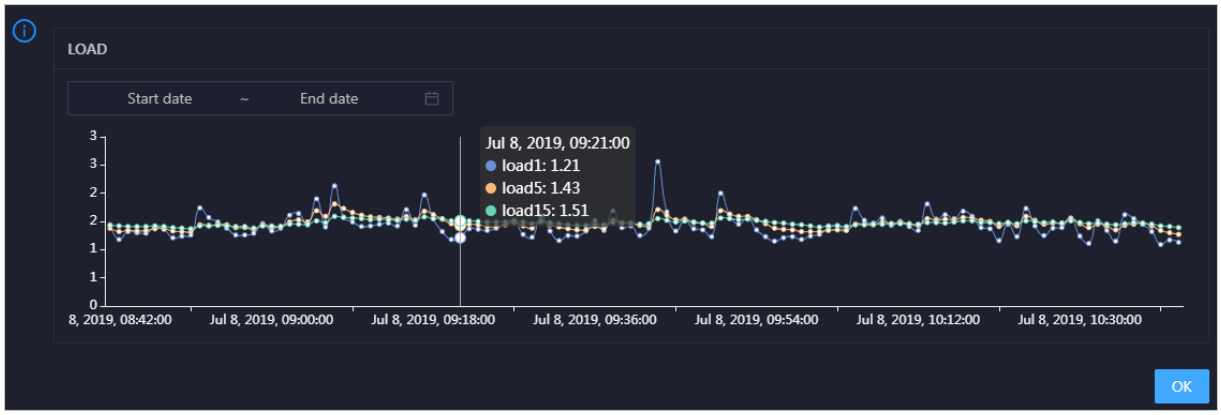

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

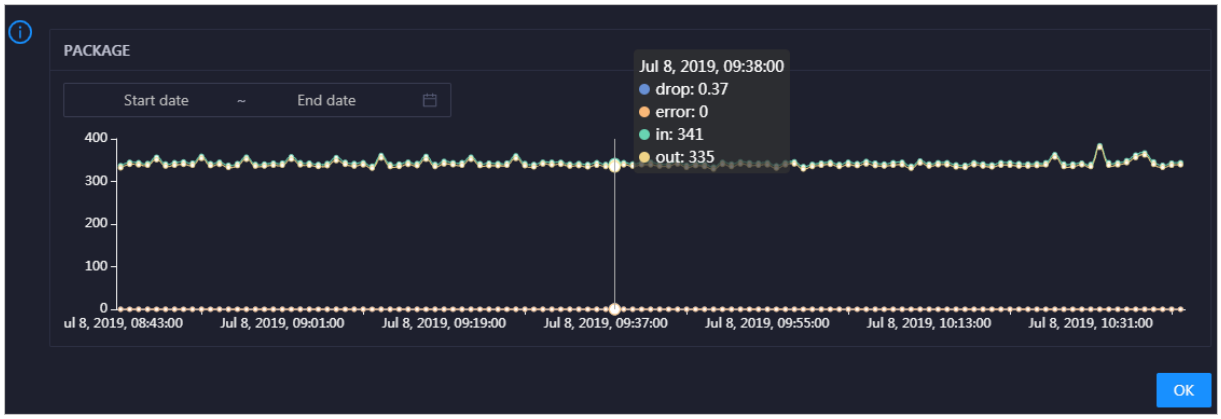

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

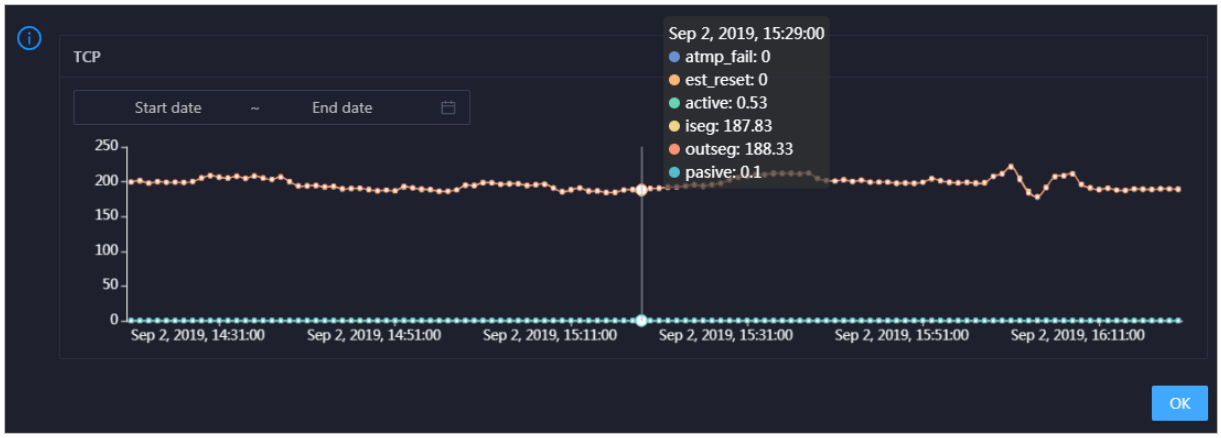

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

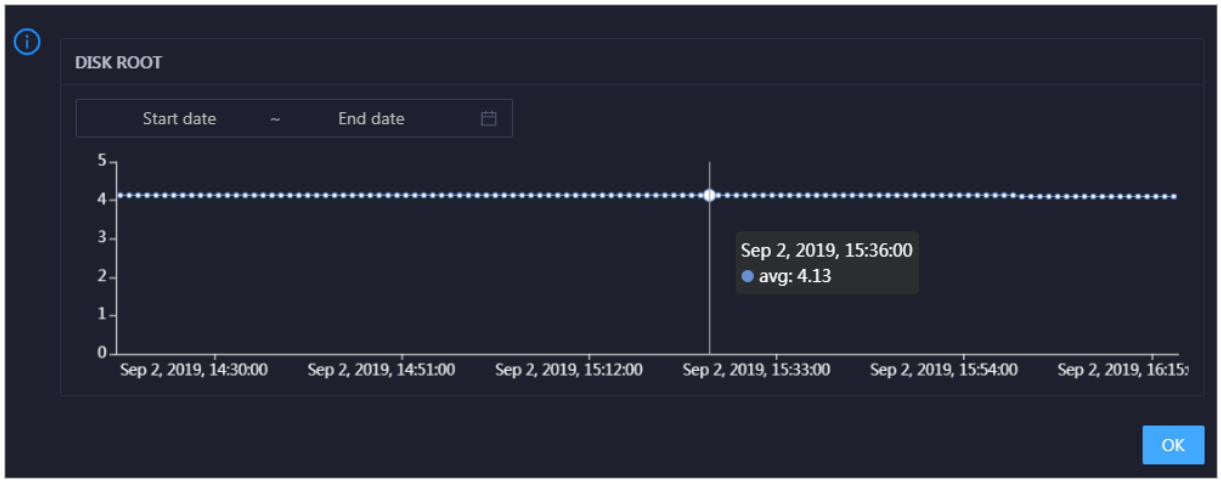

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

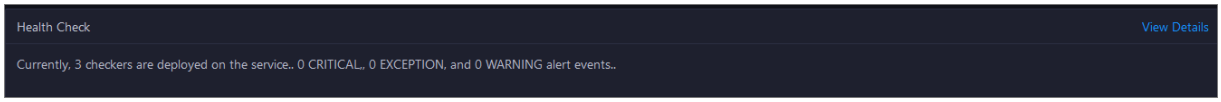

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

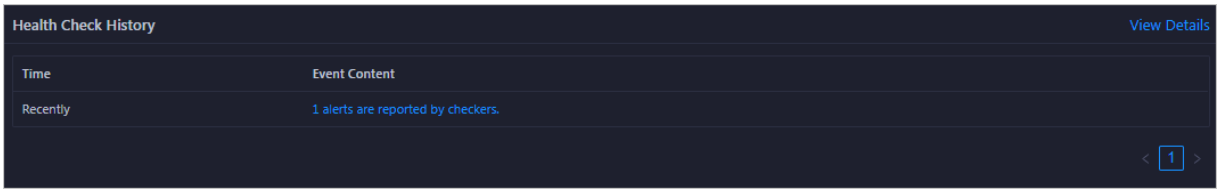

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

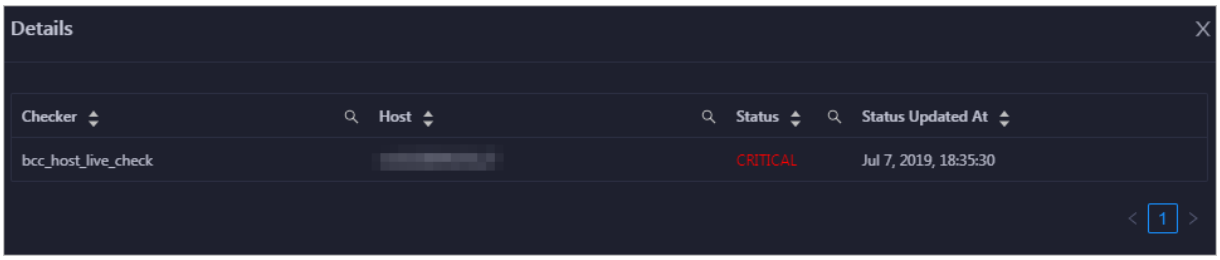

11.3.2.5.2. Host health

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.

Go to the Health Status page under Hosts

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

   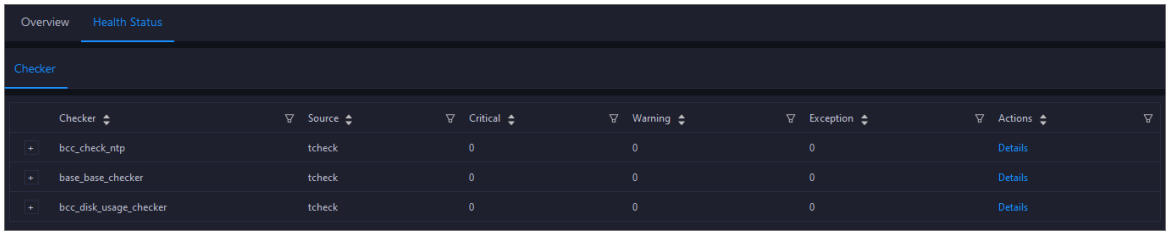

   On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

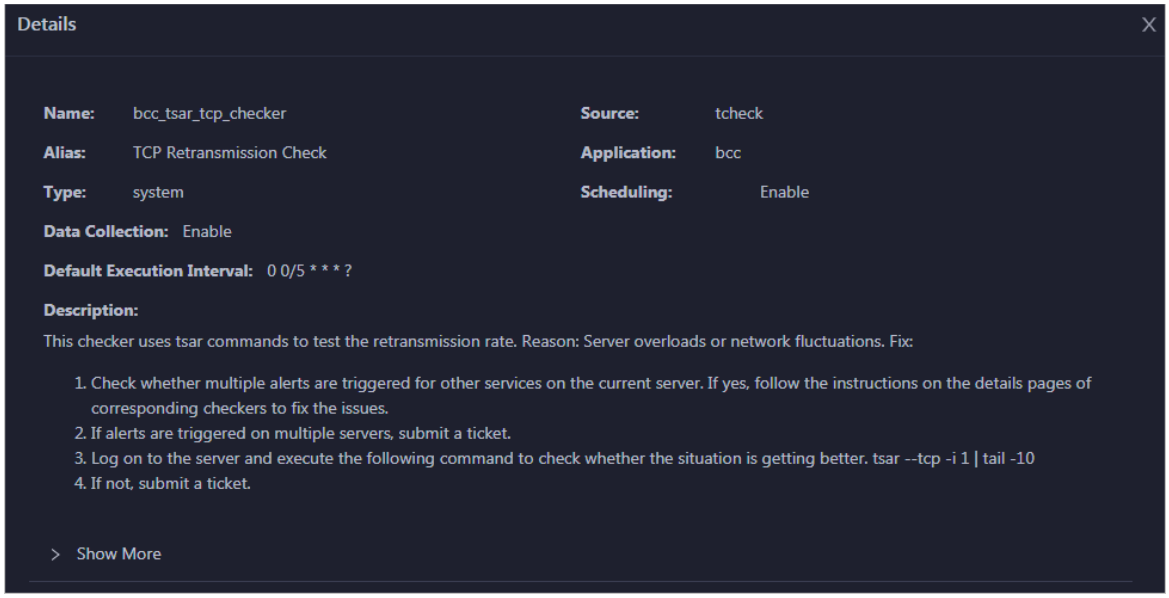

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

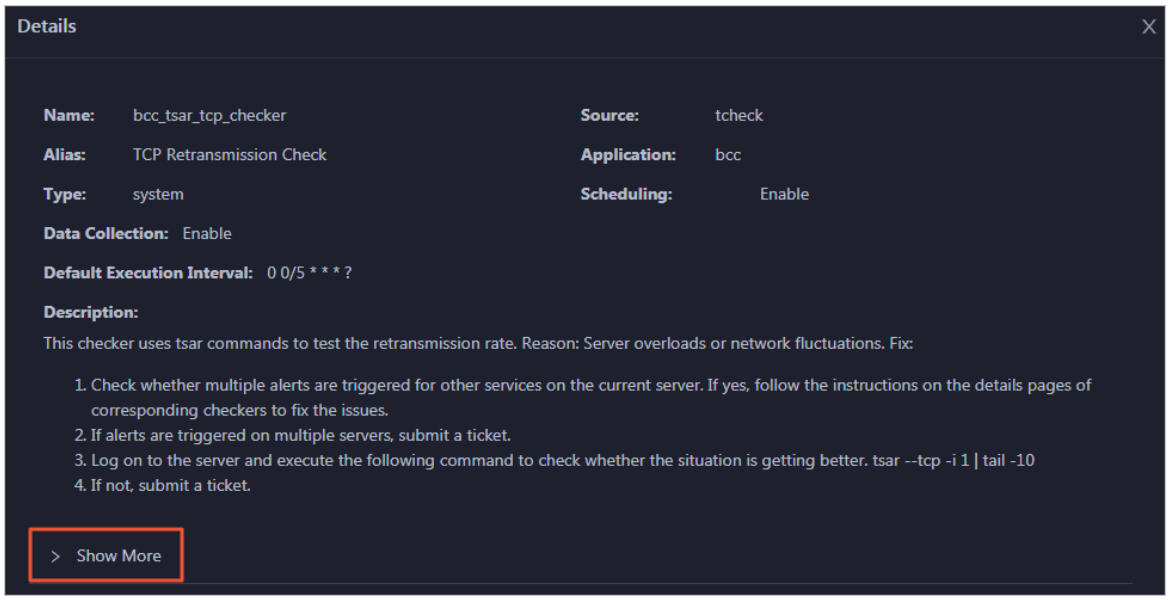

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

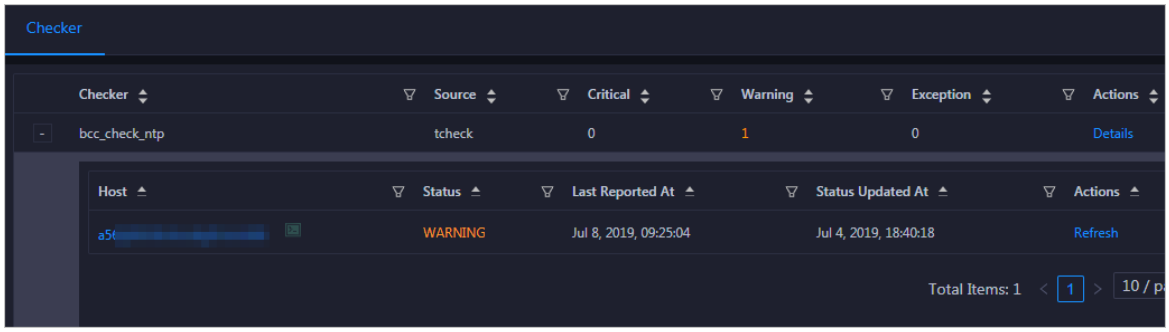

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

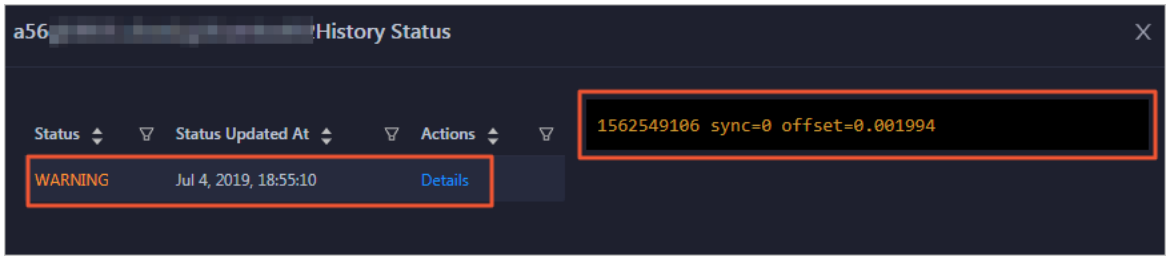

Clear alerts

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.

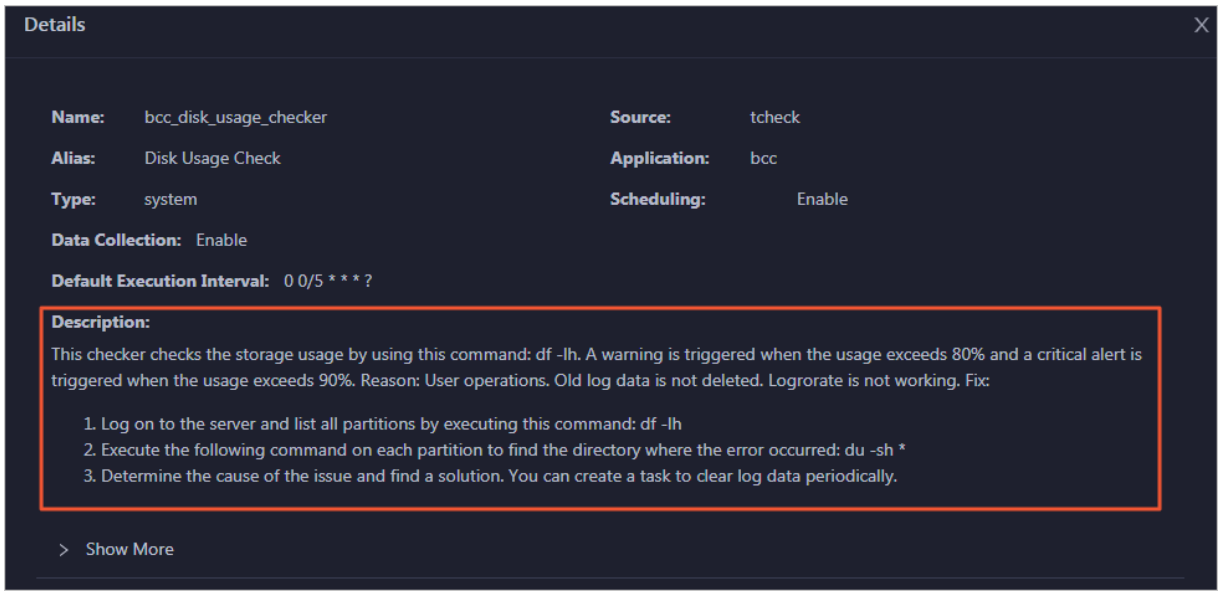

Log on to a host

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.

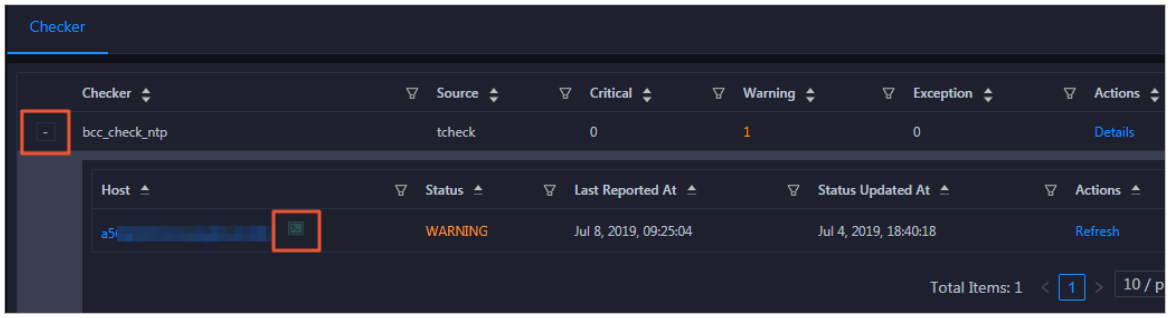

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

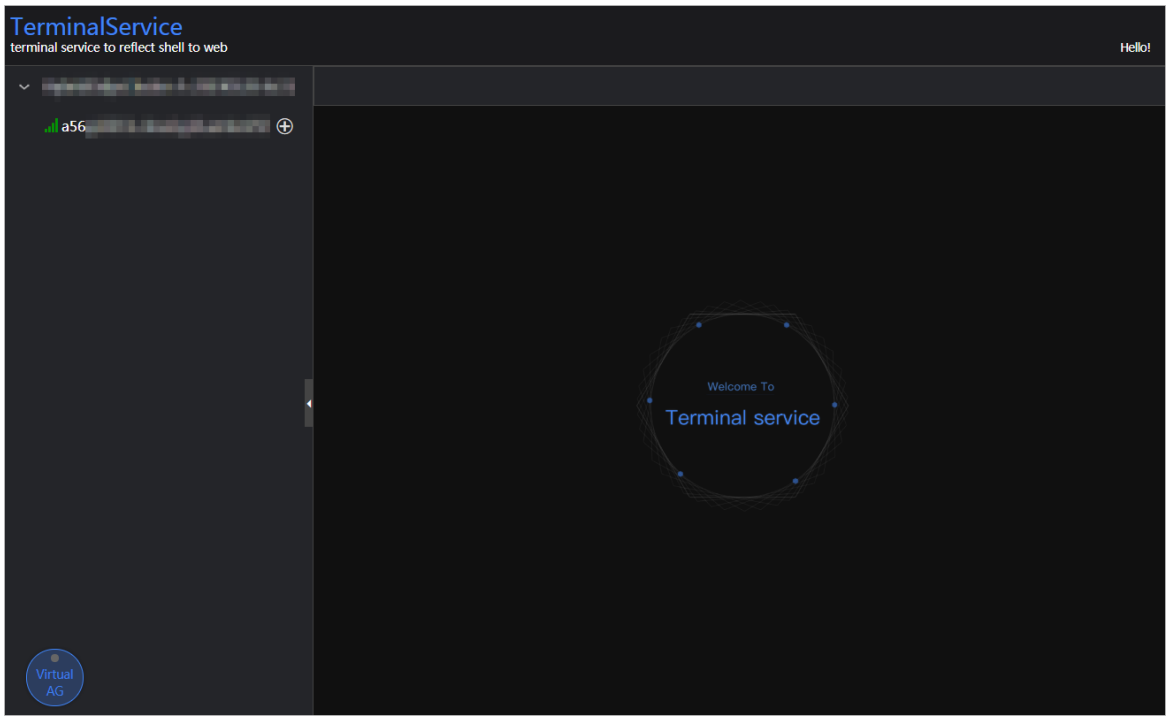

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

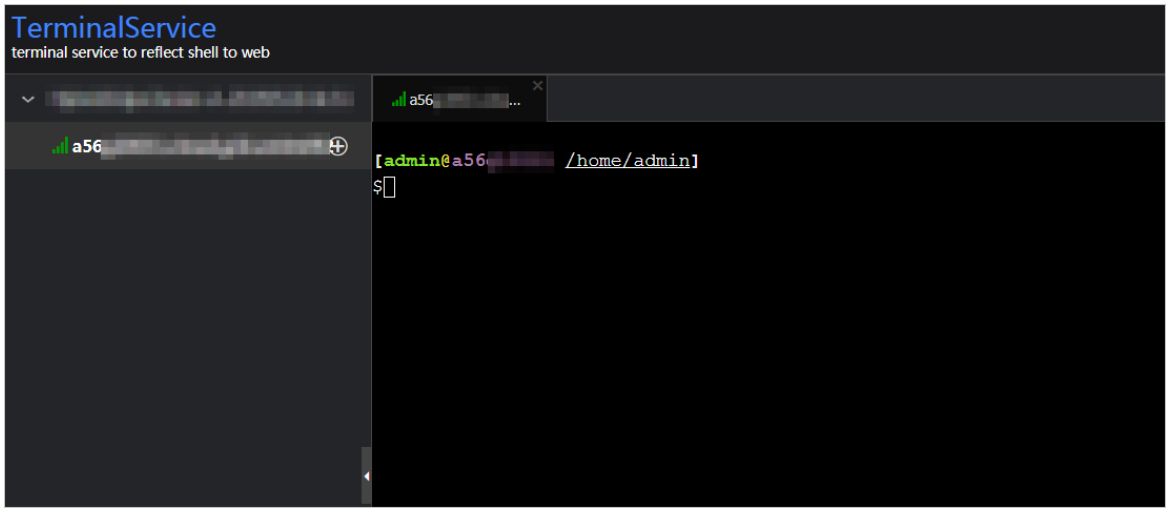

Run a checker again

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.

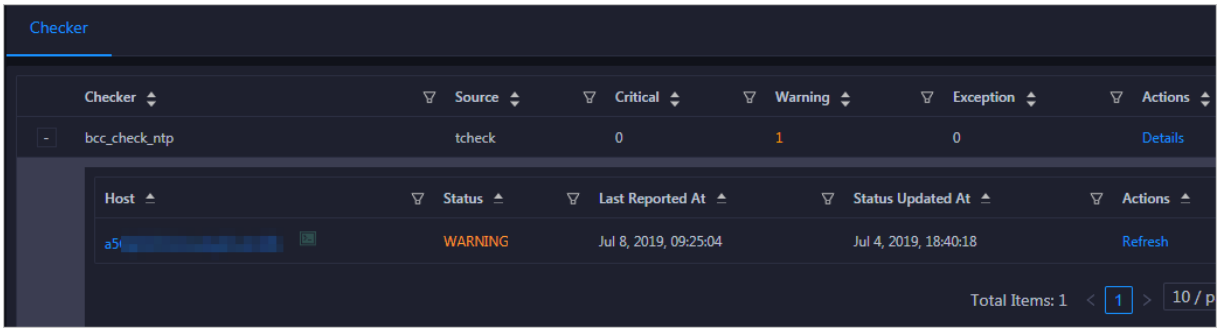

11.3.3. Common administration tools and commands

11.3.3.1. Find the host where a service resides

This topic describes how to find the host where a service resides.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.

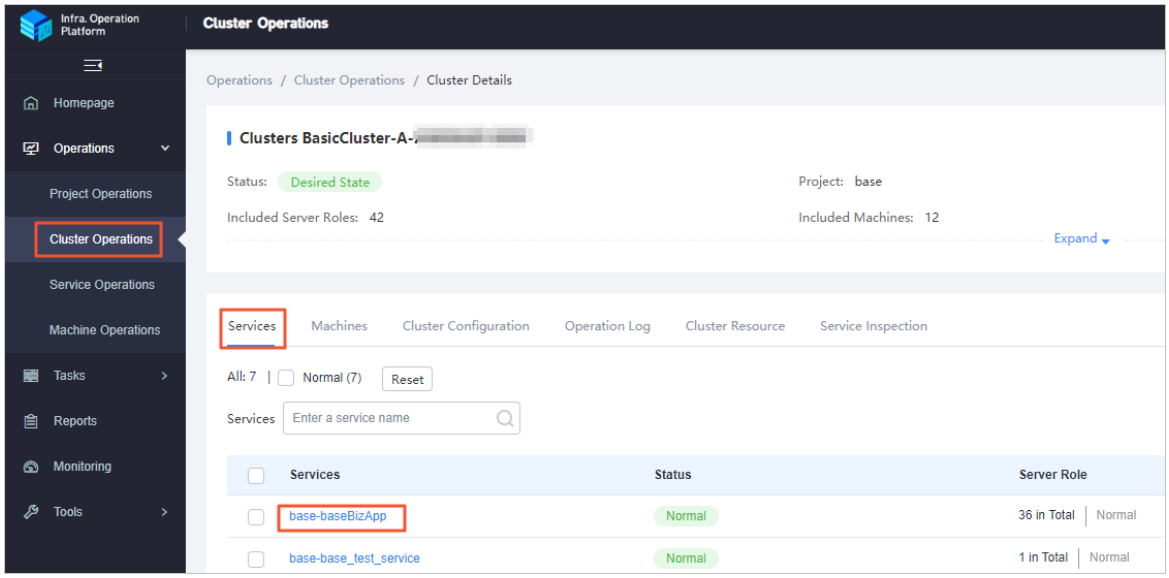

5. Click the **base-baseBizApp** service. On the **Machine Details** page, view information about the host where the service resides.

11.3.3.2. View cluster resources

This topic describes how to view the applications, resources, status, and version of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **Cluster Resource** tab. Then, you can filter required cluster information by **Server Role**, **Version**, **Name**, **Type**, or **Status**. Find the target application and click **Details** in the **Application Result** column to view detailed information about the application. For example, if you need to log on to the database of a specific application, you can find detailed logon information about the database in the **Application Result** message.

11.3.3.3. 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`.

11.3.3.4. View logs of a failed instance

This topic describes how to view logs of a failed instance in Operation Center.

Procedure

1. Log on to the DataWorks console.
2. On the DataStudio page, click  in the upper-left corner and choose **All Products > Operation Center**.
3. On the **Dashboard** page, click **Failed** in the **Instances** section. The **Cycle Instance** page appears. On this page, you can view all the instances that failed to run in the current workspace.

4. Click the target instance, and then right-click the failed node on the DAG that appears on the right side of the page.
5. Select **View Runtime Log**.

11.3.3.5. Rerun multiple instances at a time

You can use the batch rerun feature of DataWorks to rerun multiple instances at a time.

Procedure

1. Log on to the DataWorks console.
2. On the DataStudio page, click in the upper-left corner and choose **All Products > Operation Center**.
3. On the **Dashboard** page, click **Failed** in the **Instances** section. The **Cycle Instance** page appears. On this page, you can view all the instances that failed to run in the current workspace.
4. Select the instances to be rerun.
5. Choose **More > Rerun** in the lower-left corner.
6. In the message that appears, click **Rerun**.

11.3.3.6. Stop multiple instances at a time

You can use the batch stop feature of DataWorks to stop multiple instances at a time.

Procedure

1. Log on to the DataWorks console.
2. On the DataStudio page, click in the upper-left corner and choose **All Products > Operation Center**.
3. On the **Dashboard** page, click **Running** in the **Instances** section. The **Cycle Instance** page appears. On this page, you can view all the running instances in the current workspace.

   > **Note** You can stop only instances that are running.

4. Select the instances that you want to stop.
5. Choose **More > Stop** in the lower-left corner.
6. In the message that appears, click **Stop**.

11.3.3.7. Commonly used Linux commands

This topic describes the commonly used Linux commands.

Display workloads in the Linux system: top

View the three numbers after load average, which indicate the workload averages for the last 5, 10, and 15 minutes, respectively. If you divide one of these numbers by the quantity of logical CPUs and the result is greater than 5, the Linux system is overloaded.

List the sizes of files: du

You can run the `du -sh` command with a file name added at the end to view the size of the specified file. If you run the `du -sh *` command, you can view the sizes of all the files in the current directory.

List processes in the Linux system: ps

You can run the `ps -ef` command to view the processes that are running in the Linux system.

Search for strings: grep

To search for a string in a specified log file and display all lines that contain the string, run the command in the following format:

```
grep ["String"] [File name]
```

To search for a string in a specified file and display only the first few lines that contain the string, run the command in the following format:

```
grep -C Number of lines ["String"] [File name]
```

> **Note** The -C parameter is an uppercase letter. Set its value to a number.

To search for a string in a specified file and display only the last few lines that contain the string, run the command in the following format:

```
grep -A Number of lines ["String"] [File name]
```

Terminate processes: kill

You can run the `kill -9` command with the PID of a process added at the end to terminate the process.

docker commands

List all containers: `docker ps -a`

List the logs of a container: `docker logs [Container ID]`

Log on to a container: `docker exec -it [Container ID] bash`

11.3.3.8. View the slot usage of resource groups

This topic describes how to view the slot usage of a resource group.

Scenario: When a large number of nodes are waiting for resources in Operation Center, you can run a set of commands to view the slot usage of each resource group.

First, log on to the base-biz-alisa database. In an Apsara Stack V3 environment, select base from the Project drop-down list on the Clusters page in Apsara Infrastructure Management Framework. Locate the base-biz-alisa service whose type is db in the filtered results. Right-click the Result column and choose Show More from the shortcut menu to obtain the connection information of the database. Then, run a MySQL command to log on to the database based on the obtained information.

Run the following command to view the top 10 nodes by execution duration:

```
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 nodes by slot usage:

```
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 total number of nodes for each slot. Based on the command output, you can find out nodes that occupy a large number of slot resources.

```
select slot,count(*) from alisa_task where status=2 group by slot;
```

Run the following command to view the slot usage of each resource group:

```
select exec_target,sum(slot) from alisa_task where status=2 group by exec_target;
```

View the status of each gateway server. If the values of live and active_type are 1 for any gateway server, the gateway server fails.

```
select * from alisa_node;
```

11.3.4. Process daily administration operations

11.3.4.1. Daily check

11.3.4.1.1. Check the service status and basic server information

This topic describes how to view the basic cluster information, server status, and the number of servers in the desired state.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, click **Reports**.
3. On the **C** tab, select **base** from the **Project** drop-down list.
4. Move the pointer over next to the target server and select **Dashboard**. On the **Cluster Dashboard** page, view information in the **Basic Cluster Information**, **Machine Status**

Overview, and **Machines in Final Status** sections.

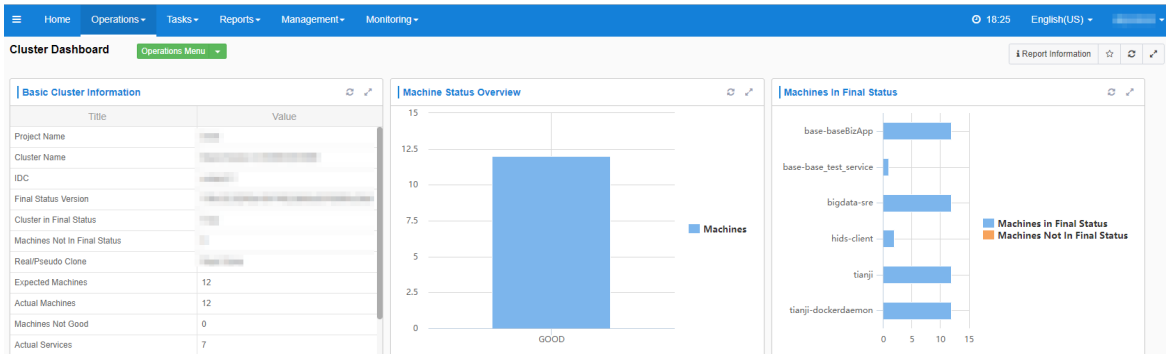

If only blue is shown in the Machine Status Overview section, all servers in the current cluster are running properly. If yellow is shown in the Machine Status Overview section, errors occur on servers.

11.3.4.1.2. Check the status of a gateway server

This topic describes how to check the status of a gateway server.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, click **Reports**.
3. On the **C** tab, select **base** from the **Project** drop-down list.
4. Move the pointer over next to the target server and select **Dashboard**.
5. In the **Cluster Resource** section, move the pointer over the **App** column and click .
6. In the dialog box that appears, enter **base-biz-alisa** in the Filter field and click **Apply Filter**.

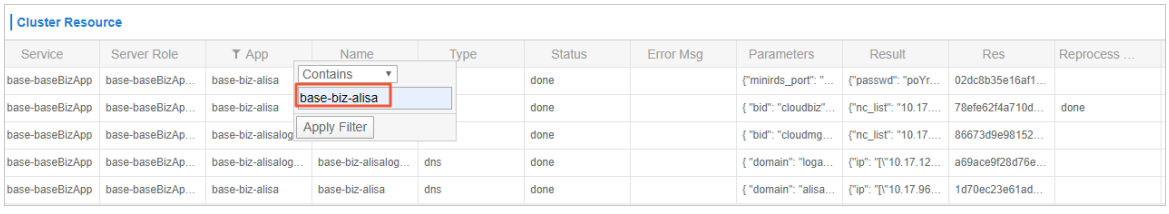

7. Filter services whose type is **db** in the same way.
8. Find the target service, right-click the information in the **Result** column, and then select **Show More** to view the endpoint, username, and password for logging on to the database.

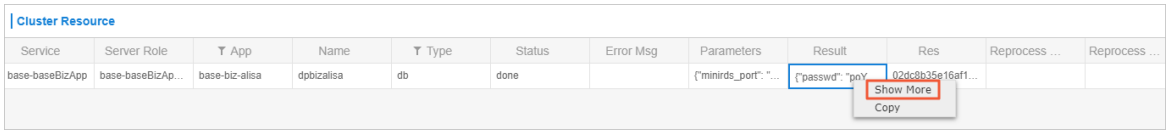

9. Connect to the database and run the following MySQL command to query the node information:

```
Select * from alisa_node;
```

If the value of the active_type or live parameter in the command output contains *-1* or *0*, the service is abnormal. Contact Alibaba Cloud technical support engineers.

11.3.4.1.3. Monitor service roles and servers

This topic describes how to view the details of monitored service roles and servers.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **Services** tab and then click the **base-baseBizApp** service.
6. On the **Service Details** page, click the target service role to view the servers where the service resides.
7. Find the target server and click **View** in the **Metric** column.
8. In the **Machine Metrics** dialog box, view the information on the **Server Role Metric** and **Machine Metrics** tabs.

11.3.4.2. 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`.

11.3.4.3. Scale out the cluster that runs the base-biz-gateway service

This topic describes how to scale out the cluster that runs the base-biz-gateway service.

Prerequisites

Before you apply a scaling change, make sure that the system is running in a status that is conducive to your change. For example, make sure that the storage space is large enough, and verify prerequisites such as the permissions on the files, the owners and paths of the files, and the software version.

- Before you scale out a BasicCluster cluster, make sure that it reaches the desired state and works as expected.
- A screenshot of the key initial configurations for the cluster is saved.
- IP addresses do not conflict. If you need 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.
- The clone_mode parameter is set to normal.

  > **Note** Apsara Infrastructure Management Framework of V3.3 and later versions support cloning protection. Before the scale-out, you must set the clone_mode parameter to normal. After the scale-out is complete, set this parameter to block.

  To set the clone_mode parameter to normal, perform the following steps:

  i. Log on to Apsara Infrastructure Management Framework.

  ii. In the left-side navigation pane, click **Reports**.

  iii. In the top navigation bar, choose **Operations > Cluster Operations**.

  iv. Click **Global Clone Switch**. In the **Global Clone Switch** dialog box, select **normal**.

  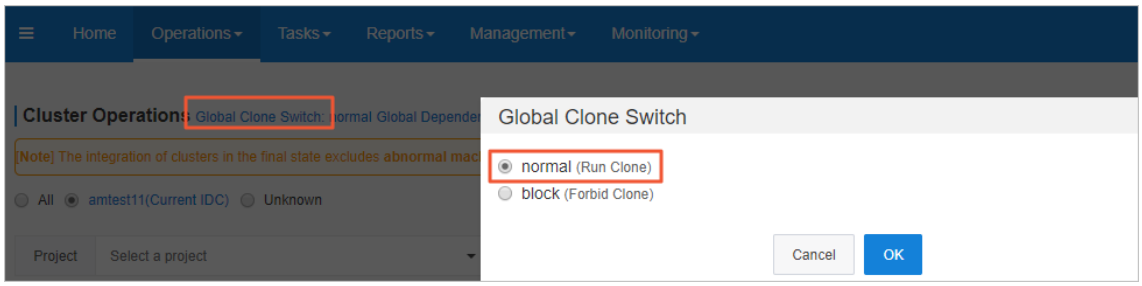

  v. Click **OK**.

Procedure

1. Create a buffer cluster. Skip this step if an existing buffer cluster has idle servers and the physical machine, memory, CPU, and disk size of the idle servers are the same as those of current servers that run the base-biz-gateway service.

   *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.

   i. 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 ${Final status in Apsara Infrastructure Management Framework} -c tianji_dest.conf
   ```

   ii. Log on to Apsara Infrastructure Management Framework. Select **buffer** from the **Project** drop-down list.

   iii. Click the buffer cluster to view the status of servers in the cluster.

iv. 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
```

2. Scale in the buffer cluster by moving idle servers to the default cluster.

> **Note** You can use the default cluster to scale out the cluster that runs base-biz-cdp and base-biz-gateway services.

i. On the **Cluster Operations** page, click the target buffer cluster.

ii. On the **Cluster Details** page, click the **Cluster Configuration** tab.

iii. Click the **machine_group.conf** file. Make sure that the value of the scalable tag value is true for the new buffer cluster.

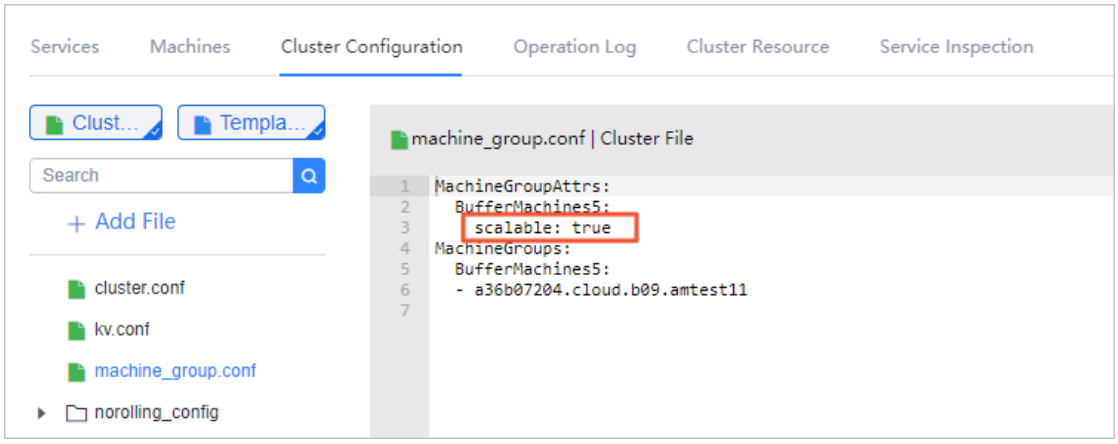

iv. Run the following commands on the OPS1 server to scale in the buffer cluster:

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.conf clt2.conf (After you run this command, if a message appears indicating that a soft link already exists, proceed with the next 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]
```

Parameter	Description
-c	The name of the buffer cluster that you scale in, which starts with buffer-cluster.
-l	The list of hostnames of servers to be removed. Separate multiple hostnames with commas (,).
-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.
-config	The tianji_clt file. **Note** These commands cannot contain Chinese characters.

v. On the **Cluster Operations** page, verify that the servers have been removed.

vi. Run the following commands on the OPS1 server to check scale-in 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 (After you run this command, if a message appears indicating that a soft link already exists, proceed with the next command.)
./tianji_clt machinestatus -c default --config clt2.conf
```

vii. Go to the details page of the new buffer cluster, and check whether the servers have been deleted from the machine_group.conf file. If the servers still exist, delete them from the machine_group.conf file and then submit a rolling task.

3. Add servers to the BasicCluster cluster and specify the name of the SRG where the servers reside.

i. Log on to Apsara Infrastructure Management Framework. In the left-side navigation pane, click **Reports**.

ii. In the top navigation bar, choose **Operations > Cluster Operations**.

iii. Right-click the target BasicCluster cluster and choose Monitoring > **Cluster Configuration**. On the page that appears, verify that **Clone Switch** is set to **Real Clone**.

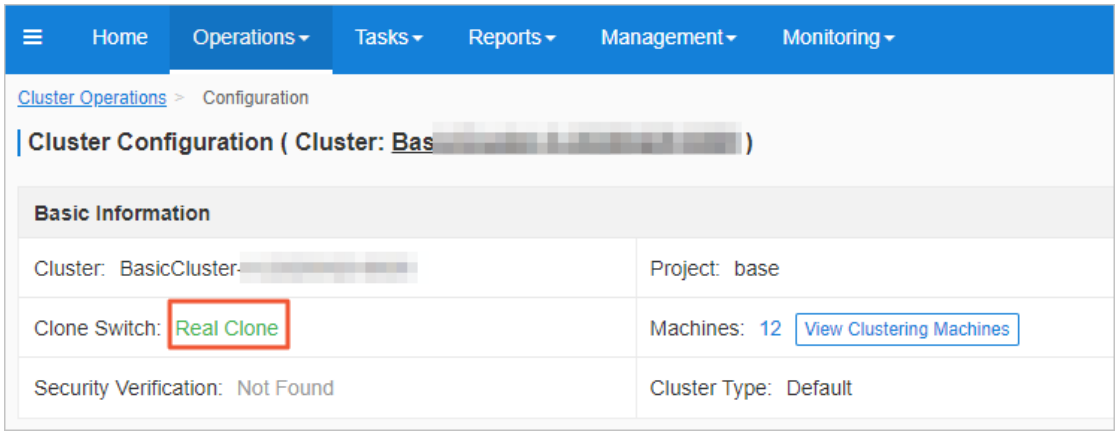

iv. Run the following commands to perform scaling. A rolling task is triggered after you run these commands.

```
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 soft link already exists, proceed with the next command.)
```

To add servers to the cluster that runs the base-biz-gateway service, run the following command:

```
./tianji_ops_tool.py expand_nc -c [BasicCluster cluster name] -s BaseGwGroup -l [machine1,machine2] --config clt2.conf
```

To add servers to the cluster that runs the base-biz-gateway service, run the following command:

```
./tianji_ops_tool.py expand_nc -c [BasicCluster cluster name] -s BaseCdpGwGroup -l [machine1,machine2] --config clt2.conf
```

v. Run the following command to call an API operation to check the scaling result:

```
curl http://127.0.0.1:7070/api/v3/column/m. *?m.id=[machine hostname]
```

vi. 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 -
```

4. 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.
5. Verify the scale-out.

i. Check the heartbeat logs of the servers. Open the terminal of each added server, log on to the gateway container, and then run the `tail -f /home/admin/alisatasknode/logs/heartbeat.log` command.

   If the logs are refreshed every 5 seconds, the heartbeat service is running as expected.

ii. Query the database information and check whether the server is online. Log on to Apsara Infrastructure Management Framework. Go to the **Cluster Details** page of the target BasicCluster cluster, click the **Cluster Resource** tab, set the Type parameter to **db**, and then find the **base-biz-alisa** service. Click **Details** in the **Application Result** column to check the database connection information.

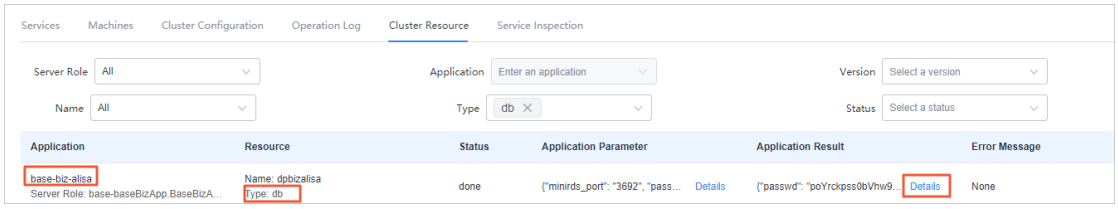

   Connect to the database by using a MySQL command, and run the `select * from alisa_node;` command. The information of all servers that run the base-biz-gateway service appears.

   Check the values of the live and active_type parameters for each added server. If both the two values are 1, the server is online.

11.3.4.4. 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/server vtl 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
   ```

11.3.4.5. Restart the base-biz-tenant service

This topic describes how to go to the Service Details page and restart the base-biz-tenant service.

Go to the Service Details page

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **base-baseBizApp** service. The **Service Details** page appears.

Restart the base-biz-tenant service

1. On the Service Details page, click the **base-baseBizApp.BaseBizTenant** service. You can also enter BaseBizTenant in the **Server Role** field to search for the target service.

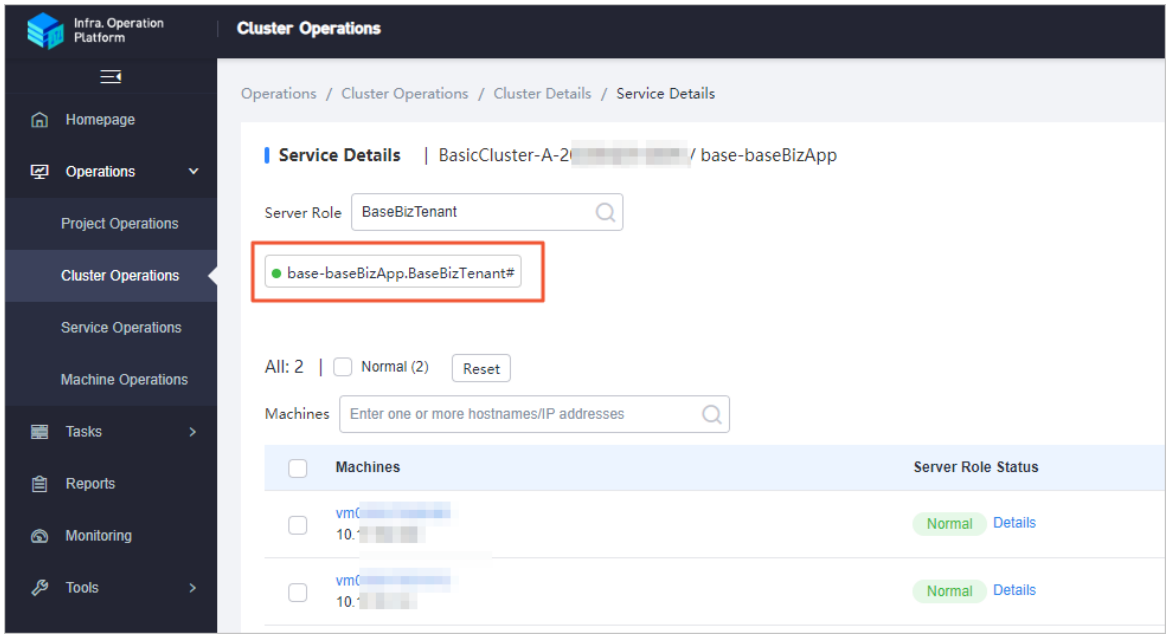

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep tenant` and press Enter to view the ID of the server.

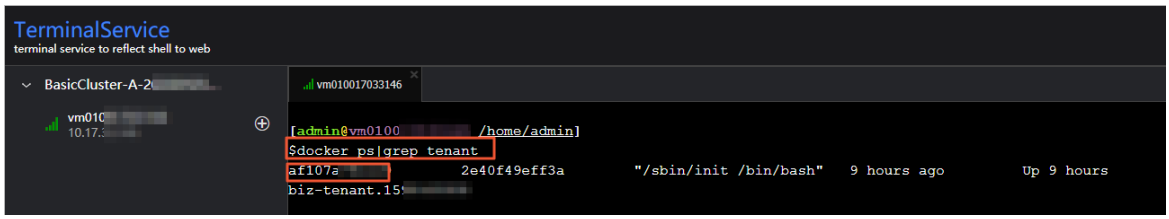

6. Enter `docker exec -it yourIDbash` and press Enter to enter the corresponding container.

> **Note** The *yourID* parameter specifies the container ID.

7. Enter `su - admin` and press Enter to switch to the admin user.
8. Enter `/home/admin/base-biz-tenant/bin/jbossctl restart` and press Enter to restart the base-biz-tenant service.

When the **OK** and **NGINX start Done** information appears, the base-biz-tenant service is restarted.

11.3.4.6. Restart the Redis services

This topic describes how to go to the Service Details page and restart the Redis services.

Go to the Service Details page

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **base-baseBizApp** service. The **Service Details** page appears.

Restart the Redis services

1. On the Service Details page, click the **base-baseBizApp.Redis1#** service.Redis services include Redis1 and redis2. This topic uses Redis1 as an example. You can also enter Redis in the **Server Role** field to search for the target service.
2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep redis` and press Enter to view the ID of the server.

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

   > **Note** The *yourID* parameter specifies the container ID.

7. Enter the following statements and press Enter to restart the Redis service:

   ```
   /etc/init.d/redis restart
   ```

   ```
   /etc/init.d/redis-sentinel restart
   ```

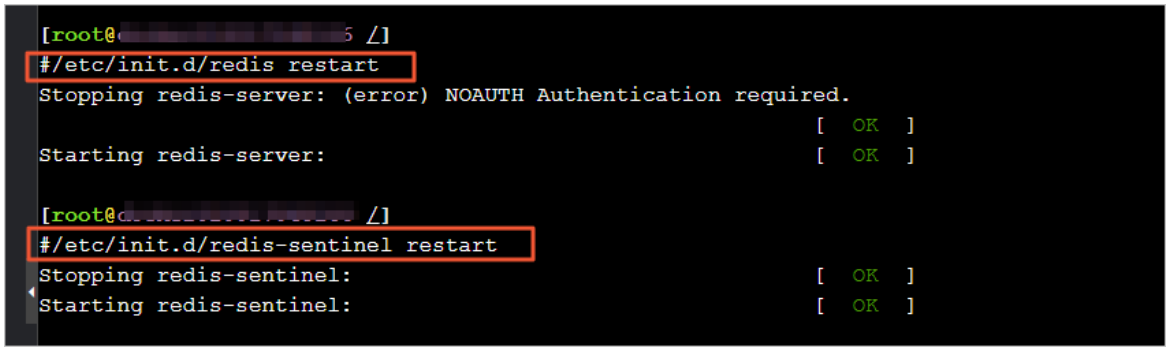

11.3.4.7. Restart the base-biz-dmc service

This topic describes how to go to the Service Details page and restart the base-biz-dmc service.

Go to the Service Details page

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **base-baseBizApp** service. The **Service Details** page appears.

Restart the base-biz-dmc service

1. On the Service Details page, click the **base-baseBizApp.BaseBizDmc** service.You can also enter BaseBizDmc in the **Server Role** field to search for the target service.

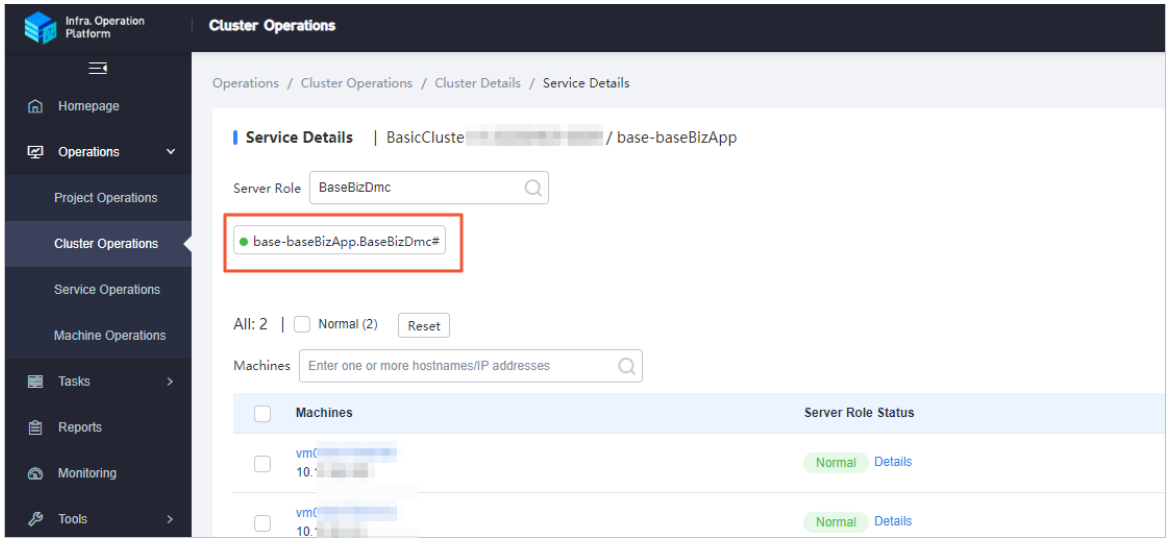

2. Click the name of the target server. The Machine Details page appears.
3. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
4. Enter `docker ps|grep dmc` and press Enter to view the ID of the server.

5. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **Note** The *yourID* parameter specifies the container ID.

6. Enter `su - admin` and press Enter to switch to the admin user.

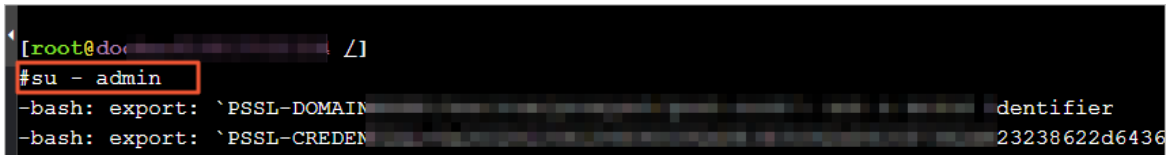

7. Enter `/home/admin/base-biz-dmc/bin/jbossctl restart` and press Enter to restart the base-biz-dmc service.

When the **OK** and **NGINX start Done** information appears, the base-biz-dmc service is restarted.

11.3.4.8. Restart the base-biz-alisa service

This topic describes how to go to the Service Details page and restart the base-biz-alisa service.

Go to the Service Details page

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **base-baseBizApp** service. The **Service Details** page appears.

Restart the base-biz-alisa service

1. On the Service Details page, click the **base-baseBizApp.BaseBizAlisa** service. You can also enter BaseBizAlisa in the **Server Role** field to search for the target service.

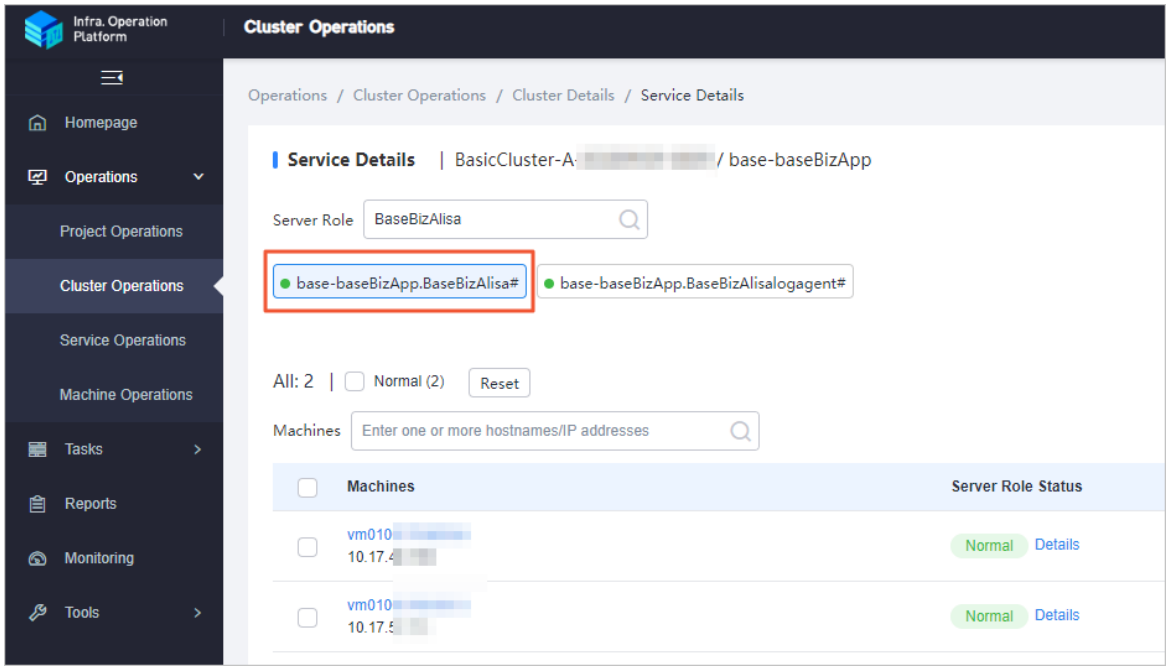

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep alisa` and press Enter to view the ID of the server.

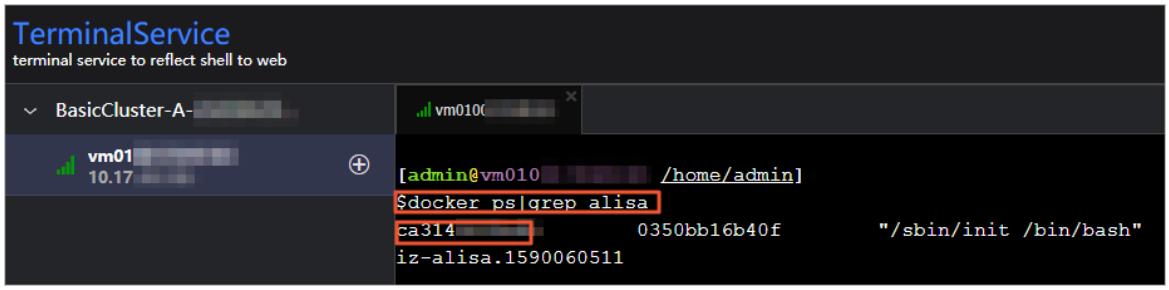

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **Note** The *yourID* parameter specifies the container ID.

7. Enter `su - admin` and press Enter to switch to the admin user.
8. Enter `/home/admin/base-biz-alisa/bin/jbossctl restart` and press Enter to restart the base-biz-alisa service. When the **OK** and **NGINX start Done** information appears, the base-biz-alisa service is started.

11.3.4.9. Restart the base-biz-phoenix service

This topic describes how to go to the Service Details page and restart the base-biz-phoenix service.

Go to the Service Details page

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **base-baseBizApp** service. The **Service Details** page appears.

Restart the base-biz-phoenix service

1. On the Service Details page, click the **base-baseBizApp.BaseBizPhoenix** service. You can also enter BaseBizPhoenix in the **Server Role** field to search for the target service.

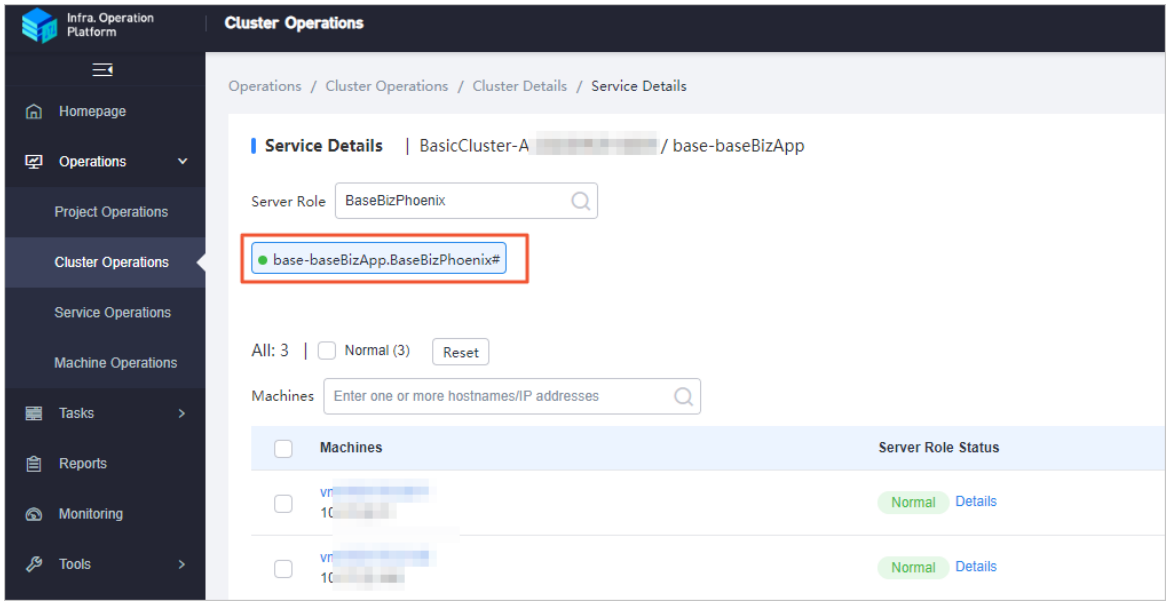

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep phoenix` and press Enter to view the ID of the server.

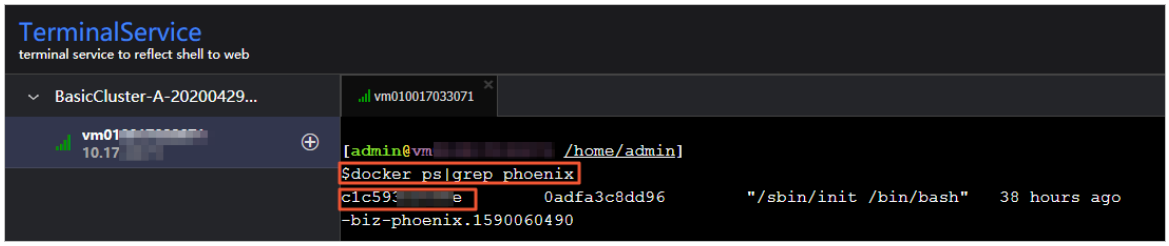

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **Note** The *yourID* parameter specifies the container ID.

7. Enter `su - admin` and press Enter to switch to the admin user.
8. Enter `/home/admin/base-biz-phoenix/bin/jbossctl restart` and press Enter to restart the base-biz-phoenix service.

```
[root@dock                    /]
#su - admin

[admin@docker(                /home/admin]
$/home/admin/base-biz-phoenix/bin/jbossctl restart

                                                    /nginx-proxy.conf -p /home/admin/cai stop

Wait Tomcat Start: 22...                                    [    ]
init /home/admin/cai/.running_conf/
copy from /opt/taobao/tengine/conf/  to  /home/admin/cai/.running_conf/
copy from /home/admin/cai/conf/   to  /home/admin/cai/.running_conf/
init /home/admin/cai/.running_conf/ done
/opt/taobao/tengine/bin/tengine -c /home/admin/cai/.running_conf/nginx-proxy.conf -p /home/admin/cai
NGINX start Done.
```

When the **OK** and **NGINX start Done** information appears, the base-biz-phoenix service is started.

11.3.4.10. Restart the base-biz-gateway service

This topic describes how to go to the Service Details page and restart the base-biz-gateway service.

Go to the Service Details page

1. Log on to Apsara Infrastructure Management Framework.
   i. Log on to the ABM console.
   ii. In the left-side navigation pane, choose **Products > Product List**.
   iii. On the page that appears, choose **Apsara Stack O&M > Basic O&M > Apsara Infrastructure Management Framework**.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the page that appears, select **base** from the **Project** drop-down list to filter clusters.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. Click the **base-baseBizApp** service. The **Service Details** page appears.

Restart the base-biz-gateway service

1. On the Service Details page, click the **base-baseBizApp.BaseBizCdpGateway** service. You can also enter BaseBizCdpGateway in the **Server Role** field to search for the target service.

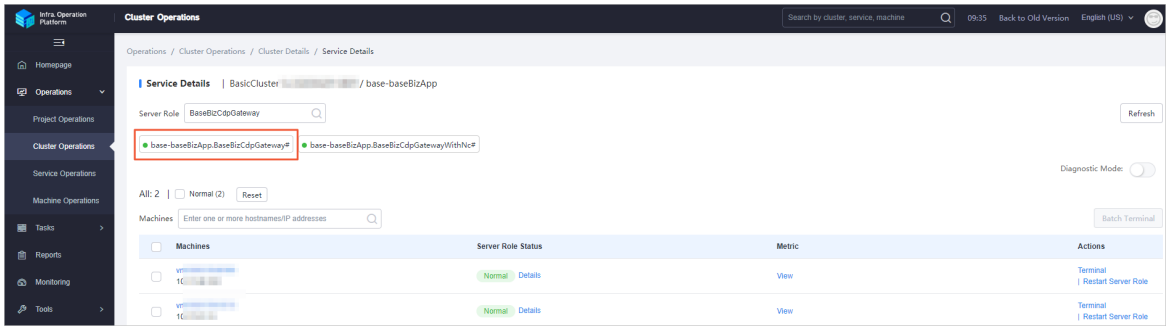

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The

configuration tab appears on the right-side of the page.

5. Enter `docker ps|grep gateway` and press Enter to view the ID of the server.

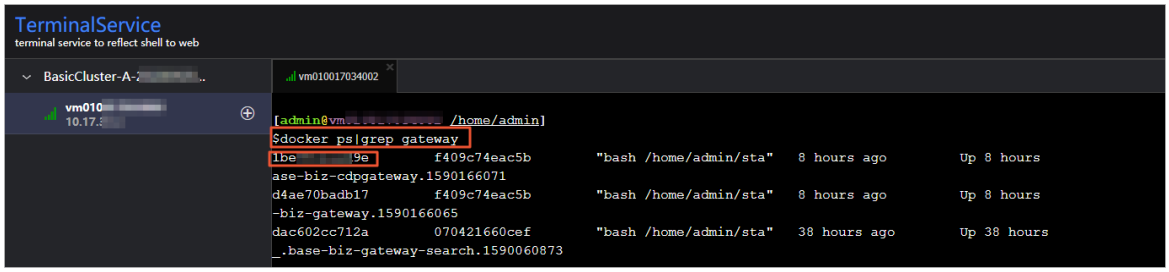

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **Note** The *yourID* parameter specifies the container ID.

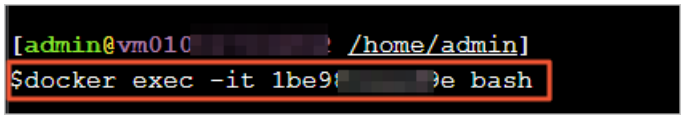

7. Enter `su - admin` and press Enter to switch to the admin user.
8. Enter `/home/admin/alisatasknode/target/alisatasknode/bin/serverctl restart` and press Enter to restart the base-biz-gateway service.

11.3.4.11. 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.

11.3.4.12. 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.

11.3.5. Common issues and solutions

11.3.5.1. 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.
  i. 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.
  ii. 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 or
der 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.

  i. 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.

  ii. Check the system time of servers in the base cluster against the time of the database.

  iii. 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:

i. Rename the /home/admin/alisatasknode/taskinfo/20180306/phoenix directory as /home/admin/alisatasknode/taskinfo/20180306/phoenix.bak.

ii. 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.

11.3.5.2. 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';
```

11.3.5.3. 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.

11.3.5.4. 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.

11.3.5.5. 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.

11.3.5.6. 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 {} \;
   ```

11.3.5.7. 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.

11.4. Realtime Compute

11.4.1. Job status

11.4.1.1. 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.

11.4.1.2. 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.

11.4.1.3. 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.

11.4.1.4. Job instantaneous values

Job parameters

Name	Description
Consumed compute time	Indicates the computing performance of a job.
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.

Name	Description
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.

11.4.1.5. 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 View mode, to offer you a more intuitive display.

View mode

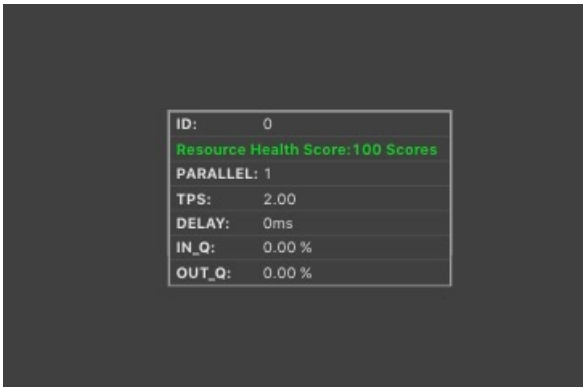

You can view the detailed information about a task by moving the pointer over the task. Parameter description describes the task parameters.

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.
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 Task details page.

Task details page

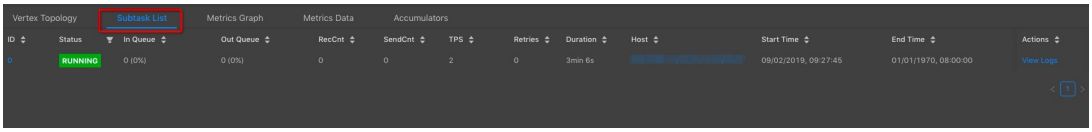

The Curve Charts tab provides curve charts to show the metrics of each task, as shown in Curve charts for task metrics.

Curve charts for task metrics

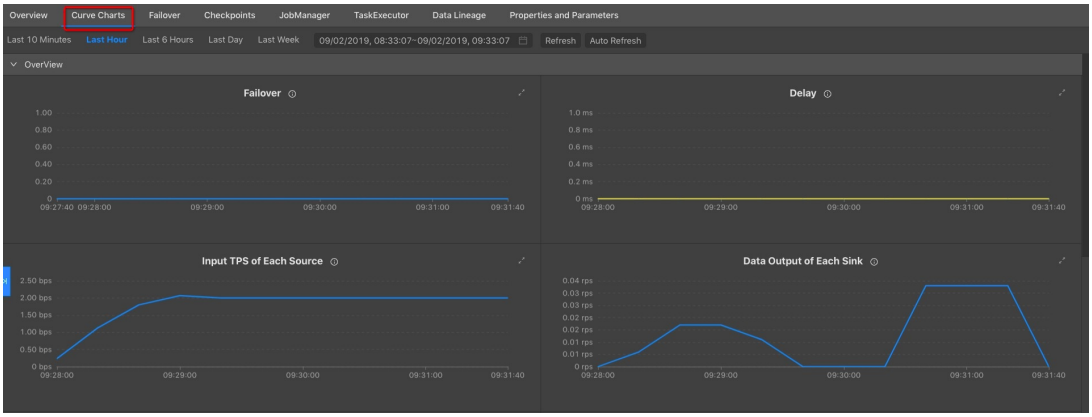

List mode

In addition to the view mode, Realtime Compute also allows you to view each task in the list mode, as shown in List mode.

List mode

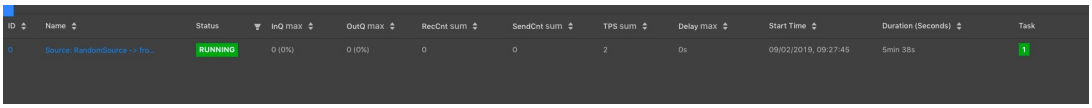

Parameter description describes the task parameters.

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.

11.4.2. Curve charts

11.4.2.1. 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.

Curve Charts tab

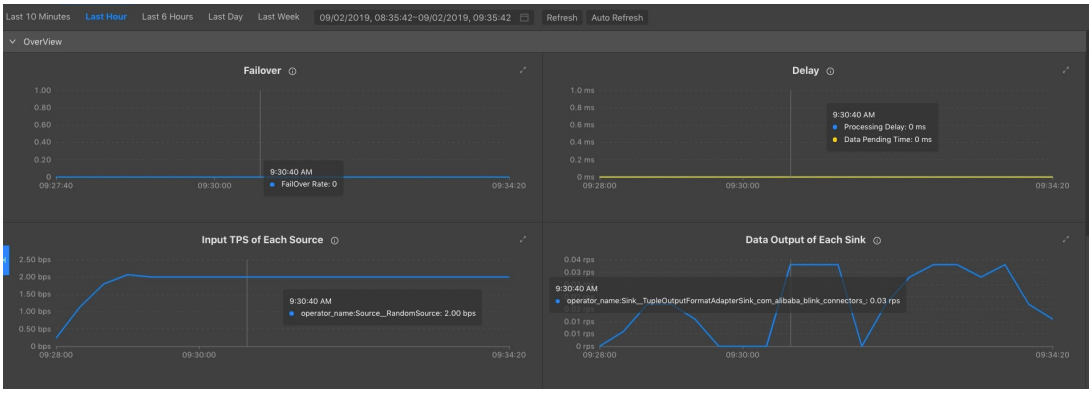

Note

- 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.

11.4.2.2. 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.

Processing delay

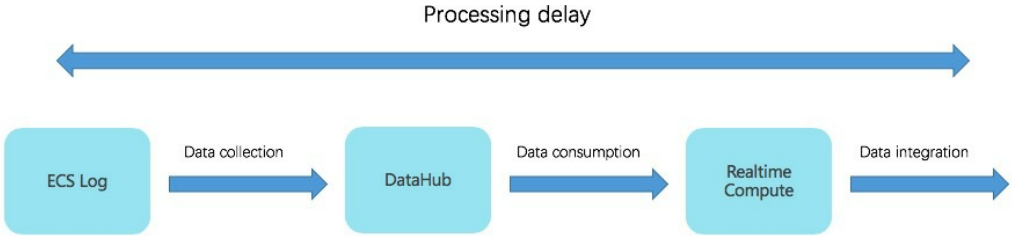

The processing delay can be categorized into the following three types:

- 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.

11.4.2.3. 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.

Barriers

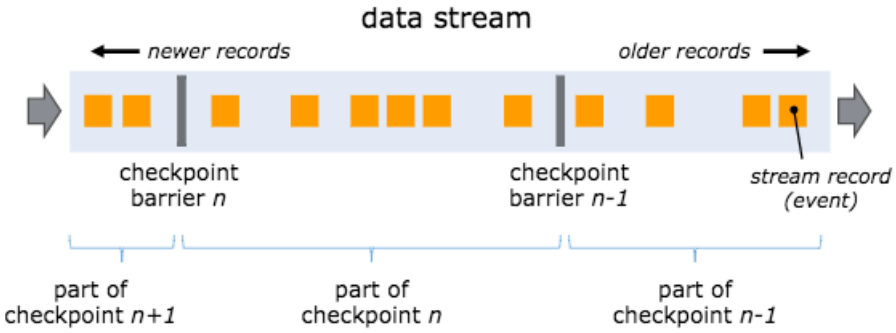

Stream barriers are injected into the dataflow at the stream sources. The point where the barrier for snapshot n is injected is the position in the source stream, up to which the snapshot covers the data. This point is indicated by Sn. The barriers then flow downstream. When an intermediate operator 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.

Barrier mechanism

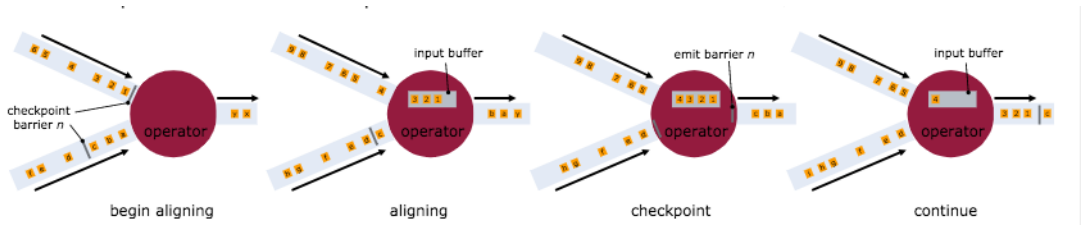

Checkpoint parameters

- **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.

11.4.2.4. Processing delay

Top 15 source subtasks with the longest processing delay

This metric describes the processing delays of each parallelism of a source.

11.4.2.5. 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.

11.4.2.6. 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.

11.4.2.7. 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.

11.4.2.8. 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.

11.4.2.9. 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.

11.4.3. 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.

11.4.4. 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. Parameter description describes the parameters for the created checkpoints.

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. Parameter description describes the parameters for the latest checkpoint.

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.

11.4.5. 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.

11.4.6. 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 TaskManager 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.

11.4.7. 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 Development 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.

11.4.8. 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. Job properties describes the basic job properties that are displayed on this tab page.

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.

11.4.9. Performance optimization by using automatic configuration

To improve user experience, Realtime Compute allows you to use automatic configuration to optimize job performance.

> **Note** Automatic configuration applies to Blink 1.0 and Blink 2.0.

Background and scope

If all the operators and both the upstream and downstream storage systems of your Realtime Compute job meet the performance requirements and remain stable, automatic configuration can help you properly adjust job configurations, such as operator resources and parallelism. It also helps optimize your job throughout the entire process to resolve performance issues such as low throughput or upstream and downstream backpressure.

In the following scenarios, you can use this feature to optimize job performance but cannot eliminate job performance bottlenecks. To eliminate the performance bottlenecks, manually configure the resources or contact the Realtime Compute support team.

- Performance issues exist in the upstream or downstream storage systems of a Realtime Compute job.
  - Performance issues in the data source, such as insufficient DataHub partitions or Message Queue (MQ) throughput. In this case, you must increase the partitions of the relevant source table.
  - Performance issues in the data sink, such as a deadlock in ApsaraDB for RDS.
- Performance issues of user-defined extensions (UDXs) such as the UDFs, UDAFs, and UDTFs in your Realtime Compute job.

Operations

- New jobs
  i. Publish a job.

a. After you complete SQL development and syntax check on the **Development** page, click **Publish**. The **Publish New Version** dialog box appears.

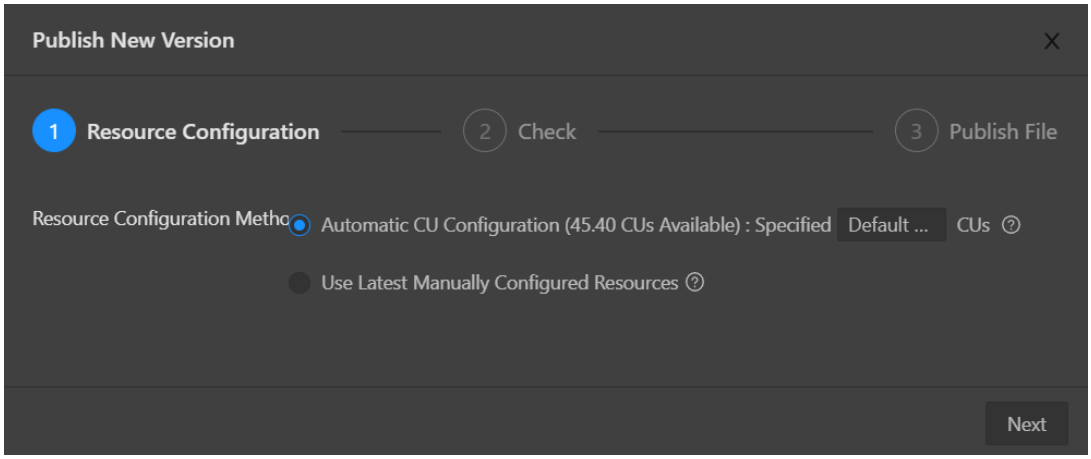

b. Specify **Resource Configuration Method**.

- **Automatic CU Configuration (2.25 CUs Available)**: If you select this option, you can specify the number of compute units (CUs). The automatic configuration algorithm generates an optimized resource configuration and assigns a value for the number of CUs based on the default configuration. If you use automatic CU configuration for the first time, the default number of CUs is used. This algorithm generates an initial configuration based on empirical data when you use automatic CU configuration for the first time. We recommend that you select Automatic CU Configuration (2.25 CUs Available) if your job has been running for 5 to 10 minutes and its metrics, such as source RPS, remain stable for 2 to 3 minutes. You can obtain the optimal configuration after you repeat the optimization process for three to five times.
- **Use Latest Manually Configured Resources**: The latest saved resource configuration is used. If the latest resource configuration is generated based on automatic CU configuration, the latest resource configuration is used. If the latest resource configuration is obtained based on the manual configuration, the manual configuration is used.

ii. Use the default configuration to start the job.

a. Use the default configuration to start the job, as shown in the following figure.

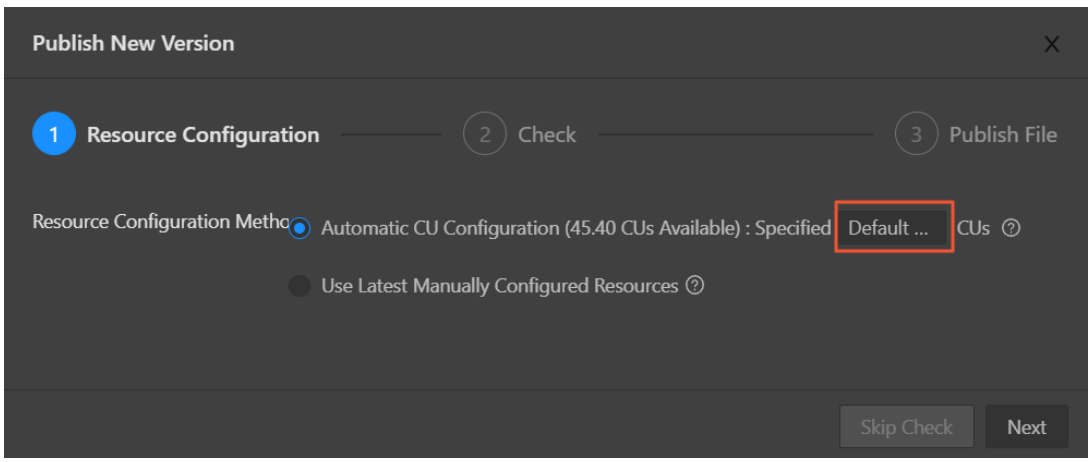

b. On the Administration page, find the job and click **Start** in the Actions column to start the job.

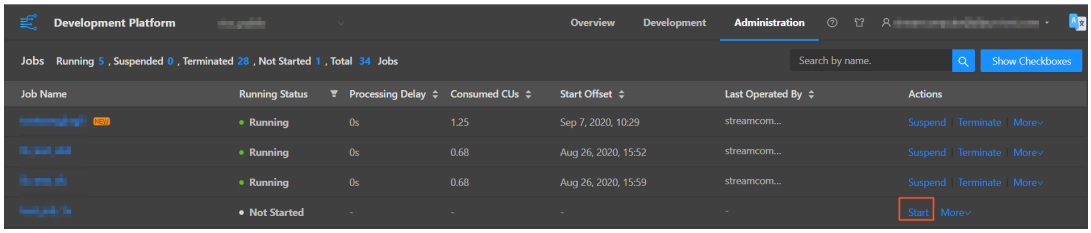

Assume that the default number of CUs generated the first time is 71.

> **Note** Make sure that your job runs longer than 10 minutes and its metrics such as source RPS remain stable for 2 to 3 minutes before you select Automatic CU Configuration (2.25 CUs Available) for Resource Configuration Method.

iii. Use the automatic CU configuration to start a job.

a. Resource performance optimization

If you select Automatic CU Configuration (2.25 CUs Available) for Resource Configuration Method and specify 40 CUs to start your job, you can change the number of CUs based on your job to optimize resource performance.

- Determine the minimum number of CUs.

  We recommend that you set the number of CUs to a value that is greater than or equal to 50% of the default value. The number of CUs cannot be less than 1. Assume that the default number of CUs for automatic CU configuration is 71. The recommended minimum number of CUs is 36, which is calculated by using the following formula: 71 CUs × 50% = 35.5 CUs.

- Increase the number of CUs.

  If the throughput of your Realtime Compute job does not meet your requirements, increase the number of CUs. We recommend that you increase the number of CUs by more than 30% of the current value. For example, if the number of CUs that you specified last time is 10 CUs, you can increase the number to 13.

- Repeat the optimization process.

  If the first optimization attempt does not meet your requirements, repeat the process until you obtain the desired results. You can change the number of CUs based on your job status after each optimization attempt.

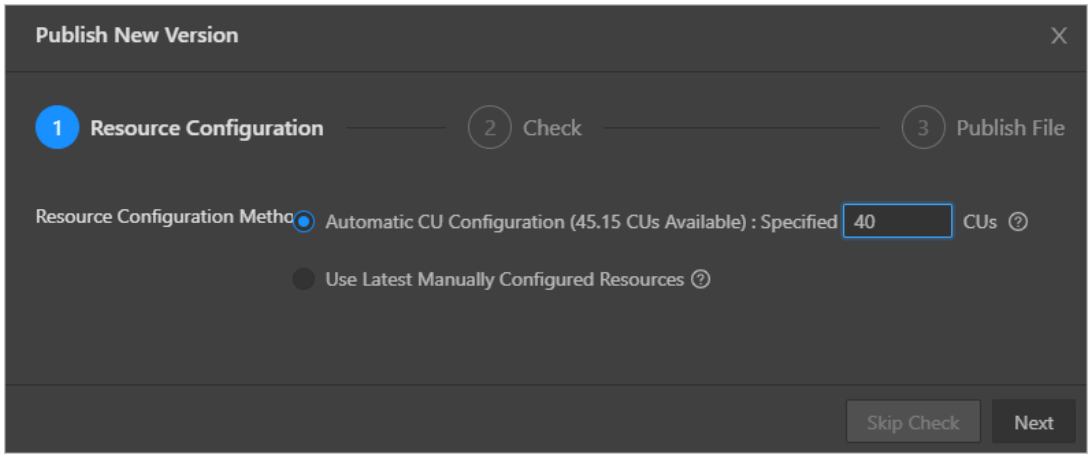

b. View the result of optimization. The following figure shows an example.

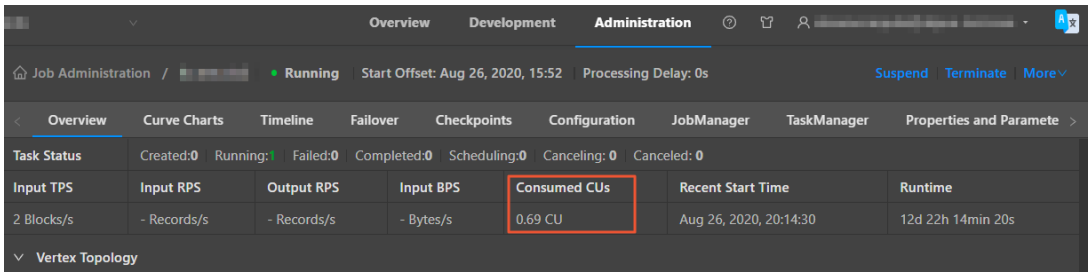

> **Note** Do not select **Use Latest Manually Configured Resources** for a new job. Otherwise, an error is returned.

- Existing jobs

- The following figure shows the optimization process of automatic configuration.

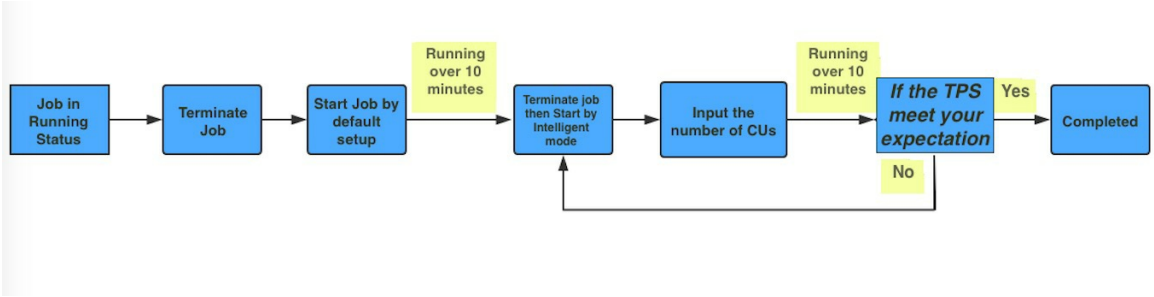

> **Note**
>
> - Before you use automatic configuration for a job that is in the running state, check whether stateful operations are involved. This is because the saved state data of a job may be cleared during the optimization process of automatic configuration.
> - If you make changes to a job, for example, modifying SQL statements or changing the Realtime Compute version, automatic configuration may fail. These changes may lead to topology changes, which results in certain issues. For example, curve charts may not be able to display the latest data, or the state data may not able to be used for fault tolerance. In this case, resource configurations cannot be optimized based on the job running history and therefore an error is returned when you perform automatic configuration. To rectify the fault, you must treat the changed job as a new job and repeat the previous operations.

- Procedure
  a. Suspend the job.

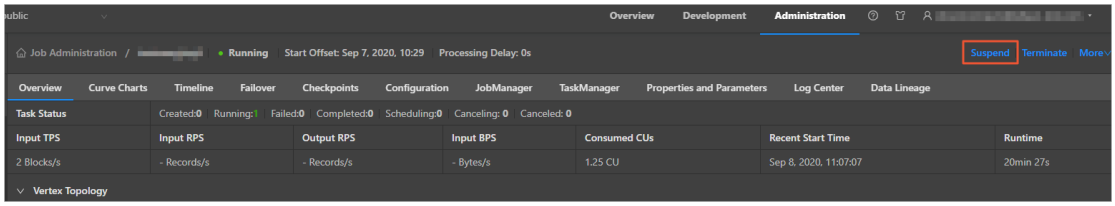

  b. Repeat the steps performed for new jobs and resume the job with the latest configuration.

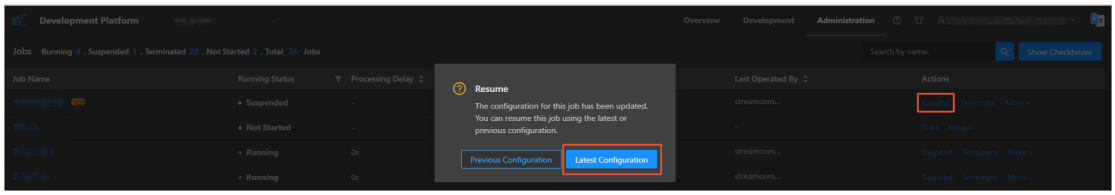

FAQ

The optimization result of automatic configuration may not be accurate in the following scenarios:

- If the job runs only for a short period of time, the data collected during data sampling is insufficient. We recommend that you increase the running duration of the job and make sure that the curves of job metrics such as source RPS remain stable for at least 2 to 3 minutes.
- A job fails. We recommend that you check and fix the failure.
- Only a small amount of data is available for a job. We recommend that you retrieve more historical data.

- The effect of automatic configuration is affected by multiple factors. Therefore, the latest configuration obtained by using automatic configuration may not be optimal. If the effect of automatic configuration does not meet your requirements, you can manually configure the resources. For more information, see Optimize performance by manual configuration.

Recommendations

- To help automatic configuration accurately collect the runtime metric information of a job, make sure that the job runs stably for more than 10 minutes before you apply automatic configuration to the job.
- Job performance can be improved after you use automatic configuration for three to five times.
- When you use automatic configuration, you can specify the start offset to retrieve historical data or even accumulate large amounts of data for a job to create backpressure to accelerate the optimization effect.

Method used to determine the effectiveness of automatic configuration

Automatic configuration of Realtime Compute is enabled based on a JSON configuration file. After you use automatic configuration to optimize a job, you can view the JSON configuration file to check whether the feature is running as expected.

- You can view the JSON configuration file by using one of the following methods:
  i. View the file on the job edit page, as shown in the following figure.

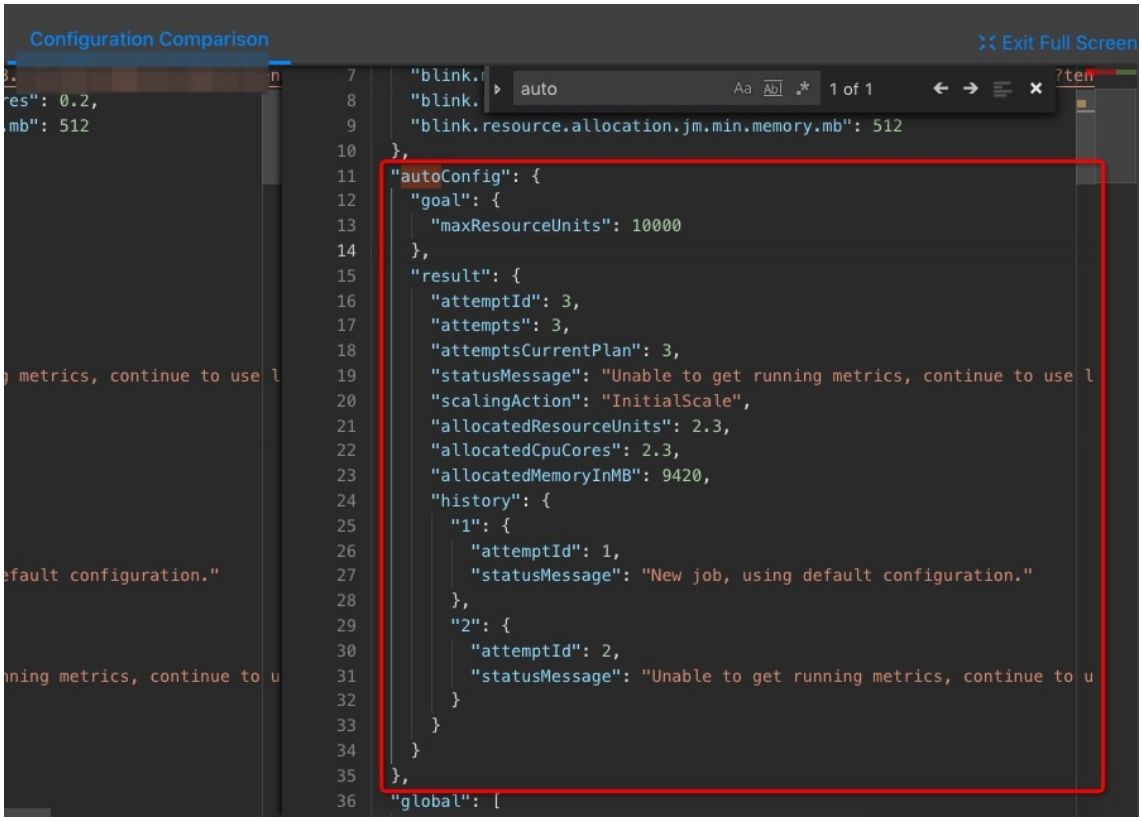

  ii. View the file on the Job Administration page, as shown in the following figure.

- JSON configuration description

```
"autoconfig" : {
  "goal": { // The goal of automatic configuration.
    "maxResourceUnits": 10000.0, // The maximum number of CUs for a Blink job. This value cannot be cha
nged. Therefore, you can ignore this item when you check whether the feature is running as expected.
    "targetResoureUnits": 20.0 // The number of CUs that you specified. The specified number of CUs is 20
.
  },
  "result" : { // The result of automatic configuration. We recommend that you pay attention to this item.
   "scalingAction" : "ScaleToTargetResource", // The action of automatic configuration. *
   "allocatedResourceUnits" : 18.5, // The total resources allocated by automatic configuration.
   "allocatedCpuCores" : 18.5,    // The total CPU cores allocated by automatic configuration.
   "allocatedMemoryInMB" : 40960   // The total memory size allocated by automatic configuration.
   "messages" : "xxxx" // We recommend that you pay attention to these messages. *
  }
}
```

  - scalingAction: If the value of this parameter is `InitialScale`, this is the first time that you use automatic configuration. If the value of this parameter is `ScaleToTargetResource`, this is not the first time that you use automatic configuration.

- If no message is displayed, automatic configuration runs properly. If some messages are displayed, you must analyze these messages. Messages are categorized into the following two types:
  - Warning: This type of message indicates that automatic configuration runs properly but you must pay attention to potential issues, such as insufficient partitions in a source table.
  - Error or exception: This type of message indicates that automatic configuration failed. The following error message is usually displayed: `Previous job statistics and configuration will be used`. The automatic configuration for a job fails in the following two scenarios:
    - The job or Blink version is modified before you use automatic configuration. In this case, the previous running information cannot be used for automatic configuration.
    - An error message that contains "exception" is reported when you use automatic configuration. In this case, you must analyze the error based on the job running information and logs. If you do not have enough information, submit a ticket.

Error messages

IllegalStateException

If the following error messages are displayed, the state data cannot be used for fault tolerance. To resolve this issue, terminate the 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, follow these steps to mitigate the negative impact of service interruption: Roll back the target job to an earlier version and specify the start offset to re-read the data during off-peak hours. To roll back the target job, click **Versions** on the right side of the **Development** page. On the page that appears, move the pointer over More in the Actions column, click Compare, and then click Roll Back to Version.

```
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.initializeOperators(StreamTask.java:870)
  at org.apache.flink.streaming.runtime.tasks.StreamTask.initializeState(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.SerializationException: Cannot serialize/deserialize the object.
  at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSecondaryState.deserializeStateEntry(AbstractRocksDBRawSecondaryState.java:167)
  at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRestoreOperation.restoreRawStateData(RocksDBIncrementalRestoreOperation.java:425)
  at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRestoreOperation.restore(RocksDBIncrementalRestoreOperation.java:119)
  at com.alibaba.blink.contrib.streaming.state.RocksDBKeyedStateBackend.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.AbstractRocksDBRawSecondaryState.deserializeStateEntry(AbstractRocksDBRawSecondaryState.java:162)
  ... 11 more
```

11.4.10. Improve performance by manual configuration

11.4.10.1. 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.

11.4.10.2. 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.

11.4.10.3. 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
   ```

11.4.10.4. 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.

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.

Object	Parameter	Description	Value
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.

11.4.10.5. 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.

11.4.10.6. 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.

11.4.11. O&M of Apsara Bigdata Manager

11.4.11.1. What is Apsara Bigdata Manager?

Apsara Bigdata Manager (ABM) provides 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 the ABM console.

Onsite Apsara Stack engineers can use ABM to easily manage big data products. They can view resource usage, check and handle alerts, and modify configurations.

For more information about the logon methods and O&M operations of Realtime Compute for Apache Flink in the ABM console, see the following topics.

11.4.11.2. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

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

Procedure

1. Open your browser.

2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

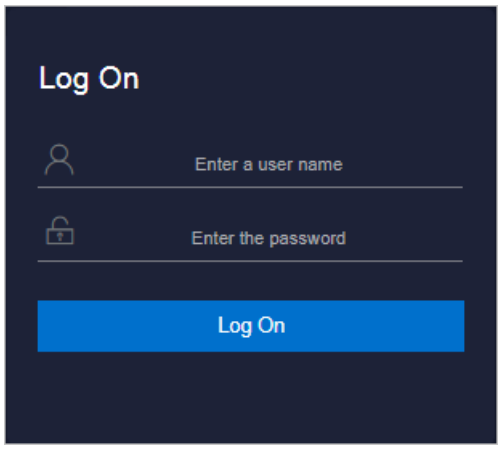

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.4.11.3. O&M overview

This topic describes the O&M features of Realtime Compute and how to access the Realtime Compute O&M page.

Modules

Realtime Compute O&M includes four modules including business O&M, service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Business O&M	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.
Service O&M	Druid	Displays the number of Druid master nodes and that of Druid worker nodes in Realtime Compute.
	Yarn	Displays information about the YARN queue APIs in Realtime Compute. Realtime Compute allocates cluster resources to YARN queue APIs. You can bind projects with these APIs to obtain the corresponding cluster resources.
Cluster O&M	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.
	Health Status	Displays the check results for a cluster. The check results are divided into Critical, Warning, Exception, and OK.
	Hosts	
Host O&M	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.
	Health Status	Displays the check results for a host. The check results are divided into Critical, Warning, Exception, and OK.

Entry

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

11.4.11.4. Business O&M

11.4.11.4.1. Projects

Apsara Bigdata Manager (ABM) allows you to view information about the projects in Realtime Compute, including the name, engine, queue, used CUs, total CUs, CU usage percentage, and number of jobs.

On the **Business** page, click **Projects** in the left-side navigation pane. The **Projects** page for Realtime Compute appears.

The **Projects** page displays the information of projects in Realtime Compute, including the name, engine, queue, used CUs, total CUs, CU usage percentage, and number of jobs.

11.4.11.4.2. Jobs

Apsara Bigdata Manager (ABM) allows you to view the information about jobs in Realtime Compute, including the name, user, project, transactions per second (TPS) in the inbound direction, latency, requested CUs, status, and start time. You can diagnose and analyze jobs to troubleshoot issues.

Jobs

On the **Business** page, click **Jobs** in the left-side navigation pane. The **Jobs** page for Realtime Compute appears.

The **Jobs** page displays the information about jobs in Realtime Compute, including the name, user, project, TPS in the inbound direction, latency, requested CUs, status, and start time.

Job analysis

Job diagnosis has two steps, namely **Failover** and **Blink Metric**. In the **Blink Metric** step, the system checks the latency, garbage collection (GC) time, TPS, the number of times of GC, data skew, and back pressure nodes of a job.

1. On the **Jobs** page, click a job name. Alternatively, click the **Job Analysis** tab at the top. The **Job Analysis** page appears.
2. Select the job to be diagnosed and analyzed from the **Select Job** drop-down list.
3. 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 complete.
4. After the diagnosis is completed, click **View Log** to view the log details if the diagnosis result appears in red.

   The metrics for job diagnosis are described as follows:

   - **Failover**
     - 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 Time**: checks whether the GC time of CMS exceeds 100 ms. This metric applies to all containers.
     - **Job TPS**: checks whether the TPS of a subtask is 0.
     - **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.

11.4.11.4.3. Queues

Apsara Bigdata Manager (ABM) allows you to view the information about the queues in Realtime Compute, including the name, status, minimum resources guaranteed, minimum resources guaranteed, maximum resources available, maximum resources available, and number of jobs.

On the **Business** page, click **Queues** in the left-side navigation pane. The **Queues** page for Realtime Compute appears on the right.

The **Queues** page displays the information about queues in Realtime Compute, including the name, status, minimum resources guaranteed, minimum resources guaranteed, maximum resources available, maximum resources available, and number of jobs.

11.4.11.5. Service O&M

11.4.11.5.1. 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.

11.4.11.5.2. 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**, **Exception**, and **OK**. They are displayed in different colors. Among them, the **Critical**, **Warning**, and **Exception** results are alerts. You need to pay special attention to them, especially the **Critical** and **Warning** results.

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.

11.4.11.5.3. 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.

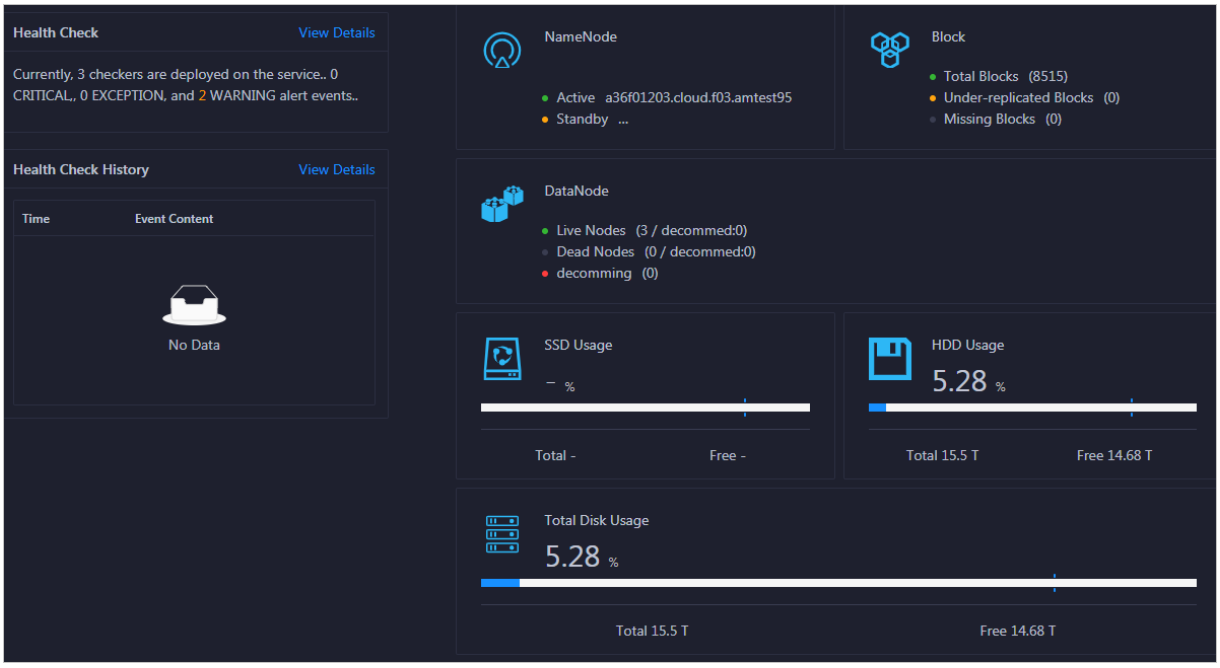

The **Overview** page displays the health check result, health check history, the information of NameNode, blocks, and DataNode, solid-state disk (SSD) usage, hard disk drive (HDD) usage, and total disk usage.

Click **View Details** in the **Health Check** section and **Health Check History** section to go to the **Health Status** page for the HDFS service. 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. The **Health Status** page for the HDFS service appears.

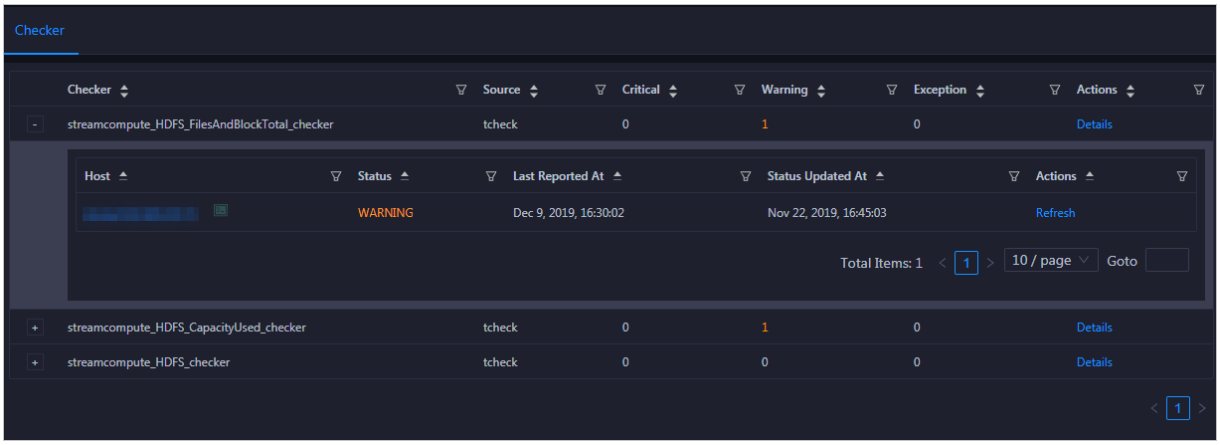

On the **Health Status** page, you can view all checkers of the HDFS service and the check results for all hosts in the cluster. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, the **Critical**, **Warning**, and **Exception** results are alerts. You need to pay attention to them, especially the **Critical** and **Warning** results.

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.

11.4.11.6. Cluster O&M

11.4.11.6.1. 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.

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 **Overview** page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. On this page, you can also view the health check result and health check history of the cluster. To view information about a cluster, select a **region** in the left-side navigation pane, and then select the **cluster** in the region.

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 it.

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.

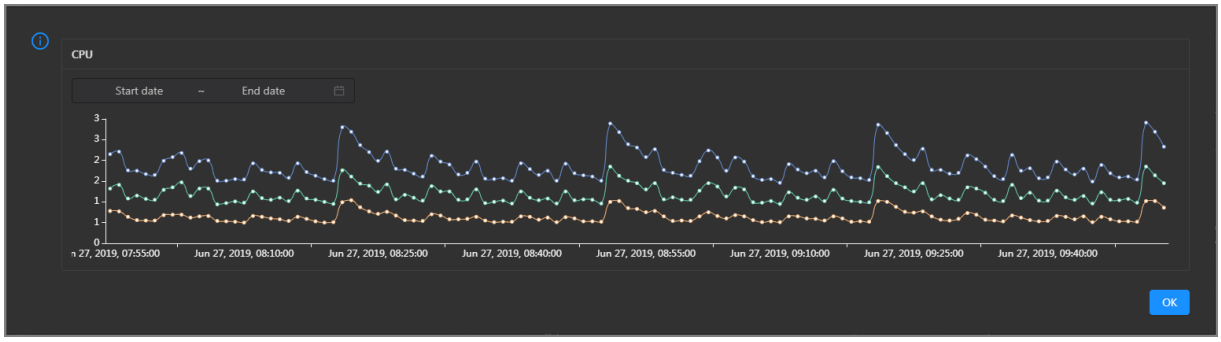

DISK

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 it.

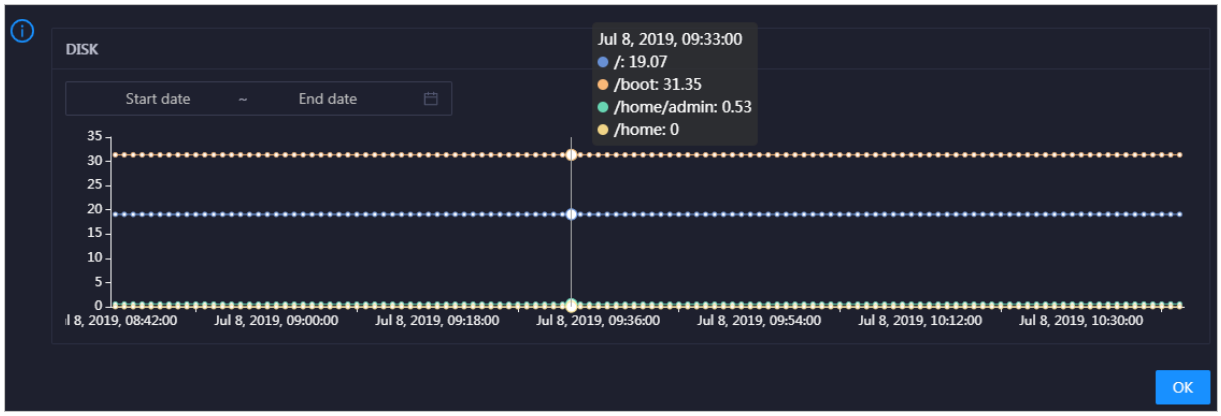

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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 it.

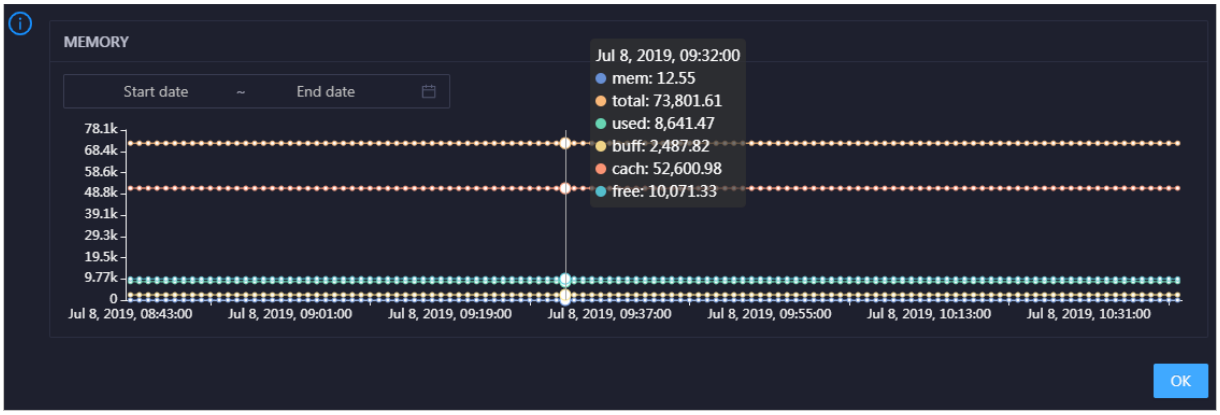

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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 it.

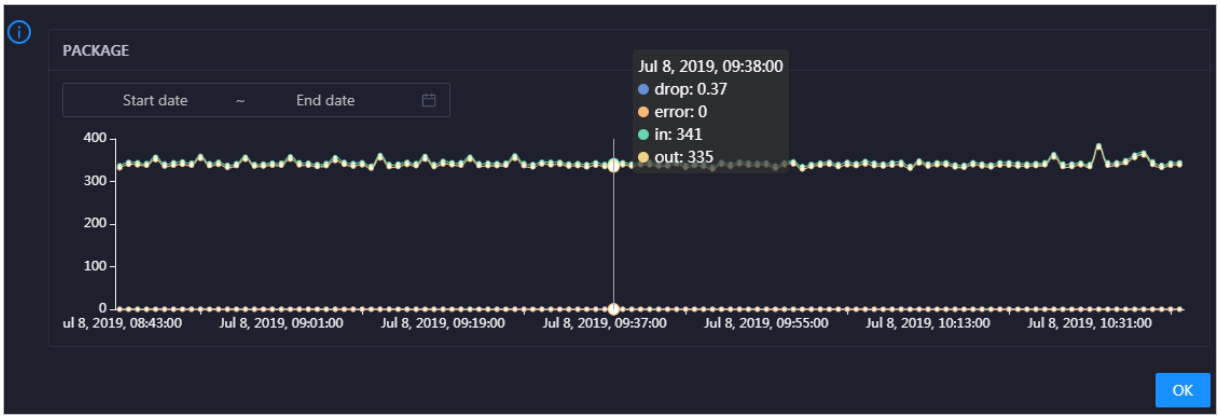

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

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 it.

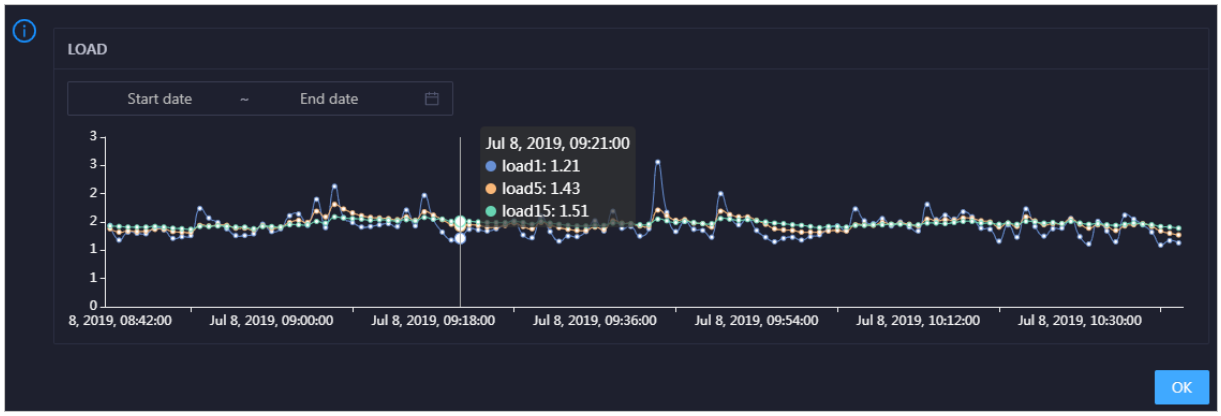

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line 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 cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

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

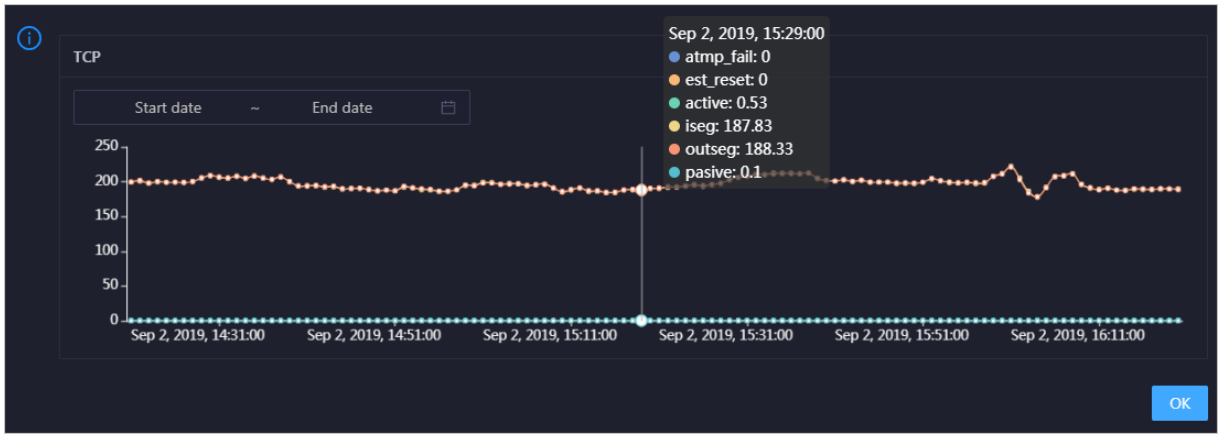

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

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

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

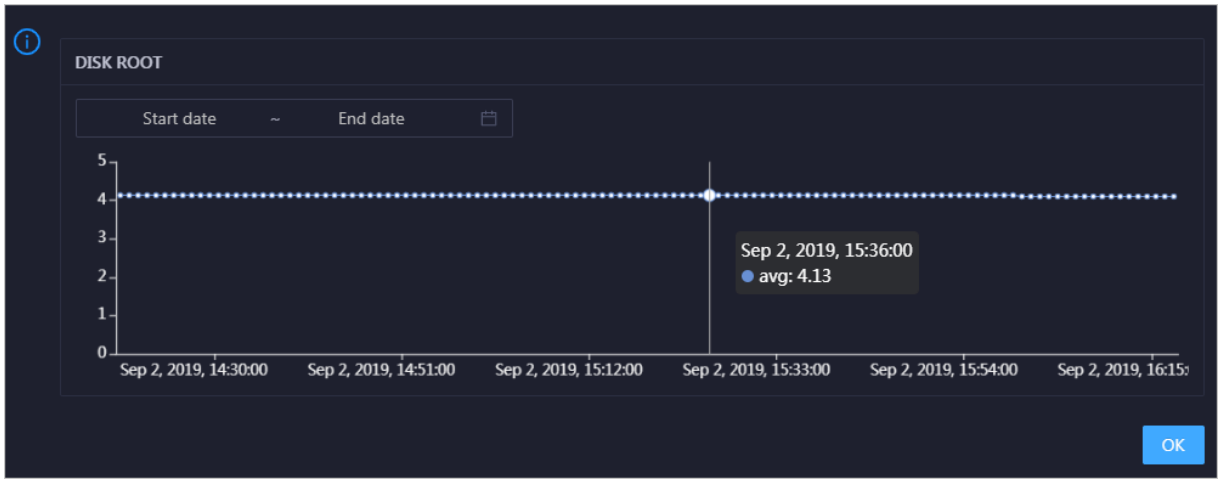

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

Health Check

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

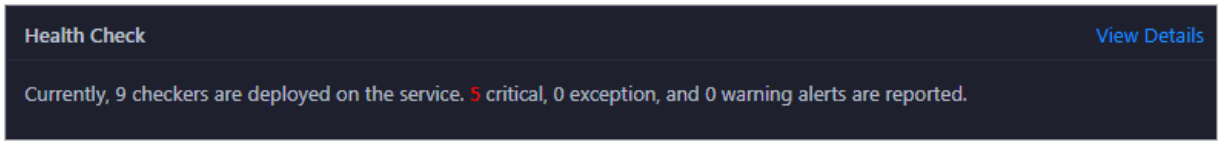

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Cluster health.

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.

11.4.11.6.2. 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 the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, **Critical**, **Warning**, and **Exception** results are alerts. You need to pay special attention to them, especially the **Critical** and **Warning** results.

View checker details

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

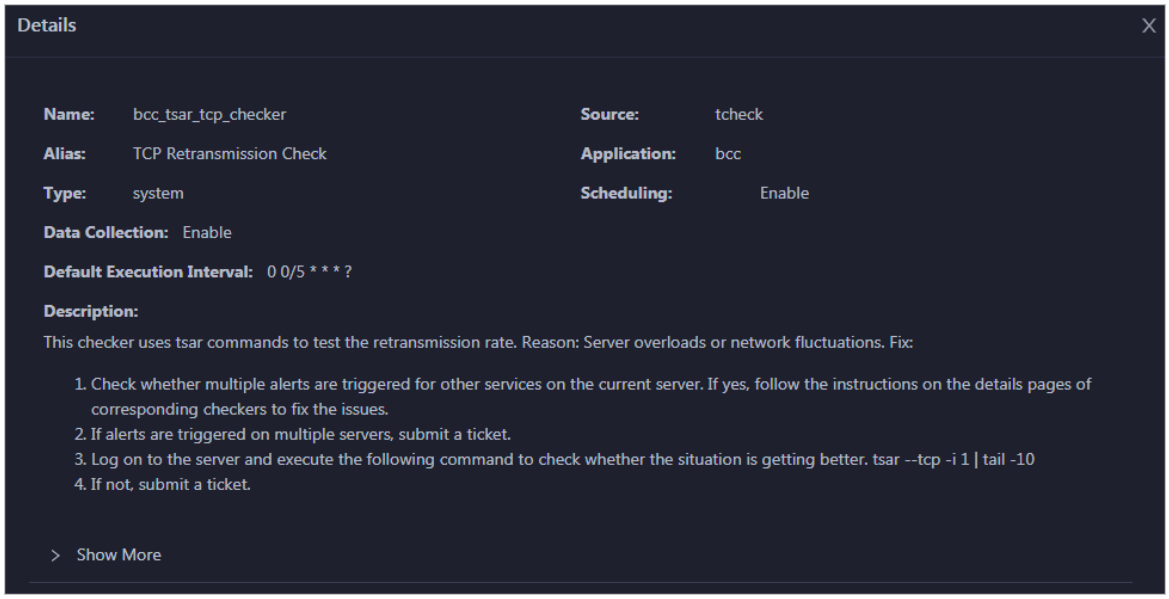

   The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

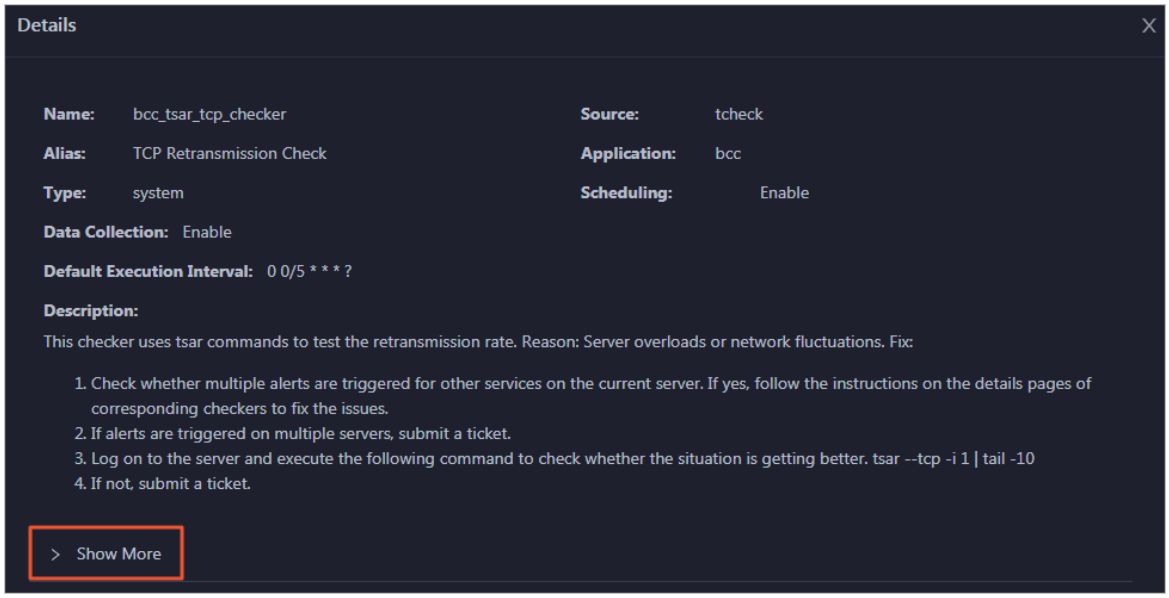

   You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

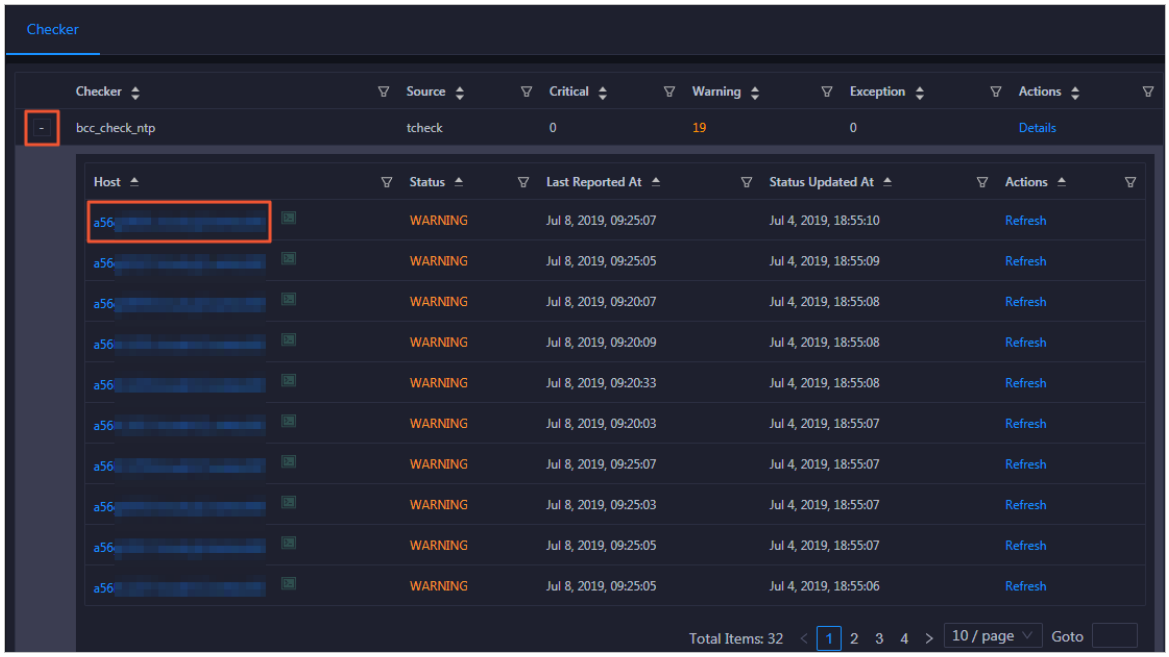

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

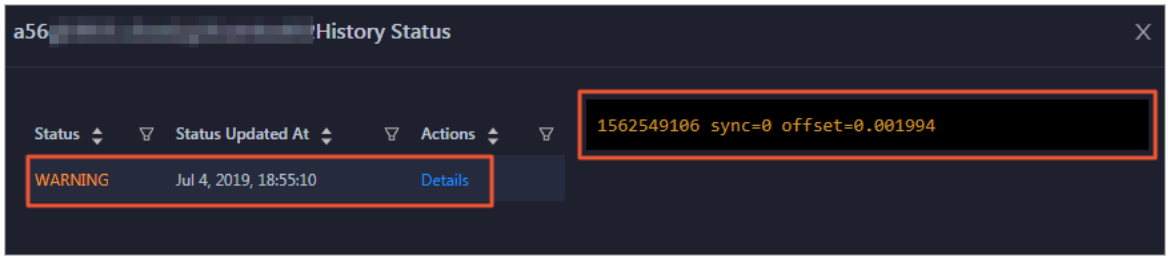

Clear alerts

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

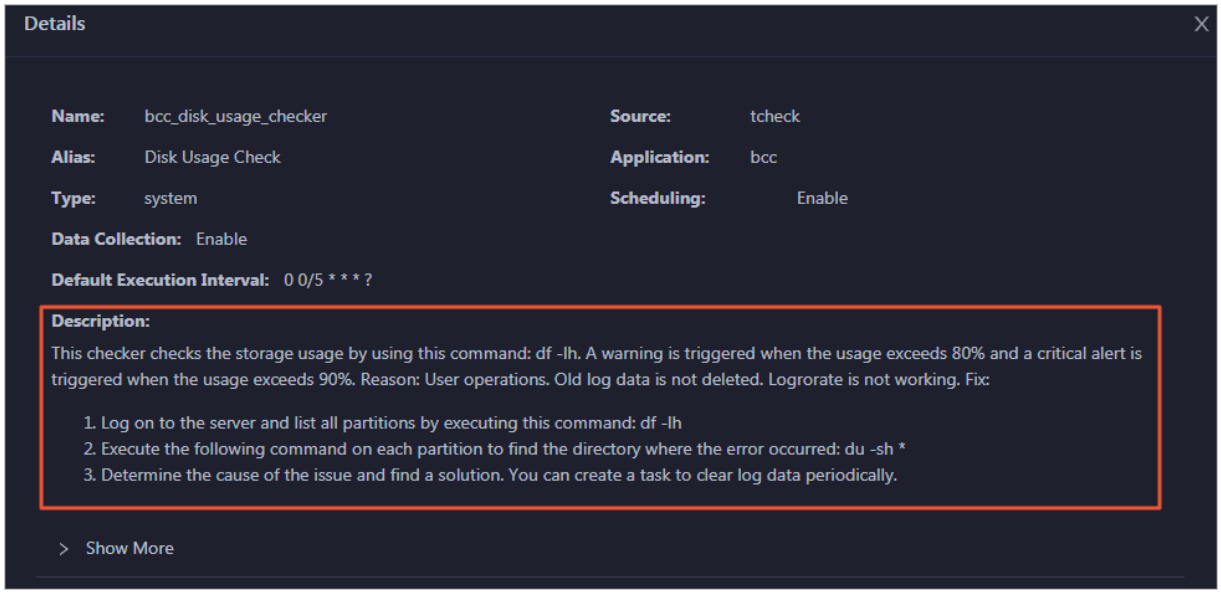

Log on to a host

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

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

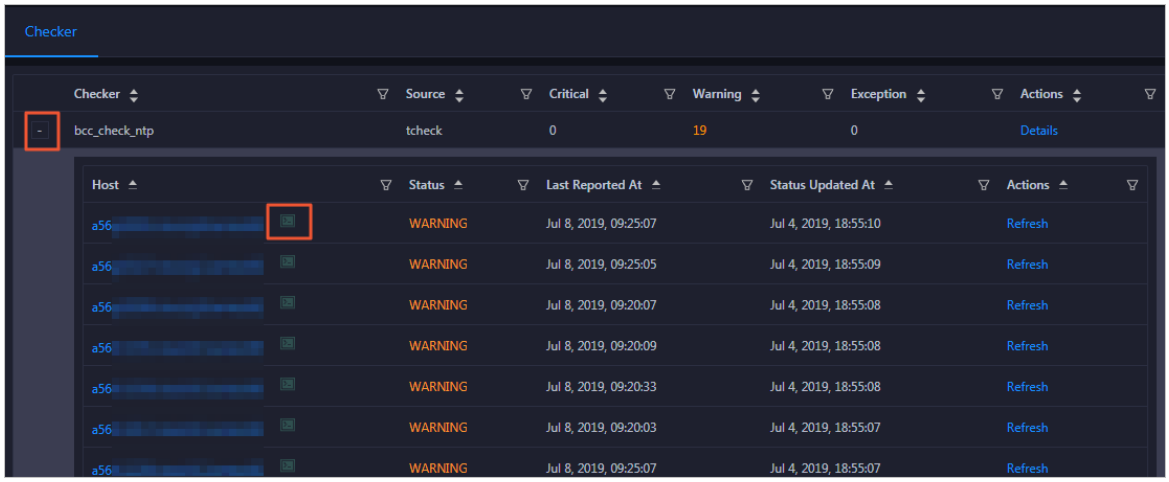

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

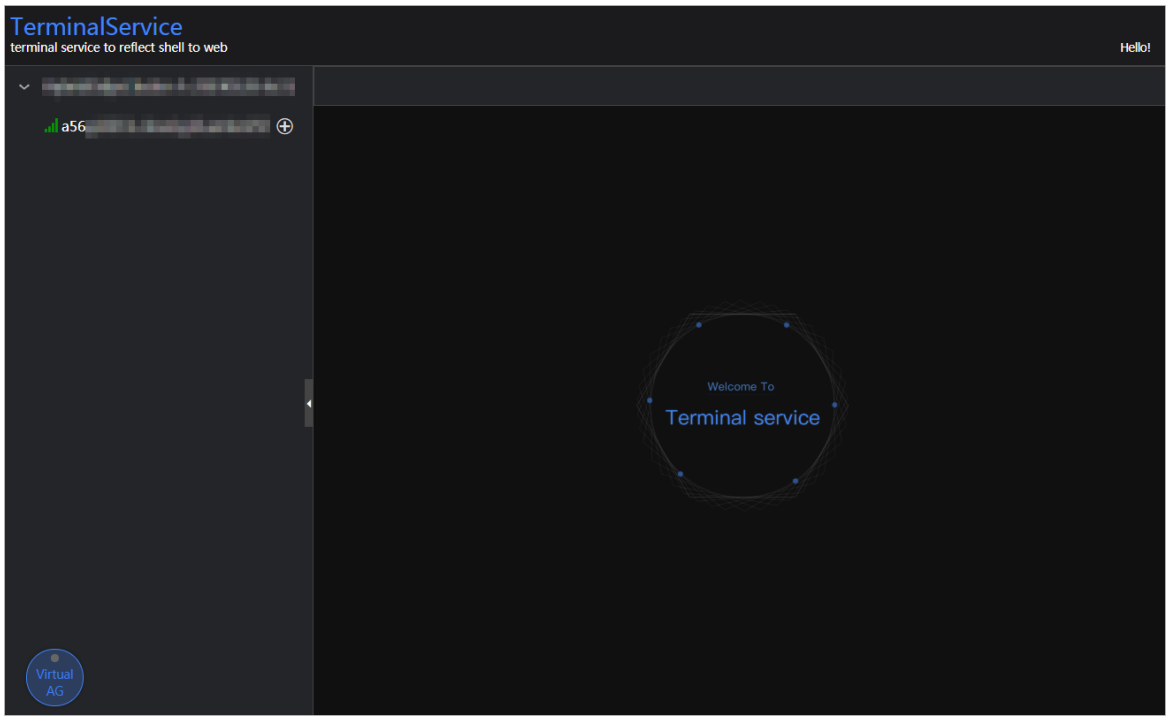

3. On the **TerminalService** page, click the hostname to log on to the host.

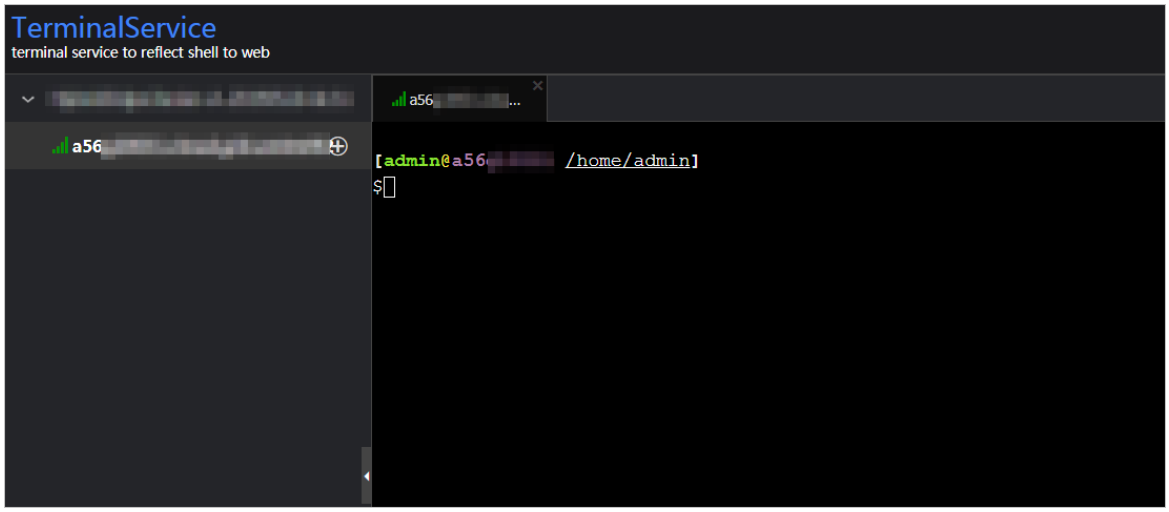

Run a checker again

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

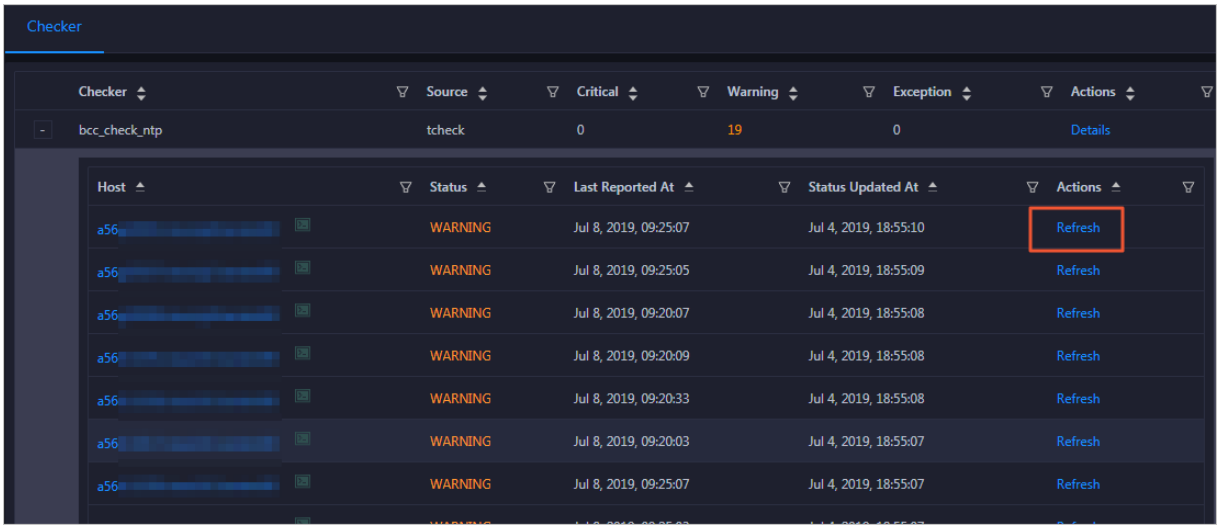

11.4.11.6.3. Hosts

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.

11.4.11.6.4. Cluster scale-out

Apsara Bigdata Manager (ABM) allows you to scale out a Realtime Compute for Apache Flink 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 for Apache Flink cluster. Currently, scale-out is only available for **worker** nodes in a Realtime Compute for Apache Flink cluster.

Prerequisites

- Your ABM account is granted the required permissions to perform O&M operations on Realtime Compute for Apache Flink.
- Hosts whose service type is **blink** are deployed in the default cluster of Apsara Infrastructure Management Framework.

Background information

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 a resource pool that can provide 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 that is to be added to a Realtime Compute for Apache Flink cluster

Before the scale-out operation, obtain the name of the available 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, 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 for Apache Flink cluster

You can add multiple hosts to a Realtime Compute for Apache Flink cluster at a time to scale out the cluster. To achieve this, you must specify an existing host as the template host. When you scale out the Realtime Compute for Apache Flink cluster, the hosts copy configurations from the template host so that they 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 **StreamCompute**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Business** page appears by default.
4. Click the **Clusters** tab. On the page that appears, click the **Hosts** tab, and select a host whose Role is **Worker** as the template host.
5. Choose **Actions > Scale out Cluster** in the upper-left corner. In the **Scale out Cluster** dialog

box, configure required 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 that you want to add to the Realtime Compute for Apache Flink 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.

6. Click **Run**. A message appears, indicating that the action has been submitted.
7. 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 scale-out operation is in progress, **SUCCESS** indicates that the scale-out operation is successful, and **FAILED** indicates that the scale-out operation fails.

Step 3: View the scale-out progress

If the status is **RUNNING**, click **Details** in the Details column to check the current step and progress of the scale-out operation.

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, scripts, and execution parameters to locate the failure cause.

11.4.11.6.5. 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 must have the required permissions to perform O&M operations on Realtime Compute.
- The current cluster has more than three **worker** nodes. 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.
- Before you scale in a cluster, check whether the resources of the cluster, including the disk, CPU, and memory, are still sufficient if the cluster is scaled in. For more information about how to check CPU usage 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.

Context

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 idle resource pool that provides resources for scaling out clusters for your business. 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

(Optional)

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.

   Click **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. (Optional)View the scale-in progress.

   If the status is **RUNNING**, click **Details** 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** to locate the failure cause.

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

11.4.11.7. Host O&M

11.4.11.7.1. 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.

Entry

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

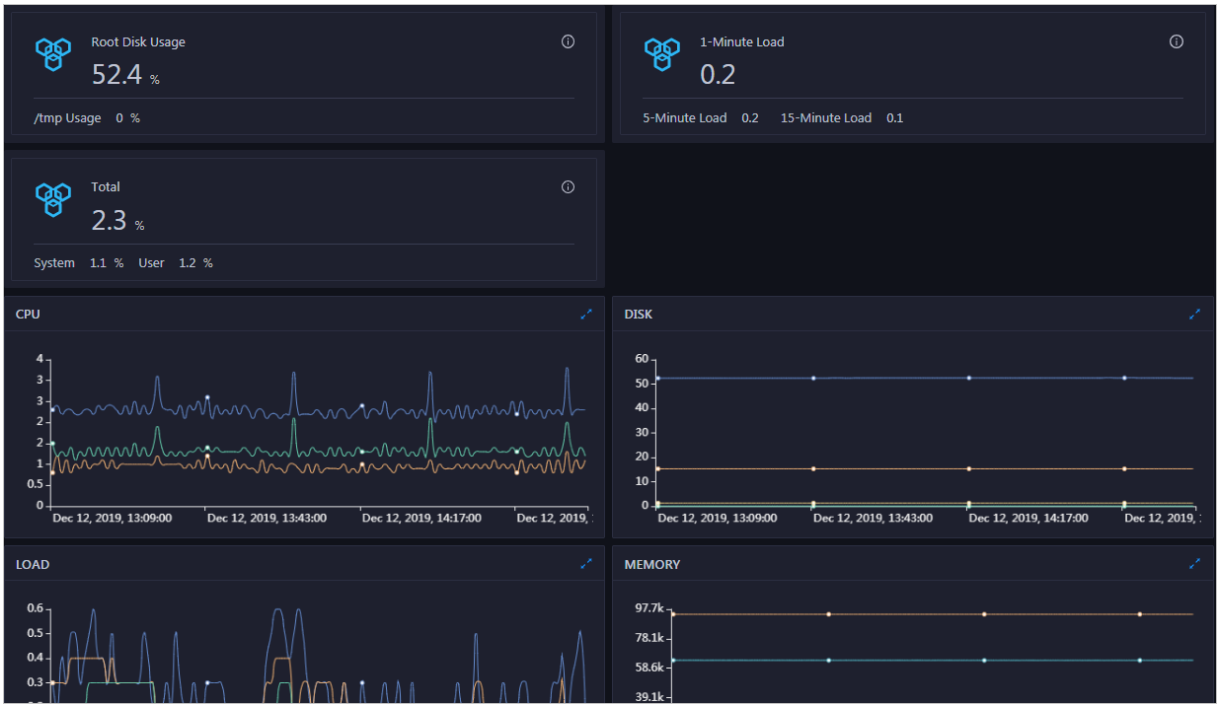

On the **Overview** page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

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.

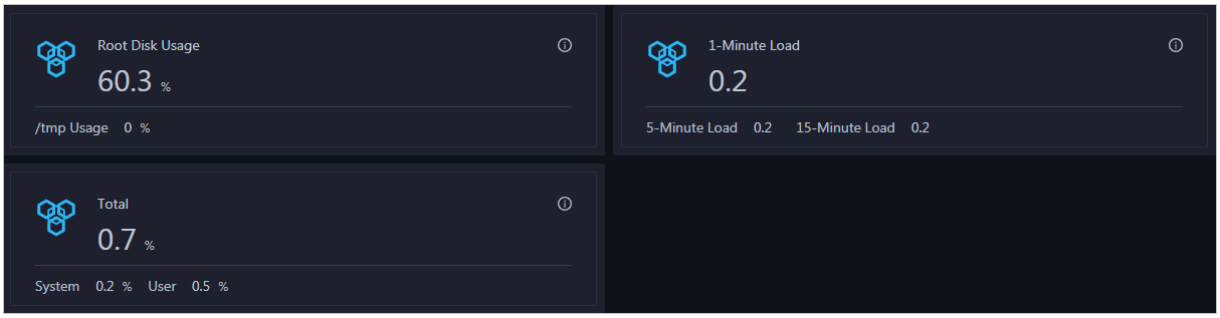

CPU

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.

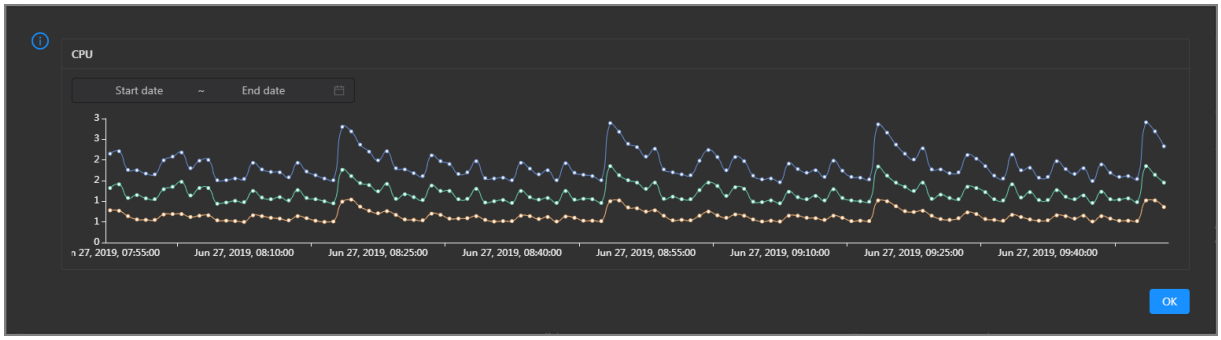

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

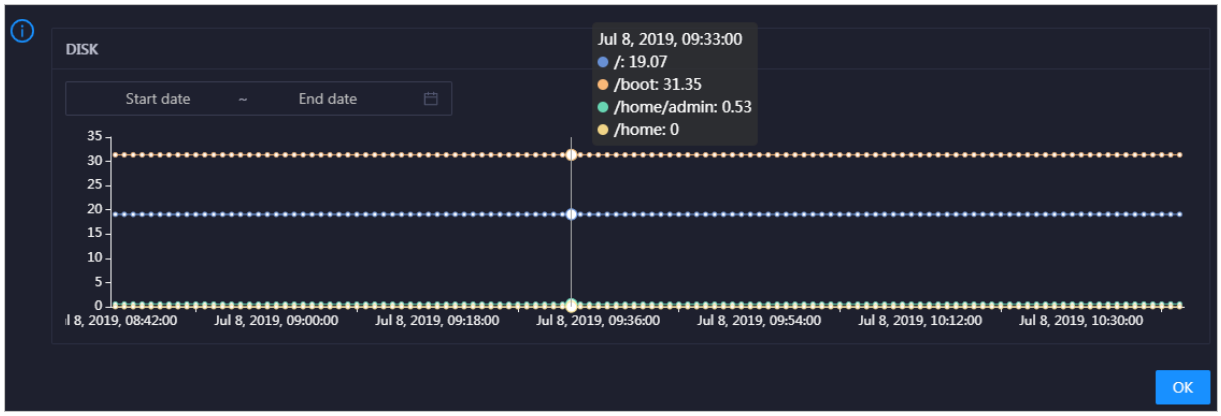

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

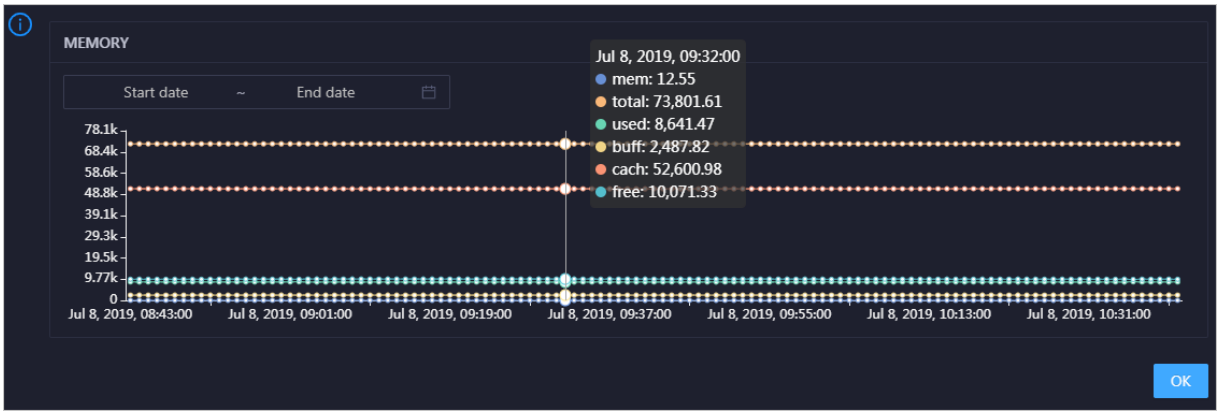

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

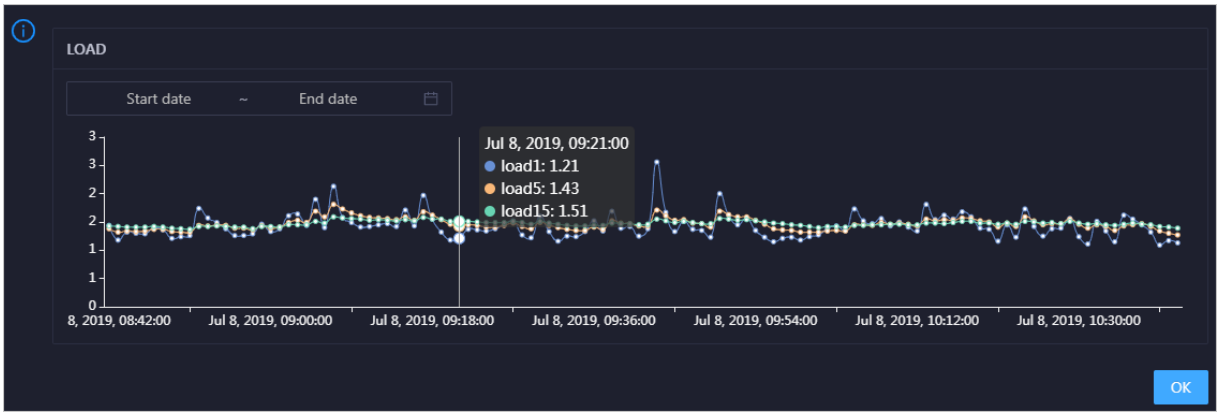

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

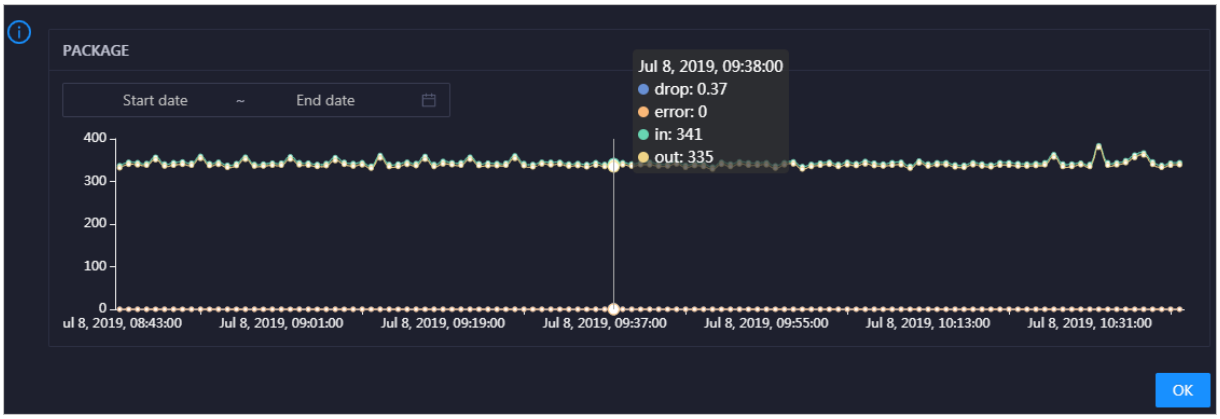

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line 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.

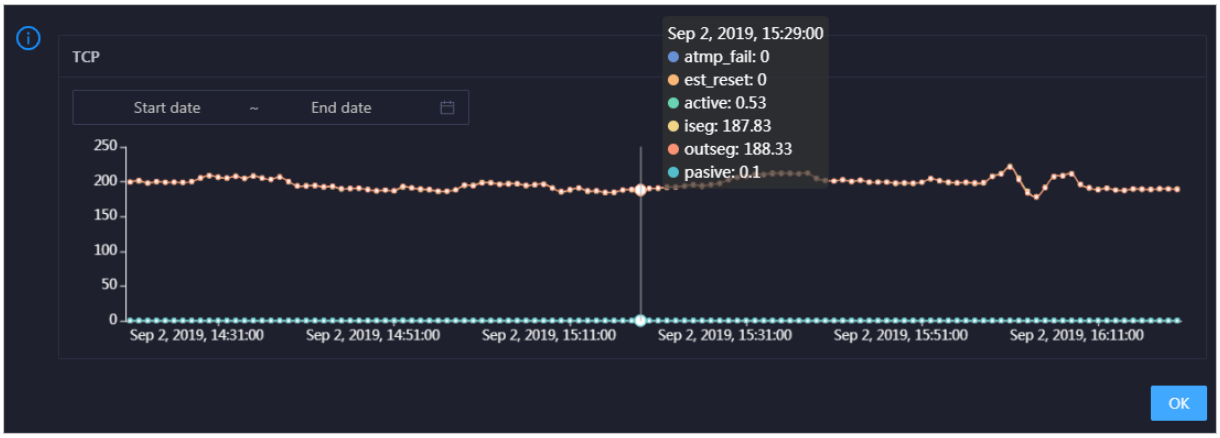

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section 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.

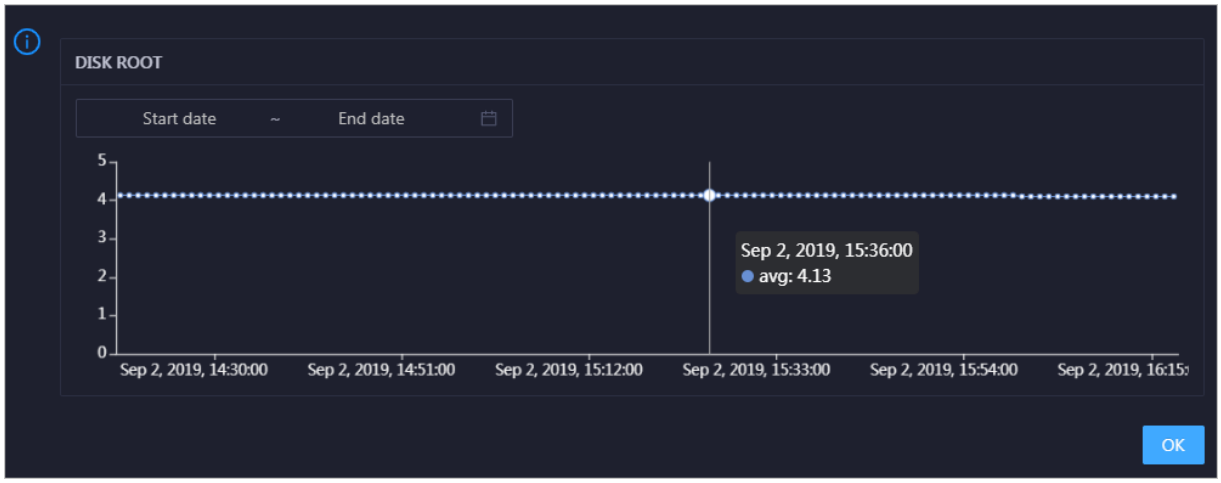

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

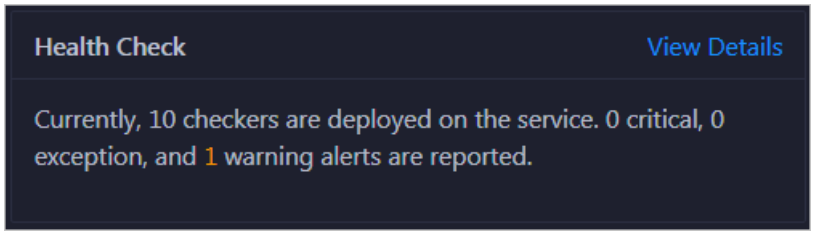

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

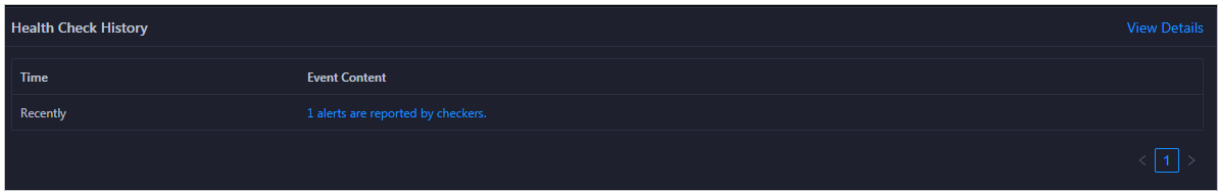

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

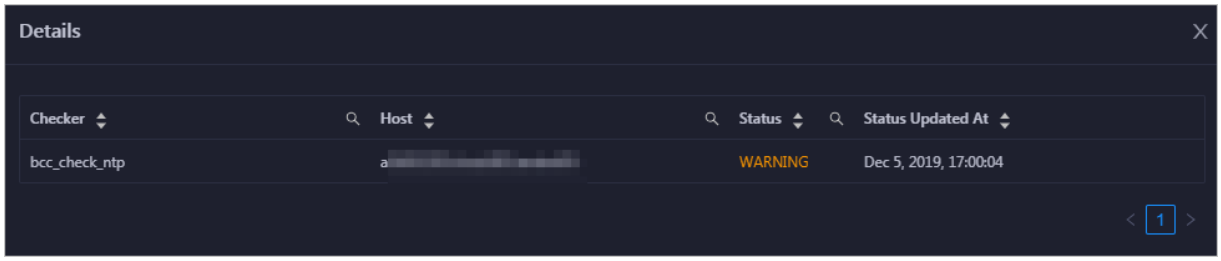

11.4.11.7.2. Host health

On the host health status page, you can view the checkers of all hosts, 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

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.

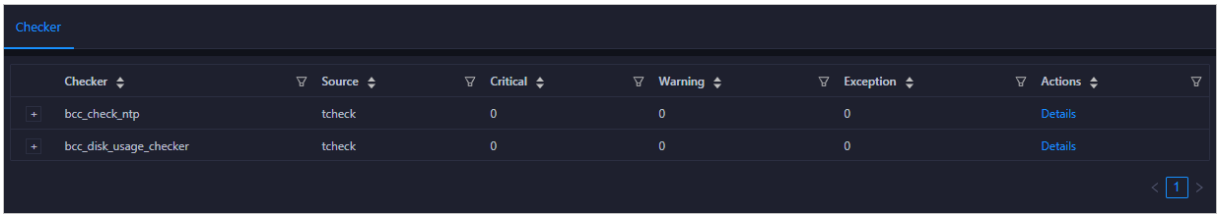

On the **Health Status** page, you can view all checkers of the host and the check results for the hosts in the host. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, **Critical**, **Warning**, and **Exception** results are alerts. You need to pay special attention to them, especially the **Critical** and **Warning** results.

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.

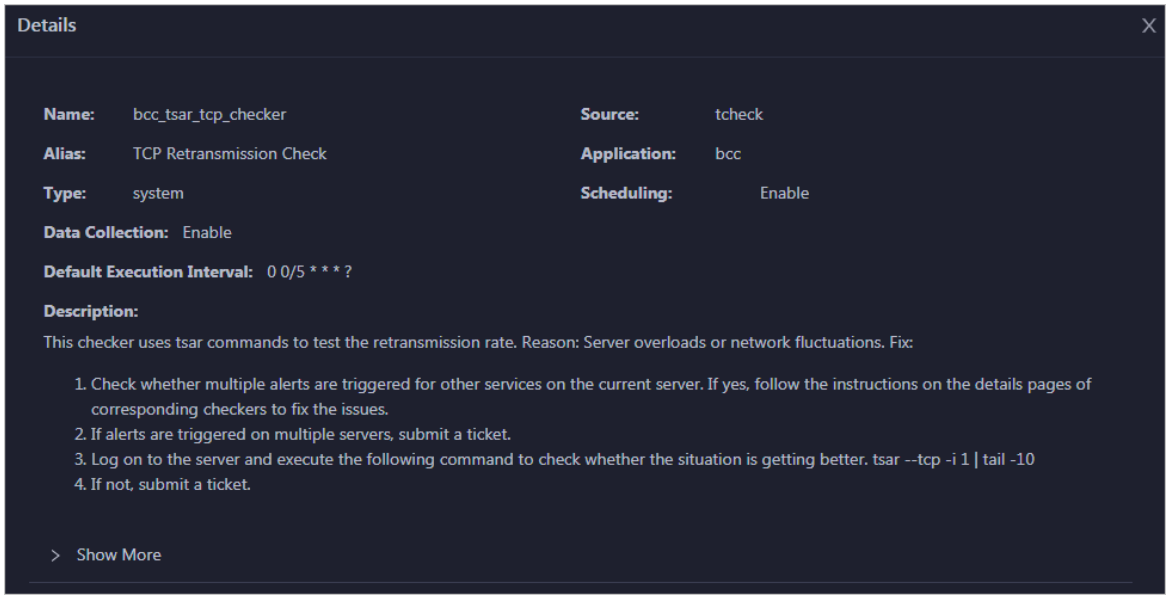

   The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

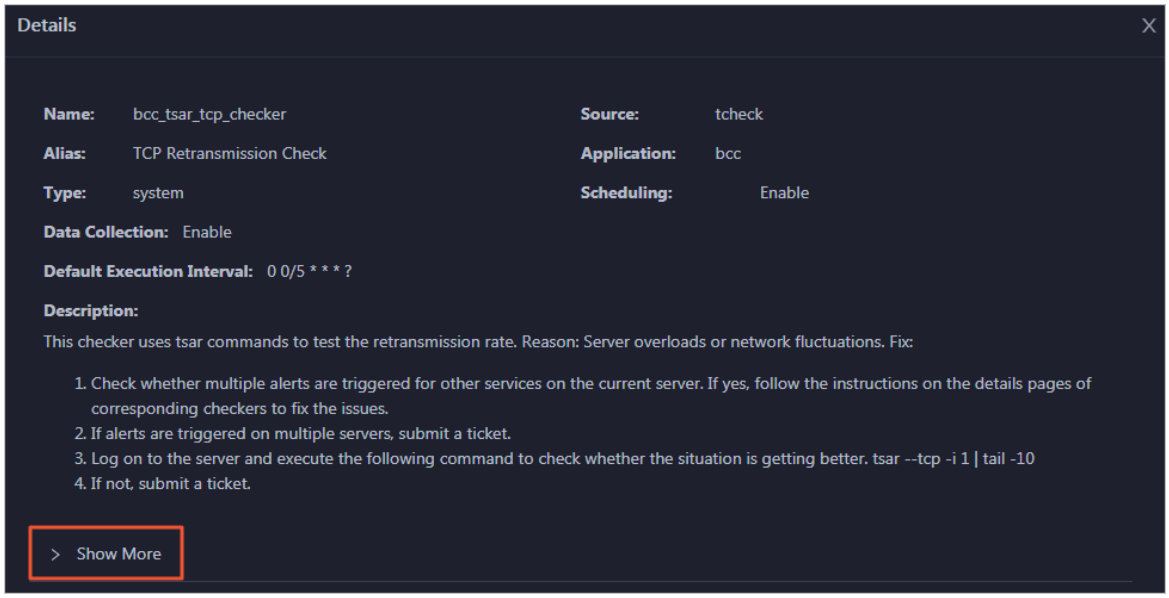

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

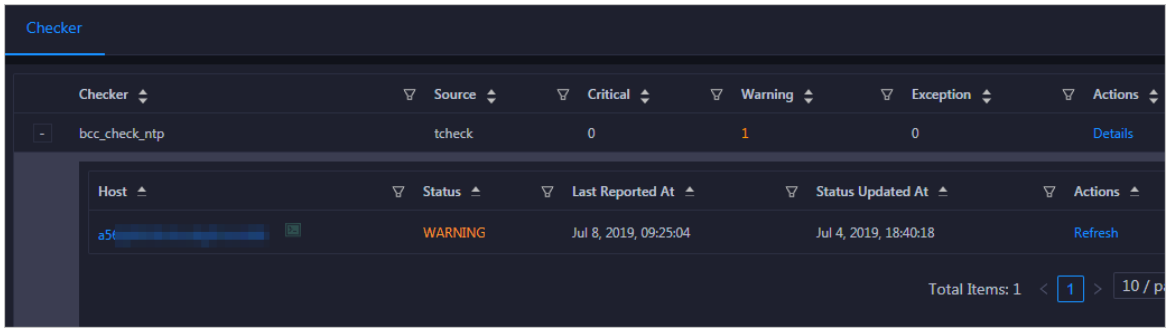

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

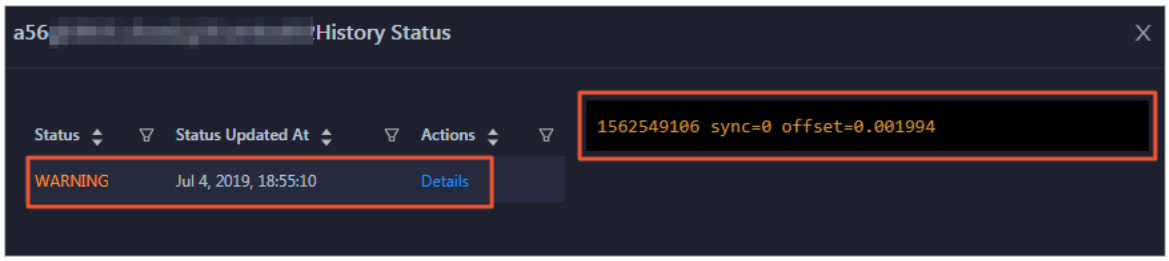

Clear alerts

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.

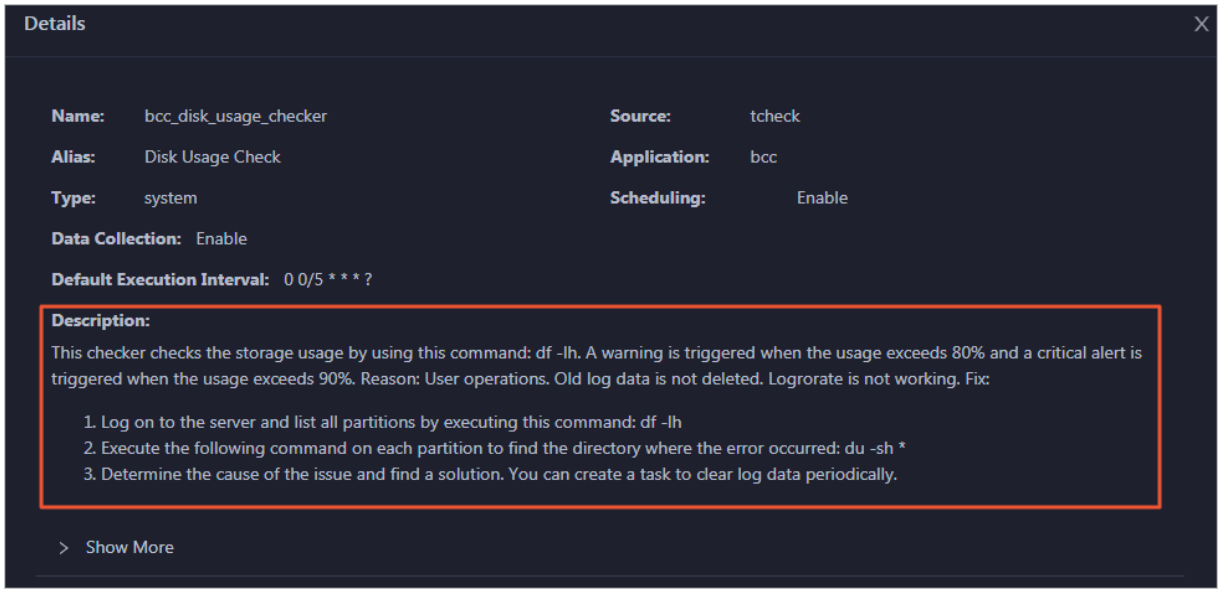

Log on to a host

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.

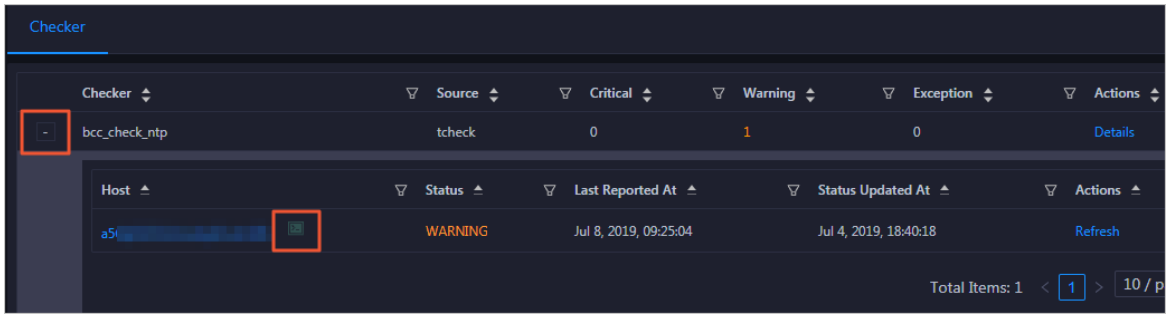

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

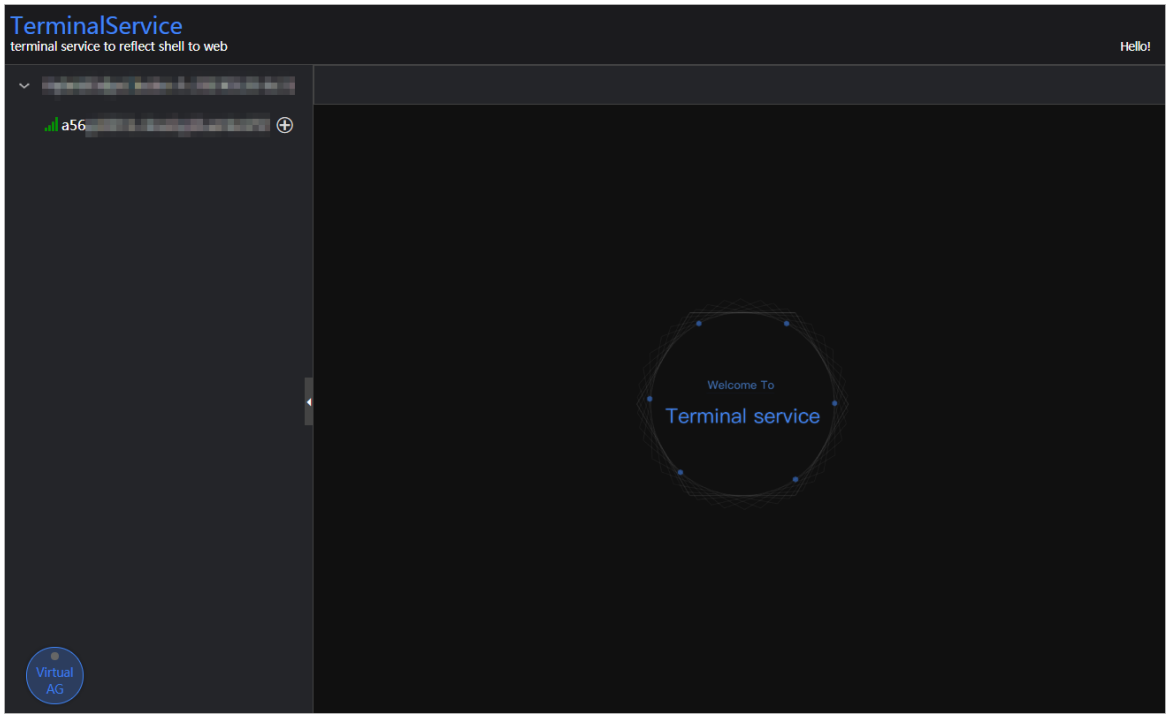

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

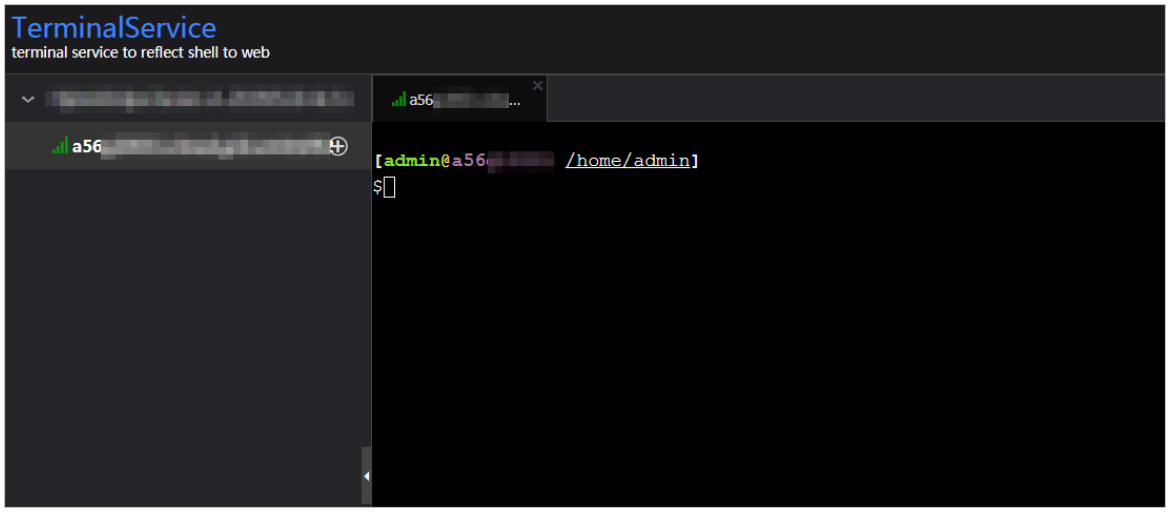

Run a checker again

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.

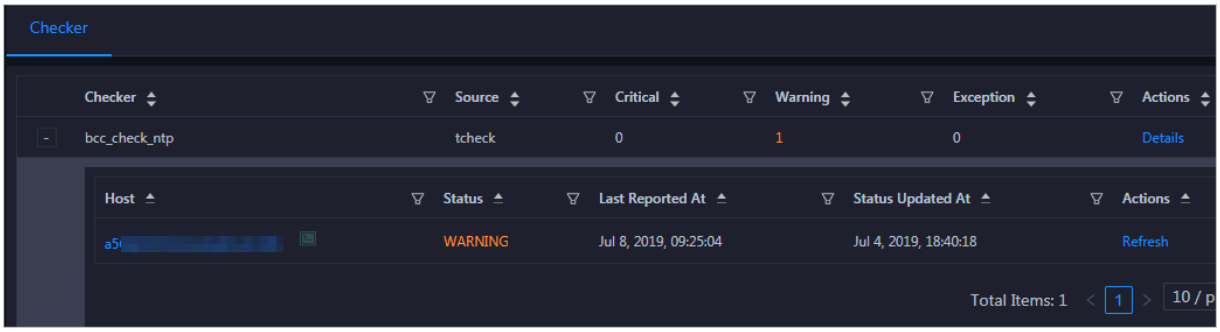

11.4.11.7.3. Host charts

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.

11.4.11.7.4. 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.

11.4.11.8. Job and queue analysis

11.4.11.8.1. 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.
	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%.

Metric	Sub-metric	Description
	Back Pressure Nodes	Checks whether each task has back pressure and finds the nodes that cause back pressure.

11.4.11.8.2. 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.

11.5. Apsara Big Data Manager (ABM)

11.5.1. Routine maintenance

11.5.1.1. 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.

11.5.1.2. 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.

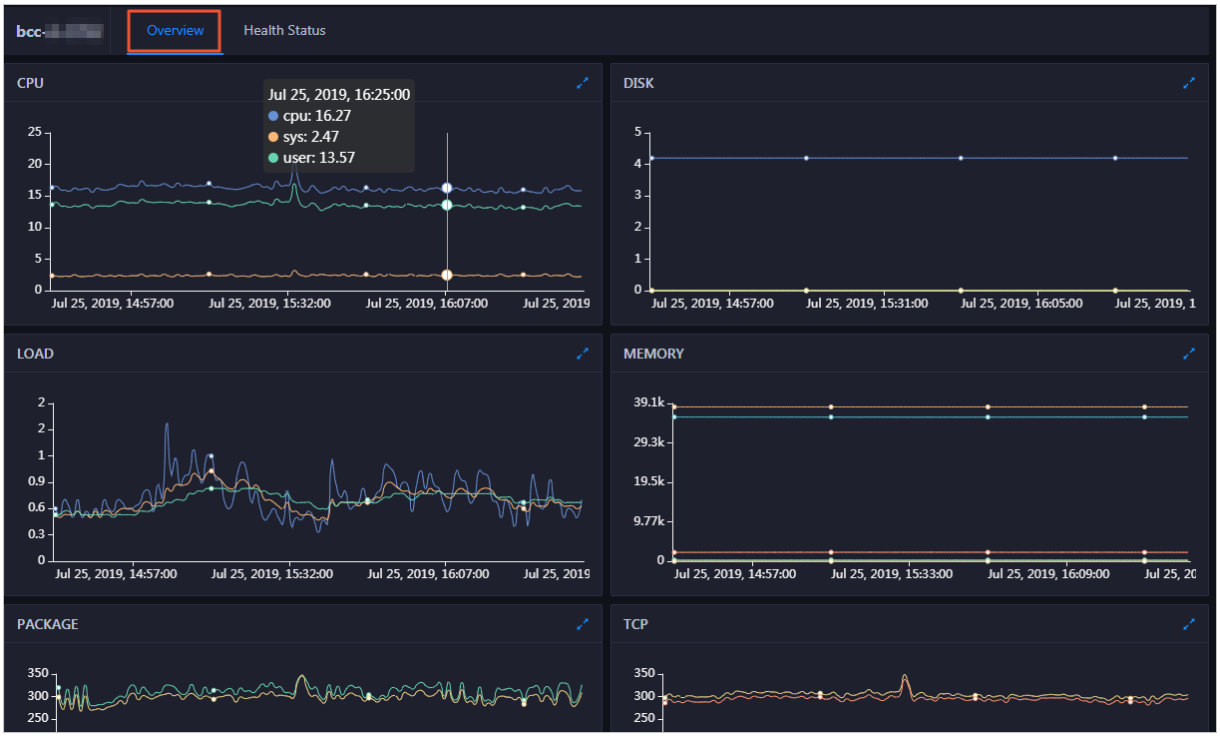

The **Overview** tab displays tendency charts for cluster metrics, including the CPU, memory, disk, load, package, TCP, and disk root directory usage. You need to regularly view and record these metrics to evaluate system downtime risks in the future.

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.

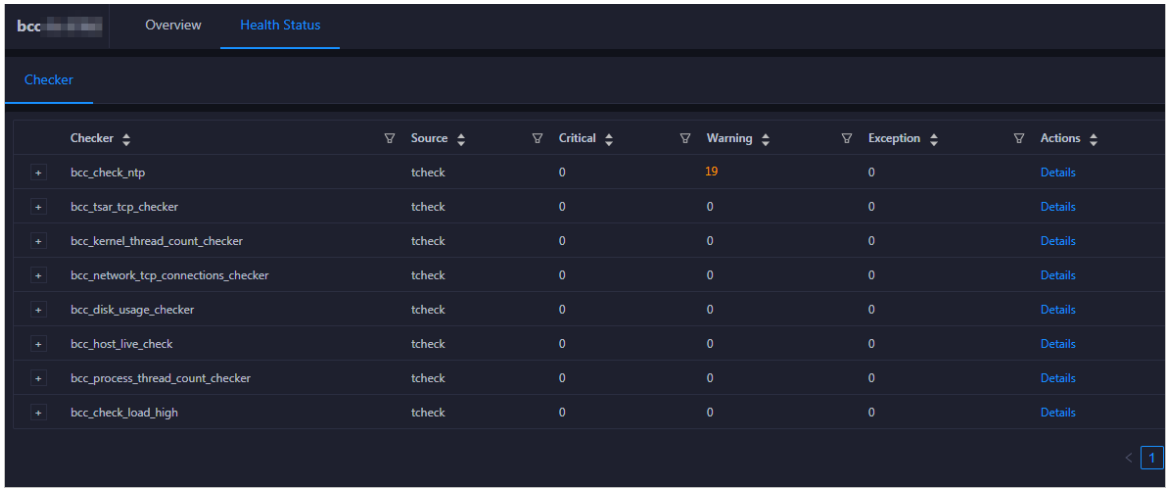

The **Health Status** tab displays all check items and the alerts that were generated during the check.

2. Click the **Fold** icon for a check item with alerts. All hosts on which the check item was performed appear.

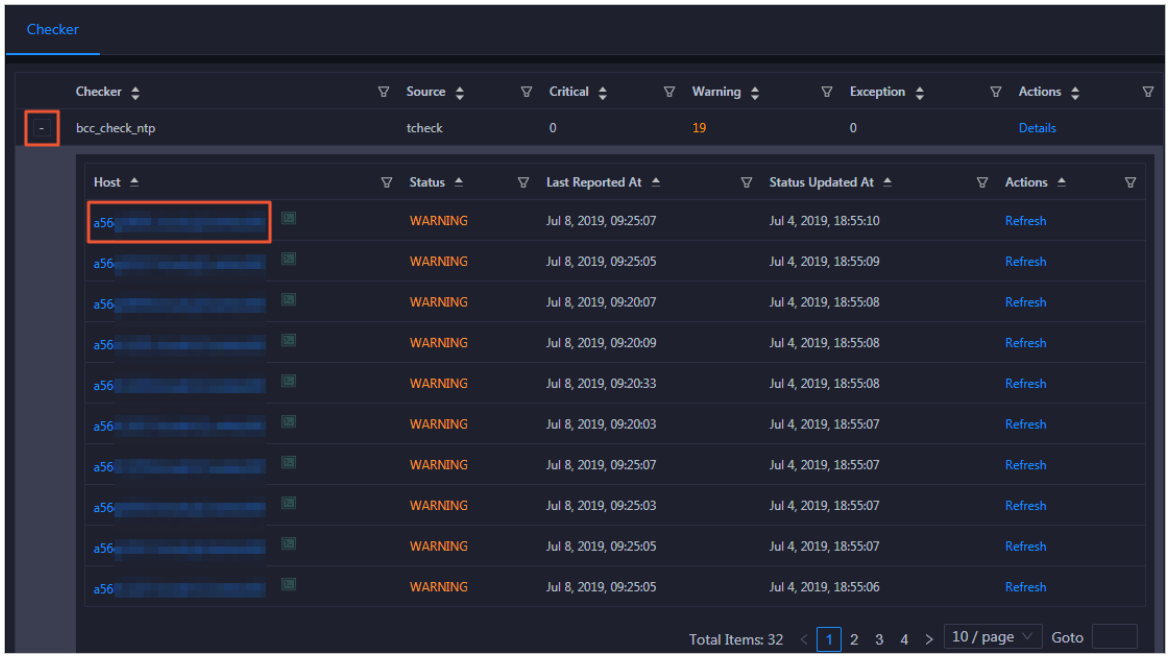

3. Click a host. In the dialog box that appears, click **Details** for an alert. The alert cause appears on the right.

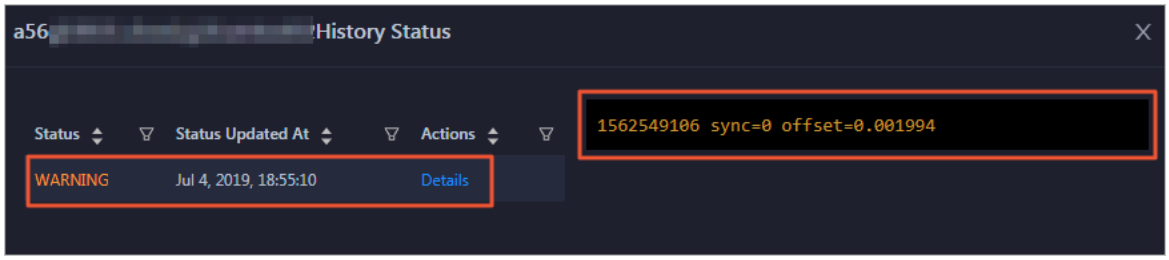

4. Click **Details** for a check item with an alert and view the fix method for the alert in the dialog box that appears.

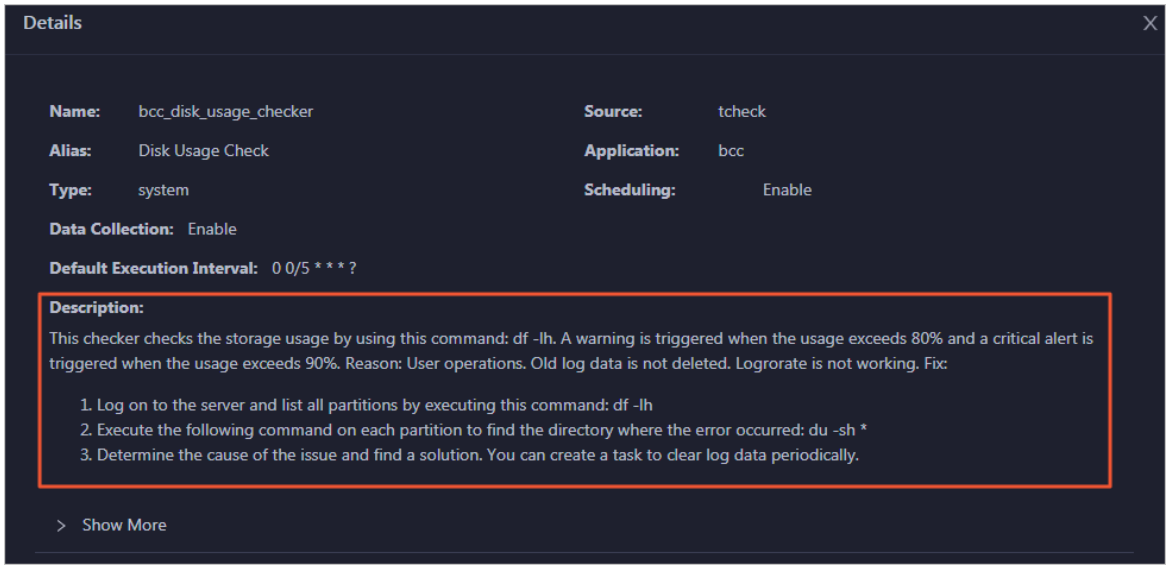

5. Handle the alert based on the fix method.

   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.

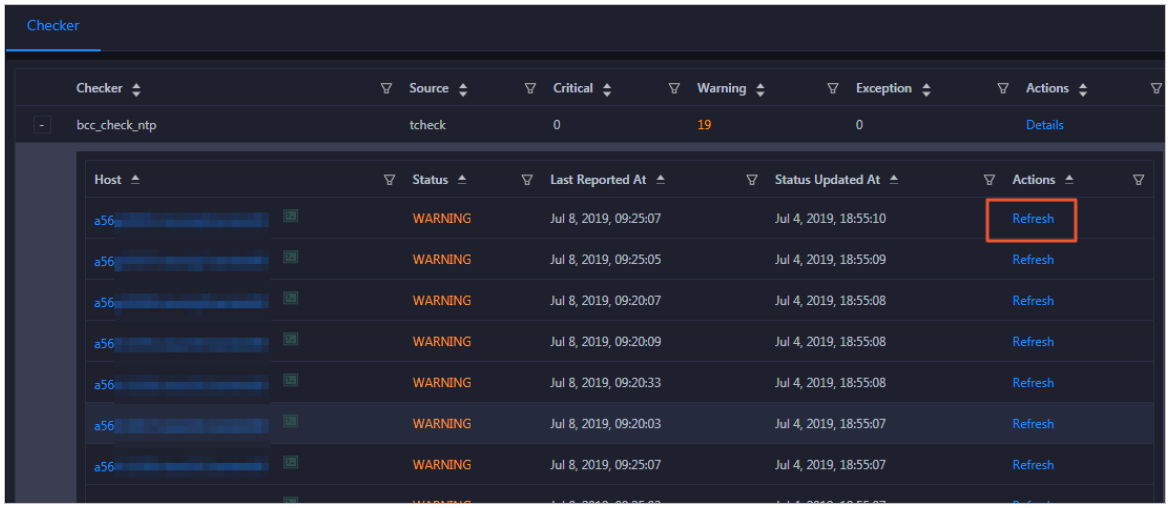

Log on to a host

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.

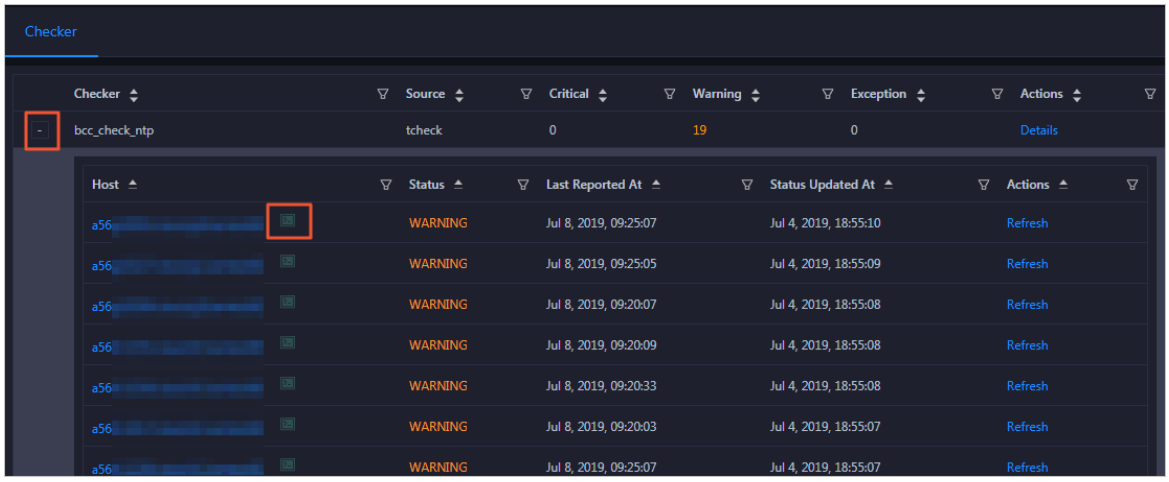

2. Click the **Logon** icon for a host. The **TerminalService** page appears.

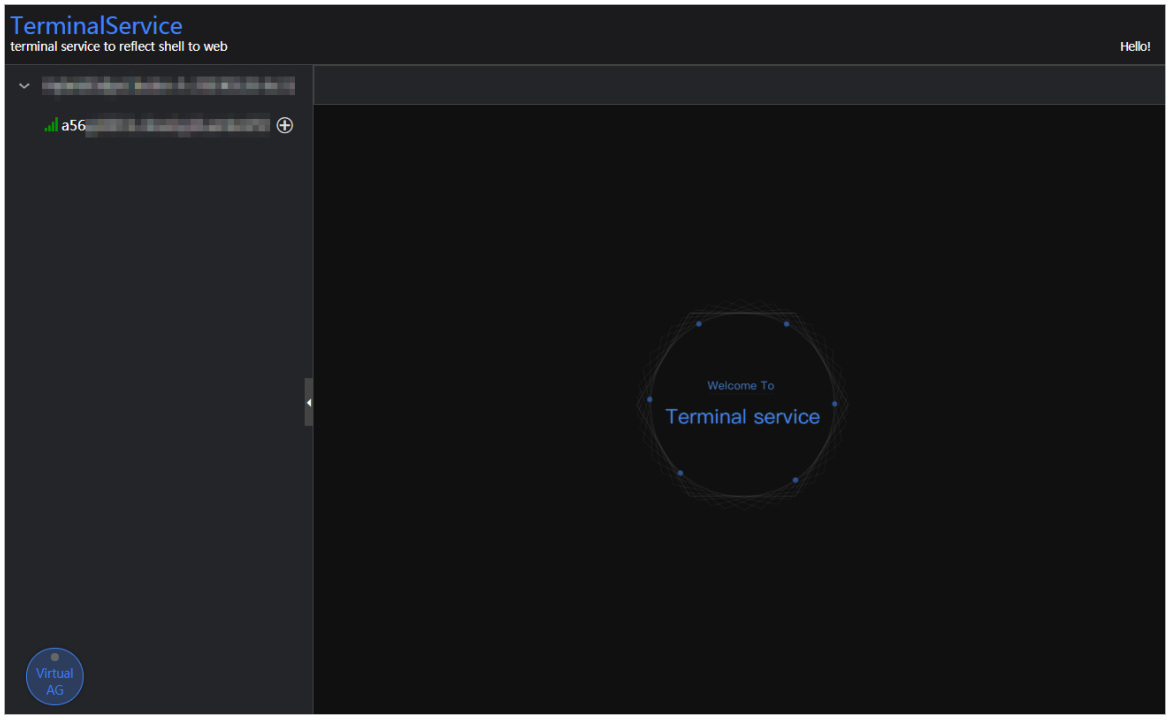

3. On the **TerminalService** page, select the host on the left to log on to it.

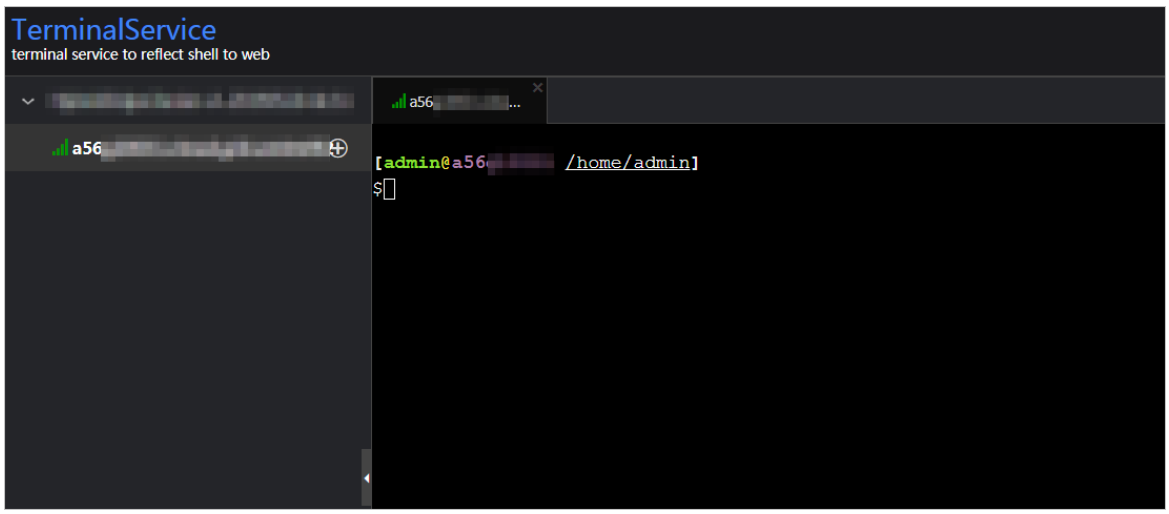

11.5.1.3. Troubleshooting

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.

11.5.2. 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.

11.6. Quick BI

11.6.1. Introduction to O&M and tools

11.6.1.1. O&M overview

Quick BI Operations and Maintenance Guide provides guidance for you to perform daily inspection, monitoring, and maintenance on Quick BI, and detect and rectify faults. These operations ensure availability, stability, and security of Quick BI.

You can use the Apsara Infrastructure Management Framework console to resolve the unavailability issues of Quick BI.

Onsite Apsara Stack engineers can use Apsara Bigdata Manager (ABM) to manage big data products. In the ABM console, they can view operation metrics, modify configurations, and check and handle alerts of big data products.

11.6.1.2. Check the Quick BI status in the Apsara Infrastructure Management Framework console

The Apsara Infrastructure Management Framework console is a tool that allows you to perform O&M on Quick BI. You can handle the unavailability issues of Quick BI in the console.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.Make sure that you have obtained the URL of the Apsara Stack Operations (ASO) console, and the username and password used to log on to the console from a deployment engineer or administrator.
   i. Open a browser, enter the URL in the address bar, and press Enter. The URL is in the format of *region-id*.aso.*intranet-domain-id*.com.

   > **Note** We recommend that you use the Google Chrome browser.
   >
   > You can select a language from the drop-down list in the upper-right corner.

   ii. Enter the correct username and password.
   - The following user roles are available by default:
     - security administrator: the user who has the permissions to manage other users or roles
     - auditor: the user who has the permissions to view audit logs
     - sysadmin: the user who has the permissions that security administrators and auditors do not have
   - When you log on to the ASO console for the first time, you must change the password of your username as prompted.

     For security concerns, your password must contain the following characters:
     - Letters
     - Digits
     - Special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%)

     The password must be 10 to 20 characters in length.

   iii. Click **Log On** to go to the **ASO** console.
   iv. In the left-side navigation pane, choose **Products > Product List**. Click **Apsara Infrastructure Management Framework**.

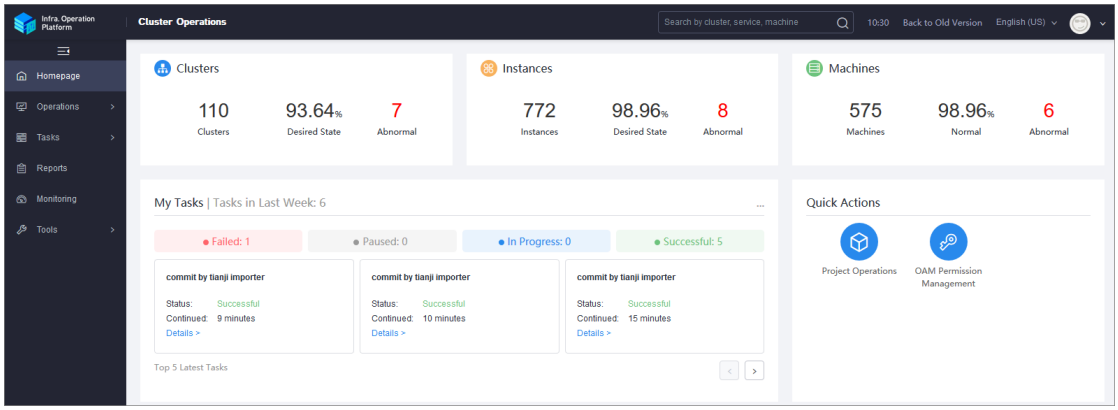

2. On the homepage of the Apsara Infrastructure Management Framework console, enter the keywords of the name of the target Quick BI cluster in the search box, select the target cluster from the drop-down list, and click **Operations** next to the cluster.
3. On the **Cluster Operations** page, check the status of the Quick BI cluster.Check whether the

cluster is at desired state:

- If the cluster is at desired state, the system is running properly.
- If the cluster is not at desired state, go to the next step.

4. Click the service name in the **Services** column or exception number in the **Service Role** column to check the details of each Quick BI service instance.
5. Click Details in the **Service Role Status** column to view the exception information.
6. (Optional)If a service instance is not at desired state, resolve the issue.Dependencies exist between server roles. If an upstream server role does not reach the desired state, its downstream server roles cannot reach the desired state. We recommend that you first troubleshoot the upstream server role. The following figure Dependencies between Quick BI server roles shows the dependencies between Quick BI server roles.

Dependencies between Quick BI server roles

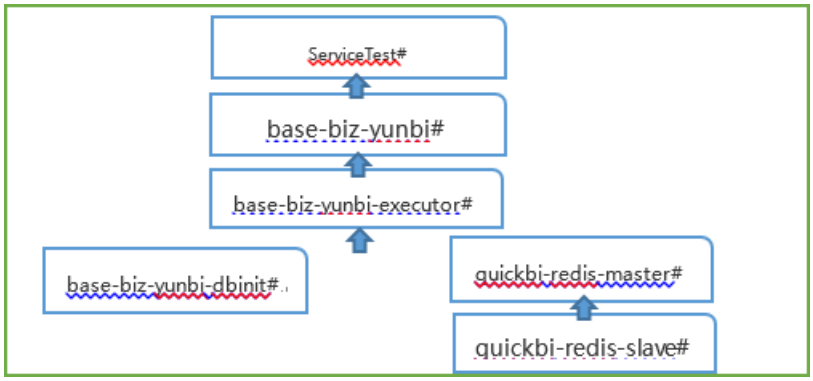

For example, if server role **base-biz-yunbi-executor#** does not reach the desired state, **base-biz-yunbi#** and **ServiceTest#** cannot reach the desired state. In this case, you must make sure that server role **base-biz-yunbi-executor#** reaches the desired state. After **base-biz-yunbi-executor#** reaches the desired state, **base-biz-yunbi#** and **ServiceTest#** will reach the desired state sequentially in normal cases.

11.6.1.3. Perform O&M on Quick BI in the ABM console

Apsara Bigdata Manager (ABM) allows you to perform O&M on big data products from the perspective of business, service, cluster, and host. You can also update big data products, customize alert configurations, and view the O&M history in the ABM console.

ABM, formally known as BCC, is an operations and management platform tailored for big data products.

For information about the ABM logon methods and Quick BI O&M operations in the ABM console, see topics under *Quick BI O&M*.

11.6.2. Routine maintenance

11.6.2.1. Introduction to Quick BI components

You can use container monitoring and inspection management to check whether server roles related to Quick BI components are at desired state, so as to perform routine maintenance. This topic describes Quick BI operations and maintenance (O&M) components, including their server roles and functions.

The following table lists Quick BI O&M components, related server roles, and component functions.

Component	Server role	Function
Database initialization component	base-biz-yunbi-dbinit#	Initializes Quick BI metadata. The server role must be at desired state. Otherwise, Quick BI cannot run properly.
Cache component	quickbi-redis-master#	Caches Quick BI data to improve query performance.
	quickbi-redis-slave#	
Runtime component	base-biz-yunbi-executor#	Retrieves table metadata and data from data sources.
Web service component	base-biz-yunbi #	Provides web services. The server role allows frontend clients to visit Quick BI web pages.
Automated testing component	ServiceTest#	Checks the overall availability of Quick BI by running multiple test cases at a time.

Note When you deploy or update Quick BI, the server roles are automatically started.

11.6.2.2. 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.

11.6.2.3. 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.

11.6.2.4. 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.

11.6.2.5. 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.

11.6.2.6. Automated testing components

This topic describes how to identify and troubleshoot issues when you perform container monitoring and inspection management on automated testing components.

Container monitoring

In the Apsara Infrastructure Management Framework console, check whether server role **ServiceTest#** 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 server role **base-biz-yunbi#**.

- The service is available but an error occurs. 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 displayed in the Description column. For more information, see *the instance status monitoring description* in Quick BI Operations and Maintenance Guide.

Solutions:

- If the service is unavailable, check whether other server roles are at desired state. If they are not, resolve the issue.
- If the service is available but an error occurs, contact Quick BI technical support and provide error information.

Inspection management

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.

11.6.3. Quick BI O&M

11.6.3.1. Log on to Apsara Bigdata Manager

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   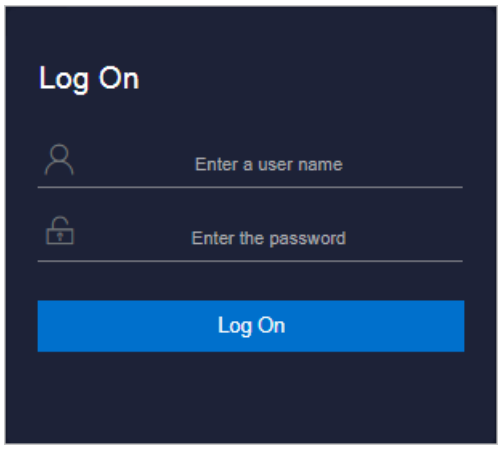

   Note You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

   Note Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.

- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.6.3.2. 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 transmission, 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 **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**.

11.6.3.3. Service O&M

11.6.3.3.1. 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.

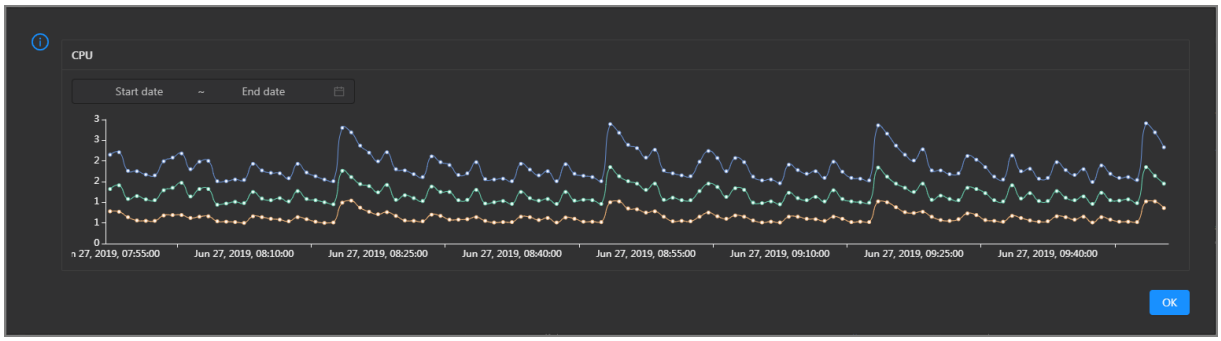

DISK

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.

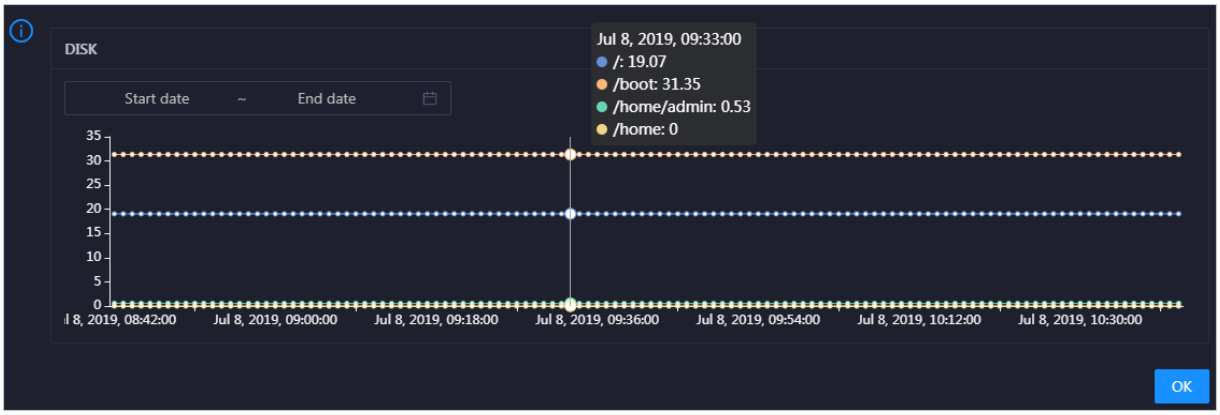

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

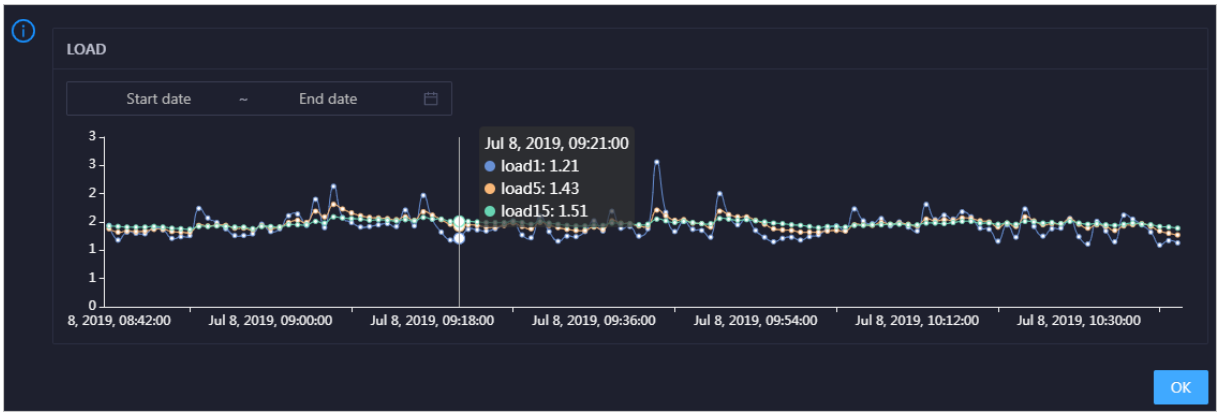

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

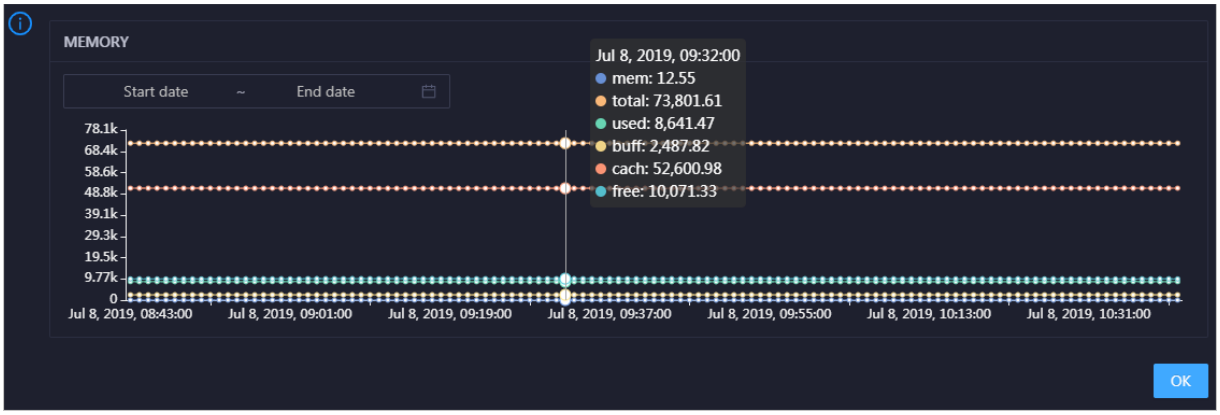

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

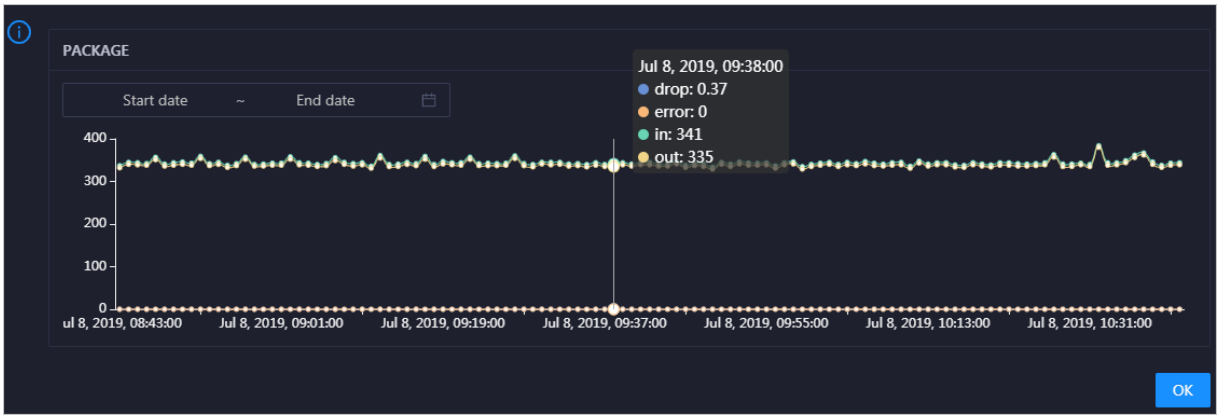

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

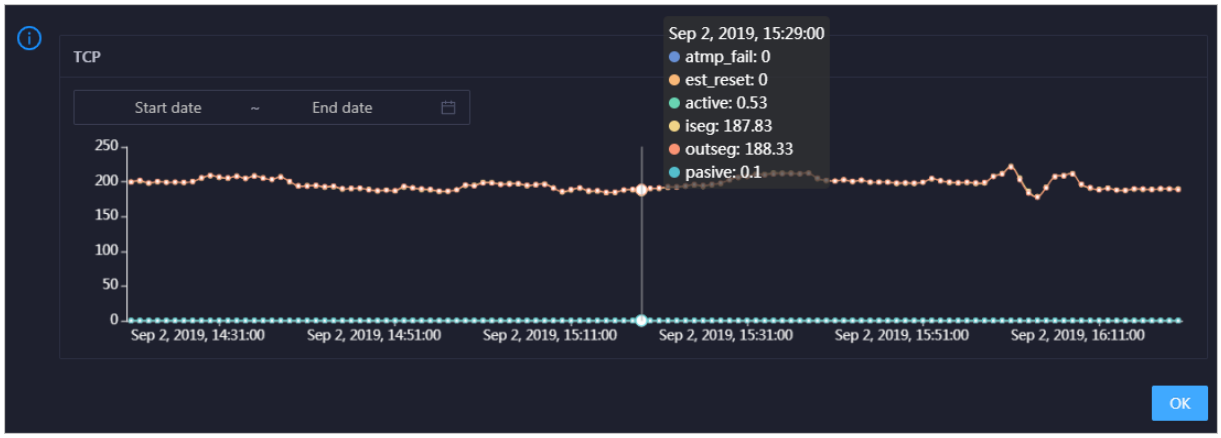

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

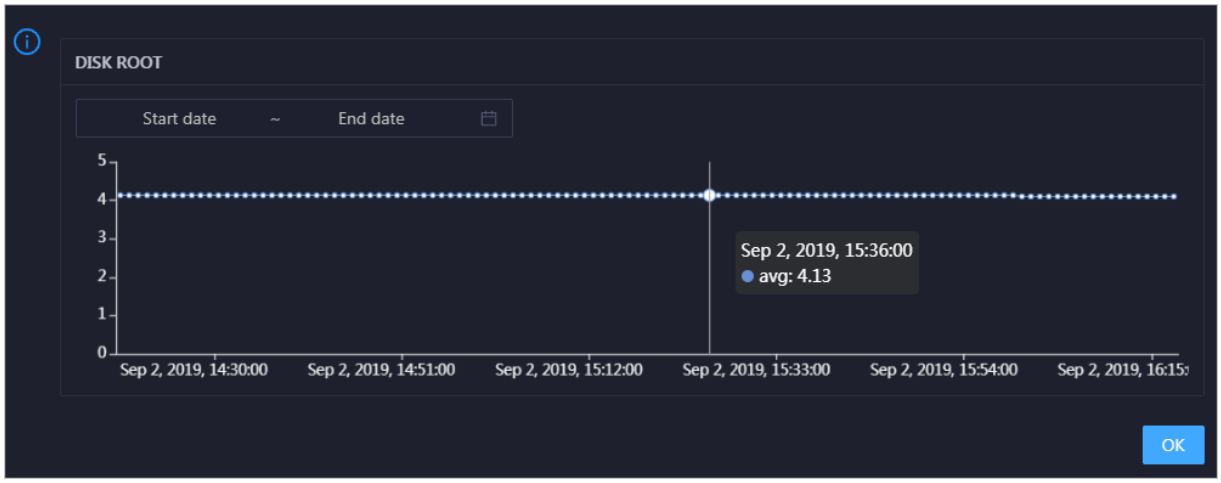

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

11.6.3.3.2. 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.

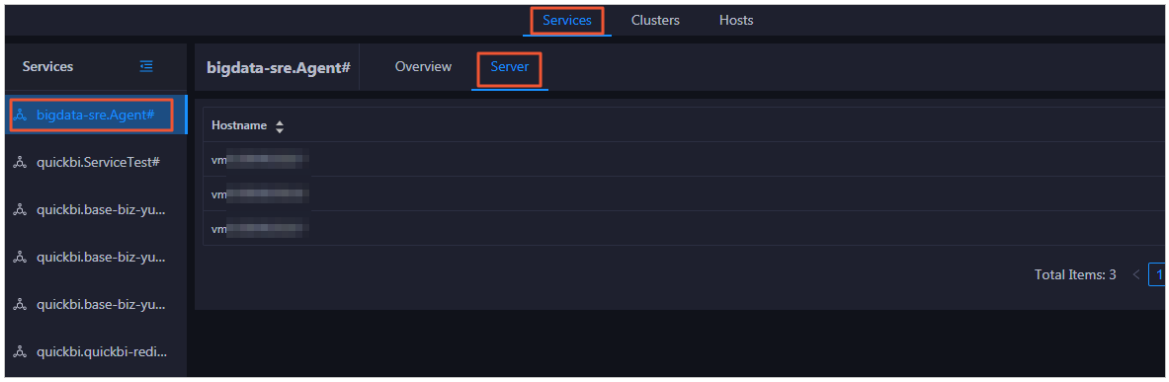

On the **Server** page, you can view the hosts where the selected service is run.

11.6.3.4. Cluster O&M

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

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

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

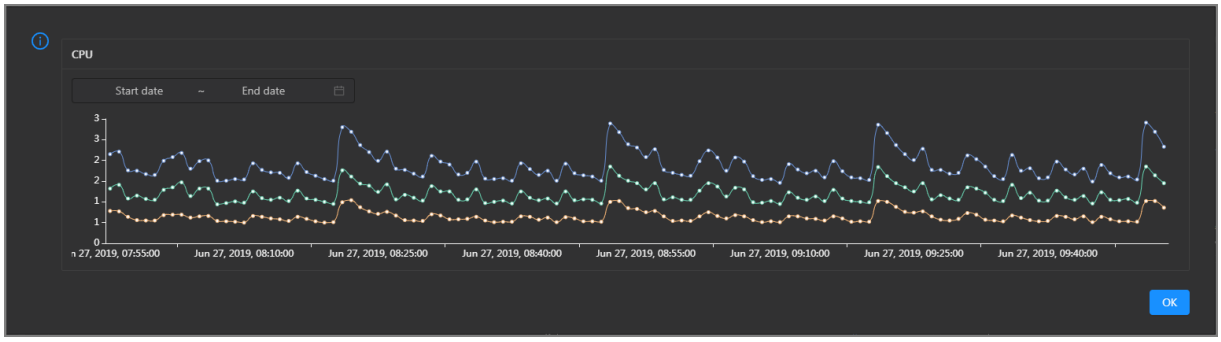

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

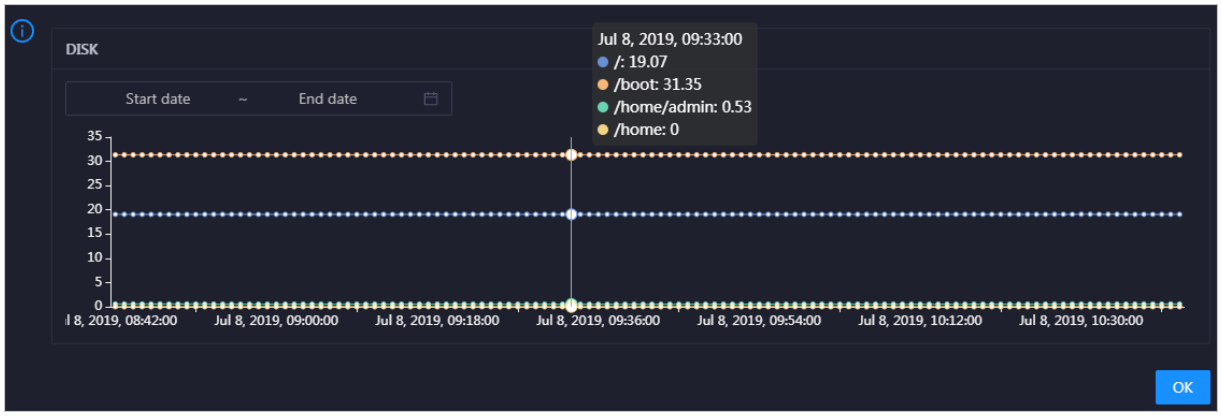

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

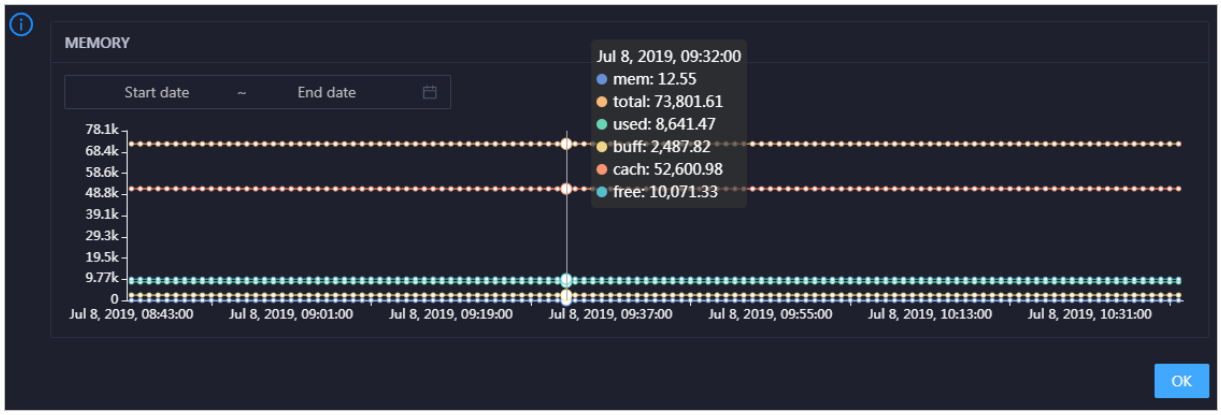

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

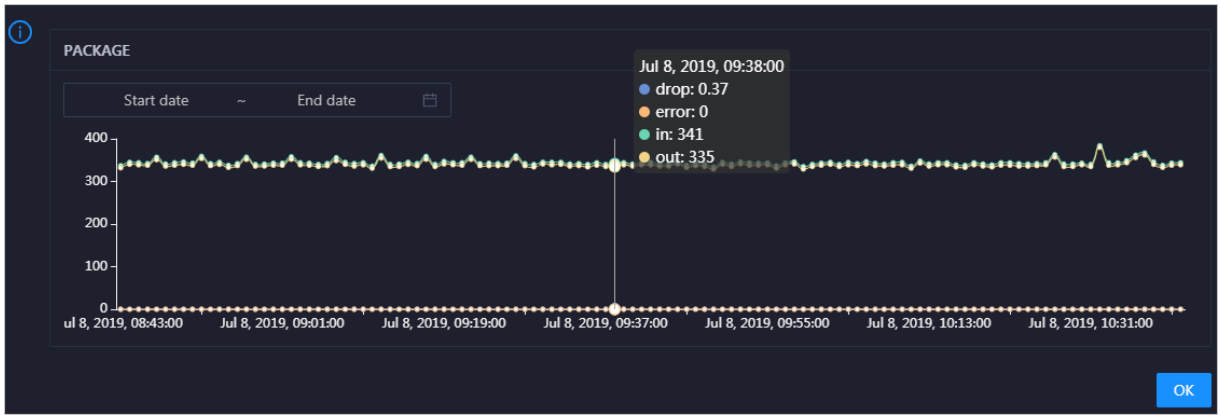

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

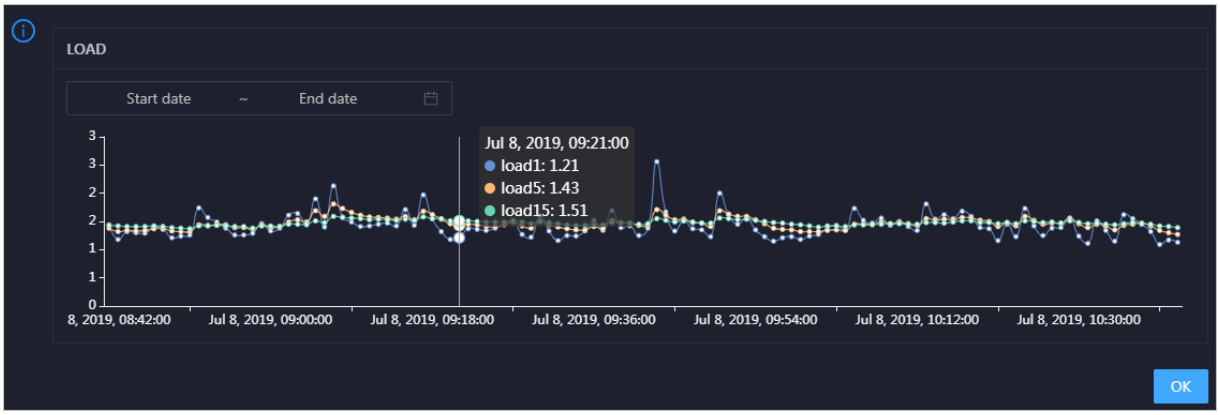

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.6.3.4.2. 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

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.

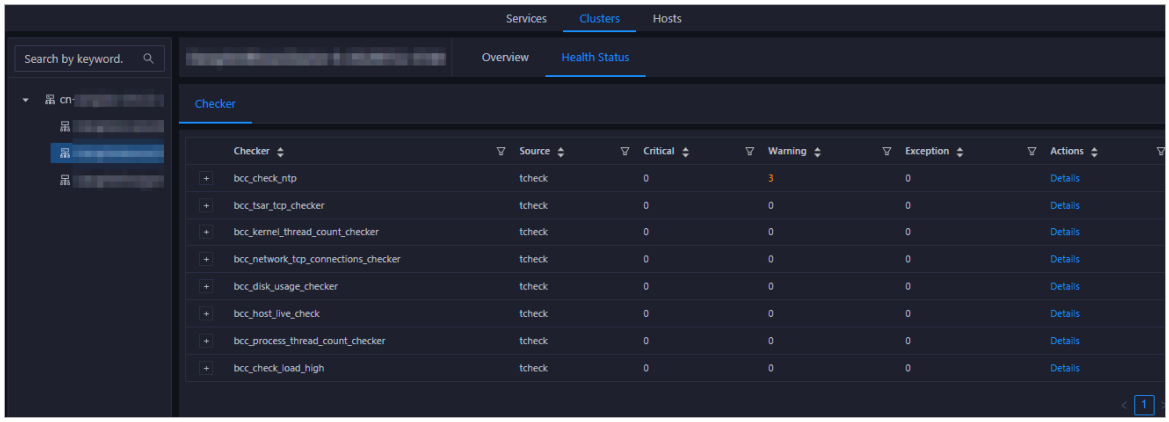

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are repesented in different colors. Handle the alerts in a timely manner,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

View checker details

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

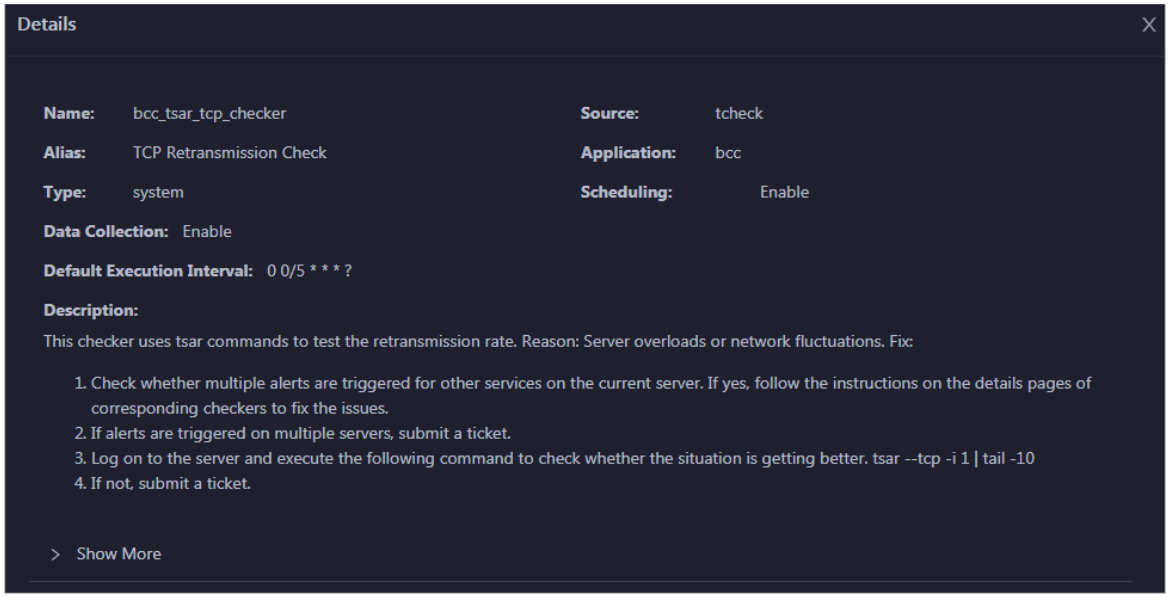

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

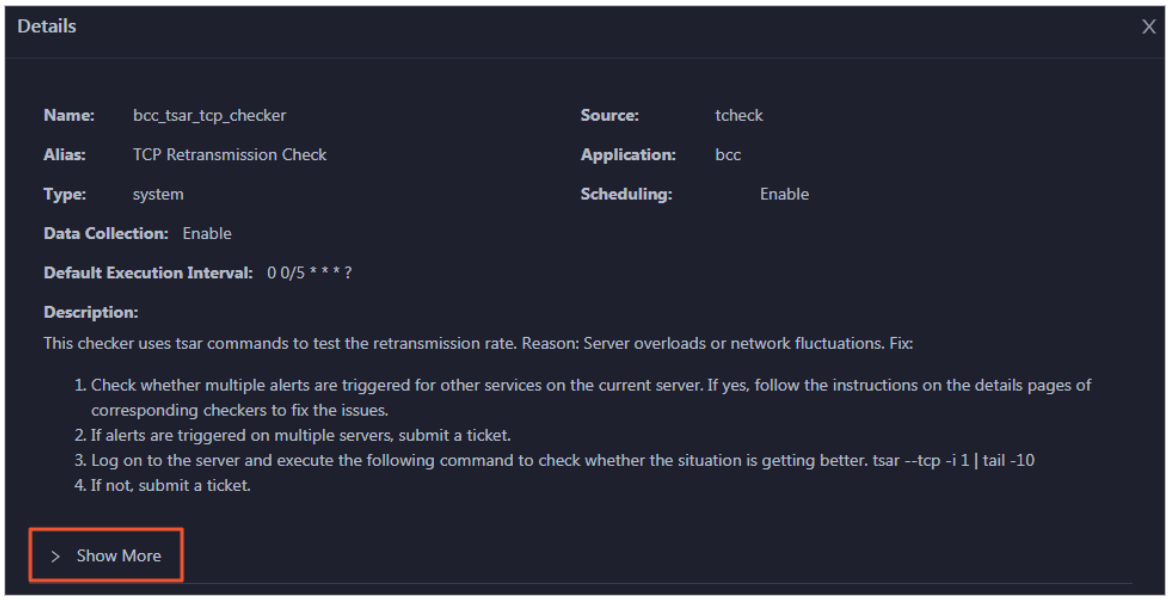

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

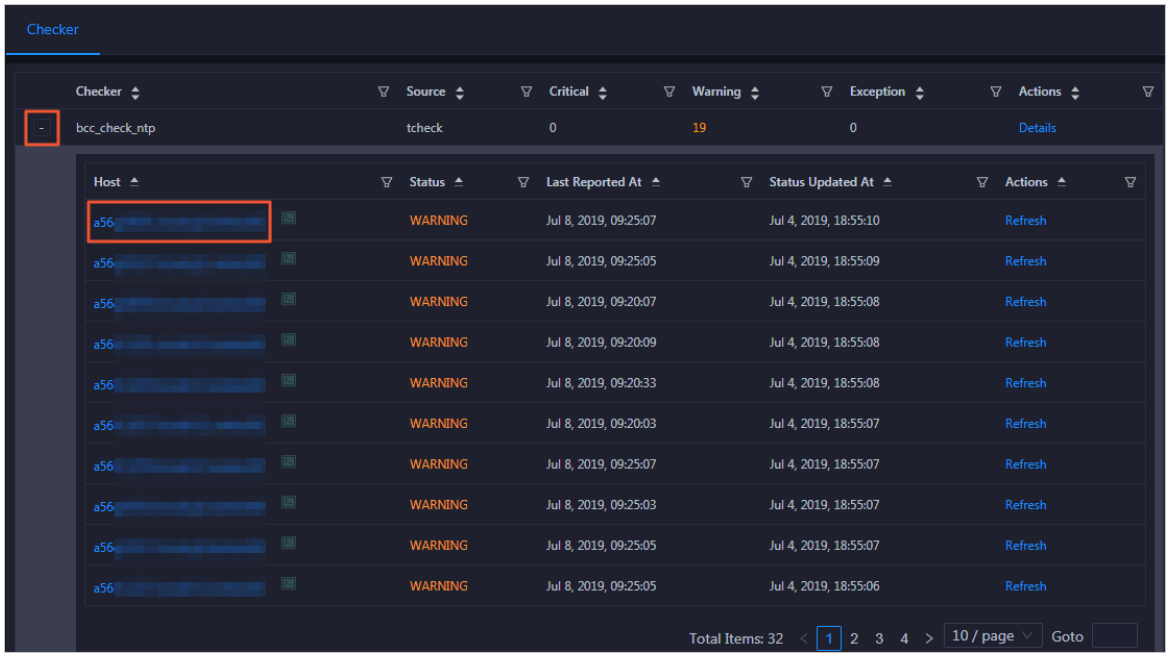

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

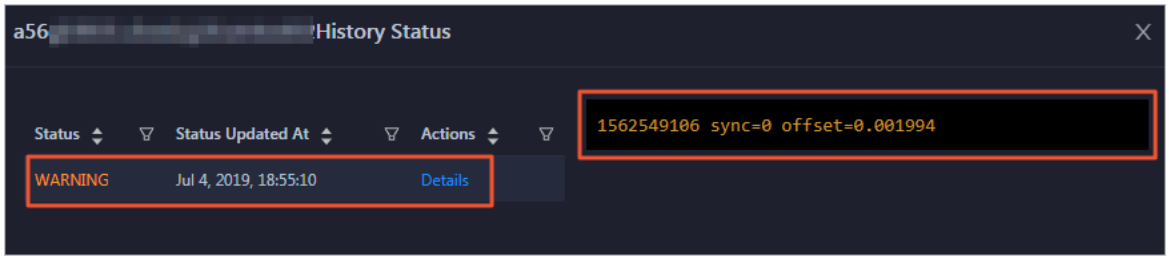

Clear alerts

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

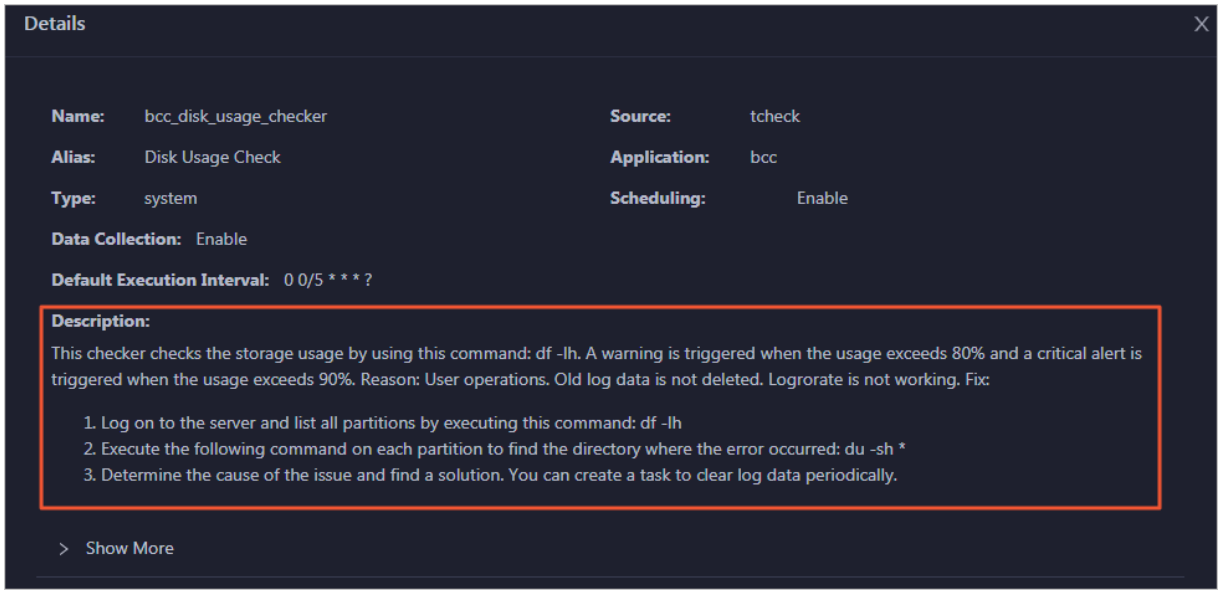

Log on to a host

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

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

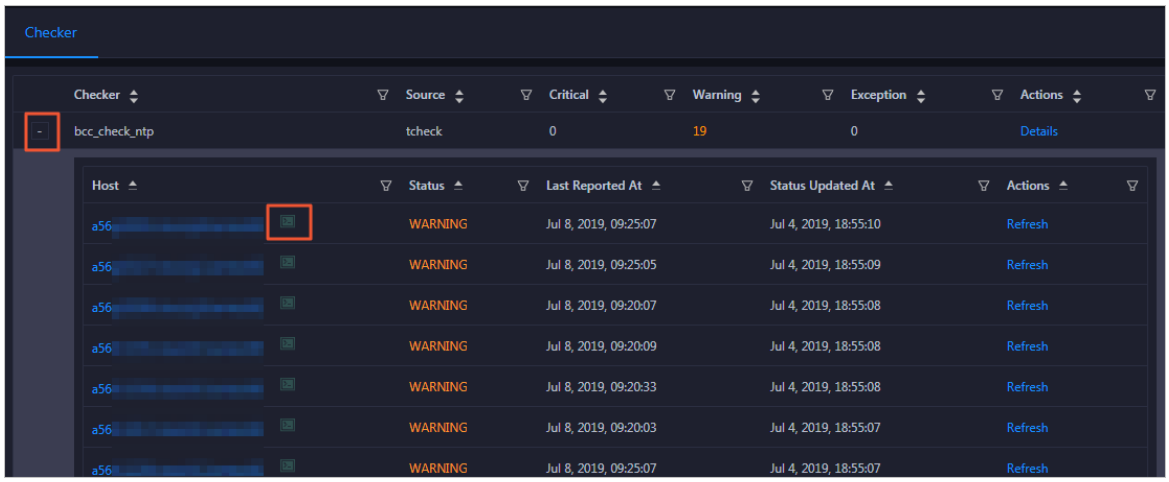

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

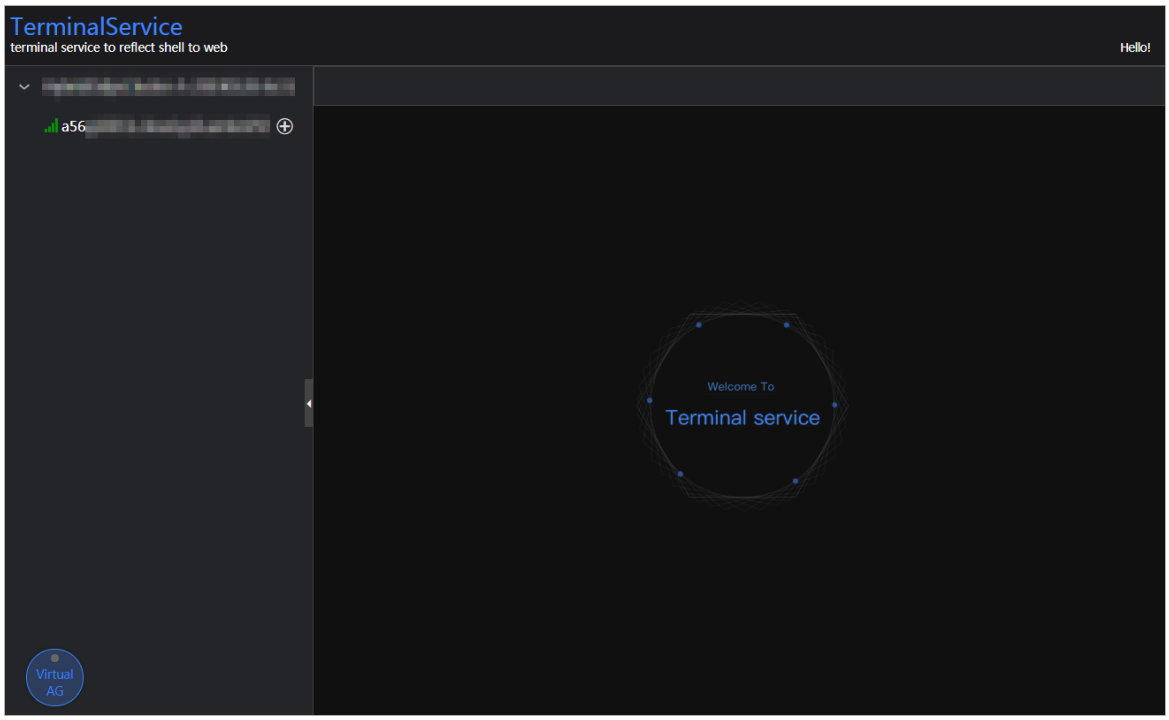

3. On the **TerminalService** page, click the hostname to log on to the host.

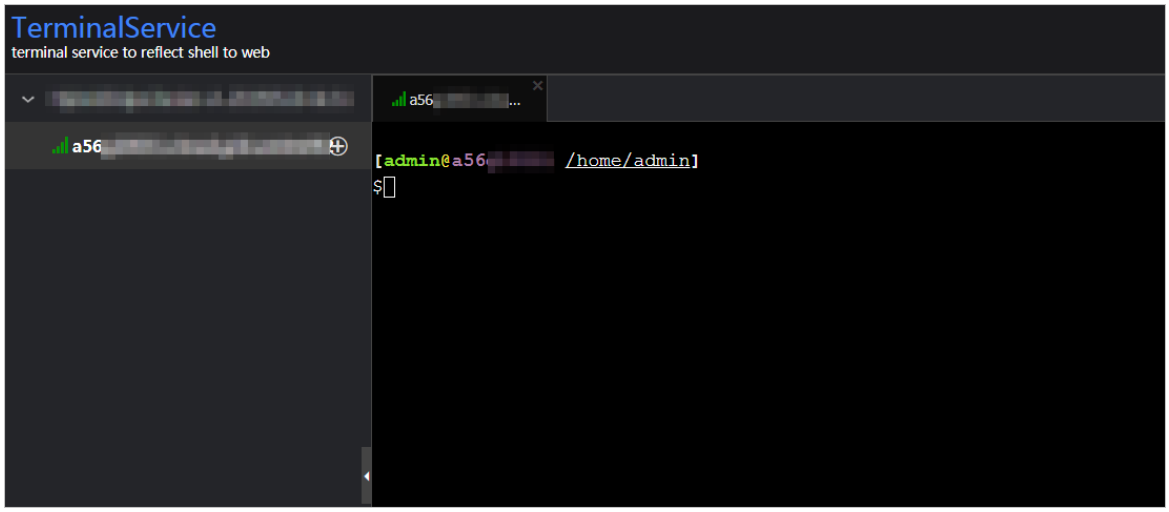

Run a checker again

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

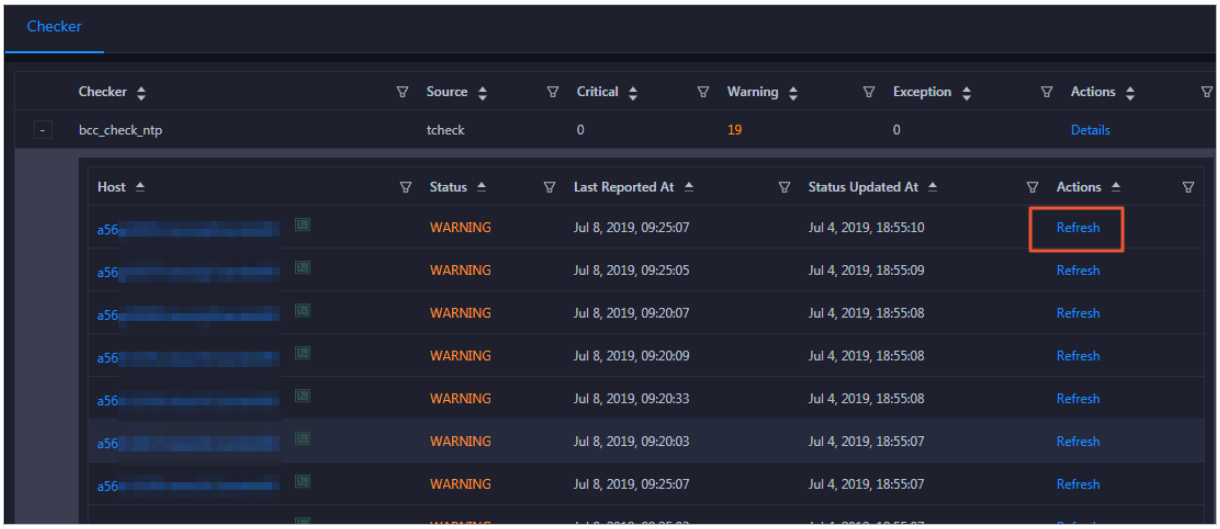

11.6.3.5. Host O&M

11.6.3.5.1. 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.

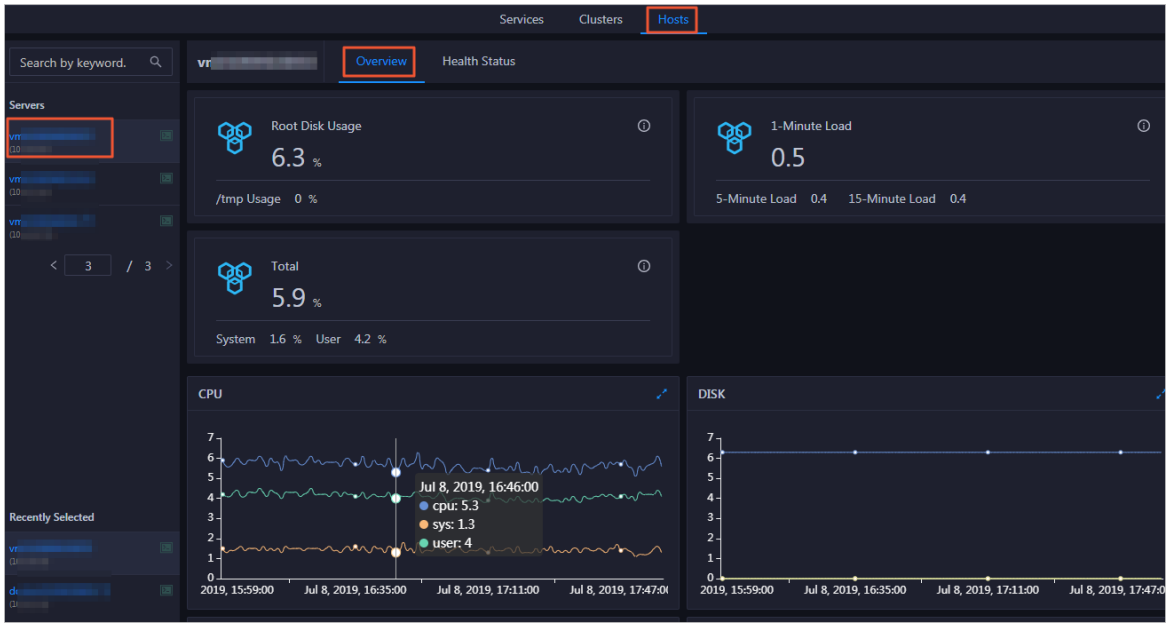

On the **Overview** page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root 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.

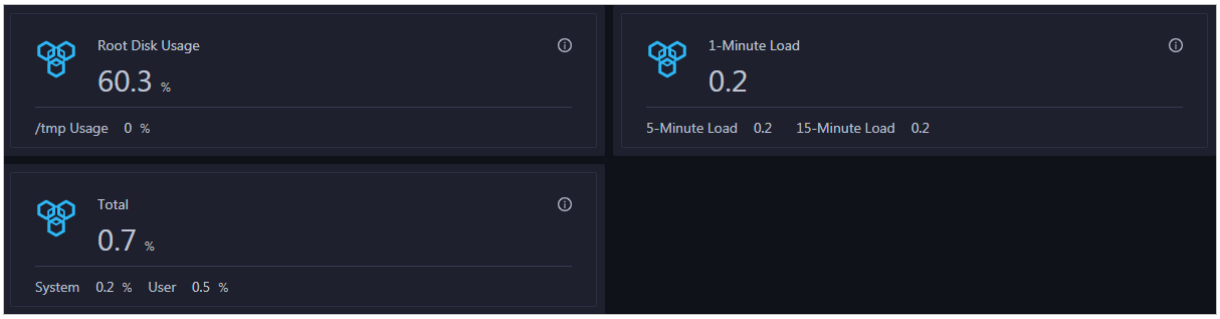

CPU

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.

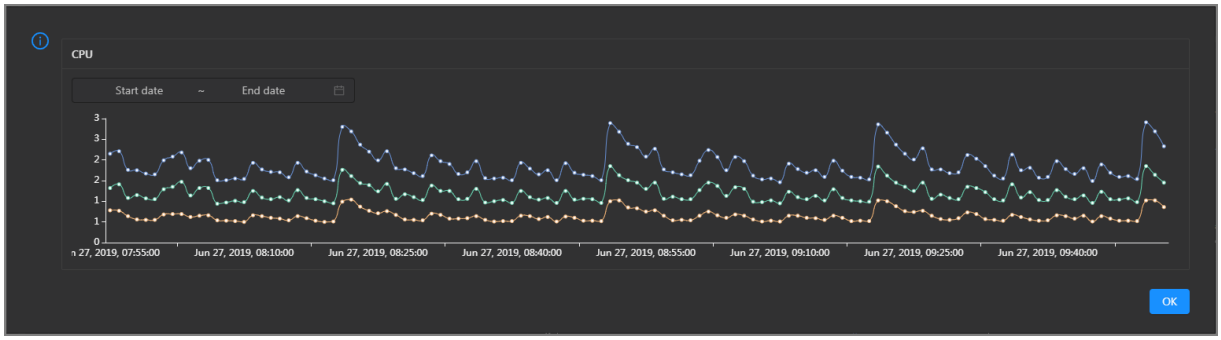

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

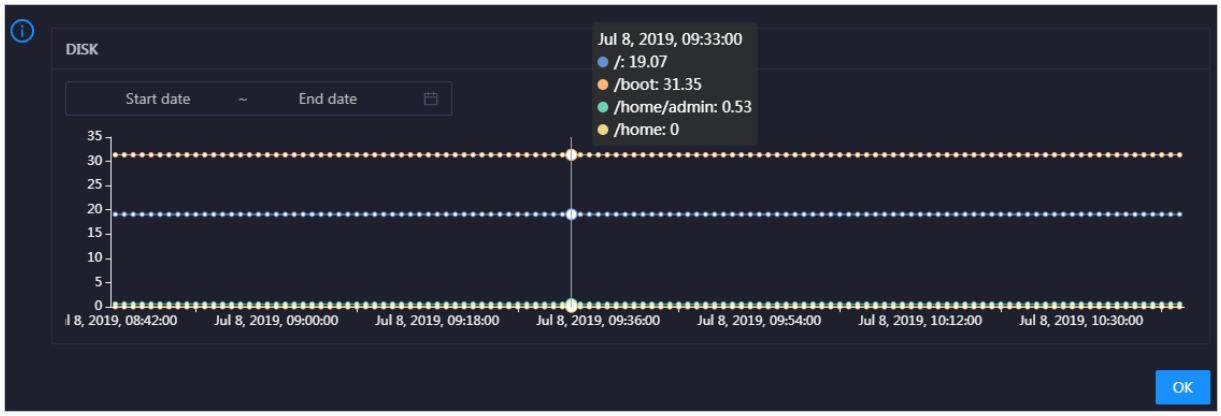

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

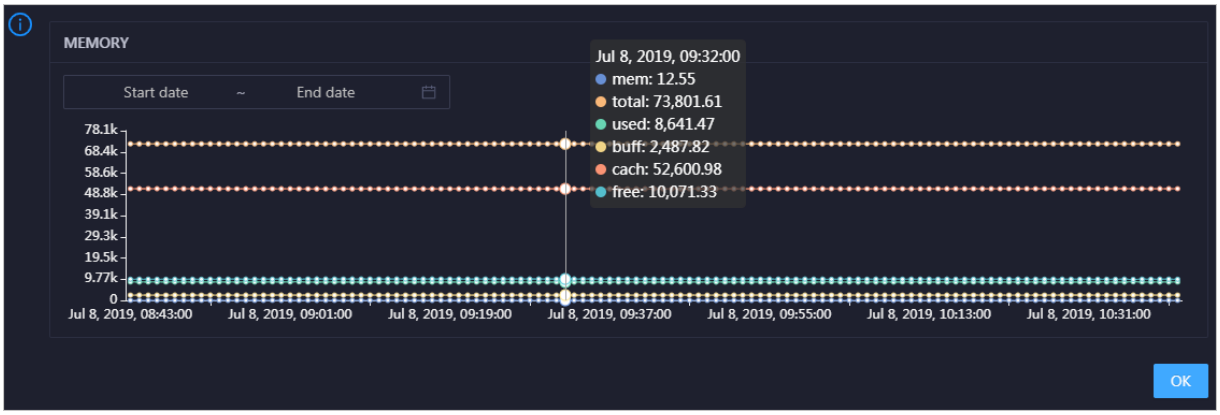

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

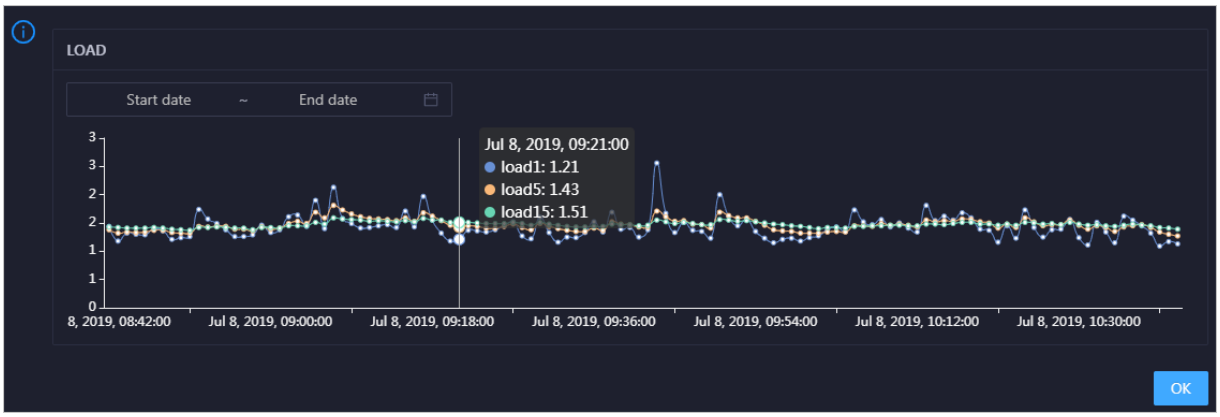

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

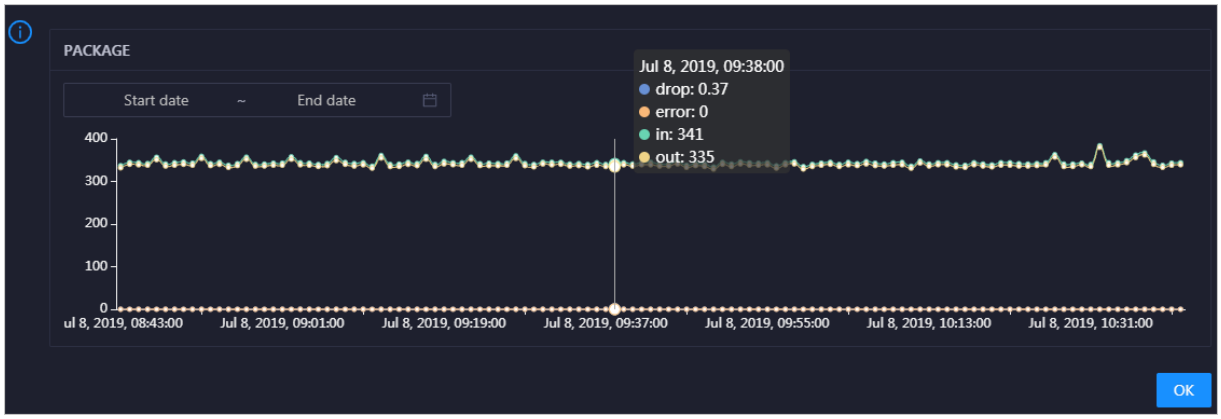

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

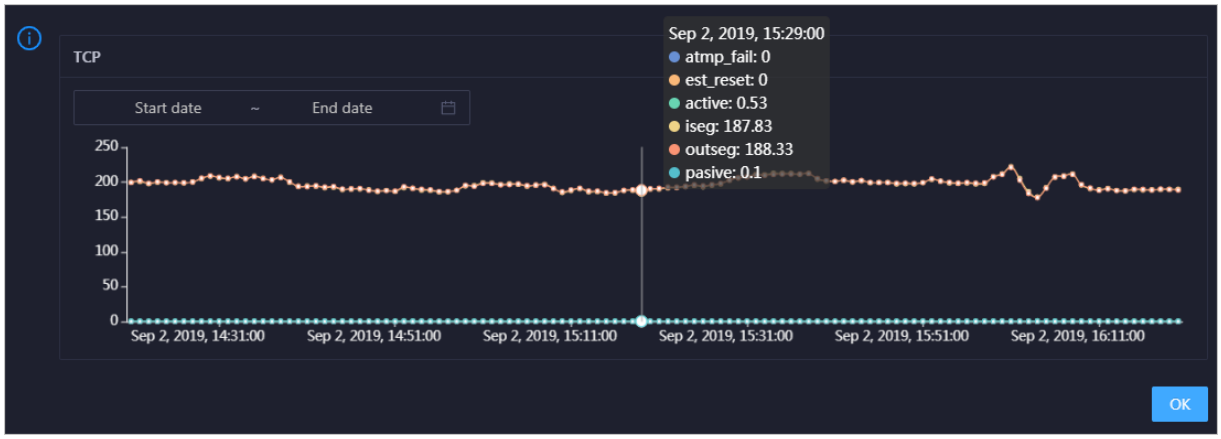

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

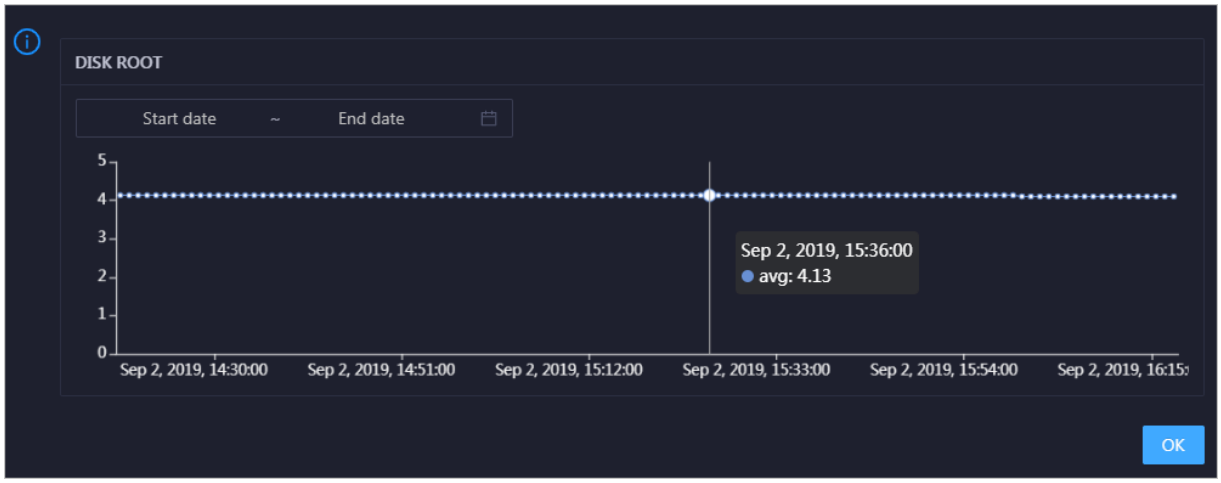

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

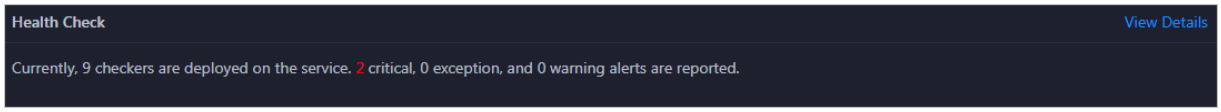

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

Health Check History

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

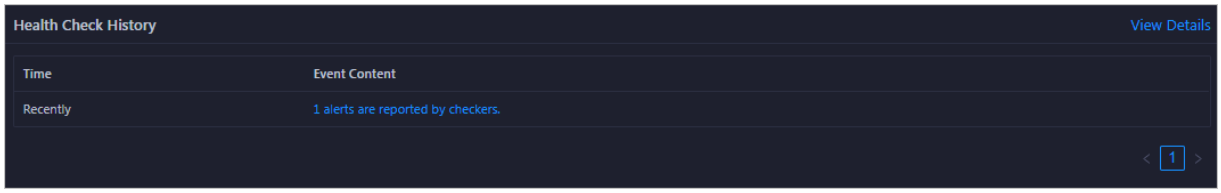

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

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

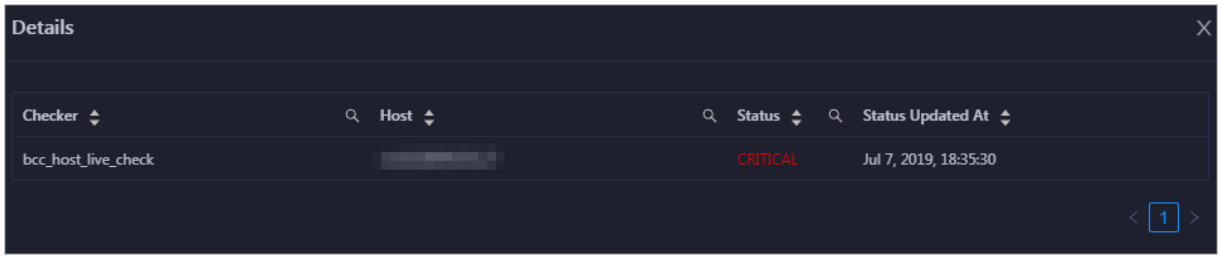

11.6.3.5.2. Host health

On the host health status page, you can view the checkers of all hosts, 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. 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 **Health Status** tab. The **Health Status** page for the host appears.

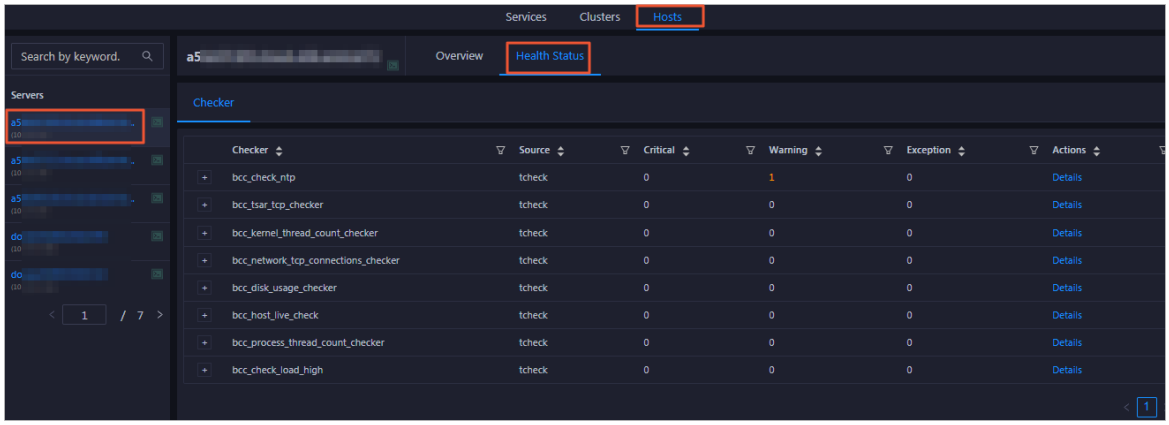

On the **Health Status** page, you can view all checkers of the host and the check results for the hosts in the host. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, **Critical**, **Warning**, and **Exception** results are alerts. You need to pay special attention to them, especially the **Critical** and **Warning** results.

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.

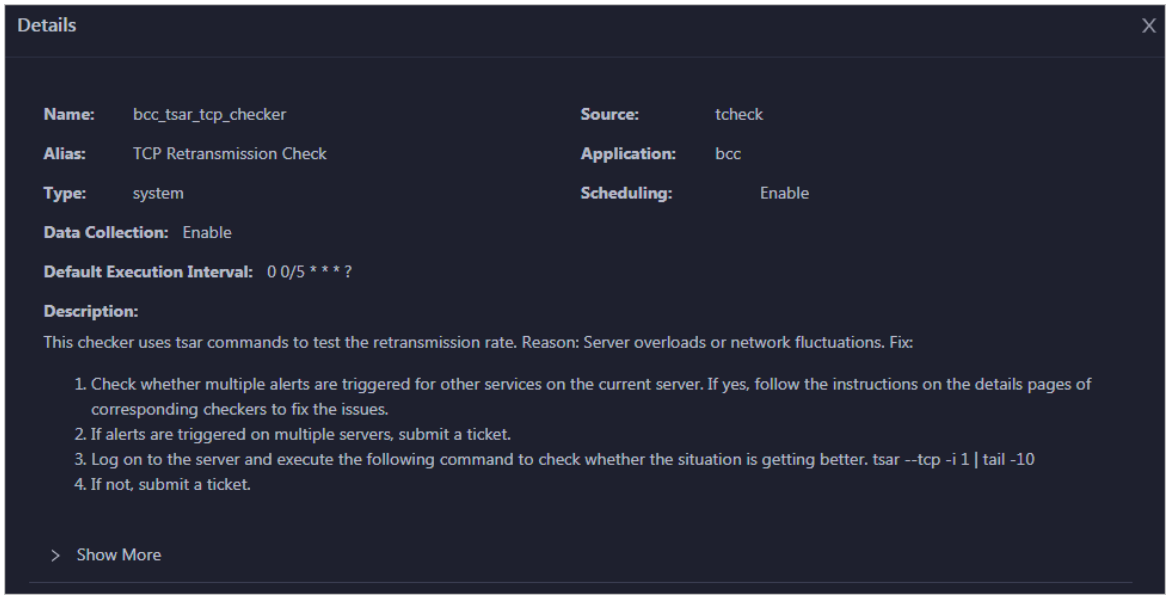

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

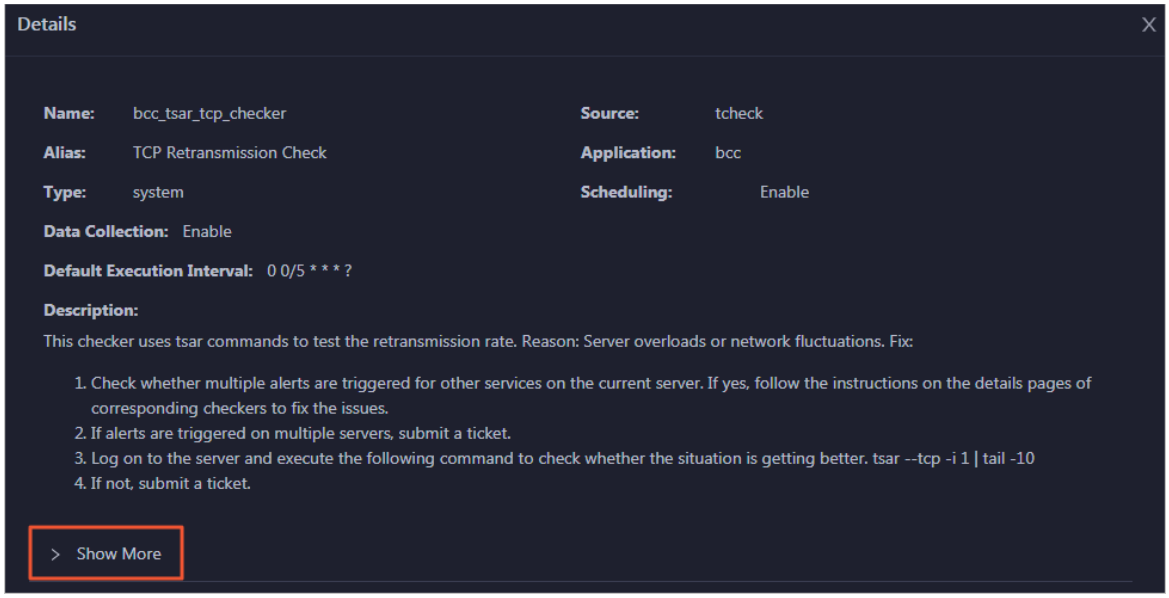

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

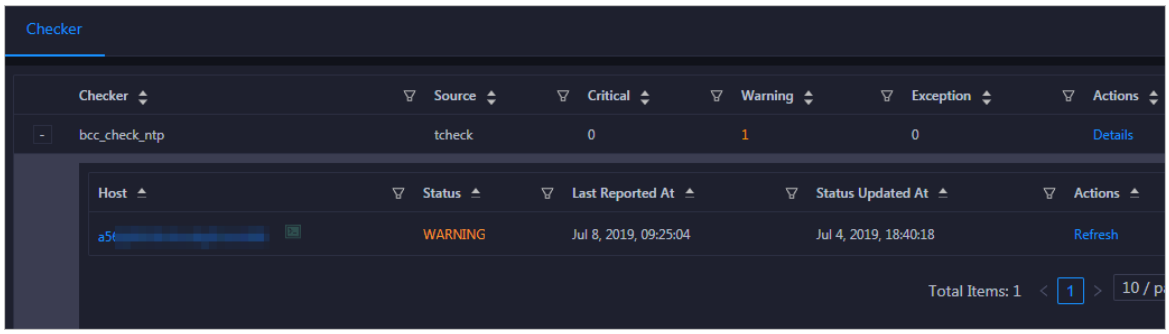

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

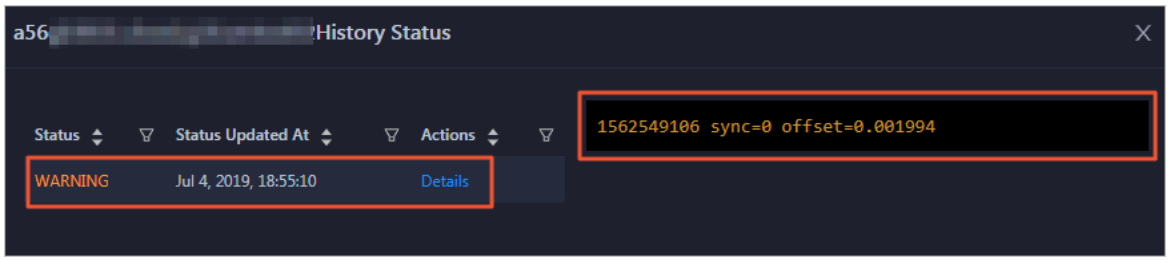

Clear alerts

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.

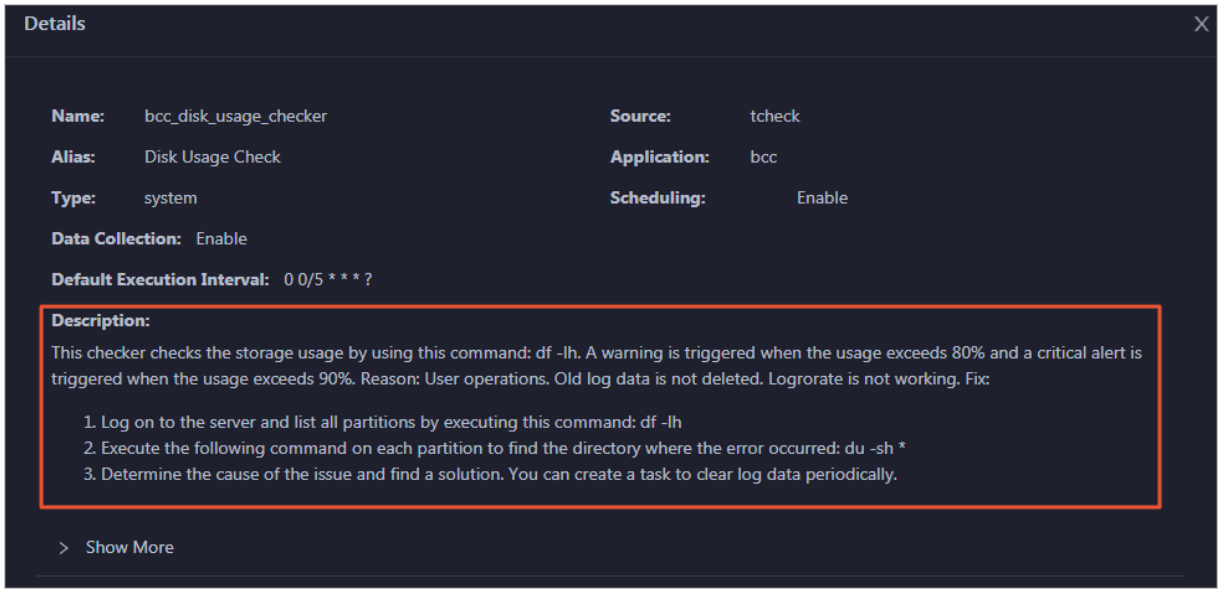

Log on to a host

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.

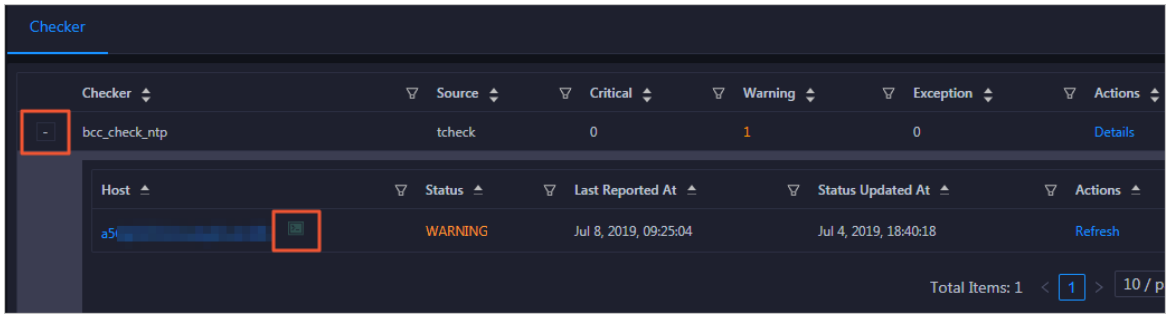

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

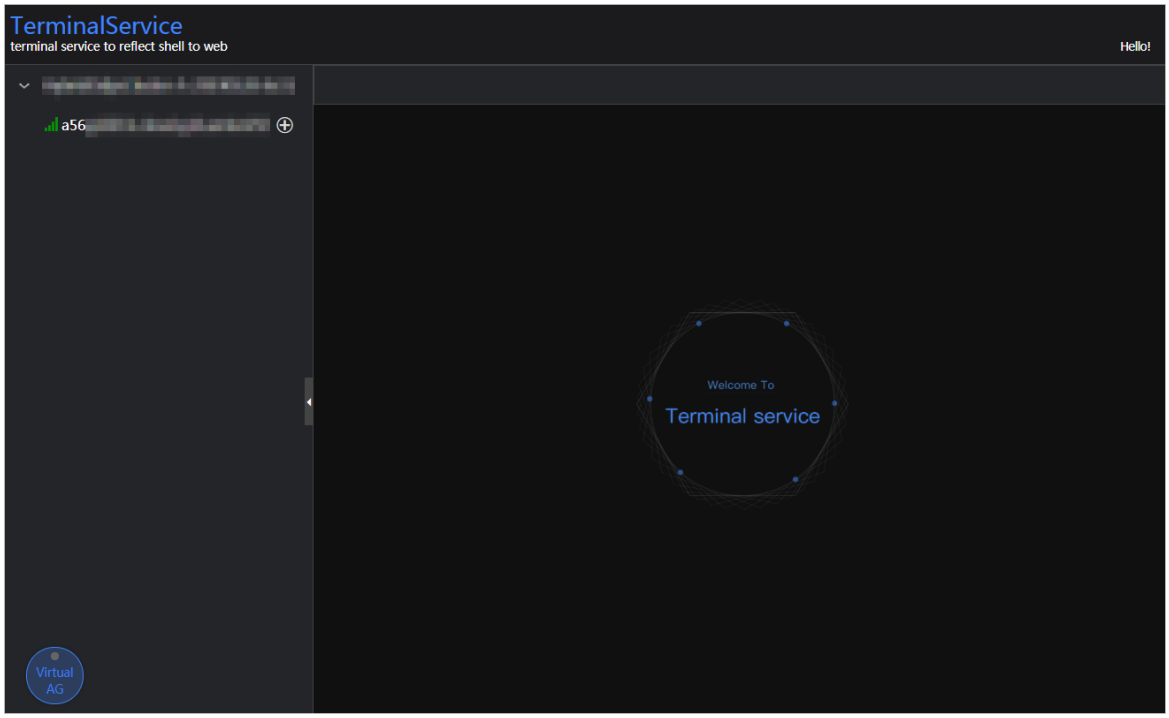

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

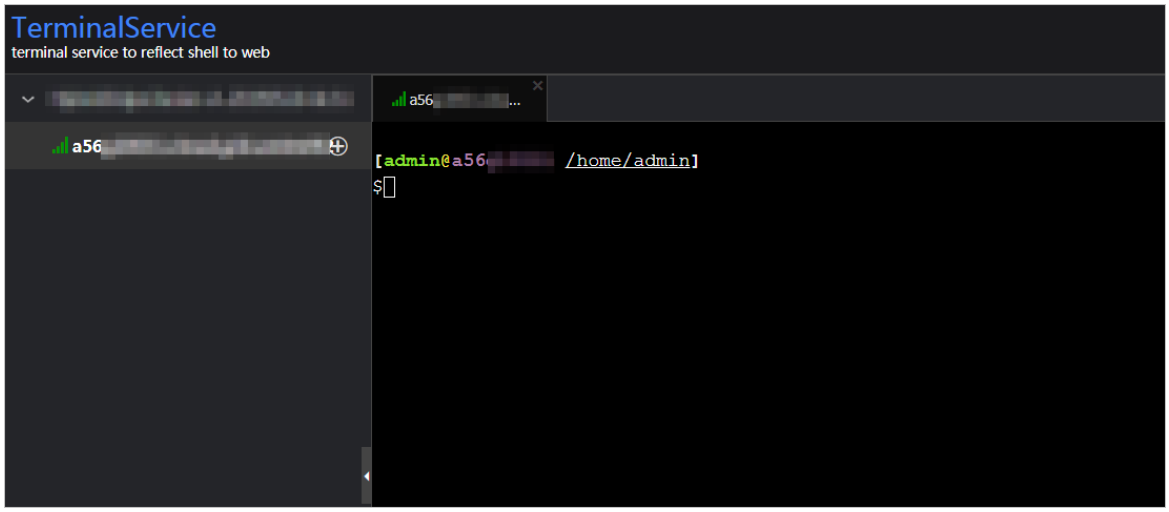

Run a checker again

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.

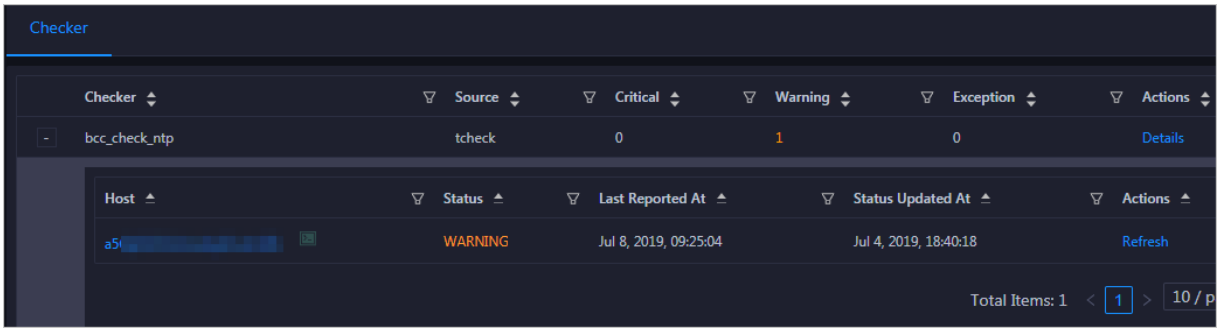

11.7. Graph Analytics

11.7.1. Operations and maintenance tools and logon methods

11.7.1.1. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   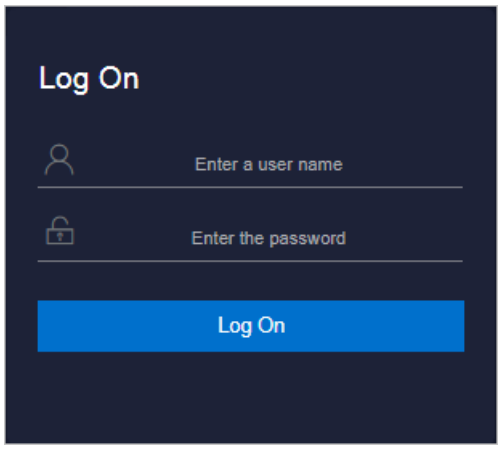

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.

- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.7.1.2. Log on to Apsara Infrastructure Management Framework

This topic describes how to log on to Apsara Infrastructure Management Framework. Apsara Infrastructure Management Framework supports operations and maintenance (O&M) management for Graph Analytics.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

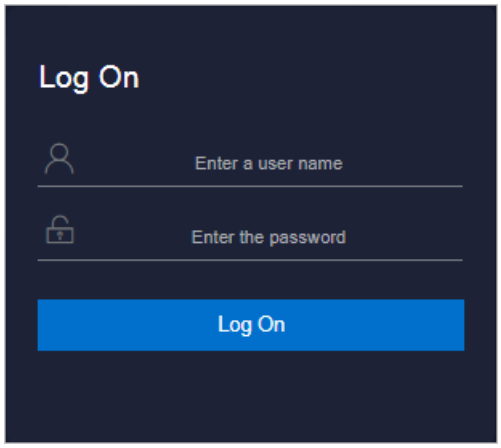

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Apsara Infrastructure Management Framework >** to redirect to the homepage of Apsara Infrastructure Management Framework.

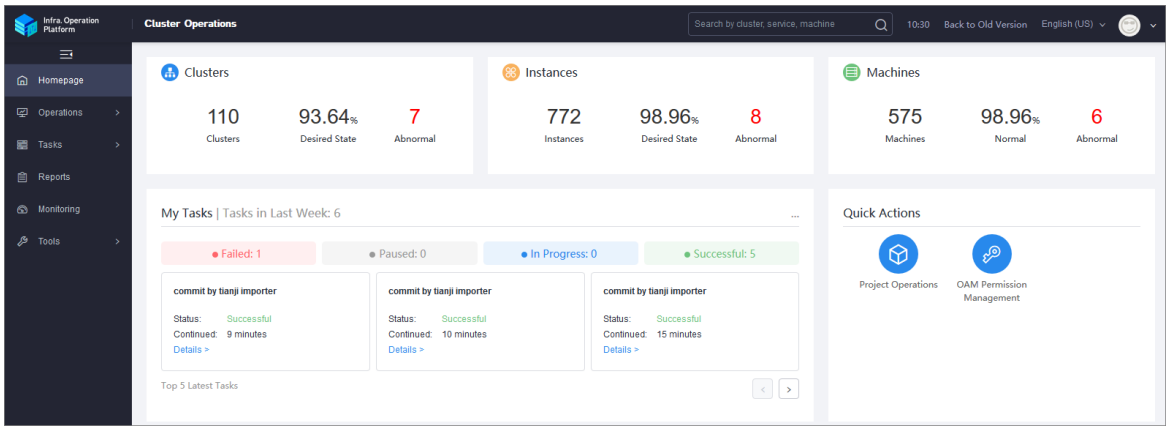

11.7.1.3. Log on to the Graph Analytics container

You can log on to the Graph Analytics container through Apsara Infrastructure Management Framework to perform operations and maintenance.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side **Project** drop-down list, enter or select *iplus* to display Graph Analytics clusters.
3. Select a Graph Analytics cluster. On the **Services** page, double click **iplus-iplus_biz > IplusBizBackend#**. Click the More icon next to **vmxxxxxxxxxxxx** and then select **Terminal** in the menu that appears. The **TerminalService** page appears.

   Operations and maintenance are typically performed on the virtual machines of **IplusBizBackend#** and **IplusBizBackendControl#**. You can use the same method to open the virtual machine where the **IplusBizBackendControl#** service is deployed.

   TerminalService page

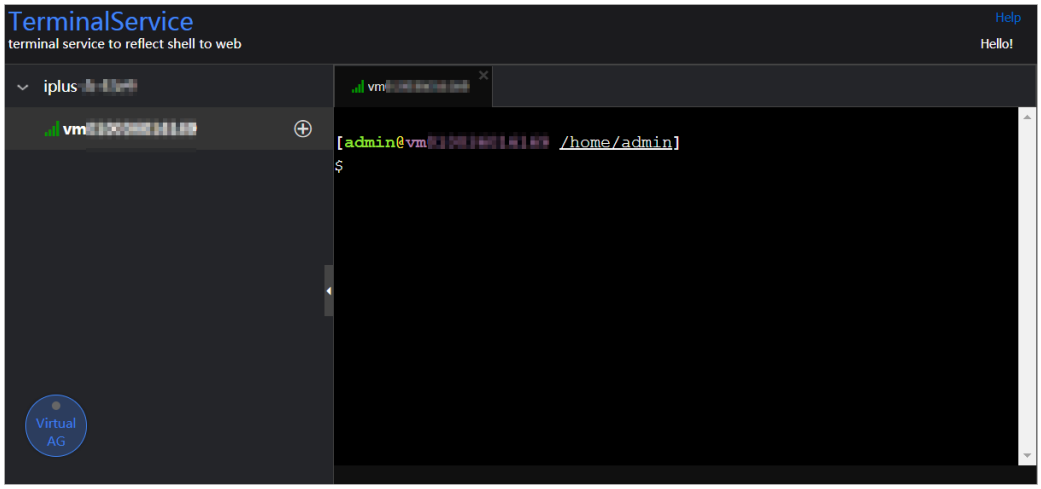

The left-side navigation pane on the **TerminalService** page displays the virtual machine selected by you (**vmxxxxxxxxxxxx**).

4. In the left-side navigation pane on the **TerminalService** page, click **vmxxxxxxxxxxxx**, and the command-line tool appears on the right-side of the page.
5. Run the **docker ps|grep** *iplus* command to query the docker ID in the Graph Analytics cluster.

   Query the docker ID

```
[admin@vm /home/admin]
$docker ps|grep iplus
bc000          reg.docker./ice_images/frontend:aeed8a997e1
508810471f302b              "/bin/sh -c 'sudo sh "   2 days ago          Up 2 days
s_biz.IplusBizFrontend__.iplus-biz-frontend.1531202126
ad97d          reg.docker./imore/imore_background_service:
c05393b231d593a656f28ccd6   "sh /home/admin/ccbin"   3 days ago          Up 3 days
s_biz.IplusBizBackend__.iplus-biz-backend.1531153975
```

   The query results of this sample display two docker IDs, which indicates that the **IplusBizBackend#** service is running on two containers.

6. Run the **docker exec -ti** *dockerID* **bash** command to log on to the docker container.

   Enter the docker ID of the container you need to log on to in *dockerID*.

   Log on to the docker container

```
[admin@vm01(  .69 /home/admin]
$docker exec -ti ad97d  bash

[root@docker01  71 /home/admin]
```

7. The root account is used by default. You can use the **su - admin** command to switch to the admin account.

   Switch to the admin account

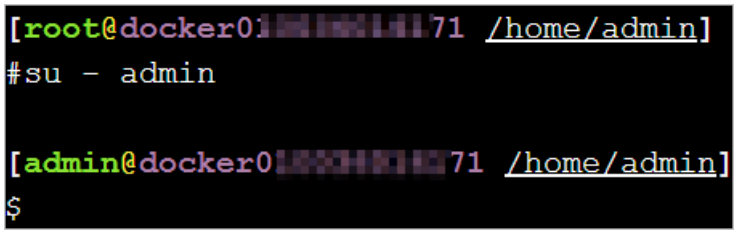

11.7.2. Operations and maintenance

11.7.2.1. Operations and maintenance based on BigData Manager

11.7.2.1.1. O&M overview

This topic describes the features of H+ O&M and how to access the H+ O&M page.

Modules

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

Module	Feature	Description
Services	Service overview	Displays the trend charts of CPU utilization, 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	Cluster overview	Displays the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	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	Host 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 utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.
	Host health	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 **I+**.
3. On the page that appears, click **O&M** in the top navigation bar. The **Services** page appears.

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

11.7.2.1.2. Service O&M

11.7.2.1.2.1. Service overview

The service overview page lists all I+ 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, search for a cluster in the search box 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.

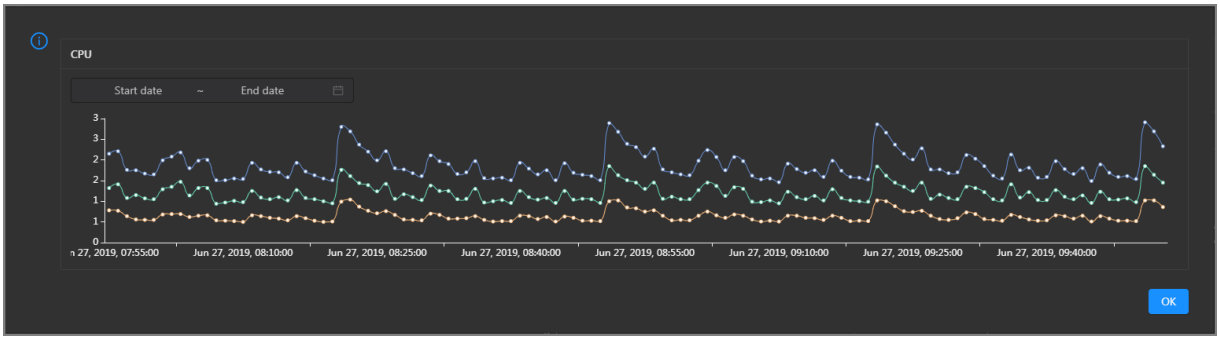

DISK

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.

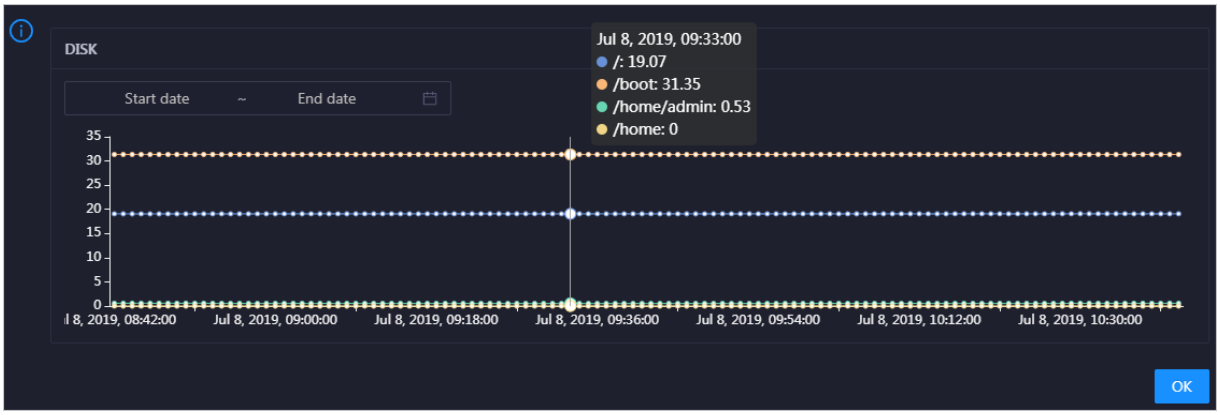

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

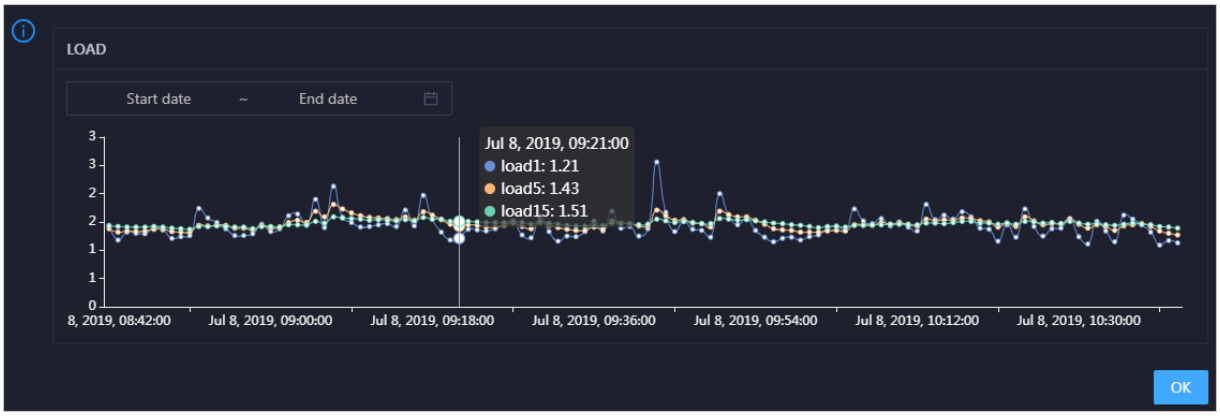

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

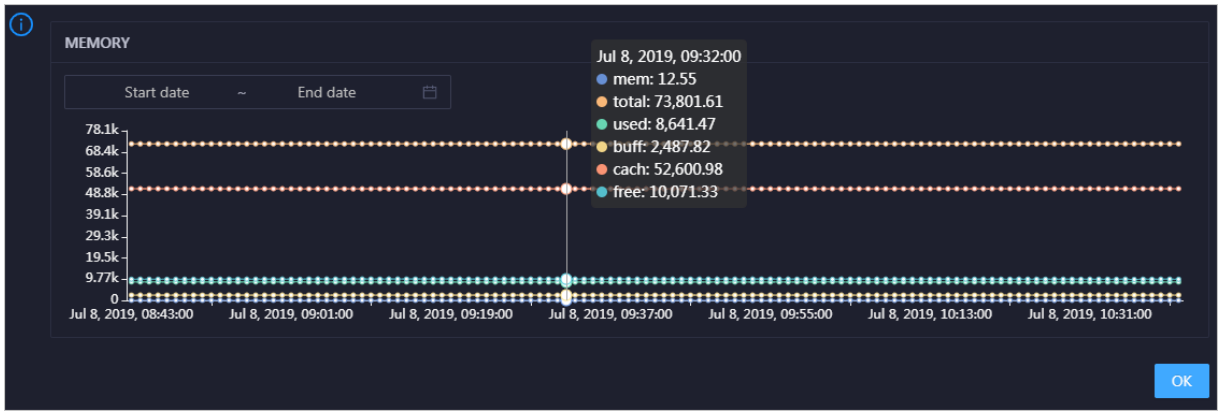

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

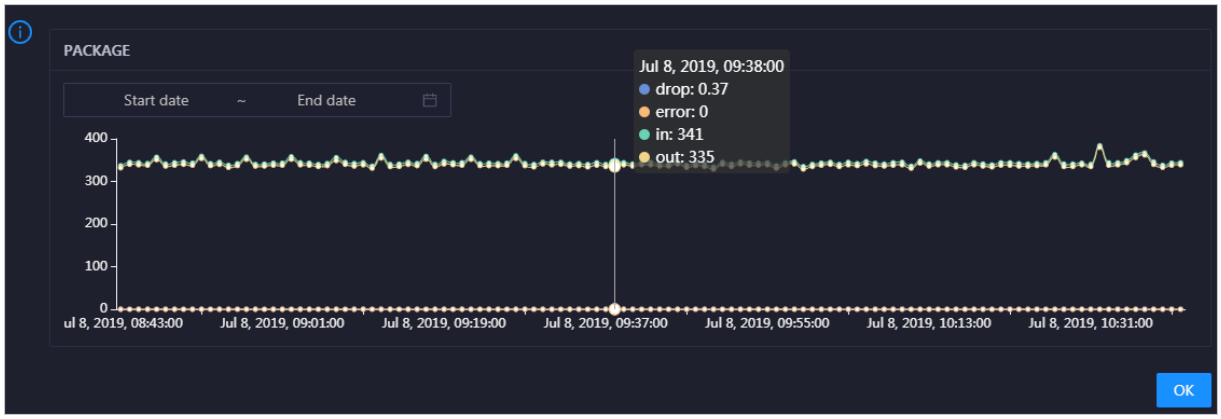

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

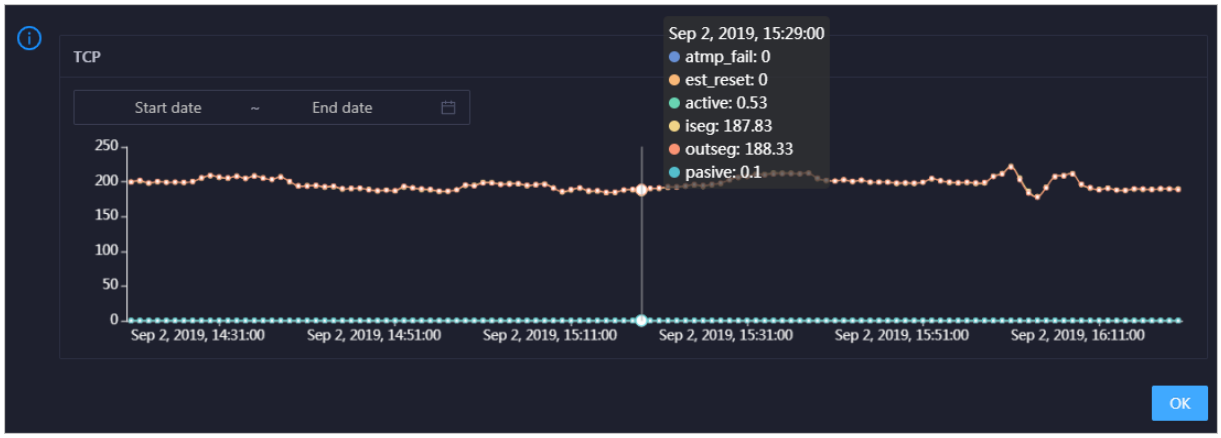

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

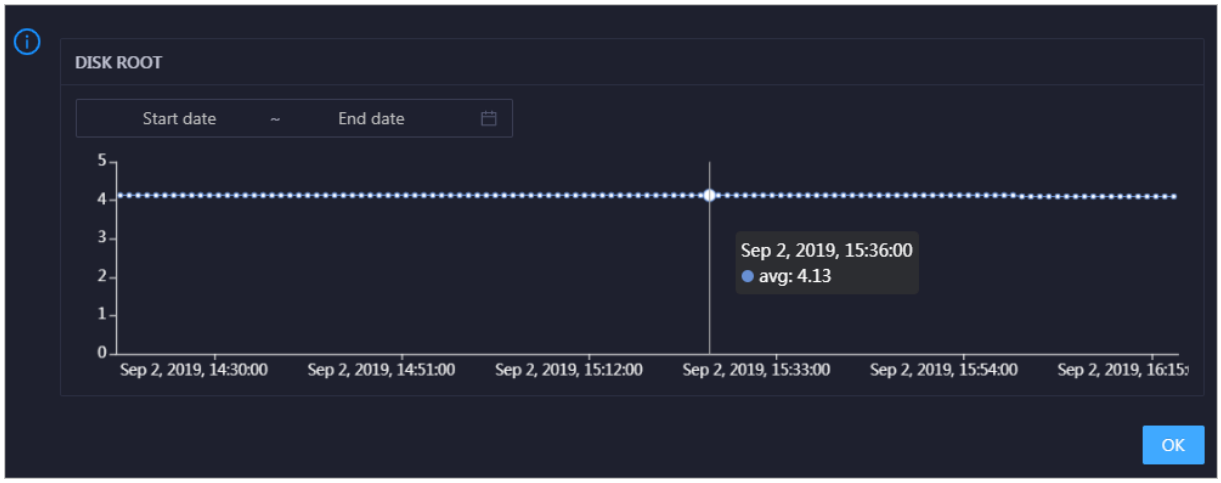

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

11.7.2.1.2.2. Service hosts

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

On the **Services** page, search for a cluster in the search box 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.

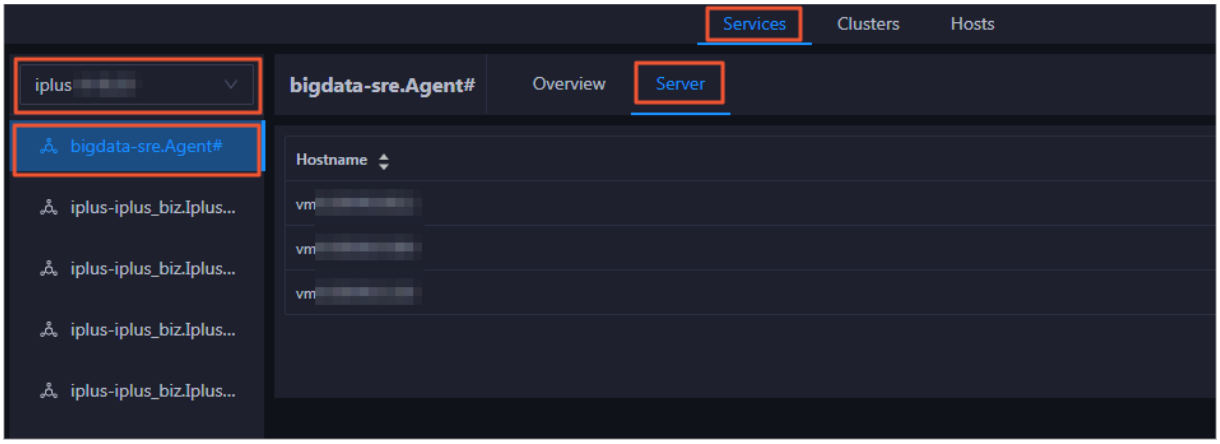

On the **Server** page, you can view the hosts where the selected service is run.

11.7.2.1.3. Cluster O&M

11.7.2.1.3.1. Cluster overview

The cluster overview page displays the trend charts of CPU utilization, 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.

CPU

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

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

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

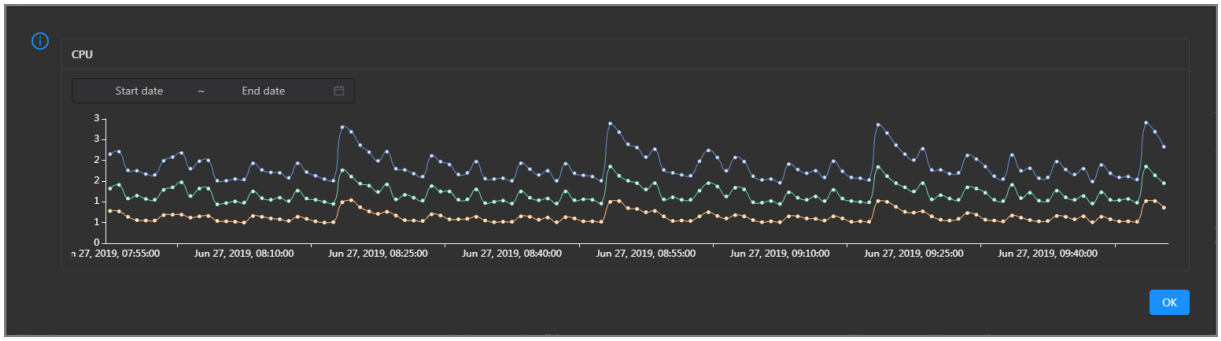

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

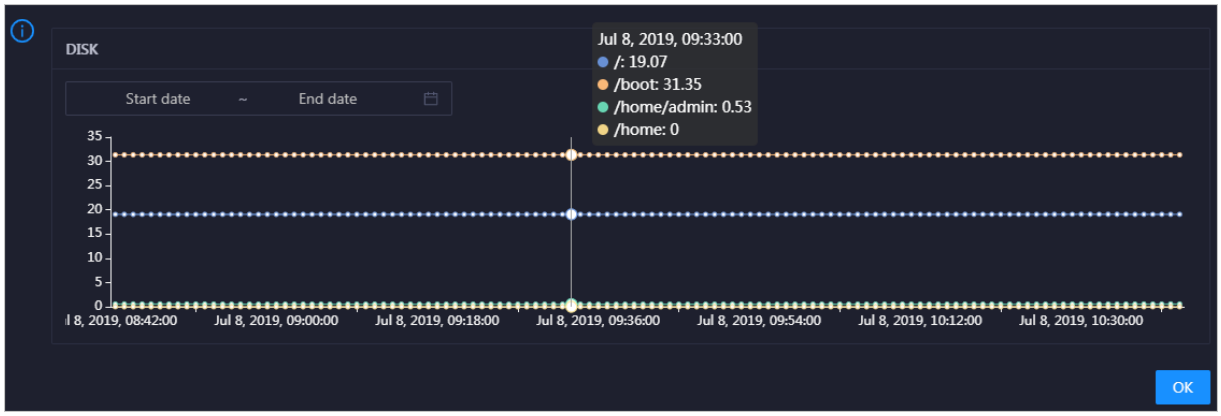

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

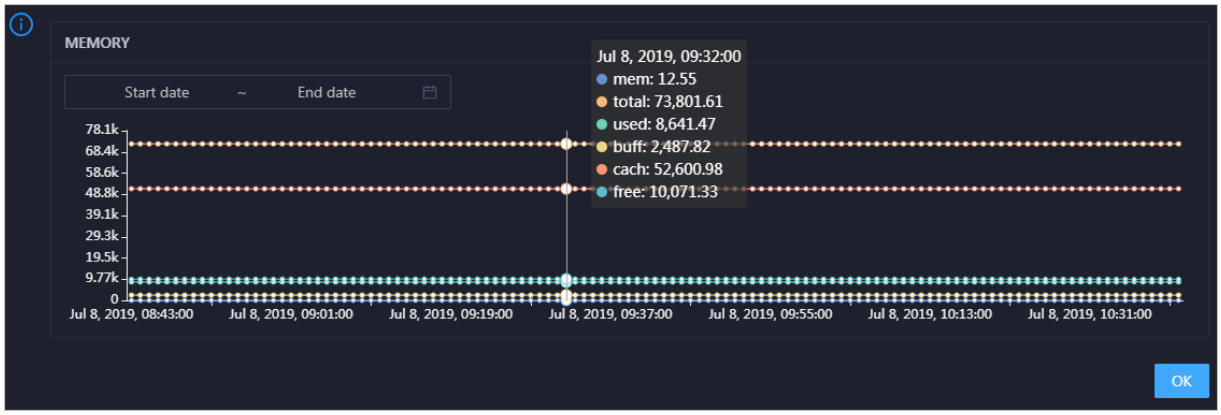

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

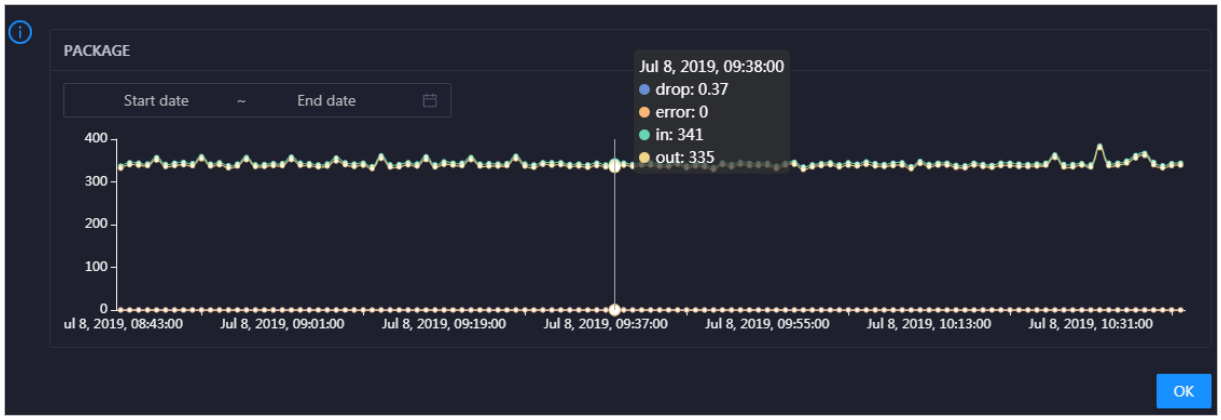

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

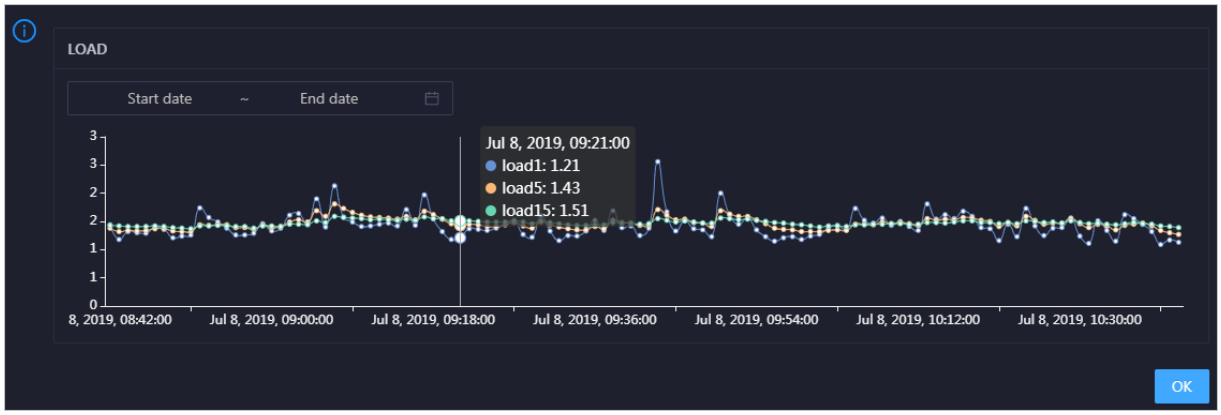

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.7.2.1.3.2. 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

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

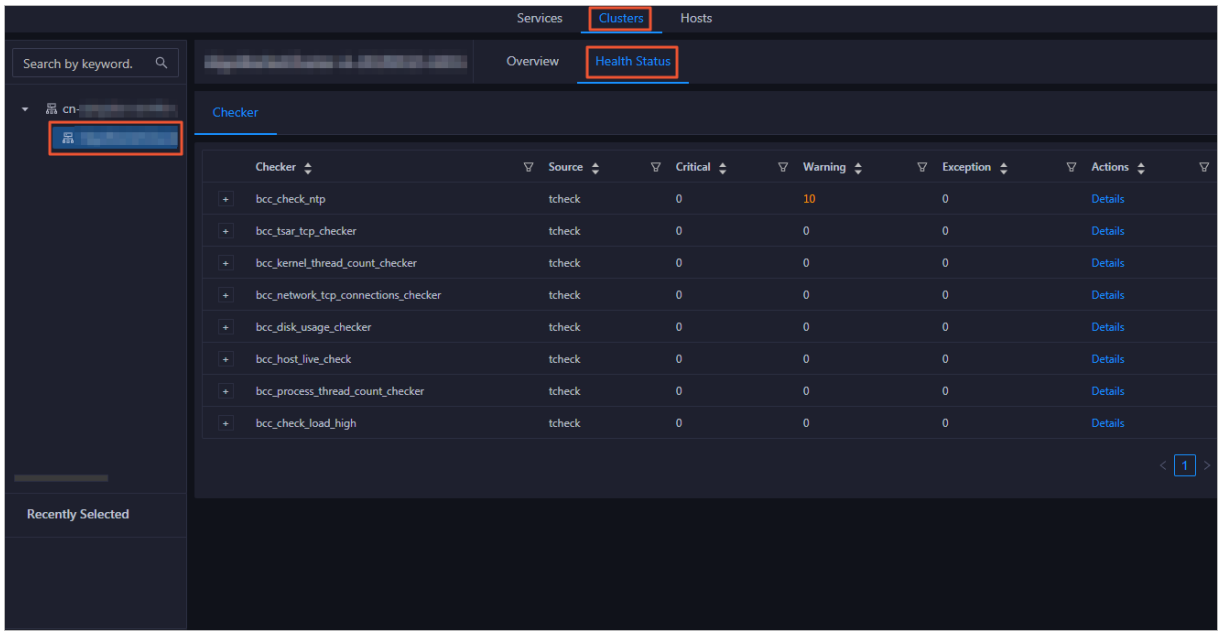

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are repesented in different colors. Handle the alerts in a timely manner,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

View checker details

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

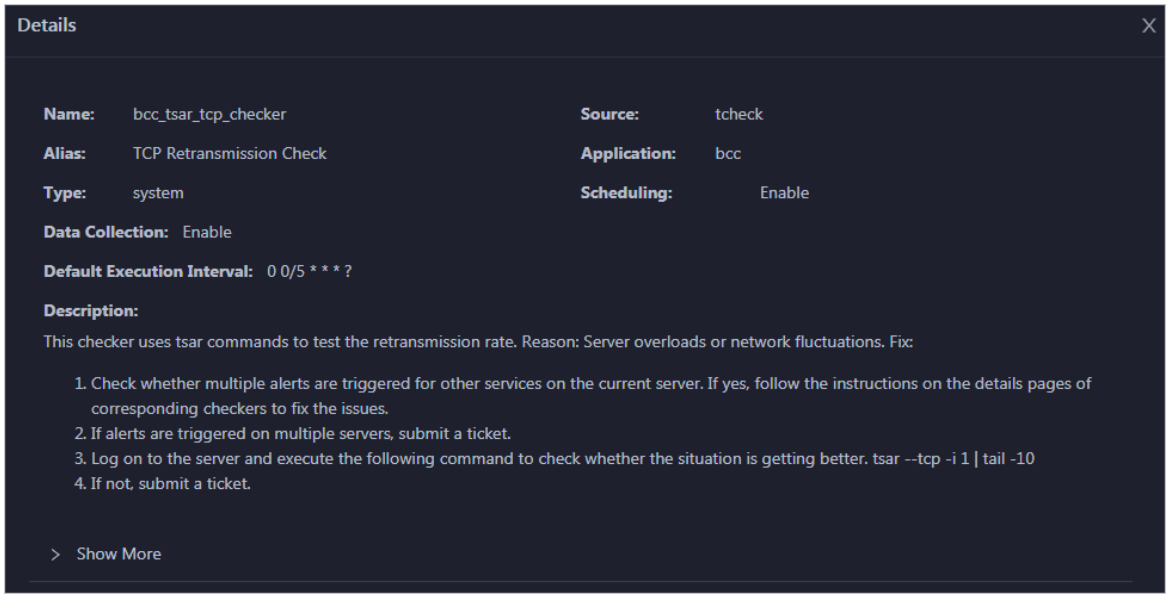

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

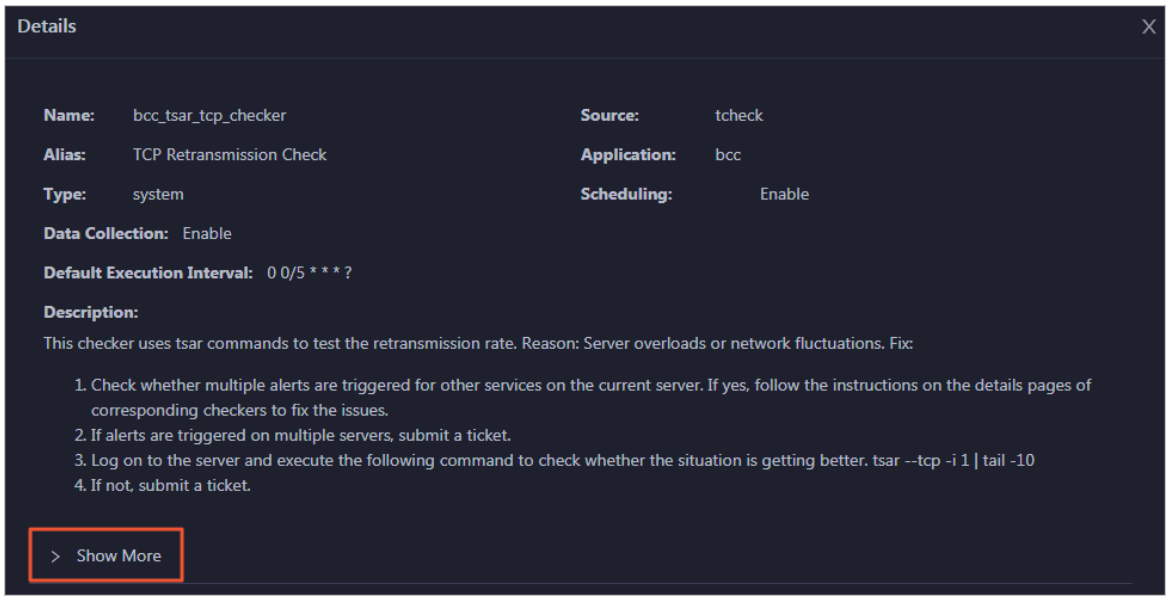

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

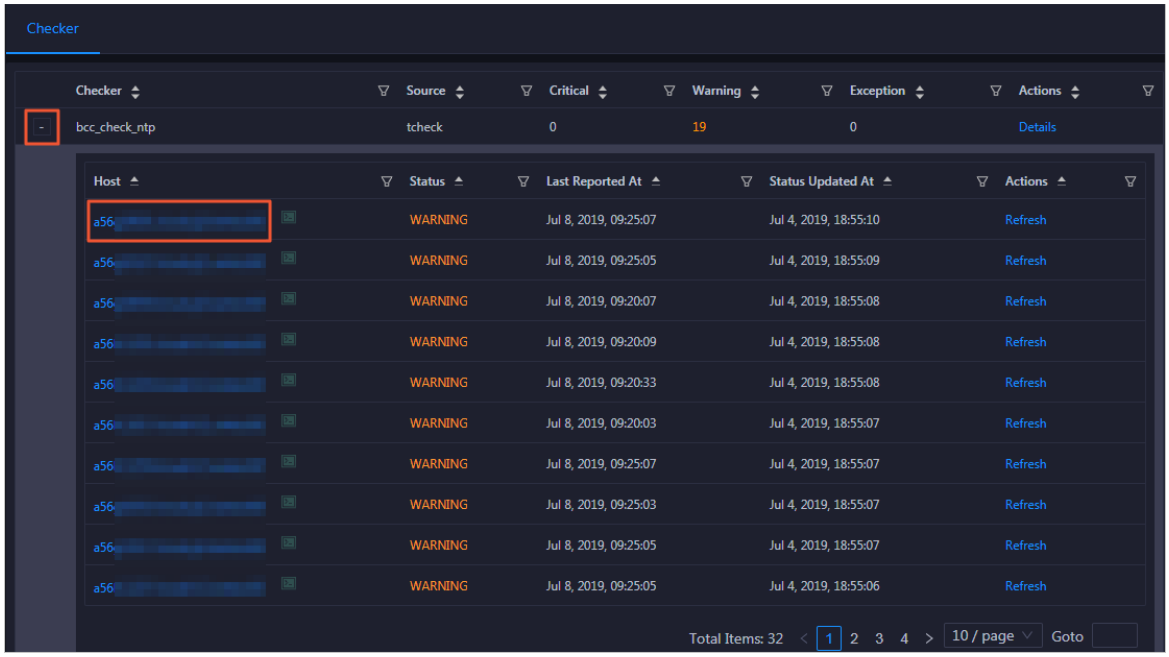

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

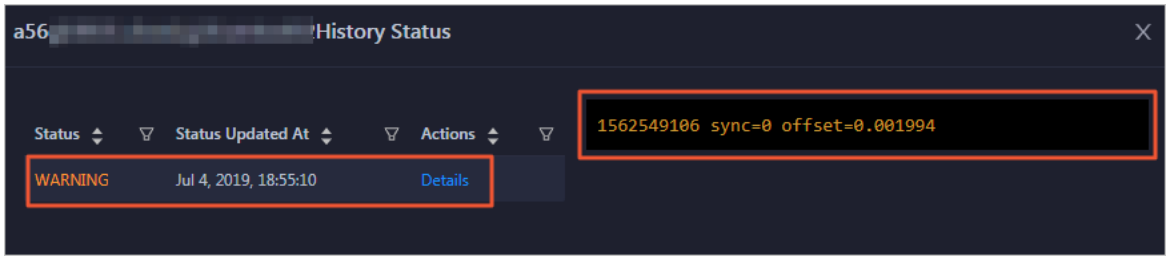

Clear alerts

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

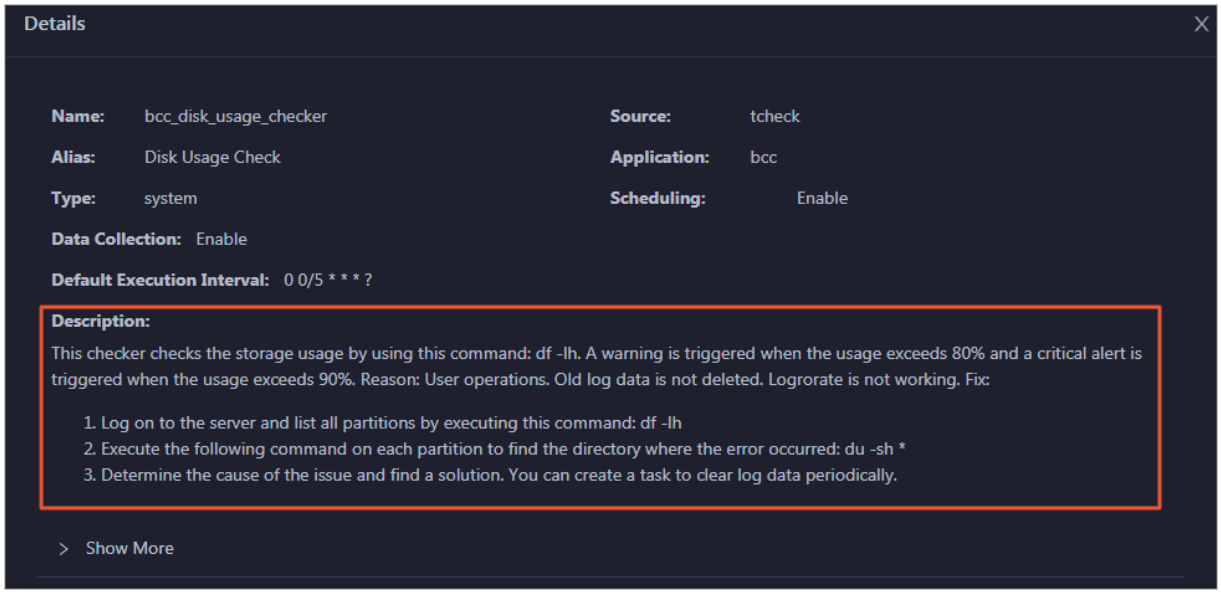

Log on to a host

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

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

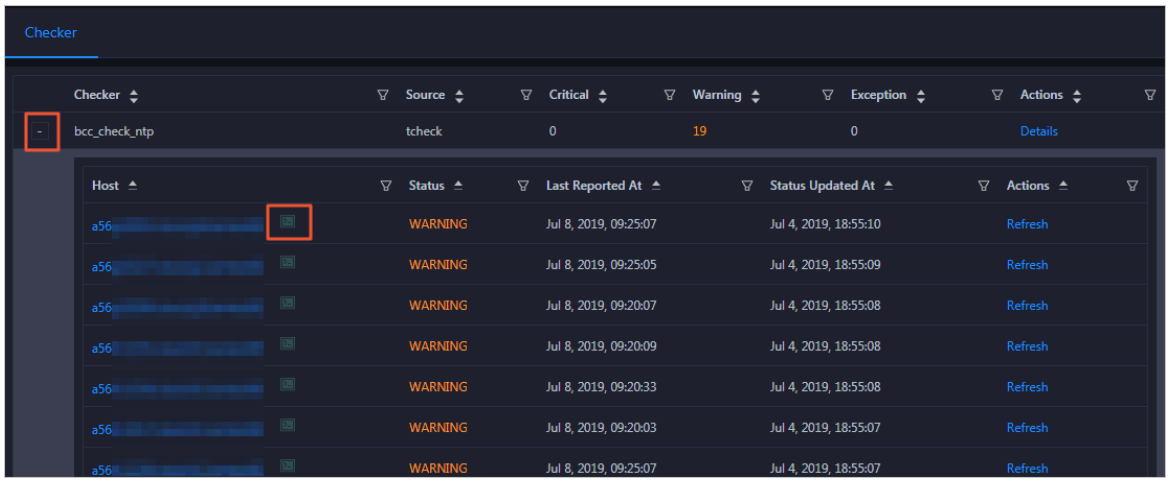

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

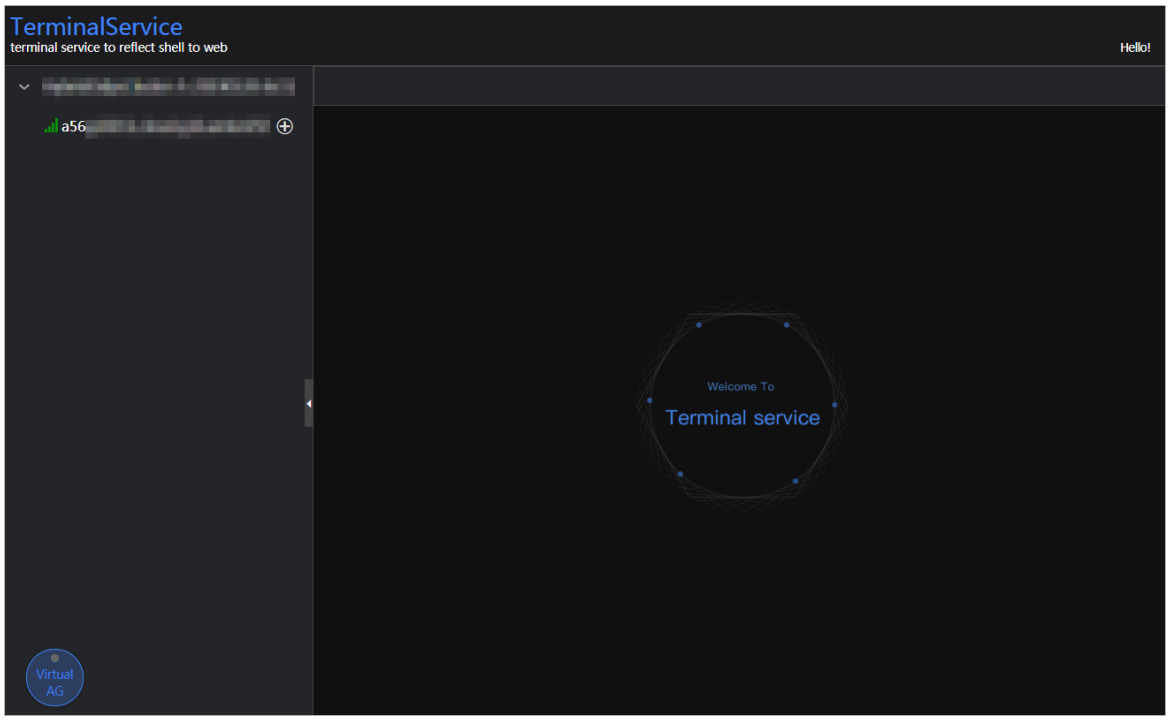

3. On the **TerminalService** page, click the hostname to log on to the host.

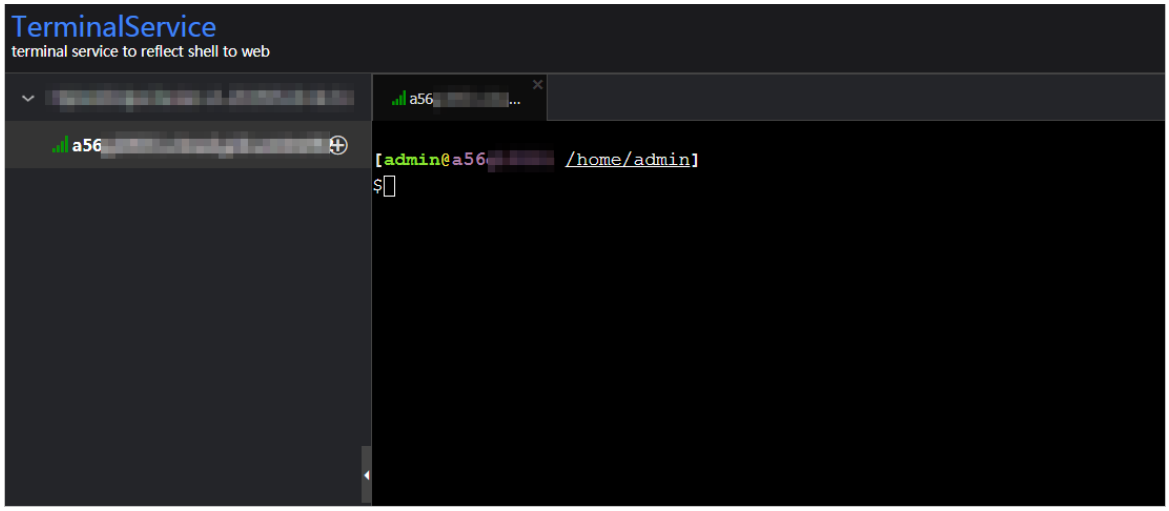

Run a checker again

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

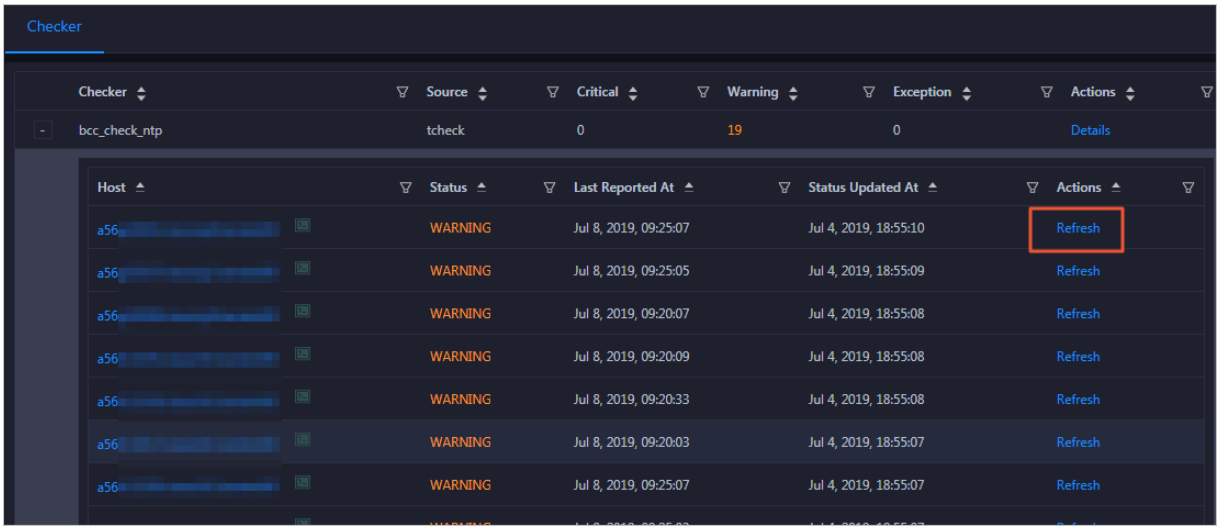

11.7.2.1.4. Host O&M

11.7.2.1.4.1. Host overview

The host overview page displays the overall running information about a host in an I+ 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

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

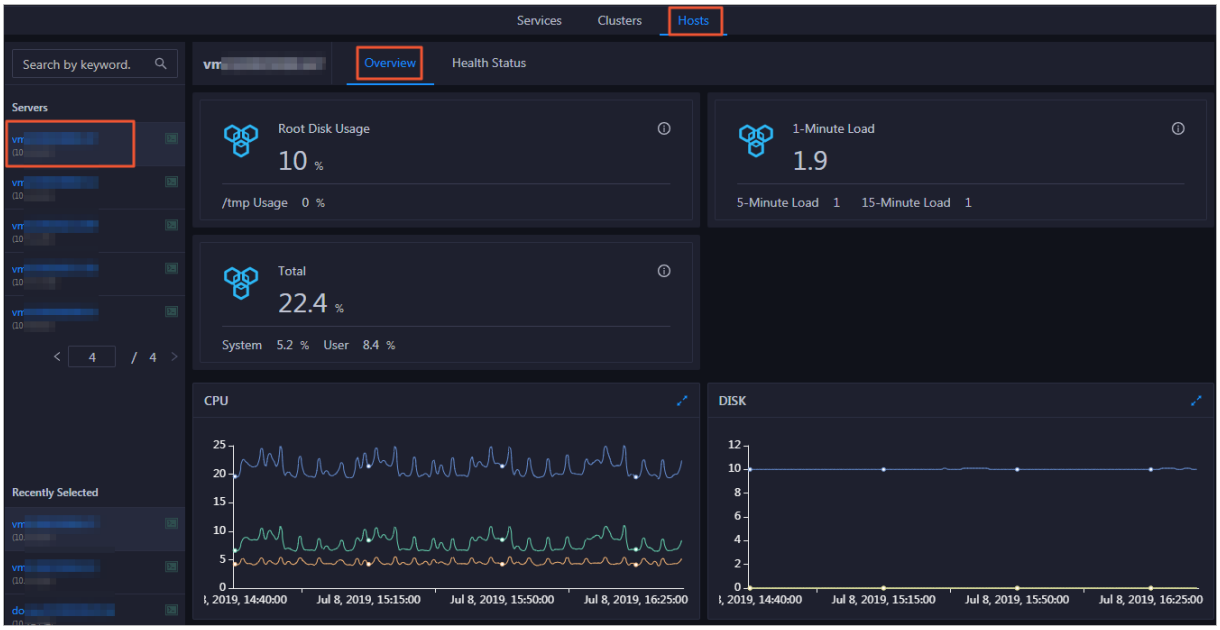

On the **Overview** page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root 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.

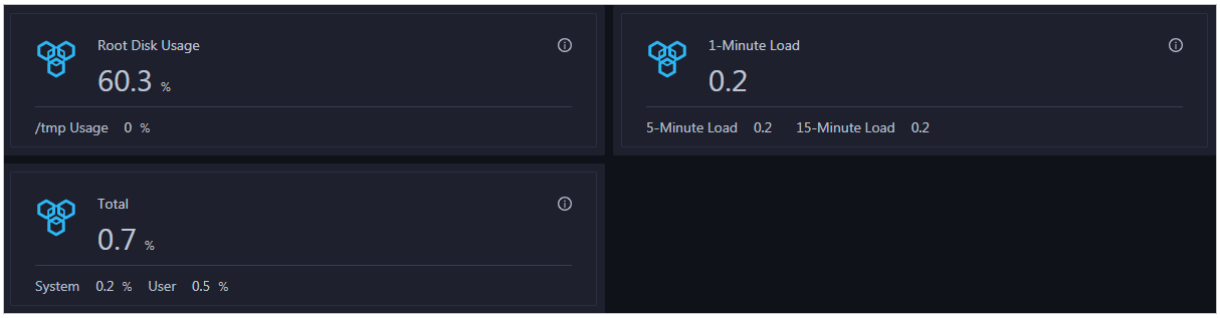

CPU

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.

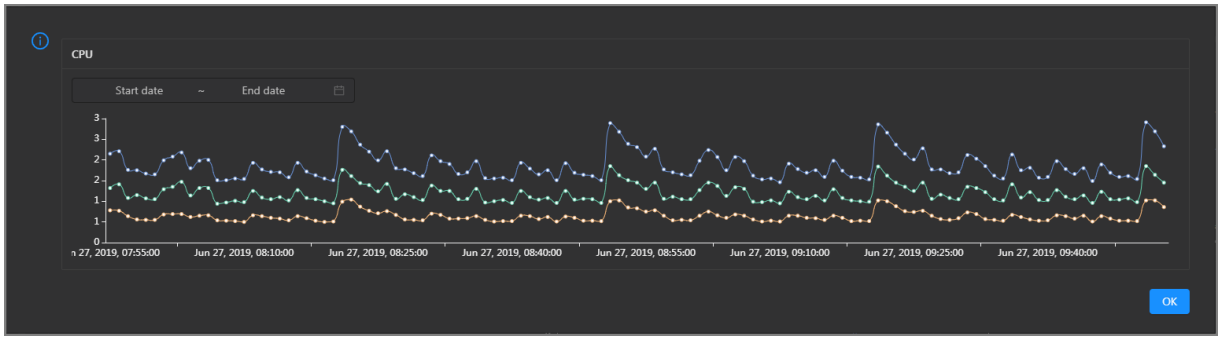

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

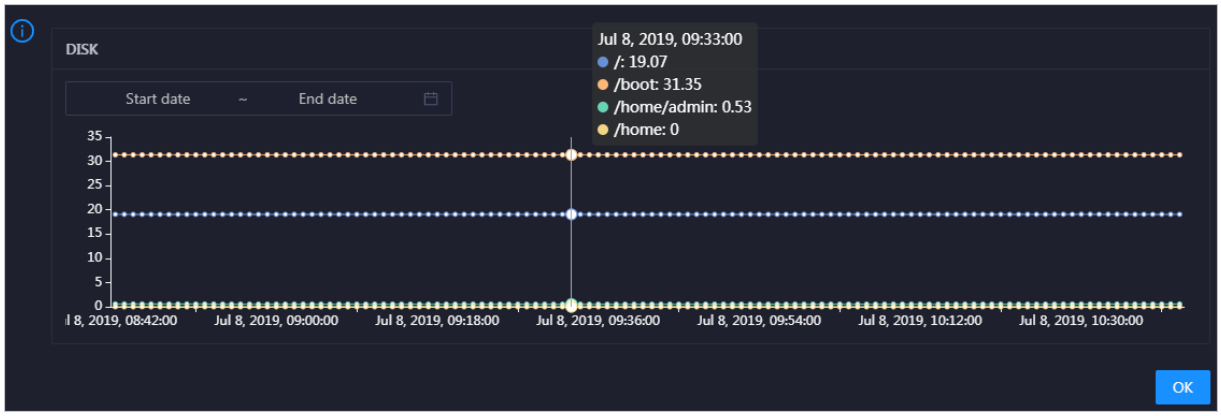

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

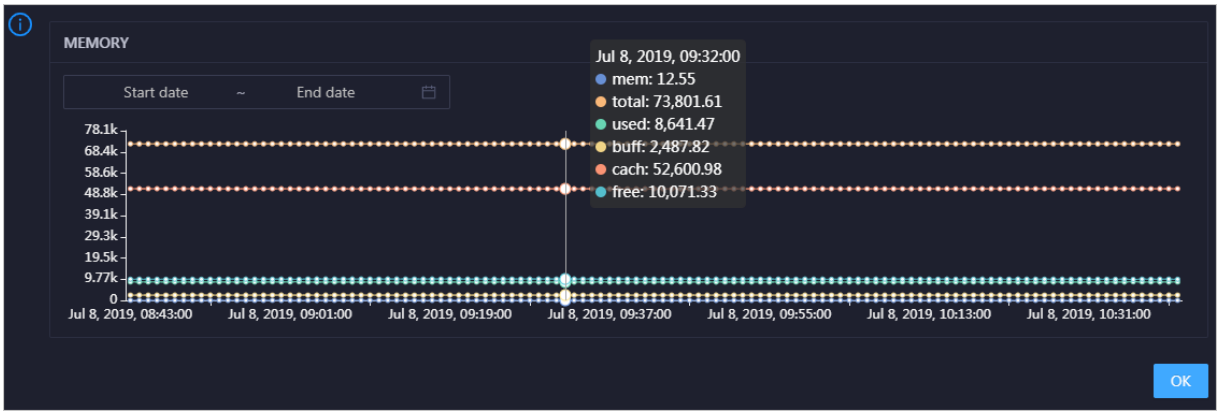

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

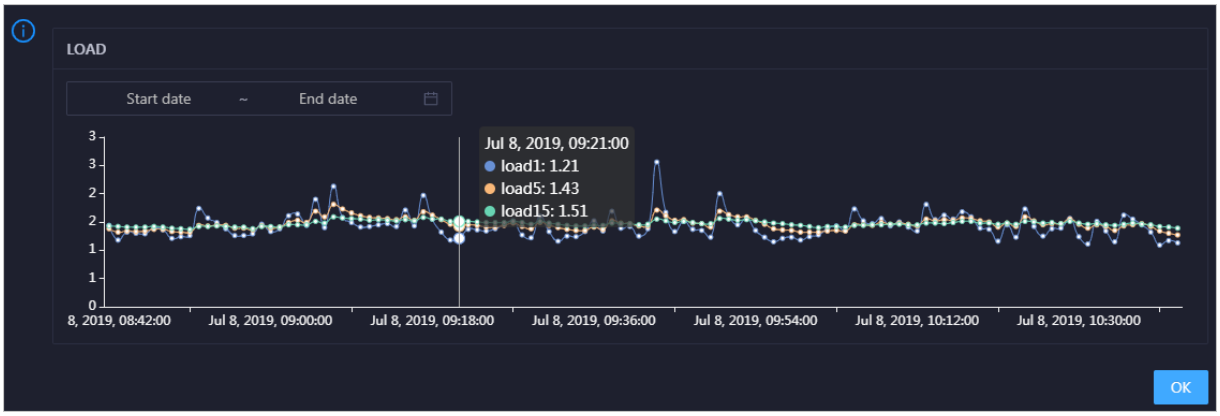

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

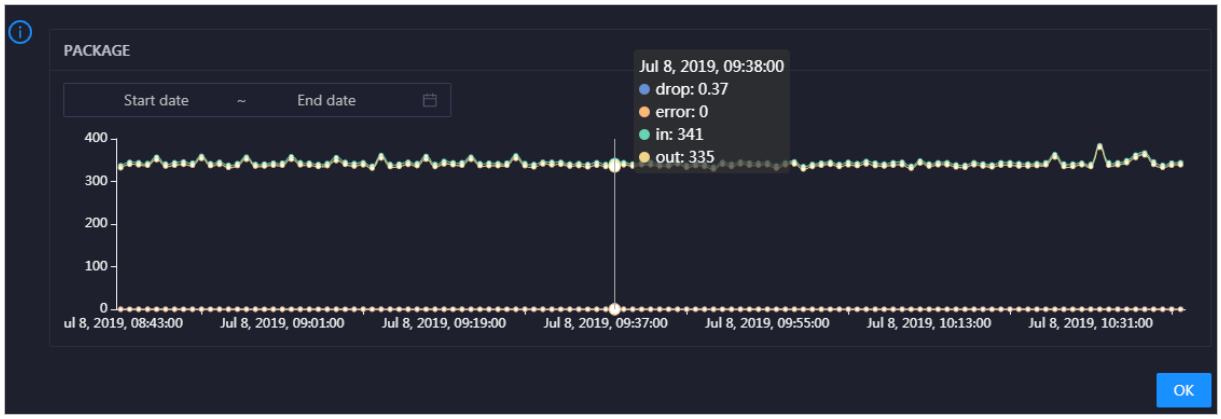

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

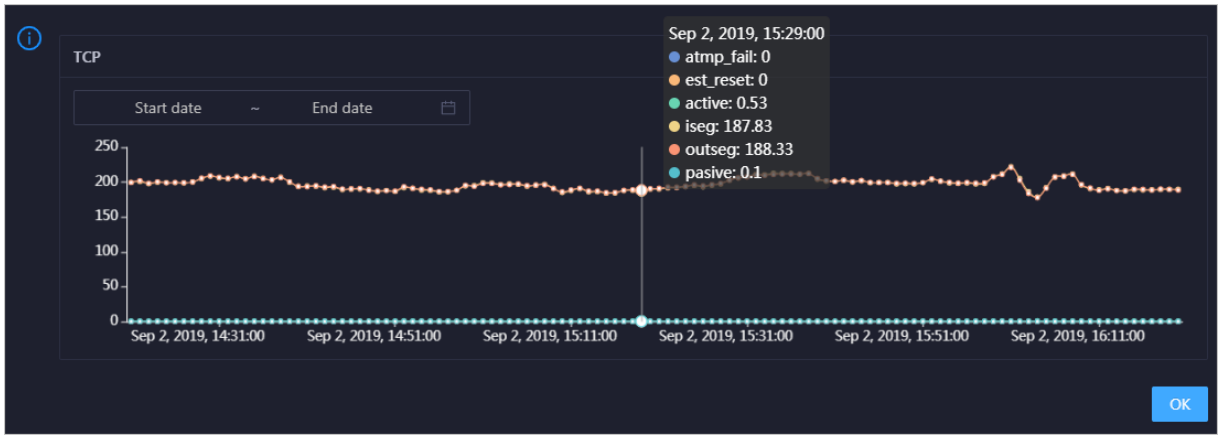

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

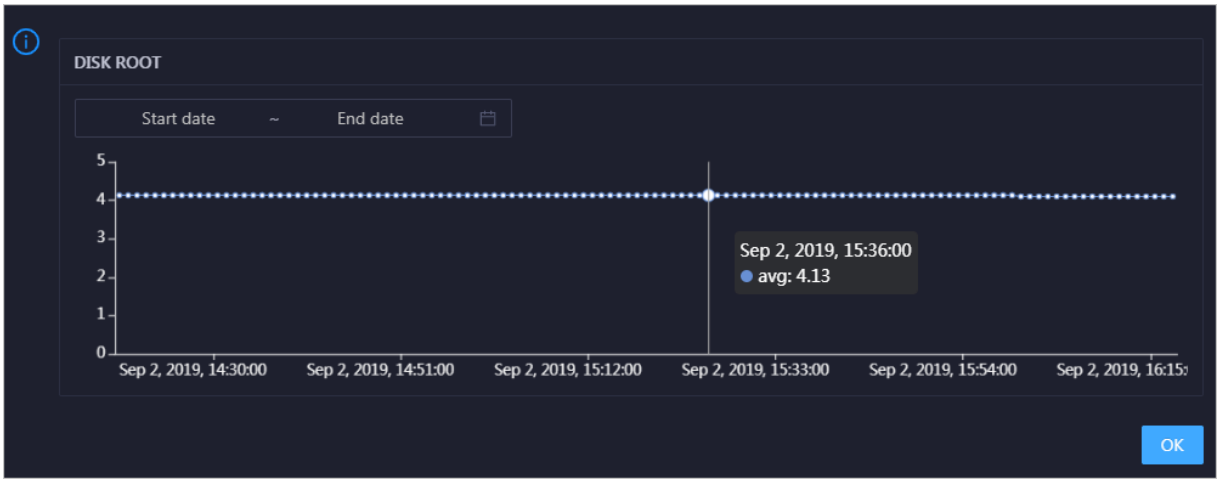

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

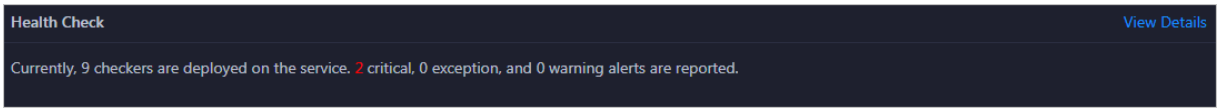

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

Health Check History

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

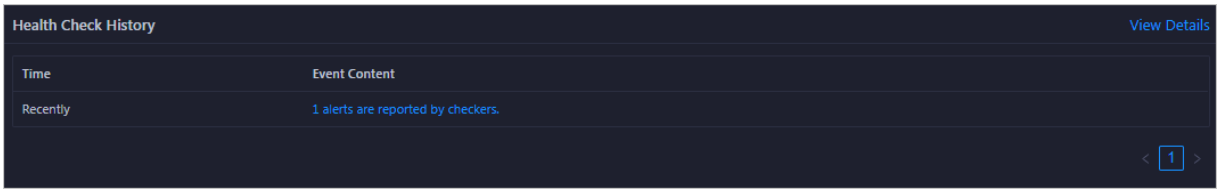

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

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

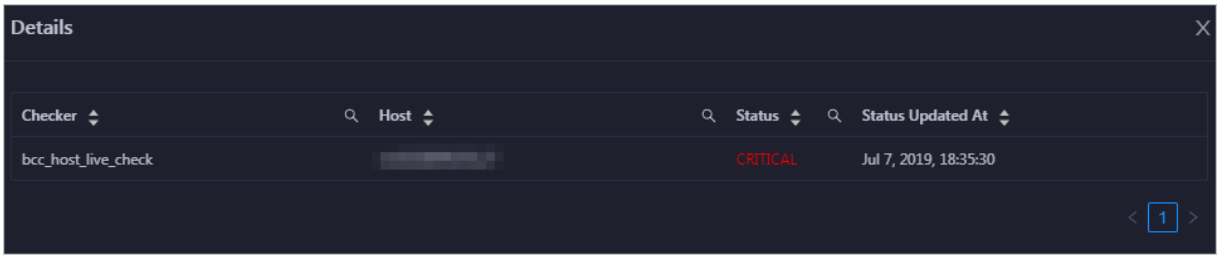

11.7.2.1.4.2. Host health

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

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.

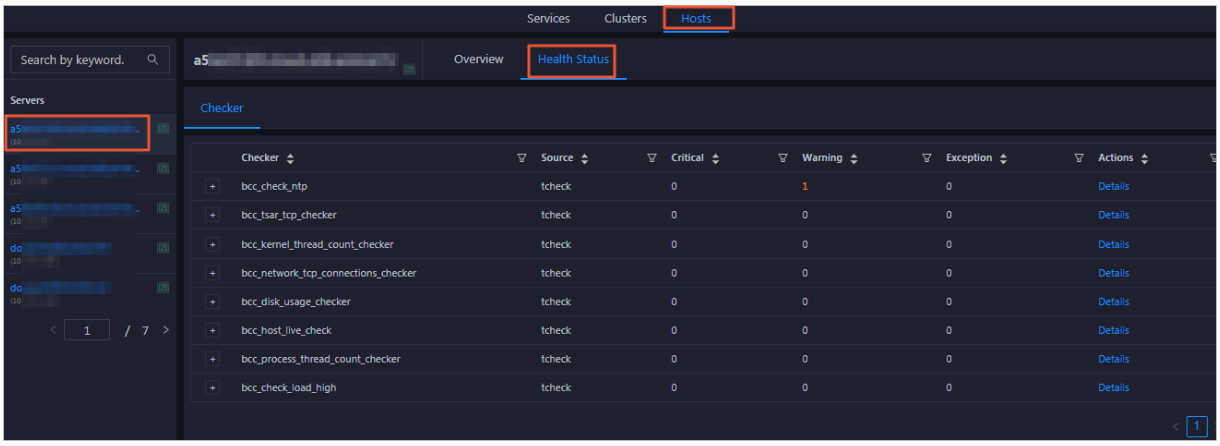

On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

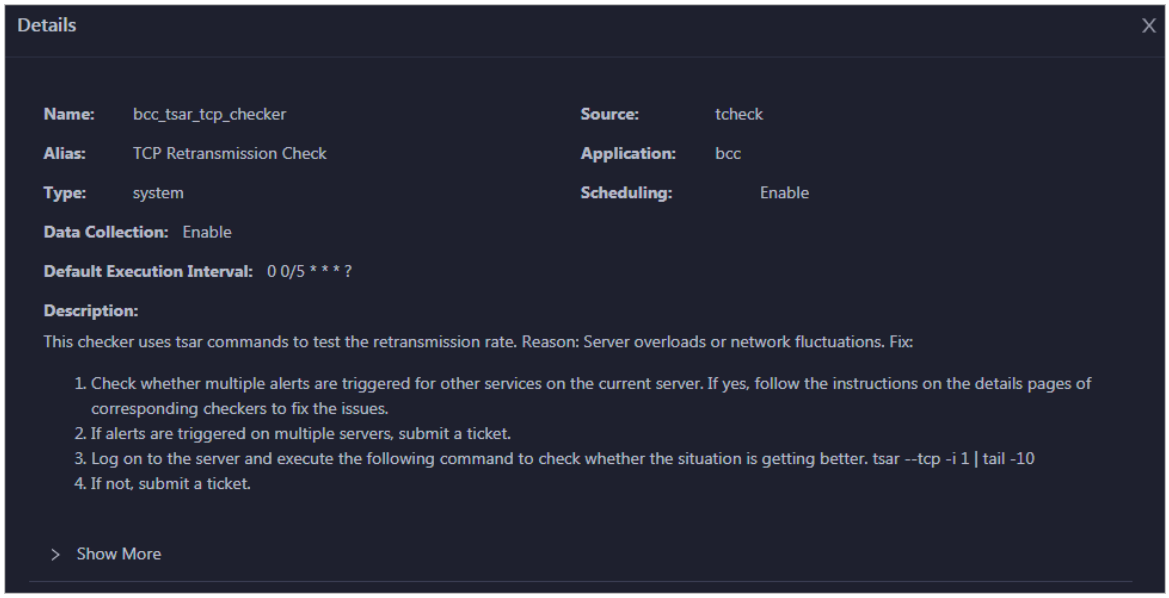

   The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

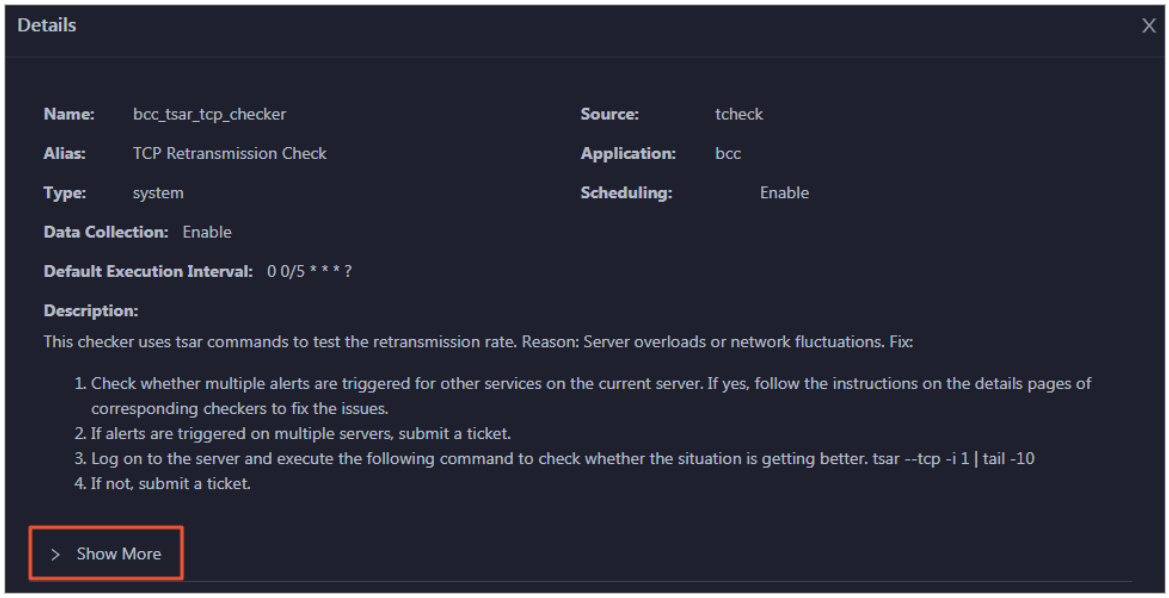

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

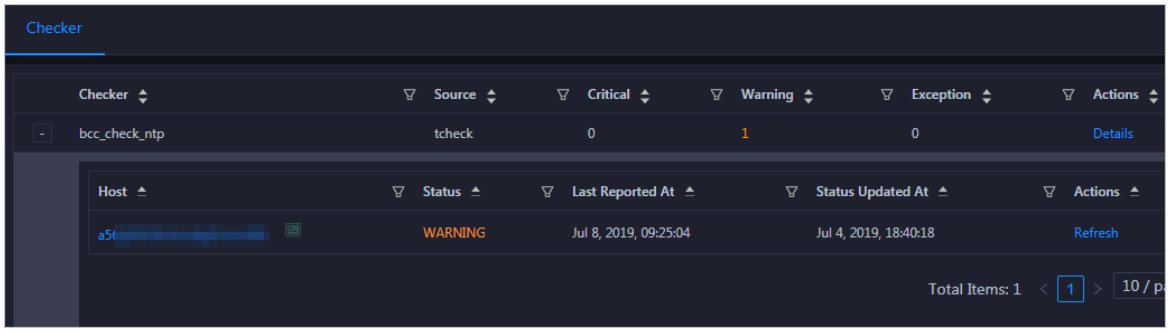

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

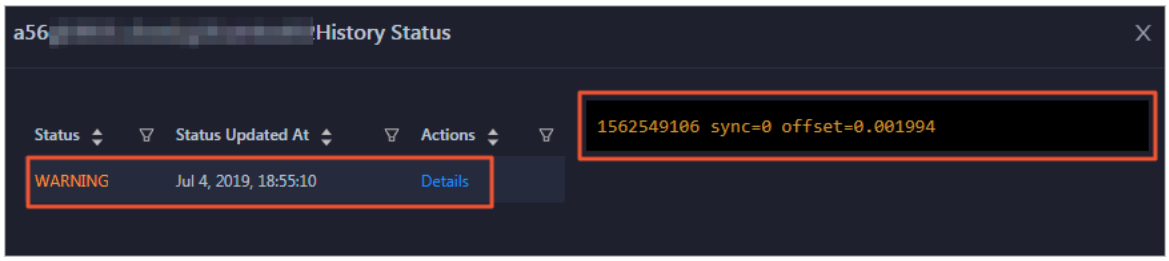

Clear alerts

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.

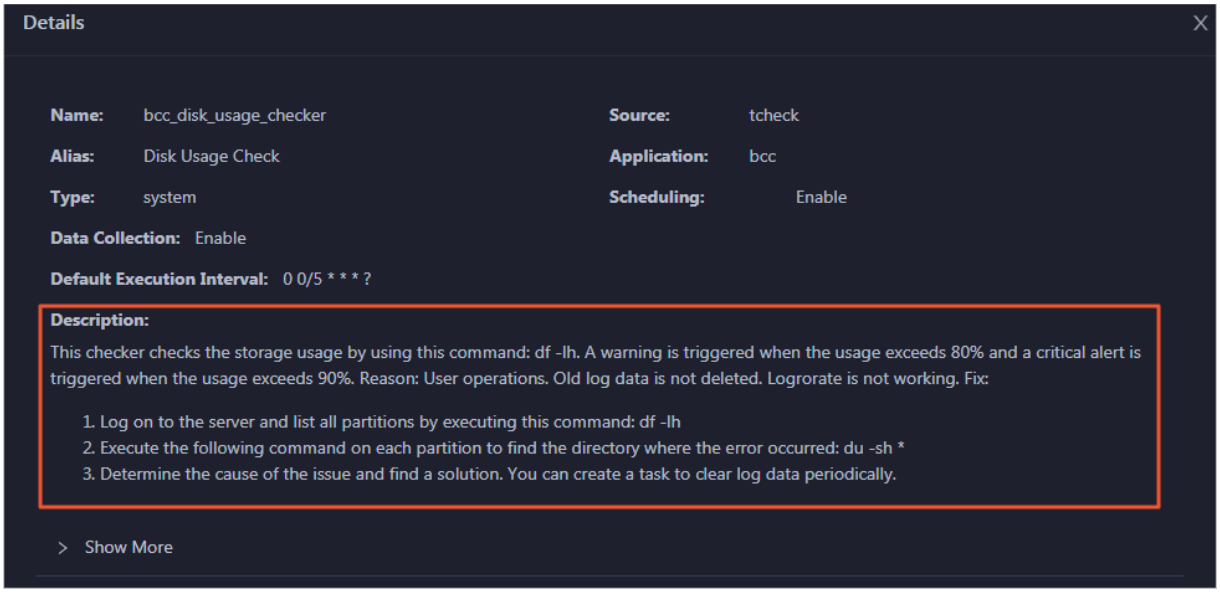

Log on to a host

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.

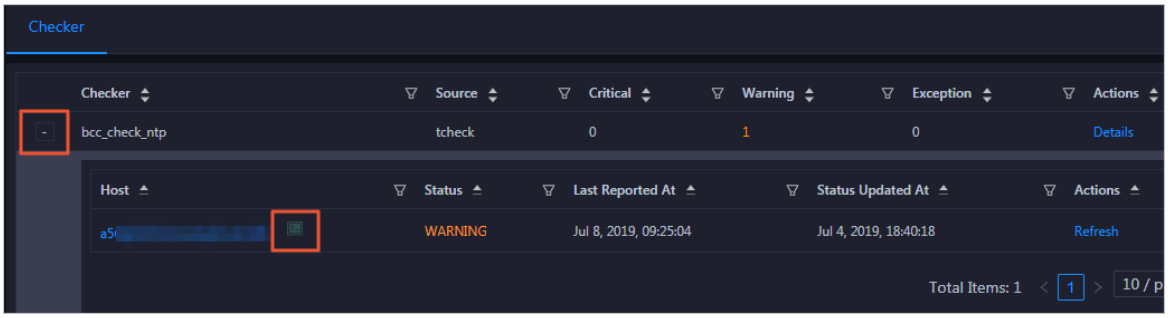

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

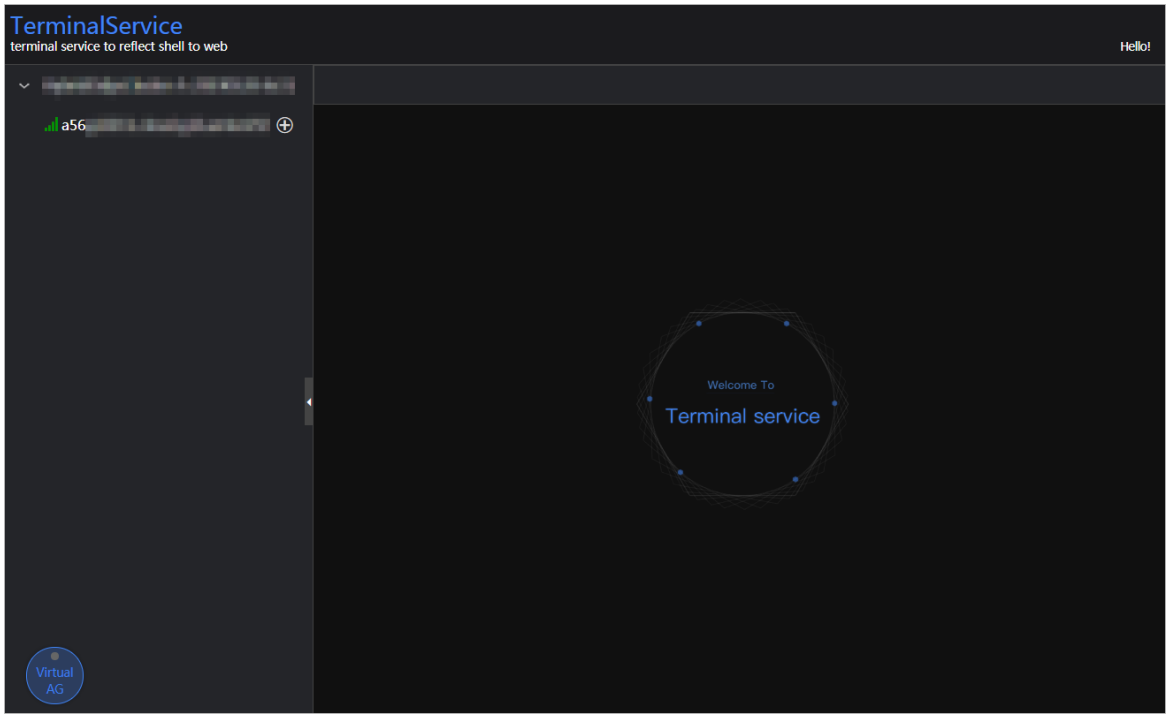

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

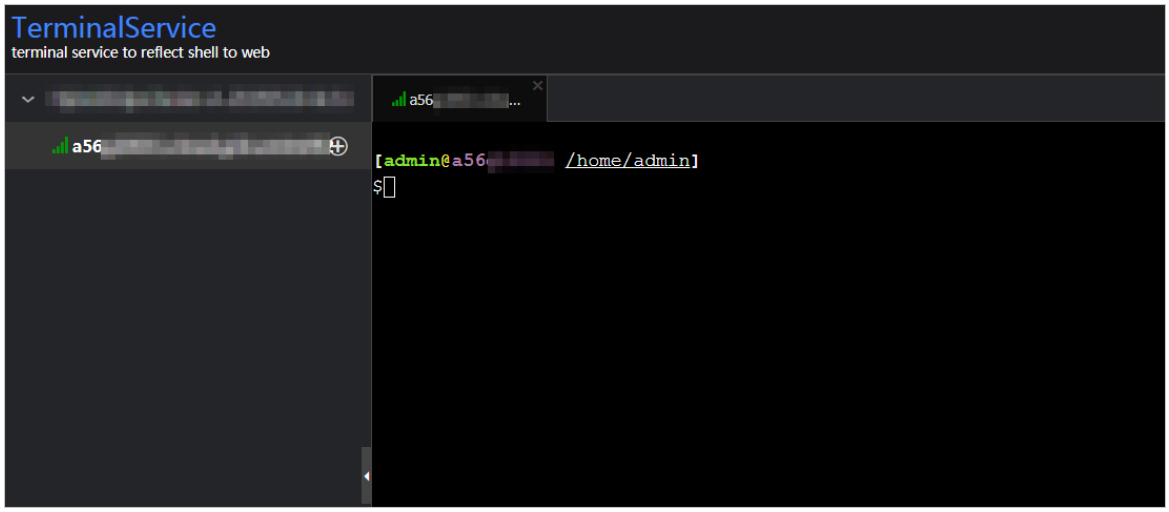

Run a checker again

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.

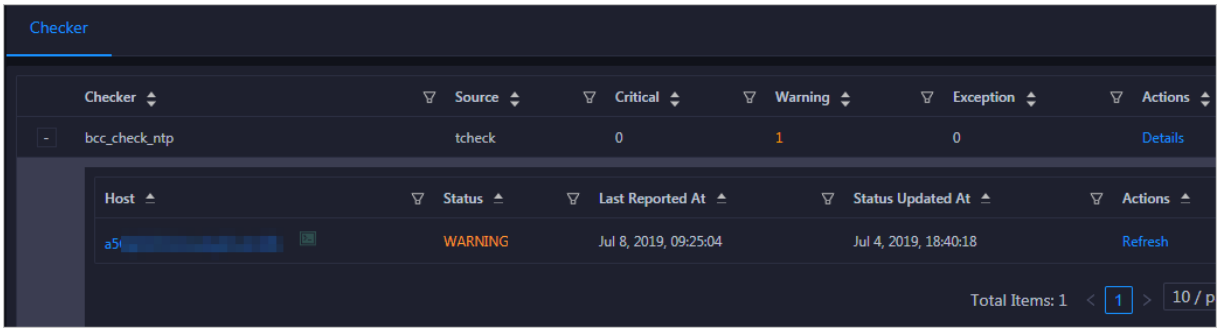

11.7.2.2. O&M on Apsara Infrastructure Management Framework

1. Log on to Apsara Infrastructure Management Framework.
2. Enter *iplus* in the search box to search for the iplus cluster, as shown in Search for the iplus cluster.

   Search for the iplus cluster

   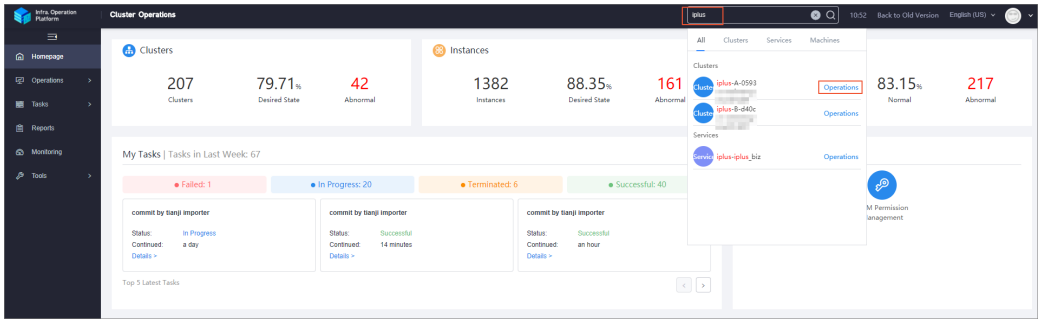

3. Click **O&M** on the right of the cluster name to redirect to the **Cluster Details** page, as shown in Cluster Details page.

   Cluster Details page

   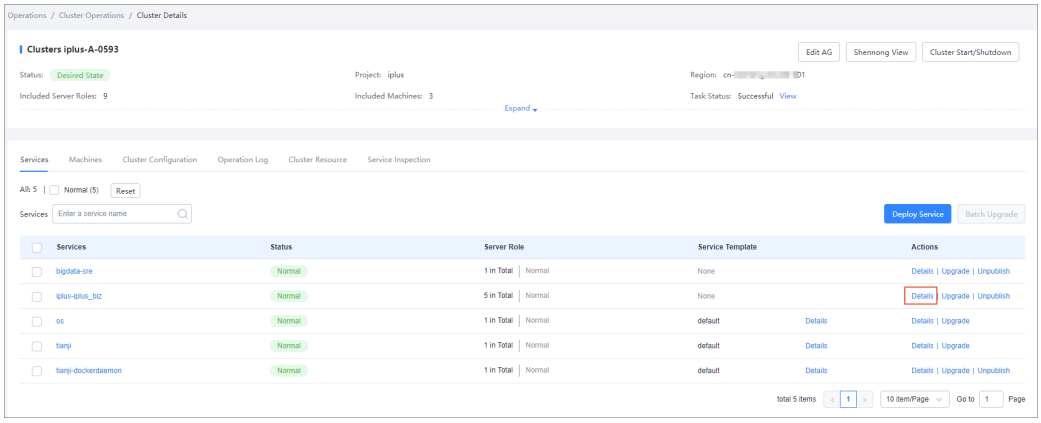

4. On the **Services** tab, find the **iplus-iplus_biz** service and click **Details** in the Actions column to view the details, as shown in Service Details page.

   Service Details page

   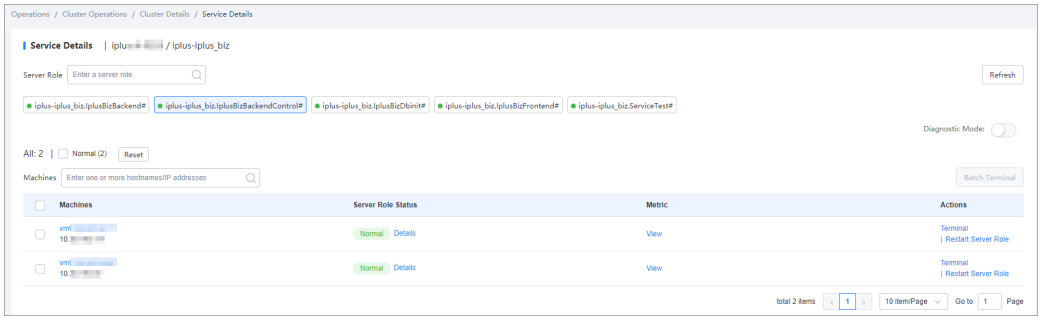

You can restart any role in the server role list, as shown in Server role list. Typically, you only need to restart the **IplusBizBackendControl#** and **IplusBizBackend#** roles.

> **Notice**
>
> You must restart the **IplusBizBackendControl#** and **IplusBizBackend#** roles in the following sequence:
>
> - Restart **IplusBizBackendControl#** first, and then **IplusBizBackend#**.
> - You must follow this sequence to restart the two roles within 10 minutes.
>
> Other roles can be restarted without strict sequence.

Server role list

5. Click the role you want to restart. In the Machines list, click **Restart** in the **Actions** column of all machines listed for this role, as shown in Restart a server role.

Restart a server role

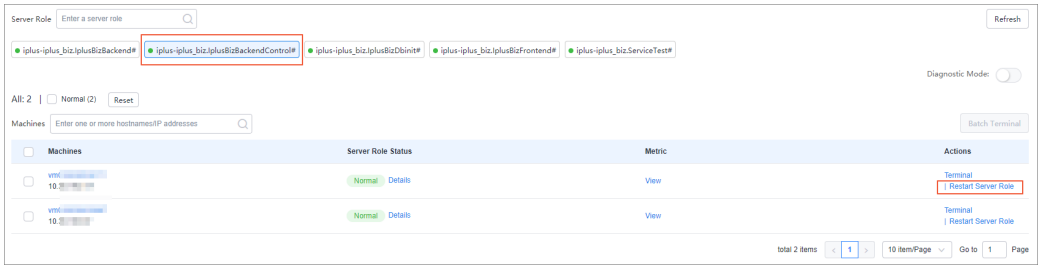

11.7.2.3. Operations and maintenance based on the Graph Analytics container

11.7.2.3.1. View instances

By viewing and examining instances, you can know the running status of instances and fix the problematic instances, for example, perform a switchover or clear logs.

View Java running instances

Log on to the Graph Analytics container, and run the `ps -ef|grep java|grep iplus` command. If the progress shown in View Java running instances exists, Administration Console is in the normal status.

View Java running instances

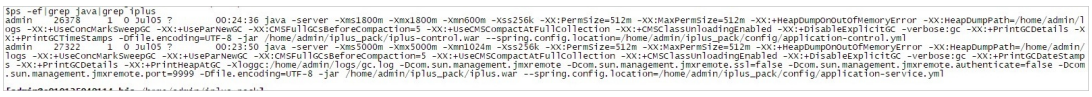

View node instances

Log on to the Graph Analytics application server, and run the `ps -ef|grep node` command. If the process shown in View node instances exists, the node service of Graph Analytics is normal.

View node instances

```
$ps -ef|grep node
admin     7974     1  0 19:12 pts/0    00:00:00 node /home/admin/i3-admin/target/i3-admin/admin-patch.js --harmony
admin     7991  7974  0 19:12 pts/0    00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.js --harmony undefined
admin     7996  7974  0 19:12 pts/0    00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.js --harmony undefined
admin     7997  7974  0 19:12 pts/0    00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.js --harmony undefined
admin     8002  7974  0 19:12 pts/0    00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.js --harmony undefined
admin    14876     1  0 Aug16 ?        00:00:00 node /home/admin/i3-web/target/i3-web/dispatch.js --harmony
admin    14887 14876  0 Aug16 ?        00:02:20 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony
admin    14892 14876  0 Aug16 ?        00:02:18 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony
admin    14893 14876  0 Aug16 ?        00:02:19 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony
admin    14898 14876  0 Aug16 ?        00:02:20 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony
```

In the preceding information, i3-web indicates that Analytics Workbench is in a normal status, and i3-admin indicates that Administration Console is in a normal status. If Administration Console is not released, the i3-admin process may not exist.

11.7.2.3.2. Log files

Graph Analytics application log:

The log files of Graph Analytics are stored in the `/home/admin/logs` directory.

A 100-GB data disk is mounted to the `/home/admin/logs` directory. Log files will increase with the execution time, which requires automatic cleanup. Two cleanup policies are available:

- Policy one: Time-based cleanup. The disk automatically deletes the log files that were created two weeks ago.
- Policy two: Cleanup based on the log size in the directory. If the log files occupy more than 80% of the total data disk space, the disk automatically deletes the earliest log files.

11.7.2.3.3. Database logs

Database logs record the execution information of i3-related programs, mainly the SQL statements. This information includes the execution time, whether the statements have been executed successfully, and whether an exception has occurred.

1. Log on to the Graph Analytics container.
2. Run the **cat** */home/admin/iplus_pack/config/application-service.yml* command to view the database information in *application-service.yml*.

   View database information

```
    datasource:
        url: jdbc:mysql://iplus[redacted].com:3177/iplus_meta?useUn
eSSL=true
        username: [redacted]
        password: [redacted]
        driver: com.mysql.jdbc.Driver
        type: com.alibaba.druid.pool.DruidDataSource
        druid:
              max-active: 50
              initial-size: 1
              min-idle: 3
              max-wait: 60000
              time-between-eviction-runs-millis: 60000
              min-evictable-idle-time-millis: 300000
              test-while-idle: true
              test-on-borrow: false
              test-on-return: false
```

3. Run the **mysql -h${db_host} -P${db_port} -u${db_user} -p${db_password} -D${db_name}** command to log on to the database.
4. Query the latest SQL statement executed by Graph Analytics and the time track.

   ```
   SELECT * from i3eye_time_trace WHERE main_time_trace_id in (
   SELECT max(main_time_trace_id) from i3eye_time_trace);
   ```

5. View the SQL statements executed within the last hour.

   ```
   select * from i3eye_time_trace where name like 'com.alibaba.iplus.common.dal.manual%' and (gmt_cr
   eate < now() and gmt_create > date_sub(now(), interval 1 hour) );
   ```

6. View the SQL statements that have errors within the last hour.

   ```
   select * from i3eye_time_trace where complete = 0 and name like 'com.alibaba.iplus.common.dal.manu
   al%' and (gmt_create < now() and gmt_create > date_sub(now(), interval 1 hour) );
   ```

11.7.2.3.4. Stop the service

Use admin Log on to the Graph Analytics container, run the start script, and run the following ps commands to view processes:

- View Java process: `ps –ef|grep java`
- View node process: `ps –ef|grep node`

You can stop a service by killing the corresponding thread.

11.7.2.3.5. Restart the service

Use admin Log on to the Graph Analytics container and run the startup script:

- Directly start iplus, i3-web, and i3-admin: `iplus-deploy.sh start`
- Start iplus only: `iplus-deploy.sh start_iplus`
- Start i3web only: `iplus-deploy.sh start_i3web`

- Start i3admin only: `iplus-deploy.sh start_i3admin`

11.7.3. Security maintenance

11.7.3.1. Network security maintenance

Network security maintenance handles the device security and the network security.

Device security

- Check network devices, and enable security management protocols and configurations of the devices.
- Check for new versions of the network device software and update to a more secure version in a timely manner.
- For more information about the security maintenance methods, see the product documentation of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network situations to detect public and internal network traffic and protect the network against attacks and unusual activities.

11.7.3.2. Account password maintenance

Account passwords include the Graph Analytics system password and the device password.

To ensure account security, you must change the system and device passwords periodically, and use passwords with high complexity.

11.7.4. Troubleshooting

11.7.4.1. Fault response mechanism

The IT administrator must establish a fault emergency response mechanism, so that the service can be recovered quickly after a fault or security accident occurs.

11.7.4.2. Troubleshooting methods

After a system fault is detected during routine maintenance, the IT administrator can read the Operations and Maintenance part of this documentation for reference.

If the fault cannot be fixed, collect the fault information, including the system information and fault symptoms, contact Alibaba Cloud technical support engineers, and troubleshoot the fault under the guidance of the engineers.

After the fault is fixed, the IT administrator must analyze the causes, review the troubleshooting process, and make improvements.

11.7.4.3. Common failure troubleshooting

Insufficient disk space

Insufficient disk space

Possible cause: The log size in the Graph Analytics system is too large.

Solution: Monitoring logs are usually stored in the `/home/admin/logs` directory. You can delete earlier logs to free up space.

Machine maintenance or downtime

Possible cause: The hardware is damaged or the warranty of the machine is expired.

Solution: Reinstall Graph Analytics.

Suspicious processes

Possible cause: If the process fails to start automatically or is terminated unexpectedly, view the logs in `/home/admin/logs` to identify the cause.

Solution: Restart Graph Analytics.

11.7.4.4. Hardware troubleshooting

Disk failure

Solution: Graph Analytics supports cluster deployment. Therefore, you can directly end all Graph Analytics threads, replace the hard drive, and then start the threads again.

Failures requiring server shutdown, including memory, MPU, CPU, and power supply failures

Solution:

Repairs involving server shutdown:

- If you can access the system, you can follow the service stop procedure to disable the Graph Analytics service on the server.
- If you cannot access the system, you must force the server to shut down.

11.8. Machine Learning Platform for AI

11.8.1. Query server and application information

11.8.1.1. Apsara Stack Machine Learning Platform for AI

11.8.1.1.1. 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.
   i. Find the service instance and click **Details**. The instance detail page appears.
   ii. Find the role list and click **Details**. The role detail page appears.
   iii. 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.

11.8.1.1.2. 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.

11.8.1.1.3. 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/BOOT-INF/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.

   i. 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.

   ii. 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
   ```

11.8.1.1.4. 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
```

11.8.1.2. Online model service

11.8.1.2.1. 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.

11.8.1.2.2. 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`

11.8.1.2.3. 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}`

11.8.1.3. GPU cluster and task information

11.8.1.3.1. 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.

11.8.1.3.2. 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 `ral` command to view the running tasks.
3. Run the `rwwl WorkItemName` command to view the status of a task and the allocated resources. `WorkItemName`: specifies the values in the first column displayed by the `ral` command.
4. Run the `rcru` 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 `rjstop WorkItemName` command to stop a Fuxi task. `WorkItemName`: specifies the values in the first column displayed by the `ral` command.

11.8.2. Maintenance and troubleshooting

11.8.2.1. Machine Learning Platform for AI maintenance

11.8.2.1.1. 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.

   ServiceTest instance

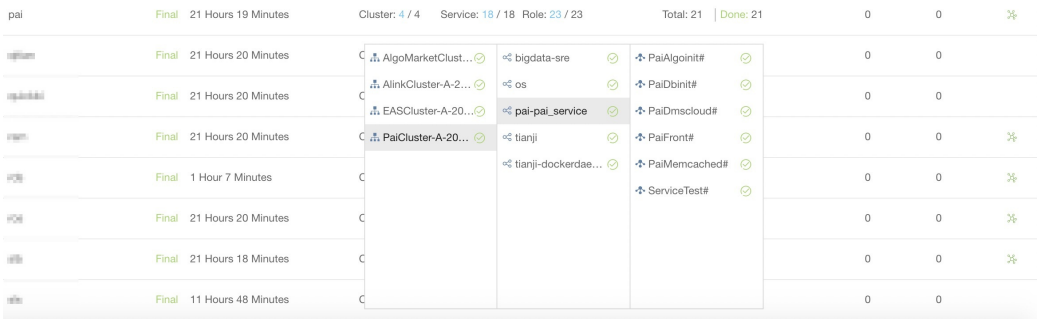

6. Run the `sudo docker restart e90f70353031` command to restart the ServiceTest service, as shown in the following figure.

   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.

   Testing results

   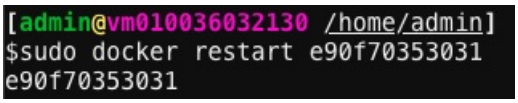

   - PASS: The algorithm is running properly.
   - SKIP or FAIL: The algorithm fails.

11.8.2.1.2. Common faults and solutions

11.8.2.2. 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.

11.8.2.3. 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 `rttrl` 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.

11.9. DataHub

11.9.1. Concepts and architecture

11.9.1.1. Terms

Project

A project is an organizational unit in DataHub and contains one or more topics. DataHub projects and MaxCompute projects are independent of each other. Projects that you create in MaxCompute cannot be used in DataHub.

Topic

The smallest unit for data subscription and publishing. You can use topics to distinguish different types of streaming data. For more information about projects and topics, see Limits in *Product Introduction*.

Topic lifecycle

The period that each record can be retained in the topic. Unit: day. Valid values: 1 to 7.

Shard

A shard in a topic. Shards ensure the concurrent data transmission of a topic. Each shard has a unique ID. A shard can be in different states. For more information about shard status, see the following table. Each active shard consumes server resources. We recommended that you create shards as needed.

Note

Shard status

Status	Description
Activating	All shards in a topic are in the Activating state when the topic is created. You cannot perform read or write operations on shards because they are being activated.
Active	Read and write operations are allowed when a shard is in the Active state.
Deactivating	A shard is in the Deactivating state when it is being split or merged with another shard. You cannot perform read or write operations on the shard because it is being deactivated.
Deactivated	A shard is in the Deactivated state when the split or merge operation is complete. The shard is read-only when it is in the Deactivated state.

Hash key range

The range of hash key values for a shard, which is in the format of [Starting hash key,Ending hash key). The hashing mechanism ensures that all records with the same partition key are written to the same shard.

Shard merge

The operation that merges two adjacent shards. Two shards are considered adjacent if the hash key ranges for the two shards form a contiguous set with no gaps.

Shard split

The operation that splits one shard into two adjacent shards.

Record

A unit of data that is written into DataHub.

Record type

The data type of records in a topic. Tuple and blob are supported. A tuple is a sequence of immutable objects. A blob is a chunk of binary data stored as a single entity.

Note

- The following table describes the data types that are supported in a tuple topic.

  Tuple data types

Data type	Description	Valid values
Bigint	An 8-byte signed integer. Note Do not use the minimum value, which is -9223372036854775808, because this is a system reserved value.	-9223372036854775807 to 9223372036854775807
String	A string. Only UTF-8 encoding is supported.	A string whose size is no greater than 1 MB
Boolean	One of two possible values.	True and False, true and false, or 0 and 1
Double	A double-precision floating-point number. It is 8 bytes in length.	-1.0 10^{308} to 1.0 10^{308}
TimeStamp	A timestamp.	A timestamp that is accurate to microseconds

- In a blob topic, a chunk of binary data is stored as a record. Records written to DataHub are Base64 encoded.

Service roles

Available service roles in DataHub

Service	Service role	Description
DataHub	Xstream	Receives read and write requests from the frontend server and forwards the requests to Apsara Distributed File System.
	Shipper/Connector	Synchronizes data from DataHub to other Apsara Stack services, including MaxCompute, ApsaraDB RDS for MySQL, and Object Storage Service (OSS).
	Coordinator	Saves consumption offsets for applications. You can resume data consumption from a saved consumption offset.
	Frontend	Receives all the read and write requests.

Run the following command on the admin gateway of a cluster to query the services deployed on the cluster:

```
r al
```

Services deployed on the cluster

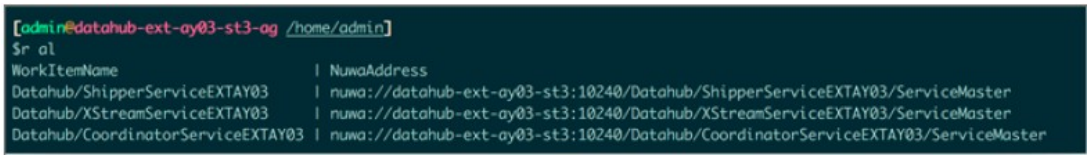

Run the following command on the admin gateway of the cluster to query the service role and the hosts where the service is running:

```
r wwl $WorkItemName
```

Service role and hosts where the service is running

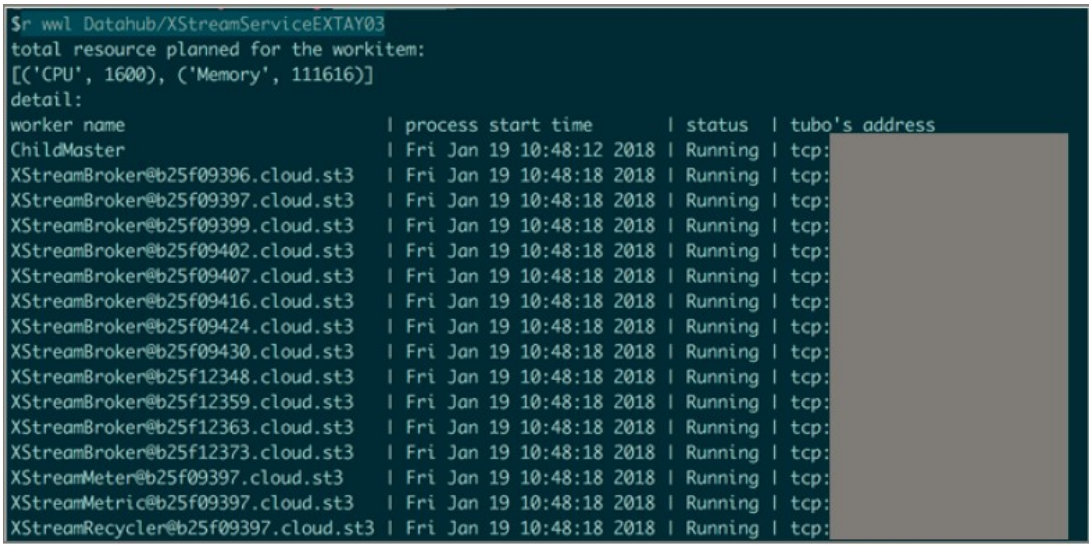

11.9.1.2. Architecture

11.9.1.2.1. Architecture

Architecture shows the architecture of DataHub.

Architecture

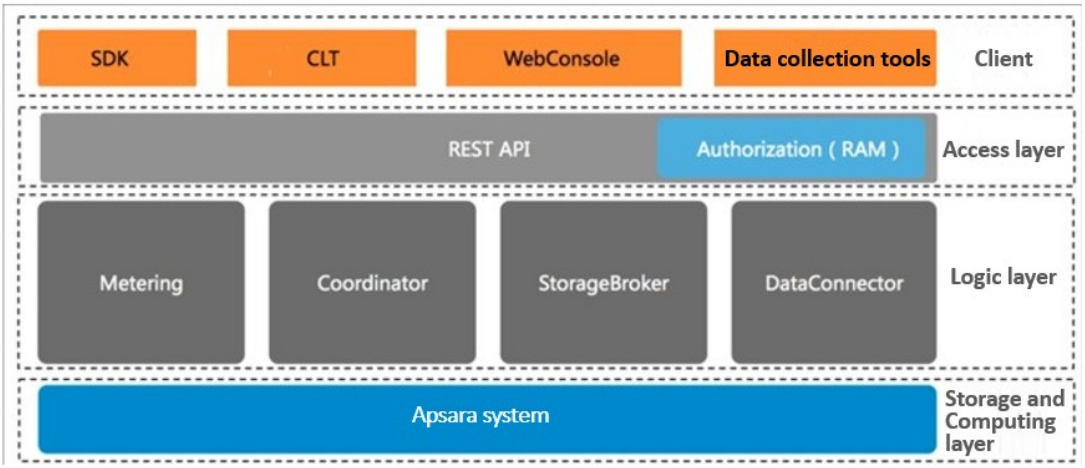

The architecture of DataHub consists of four layers: **clients**, **access layer**, **logic layer**, and **storage and scheduling layer**.

Clients

DataHub supports the following types of clients:

- SDKs: DataHub provides SDKs in a variety of languages such as C++, Java, Python, Ruby, and Go.
- Command-line tools (CLTs): You can run commands in Windows, Linux, or Mac operating systems to manage projects and topics.
- Console: In the console, you can manage projects and topics, create subscriptions, view the shard status, monitor topic performance, and manage DataConnectors.
- Data collection tools: You can use Logstash, Fluentd, and Oracle GoldenGate (OGG) to collect data to DataHub.

Access layer

You can access DataHub by using HTTP and HTTPS. DataHub supports Resource Access Management (RAM) authorization and horizontal scaling of topic performance.

Logic layer

The logic layer handles the key features of DataHub, including project and topic management, data read and write, offset-based data consumption, traffic statistics, and data synchronization. Based on these key features, the logic layer is composed of the following modules: StorageBroker, Metering, Coordinator, and DataConnector.

- StorageBroker: provides data reads and writes in DataHub. This module adopts the log file storage model of Apsara Distributed File System, halving the read/write volume compared with the conventional write-ahead logging (WAL) model. This module stores three copies of data to ensure that no data is lost if a server fault occurs, and supports disaster recovery between data centers. It supports real-time data caching to ensure efficient consumption of real-time data and supports an independent read cache of historical data to enable concurrent consumption of historical data.
- Metering: supports shard-level billing based on the consumption period.
- Coordinator: supports offset-based data consumption and horizontal scaling of the processing capacity. It supports up to 150,000 QPS on a single node.
- DataConnector: supports automatic data synchronization from DataHub to other Apsara Stack services, including MaxCompute, OSS, AnalyticDB, ApsaraDB RDS for MySQL, Tablestore, and

Elasticsearch.

Storage and scheduling layer

- Storage: Based on the log file storage model of Apsara Distributed File System, DataHub supports append operations and solid state drive (SSD) storage. Data in each shard is stored in a separate file based on the timestamp of the data.
- Scheduling: Based on Job Scheduler, DataHub assigns shards to nodes based on the traffic on each shard. This ensures that the shards do not occupy the CPU or memory of Job Scheduler. The number of partitions on a single node has no upper limit. DataHub supports failovers within milliseconds and hot upgrades.

11.9.1.2.2. Technical architecture

Technical architecture of DataHub shows the technical architecture of DataHub.

Technical architecture of DataHub

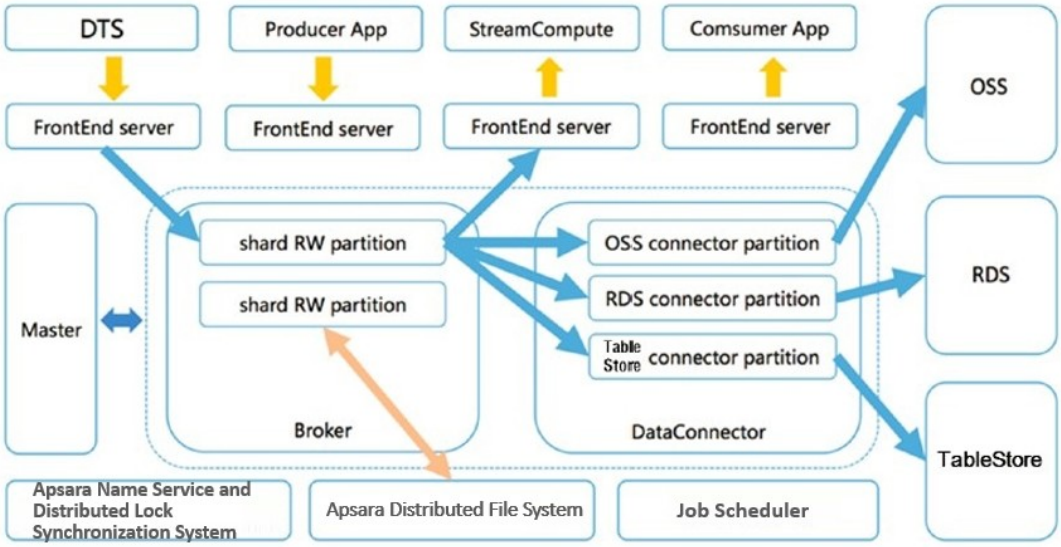

The figure shows the process from data ingestion to consumption.

1. A shard is the smallest unit of data management in DataHub, and is a first-in, first-out (FIFO) collection of records.
2. Data in each shard is stored in a set of log files in Apsara Distributed File System.
3. The master distributes each shard to a StorageBroker. Each StorageBroker is responsible for the read and write operations on multiple shards.
4. The frontend server finds a StorageBroker based on the project, topic, and shard information specified in the request and forwards the request to the StorageBroker.
5. DataConnectors read data from the StorageBroker and forward the data to other Apsara Stack services.

Data collection

You can write data to DataHub from applications developed by using SDKs and data collection tools such as Logstash, Fluentd, and OGG. You can also write data by using Data Transmission Service (DTS) and Realtime Compute.

Frontend server

Frontend servers constitute the access layer and support horizontal scaling. You can call RESTful API operations to access DataHub. RAM authorization is supported.

Master

The master handles metadata management and shard scheduling. It supports create, read, update, and delete operations on projects and topics. The master also supports split and merge operations on shards.

StorageBroker

StorageBrokers handle read and write operations on each shard including data indexing, caching, and file organization and management.

DataConnector

DataConnectors forward data in DataHub to other Apsara Stack services. DataConnectors provide different features for various destination services. These features include automatically creating partitions in MaxCompute and converting data streams into files stored in OSS.

11.9.2. Commands and tools

11.9.2.1. Common commands for the Apsara system

DataHub is built based on the Apsara system. Both DataHub and the Apsara system including Job Scheduler, Apsara Distributed File System, and Apsara Name Service and Distributed Lock Synchronization System are hosted by Apsara Infrastructure Management Framework.

- Run the following command to view the server roles that are installed on the server:

  ```
  tj_show
  ```

- Run the following command to view all server roles:

  ```
  tj_show –l
  ```

- Run the following command to retrieve a list of servers that the pangu_chunkserver server role is installed on:

  ```
  tj_show -r pangu.PanguChunkserver# //The hostnames of the servers are returned.
  tj_show -r pangu.PanguChunkserver# -ip //The IP addresses of the servers are returned.
  ```

- Run the following command to retrieve a list of servers that the FrontEnd server role is installed on:

  ```
  tj_show –r datahub-frontend.Frontend#
  ```

- Run the following command to retrieve a list of servers that the WebConsole server role is installed on:

  ```
  tj_show –r datahub-webconsole.WebConsole#
  ```

11.9.2.2. Common commands for Apsara Distributed File System

Commands for Apsara Distributed File System start with pu or puadmin. To view the complete description of a command, enter the command followed by --**help** and press enter.

- Run the following command similar to the ls command used in Linux to retrieve the file content in a specific directory:

  ```
  pu ls
  ```

- Run the following command to upload local files to Apsara Distributed File System:

  ```
  pu put
  ```

- Run the following command to retrieve metadata:

  ```
  pu meta
  ```

- Run the following command to retrieve details about all masters in Apsara Distributed File System:

  ```
  puadmin gems
  ```

- Run the following command to retrieve details about all chunk servers:

  ```
  puadmin lscs
  ```

- Run the following command to view version information:

  ```
  puadmin --buildinfo
  ```

- Before maintaining a chunk server, remove the chunk server from the cluster. Perform the following operations:

  i. Run the following command to retrieve the current status of a chunk server:

  ```
  pyadmin cs -stat tcp://x.x.x.x:10260
  ```

  ii. Run the following command to remove the chunk server from the cluster by setting its status to shutdown:

  ```
  pyadmin cs -stat tcp://x.x.x.x:10260 --set=shutdown
  ```

  iii. After the maintenance is completed, run the following command to add the chunk server back to the cluster:

  ```
  pyadmin cs -stat tcp://x.x.x.x:10260 --set=normal
  ```

11.9.2.3. Common commands for Job Scheduler

The commands for Job Scheduler start with `r`, which is encapsulation of rpc.sh.

```
alias r='sh /apsara/deploy/rpc_wrapper/rpc.sh'
```

- Run the following command to retrieve all services and service jobs:

  ```
  r al
  ```

  > **Note** Typically, service jobs are deployed on the DataHub cluster. The list returned has many entries.

- Run the following command to retrieve the status of a service:

  ```
  r wwl $servicename
  ```

- Run the following command to terminate a service:

  ```
  r sstop $servicename
  ```

- Run the following command to start a service:

  ```
  r sstart $servicename
  ```

- Run the following command to retrieve a list of all resources in the cluster:

  ```
  r ttrl
  ```

- Run the following command to retrieve a list of idle resources in the cluster:

  ```
  r tfrl
  ```

  You can run other commands for scheduling purposes as needed.

11.9.2.4. Xstream

You can run commands on a service terminal by using Xstream for maintenance purposes. To access the target service terminal, perform the following operations:

> **Note** To use Xstream, you must log on as the administrator.

1. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose **Operations > Cluster Operations**. On the Cluster Operations page, enter datahub in the **Clusters** search box in the upper-right corner.
2. Click the name of the cluster in the search result. The **Services** tab on the Cluster Details page appears. In the **Service** search box, enter **datahub-webconsole**. Click datahub-webconsole in the search result.
3. The server role **datahub-webconsole.WebConsole#** appears. Click **Terminal** in the Actions column of the host to go to the TerminalService page.

On the TerminalService page, you can use Xstream to run commands for maintenance purposes.

1. Run the following command and find the IP address of ChildMaster. Log on to the host where ChildMaster is running by using Secure Shell (SSH).

```
r wwl Datahub/XStreamServicex
```

Find the IP address of ChildMaster

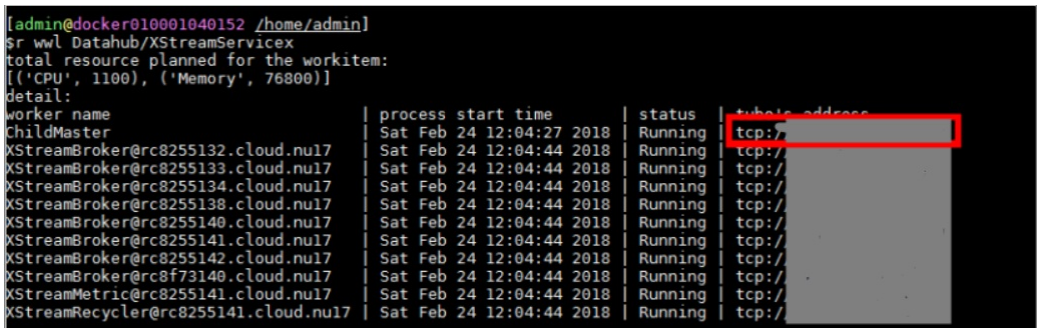

2. Run the following command to go to the specified directory:

```
cd /apsara/tubo/TempRoot/Datahub/XStreamServicex/tool
```

3. Run the following command to configure environment variables:

```
export LD_LIBRARY_PATH=/apsara/lib64/:../lib/
```

4. Run the following command to view resources:

```
./xstream_tool -x x mo
```

View resources

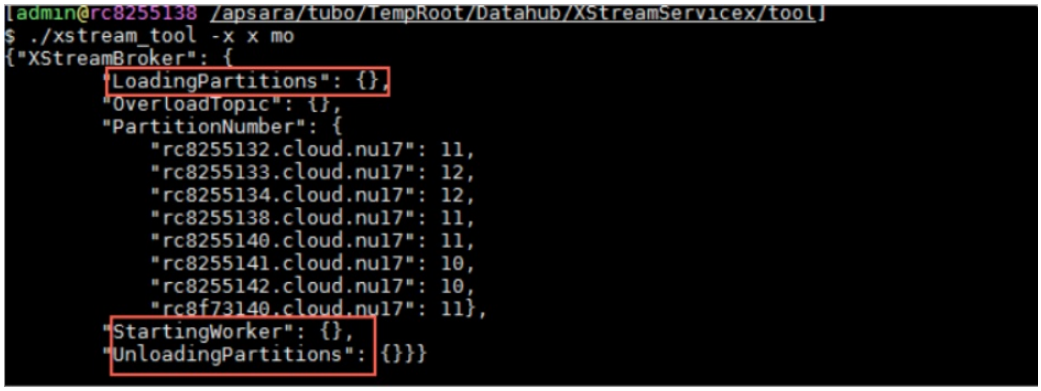

If **LoadingPartitions**, **UnloadingPartitions**, and **StartingWorker** are returned with values, run the command again. If these parameters are repeatedly returned with values, an error may occur when the shards are being activated or deactivated.

5. Run the following command to check the status of all StorageBrokers:

```
./xstream_tool gws -x x -r broker
```

Check the status of all StorageBrokers

```
[admin@rc8255138 /apsara/tubo/TempRoot/Datahub/XStreamServicex/tool]
$./xstream_tool gws -x x -r broker
Machine Name        | Requirement | Assignment | LoadedPartition | UnloadedPartition | UnconnectedWorker
rc8255132.cloud.nu17| 1           | 1          | 11              | 0                 | 0
rc8255133.cloud.nu17| 1           | 1          | 12              | 0                 | 0
rc8255134.cloud.nu17| 1           | 1          | 12              | 0                 | 0
rc8255138.cloud.nu17| 1           | 1          | 11              | 0                 | 0
rc8255140.cloud.nu17| 1           | 1          | 11              | 0                 | 0
rc8255141.cloud.nu17| 1           | 1          | 10              | 0                 | 0
rc8255142.cloud.nu17| 1           | 1          | 10              | 0                 | 0
rc8f73140.cloud.nu17| 1           | 1          | 11              | 0                 | 0
8                   | 8           | 8          | 88              | 0                 | 0
```

When 0 is returned for **UnloadedPartition** and **UnconnectedWorker**, the StorageBrokers are functioning properly.

6. Run the following command to check the status of all shards in the topic:

```
./xstream_tool -x x lsw -p $project -t $topic -r broker
```

Check the status of all shards in the topic

```
$ ./xstream_tool -x x lsw -p smoke_test_project -t datahub_to_datahub_input_1 -r broker
err_code: 0
err_msg: "Success"
workers {
  key: 3
  value: "Datahub/XStreamServicex/XStreamBroker@rc8255140.cloud.nu17"
}
workers {
  key: 5
  value: "Datahub/XStreamServicex/XStreamBroker@rc8255142.cloud.nu17"
}
workers {
  key: 2
  value: "Datahub/XStreamServicex/XStreamBroker@rc8255138.cloud.nu17"
}
workers {
  key: 7
  value: "Datahub/XStreamServicex/XStreamBroker@rc8255132.cloud.nu17"
}
workers {
  key: 4
  value: "Datahub/XStreamServicex/XStreamBroker@rc8255141.cloud.nu17"
}
```

From the command output, you can find the anomalous shards.

Note We recommend that you do not run other commands by using Xstream except for those described in the preceding example. If you need to run other commands, contact an operations engineer.

11.9.2.5. DataHub console

In the DataHub console, you can obtain performance statistics to facilitate O&M.

For more information about how to log on to the DataHub console, see Log on to the DataHub console in *User Guide*.

To check the performance statistics in the DataHub console, perform the following steps:

1. Log on to the DataHub console. In the left-side navigation pane, click **Project Manager**. On the Project List page, find the target project that you want to view performance statistics and click **View** in the Operate column. The project details page appears.
2. On the project details page, find the target topic and click **View** in the Operate column.
3. On the topic details page that appears, click the **Metric Statistics** tab to view the charts that display the performance statistics of the selected topic.

11.9.2.6. Apsara Bigdata Manager

Apsara Bigdata Manager (ABM) provides O&M on big data services from the perspective of business, services, clusters, and hosts. You can also upgrade big data services, customize alert configurations, and view the O&M history in the ABM console.

ABM allows onsite Apsara Stack engineers to manage big data services. As an engineer, you can view resource usage, check and handle alerts, and modify configurations.

For more information about how to log on to the ABM console and the supported O&M for DataHub, see the relevant topics in *DataHub O&M*.

11.9.3. Routine maintenance

11.9.3.1. Restore data after a power outage

Prerequisites

None.

Procedure

1. DataHub stores data in Apsara Distributed File System. A power outage may cause data loss. After a power outage, run the following command in the DataHub console to check whether the data stored in Apsara Distributed File System has been lost:

   ```
   puadmin fs -abnchunk|grep NONE|awk '{print $1}'|awk -F"_" '{print $1}'|while read line;do puadmin whois $line;done|grep FileId|awk '{print $4}' |sort|uniq >/home/admin/lostfile
   -- Ignore directories that start with /deleted/ and send all other directories to an operations engineer to check the lost data.
   ```

2. Restore data based on file types.
   - If DataHub files have been lost, notify your users that they must re-create the corresponding topics.
   - If metadata has been lost, re-install the corresponding package or initialize the Docker container.
3. After the data is restored, wait until the tianji cluster is at desired state. For assistance, contact an operations engineer.

11.9.3.2. Shut down anomalous chunkserver hosts

Prerequisites

None.

Procedure

1. Configure the action and action status for the anomalous chunkserver hosts.

i. Log on to the ops1 host and set action to rma and action status to pending for the anomalous chunkserver host. In this example, the name of the anomalous chunkserver host is m1.

Run the following command to configure the action and action status for the anomalous chunkserver host:

```
curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1" -d '{"action_name":"rma", "action_status":"pending"}'
```

The following response is returned:

```
{
  "err_code": 0,
  "err_msg": "",
  "data": [
    {
      "hostname": "m1"
    }
  ]
}
```

Set action to rma and action status to pending for the anomalous chunkserver host

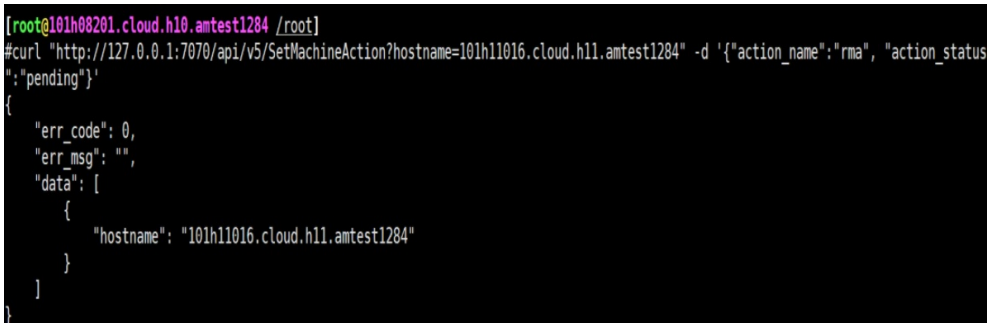

Note Replace the IP address and hostname in the sample code with those of your anomalous chunkserver host.

ii. Run the following command to configure audit logs:

```
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" }'
```

Note

- Replace the IP address and hostname in the sample code with those of your anomalous chunkserver host.
- Replace the value of the mtime parameter in the sample code with the current time.
- Run the following command to query mtime. The sample code is for your reference only.

```
curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1&attr=action_name,action_status,action_description@mtime"
```

The following response is returned:

```
{
  "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
  }
}
```

Query mtime

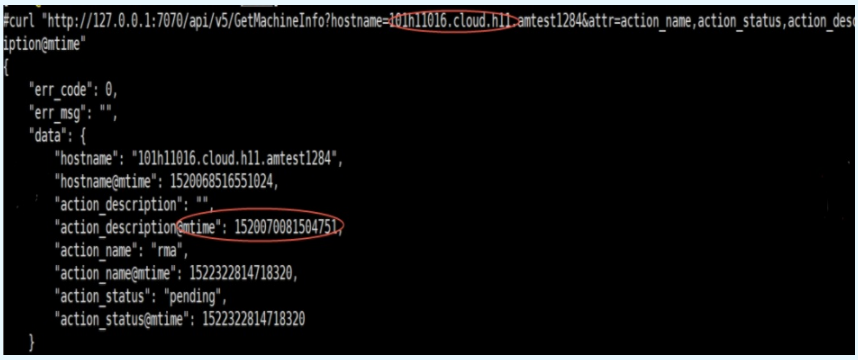

2. Wait for approval.

i. Check the action status of the host.

Run the following command to check the action status of the host:

```
curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1"
```

The response is a long list. We recommend that you search for the host by the keyword **"action_status": "pending"**.

After you verify that the action status is pending, you can approve the action in the Apsara Infrastructure Management Framework console.

ii. Check the action status of the server role. When the status is approved or done, you can shut down the host for maintenance.

Run the following command to check the action status:

```
curl http://127.0.0.1:7070/api/v5/GetMachineInfoPackage?hostname=m1&attr=sr.id,sr.action_name,sr.action_status
```

The response is a long list. We recommend that you search for the host by the keyword **"action_status": "pending"**.

3. After the action of the host changes to rma and action status changes to approved or done, shut down the host. Restart the host after the maintenance is completed.
4. After the host is restarted, run the following command to configure the action status of the host:

```
curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1&action_name=rma" -d '{"action_name":"rma","action_status":"done", "force":true}'
```

5. Check whether the cluster has reached the desired state.

11.9.3.3. Shut down a DataHub cluster

Prerequisites

None.

Procedure

1. Terminate DataHub services.

i. Log on to the webconsole host of the target cluster and run the following commands as an administrator. Ensure that no data is returned.

```
puadmin abnchunk fs -t none
puadmin abnchunk fs -t onecopy
puadmin abnchunk fs -t lessmin
```

ii. On the webconsole host, run the following commands as an administrator to terminate all services run by chunkserver hosts in the Apsara system:

```
r ttrl |grep disk |awk '{print $1}' > tubo.list
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad stop"
```

iii. On the webconsole host, run the following command as an administrator to make sure that all services in the Apsara system have been terminated:

```
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad status"
```

2. Shut down the cluster.
3. Restart DataHub services.
   i. On the webconsole host, run the following command as an administrator to restart all services run by chunkserver hosts in the Apsara system:

```
r ttrl |grep disk |awk '{print $1}' > tubo.list
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad start"
```

   ii. On the webconsole host, run the following command as an administrator to make sure that all services in the Apsara system are functioning properly:

```
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad status"
```

11.9.3.4. Replace a hard drive with a new one on the pangu_cs node

Prerequisite

Obtain the following information:

- The hostname or the IP address.
- The drive letters of the problematic drive. For example, /dev/sdk.
- The ID of the problematic drive. For example, if the path of the problematic drive in the Apsara Distributed File System is /apsarapangu/disk5, the drive ID is 5. You can also obtain the drive ID by running the following command: puadmin lscs -m

Procedure

1. Run the following command to check that the drive to be replaced is in DISK_ERROR status.

```
puadmin lscs -m
```

> **Note** If the hard drive is not in DISK_ERROR status, run the following command to change the status:
>
> ```
> puadmin cs -stat tcp://hostname or IP address:10260 -d drive ID  --set=ERROR
> ```

2. Run the following command to unmount the drive. In this example, the drive letters of the drive to be unmounted are /dev/sdk.

```
sudo umount /dev/sdk1
```

> **Note** Ignore this operation if the df command output shows that the drive is not mounted.

3. After the unmount operation is completed, replace the hard drive in hot swap mode.
4. Upload the **sudo repair_app_disk.sh** script to the server and execute the script to format the drive.
5. Run the following command to set the drive status in the Apsara Distributed File System to OK:

```
puadmin cs -stat tcp://hostname or IP address:10260 -d drive ID --set=OK
```

6. Restart the server. After the server is started up, it detects a new hard drive.

> **Note** Kill the processes running on the pangu_cs chunk server and restart the server. Restarting a chunk server does not affect the continuity of your business because DataHub adopts a distributed storage model.

7. Run the following command to check whether the drive status is DISK_OK.

```
puadmin lscs -m
```

You can log on to the server to confirm that the drive has the chunks sub-directory. For example, the chunks exists in the /apsarapangu/disk5/chunks/ directory and new chunks are written into the sub-directory.

11.9.4. DataHub O&M

11.9.4.1. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

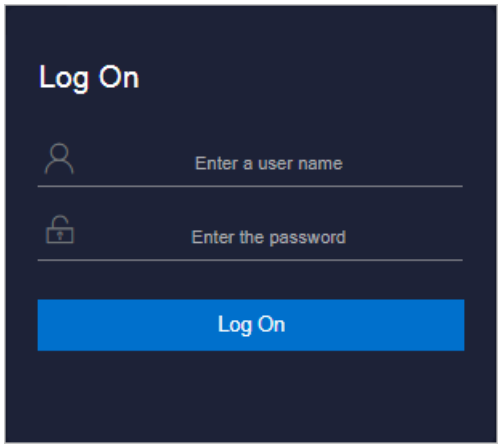

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.9.4.2. Common operations

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

Search for a project

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

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

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

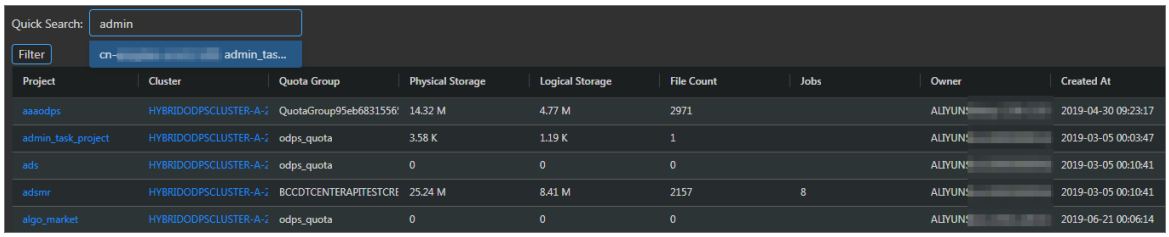

The following figure shows the search result.

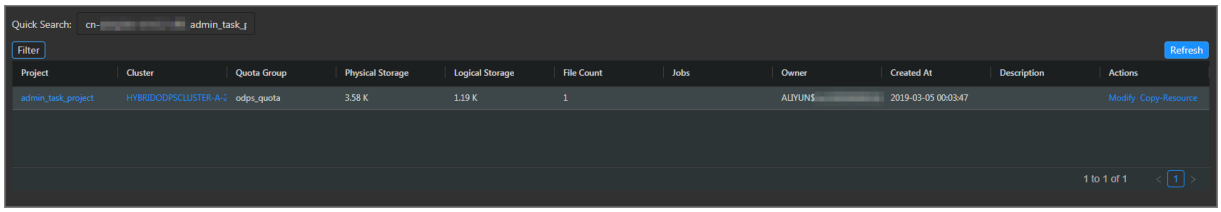

Filter projects

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

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

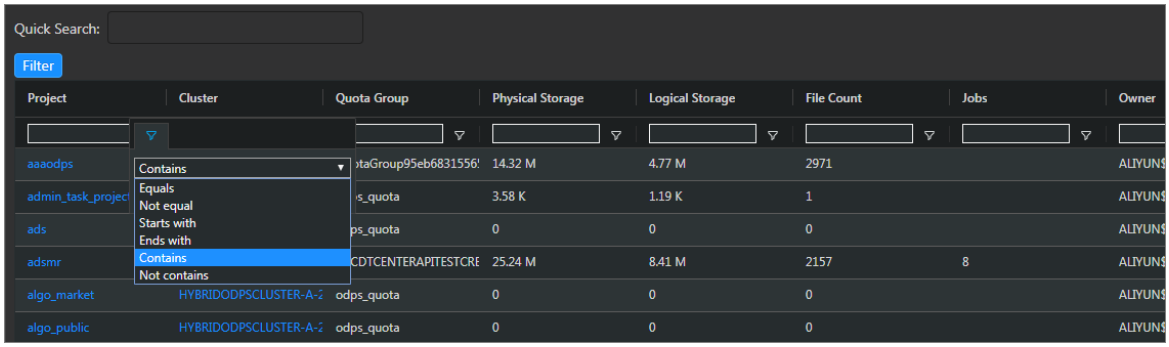

You can select one of the following filtering methods:

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

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

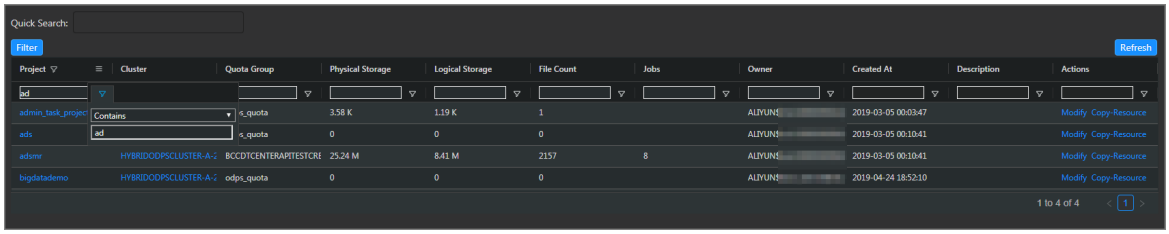

5. If the filtering result is not accurate, you can continue performing this operation on other columns.

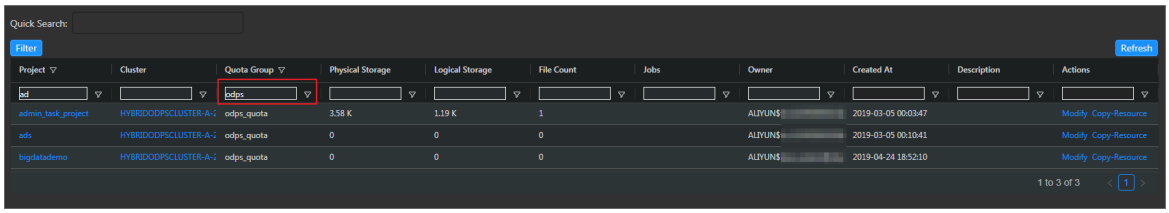

After you set the filter conditions for the projects, the **Filter** button is highlighted. If you need to cancel filtering, click the highlighted **Filter** button.

Search for an item

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

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

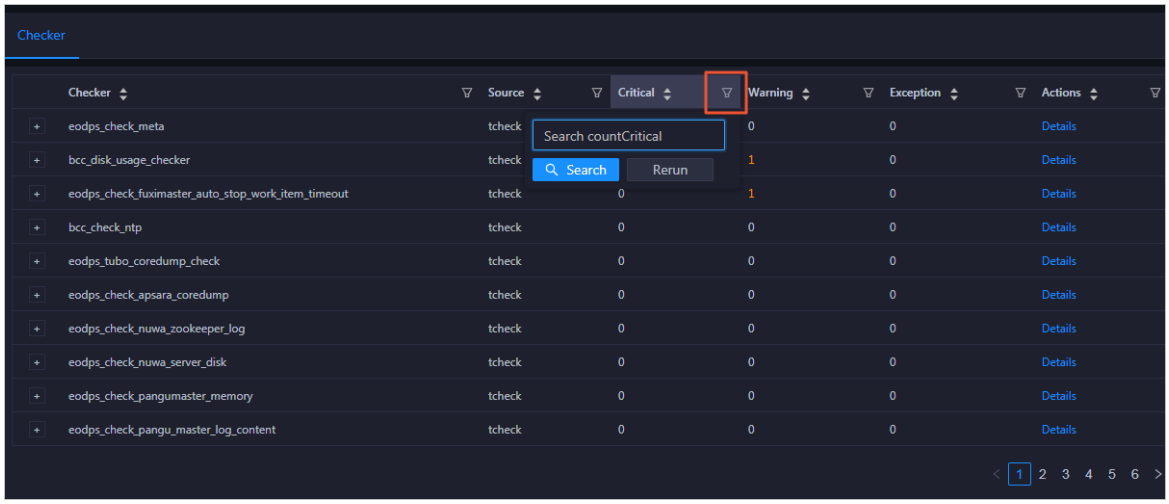

3. Click **Search**. The checkers that meet the requirements appear.
4. If the search result is not accurate, you can continue performing this operation on other columns.

Customize a column

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

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

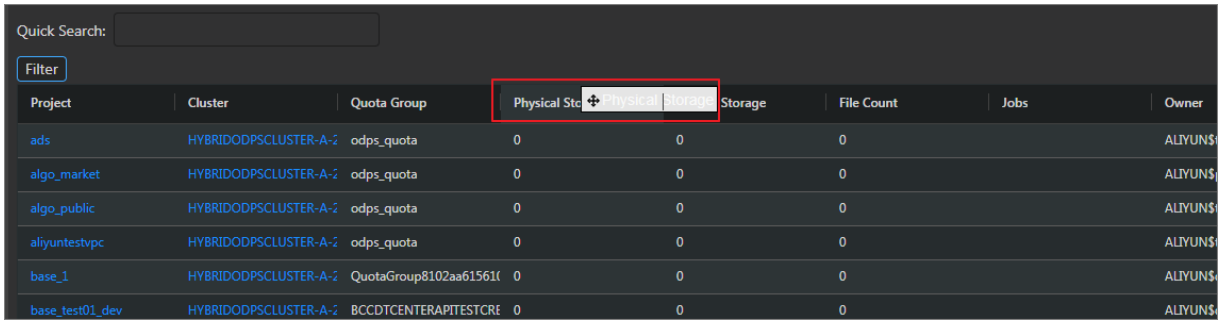

You can click in a column heading to customize the column.

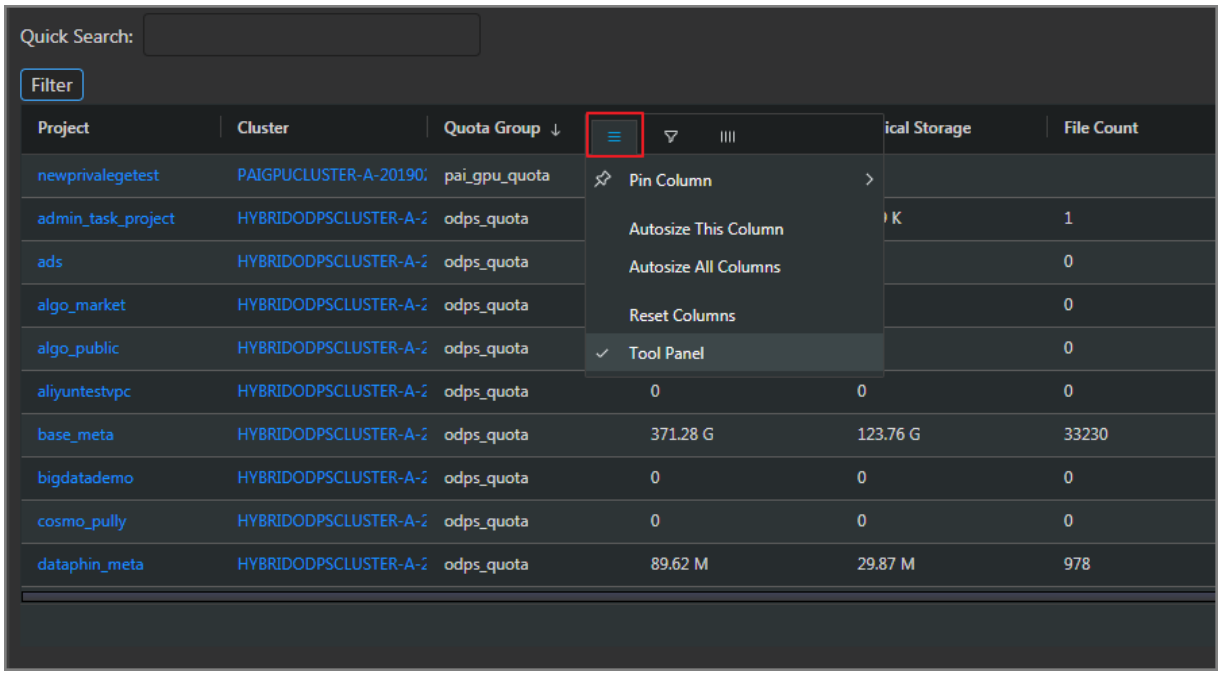

- **Pin Column**: allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.
- **Autosize This Column**: allows you to adjust the width of a column automatically.
- **Autosize All Columns**: allows you to adjust the width of all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

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

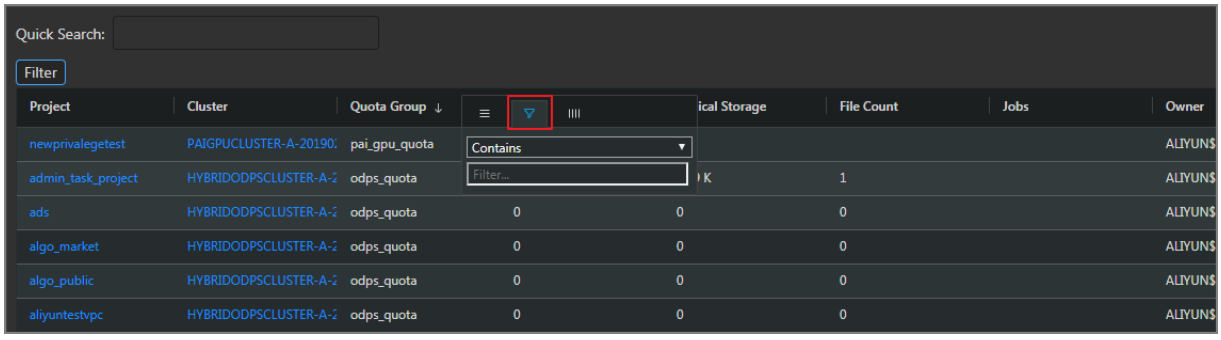

Click in a column heading and select the columns to display.

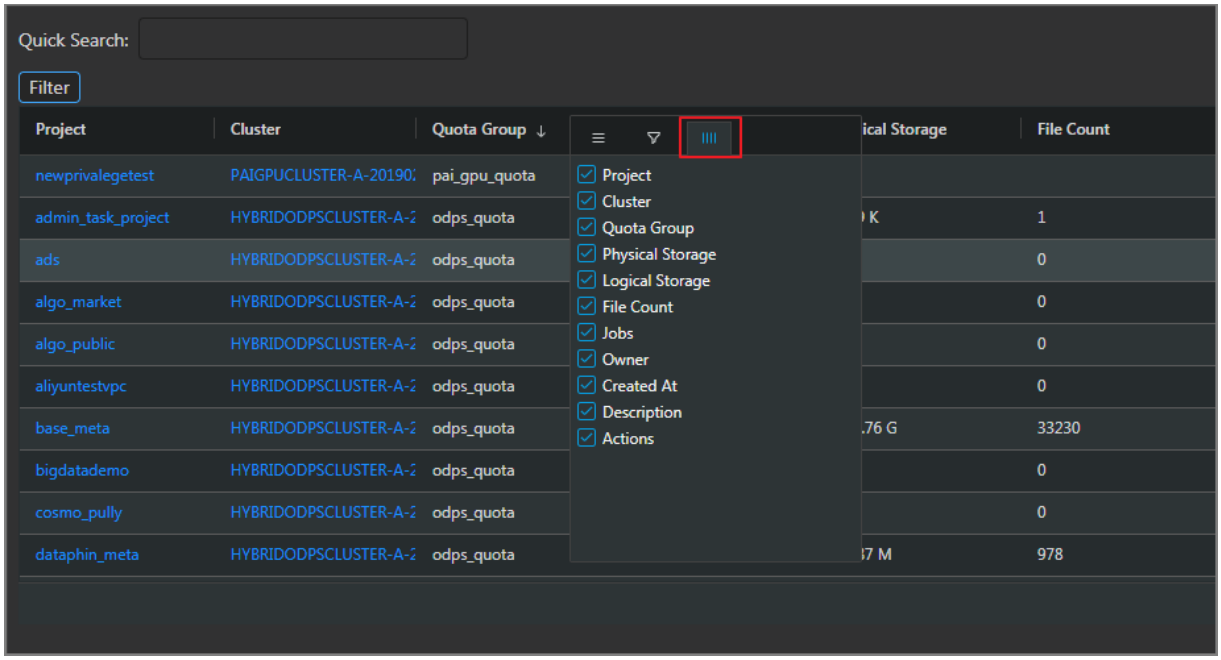

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

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

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

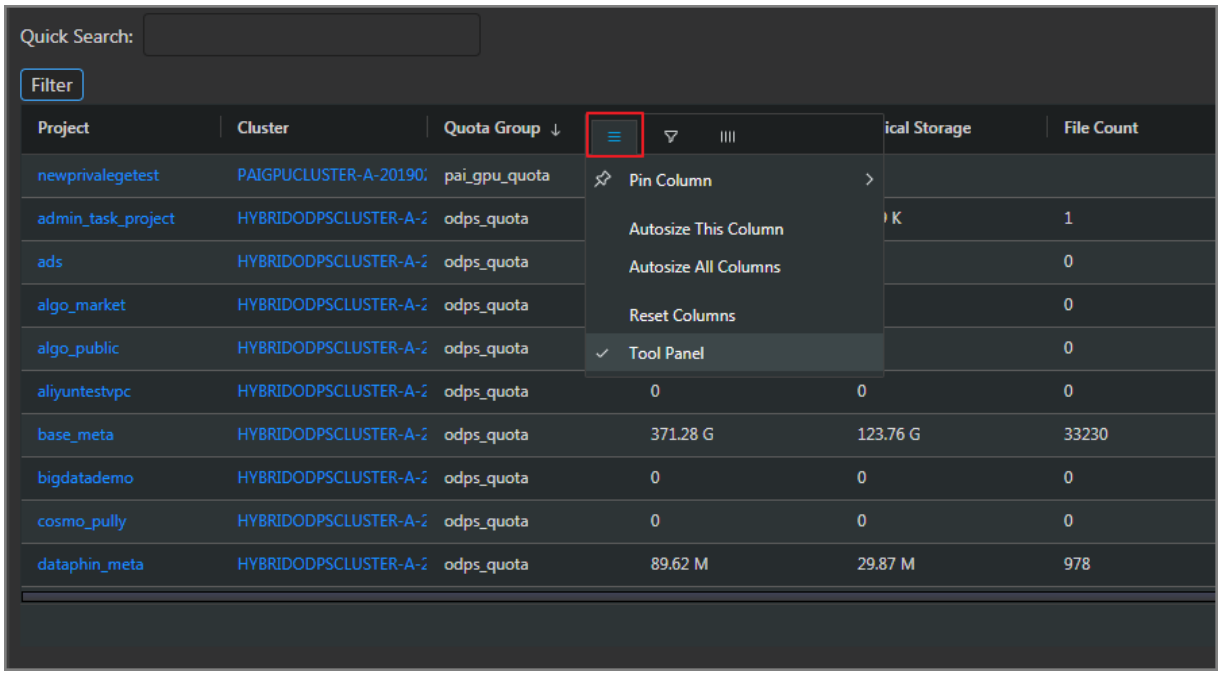

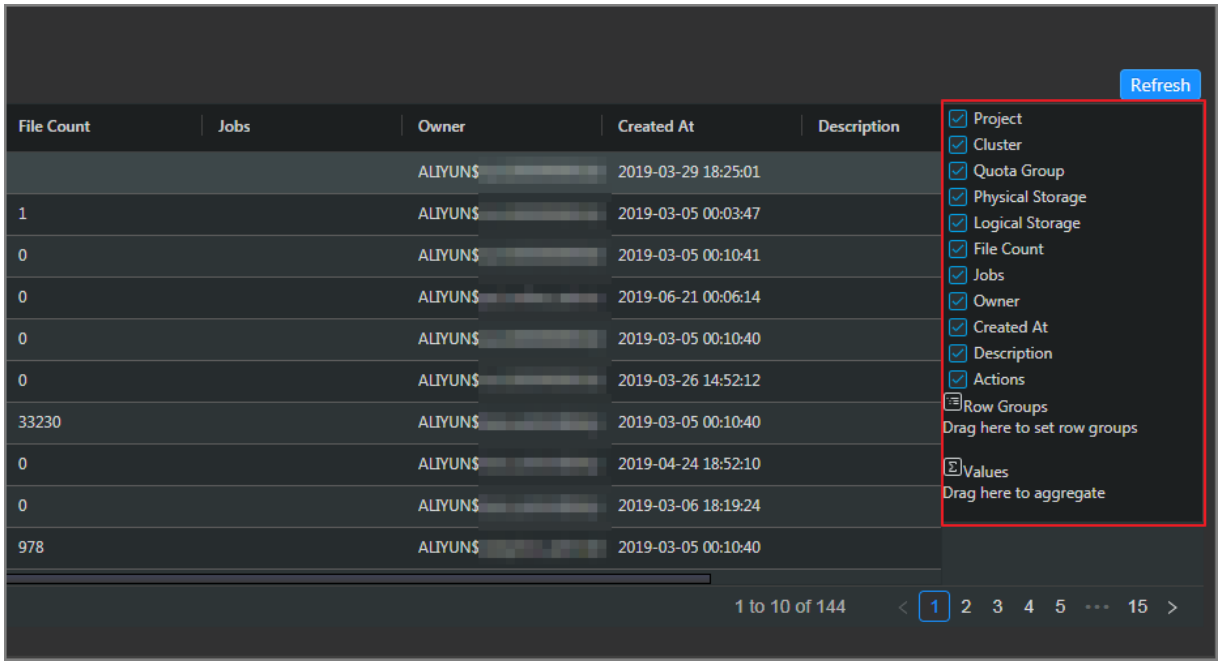

Sort projects based on a column

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

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

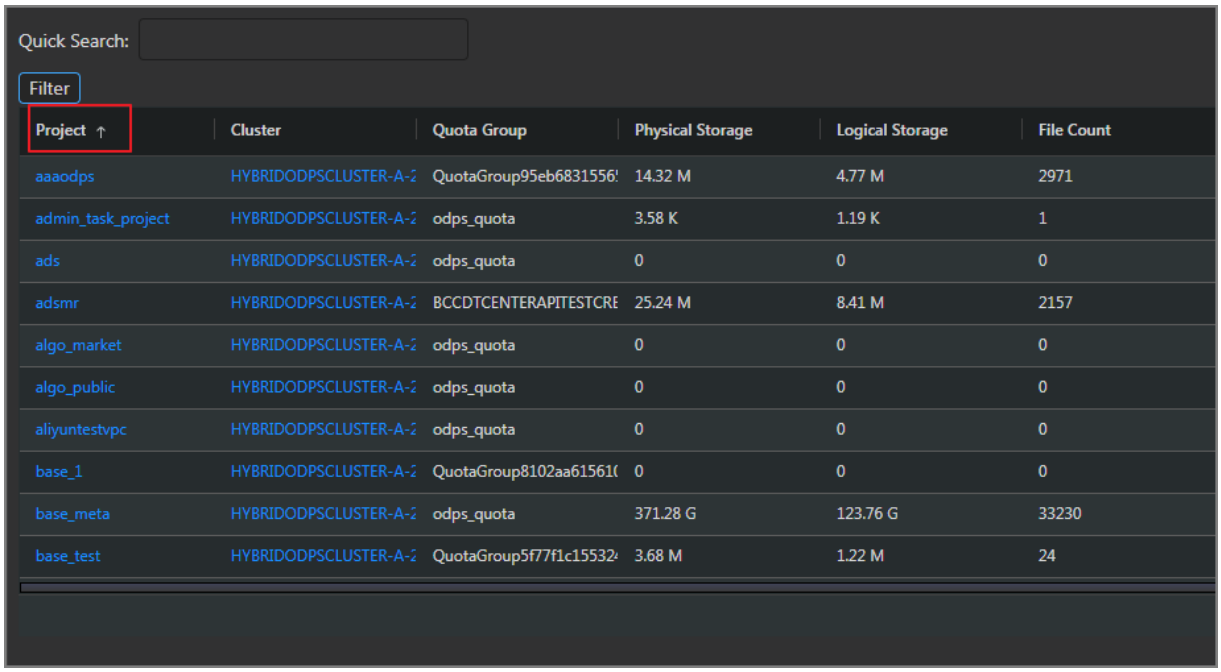

Sort items based on a column

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

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

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

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

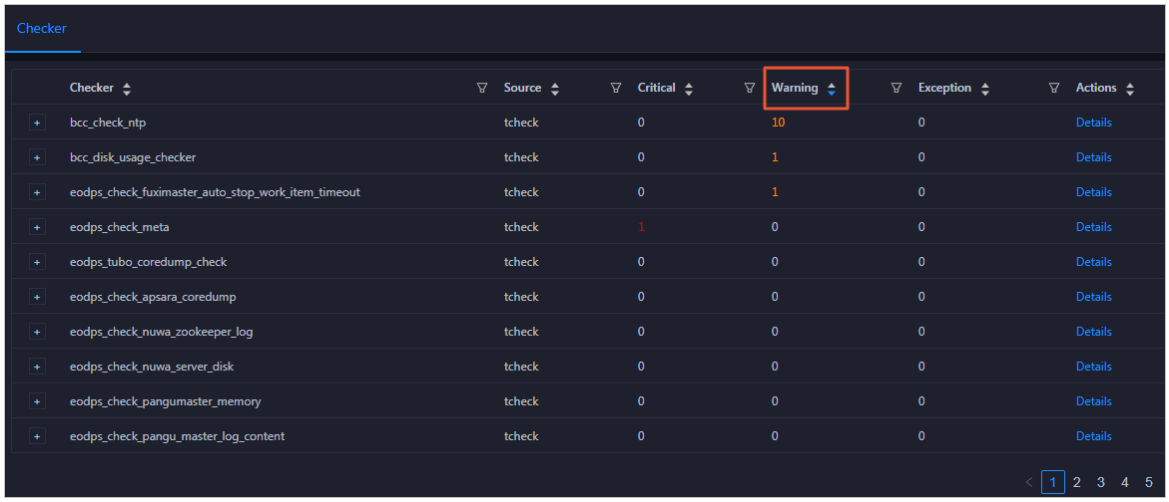

The highlighted up arrow indicates that the checkers are sorted in ascending order. The highlighted down arrow indicates that the checkers are sorted in descending order.

View the trend charts for a MaxCompute cluster

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

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

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

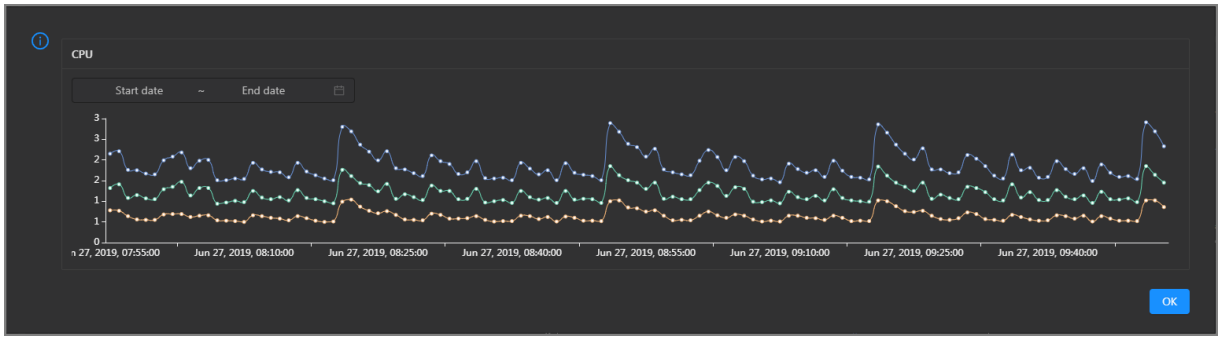

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

View the trend charts for a DataHub cluster

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

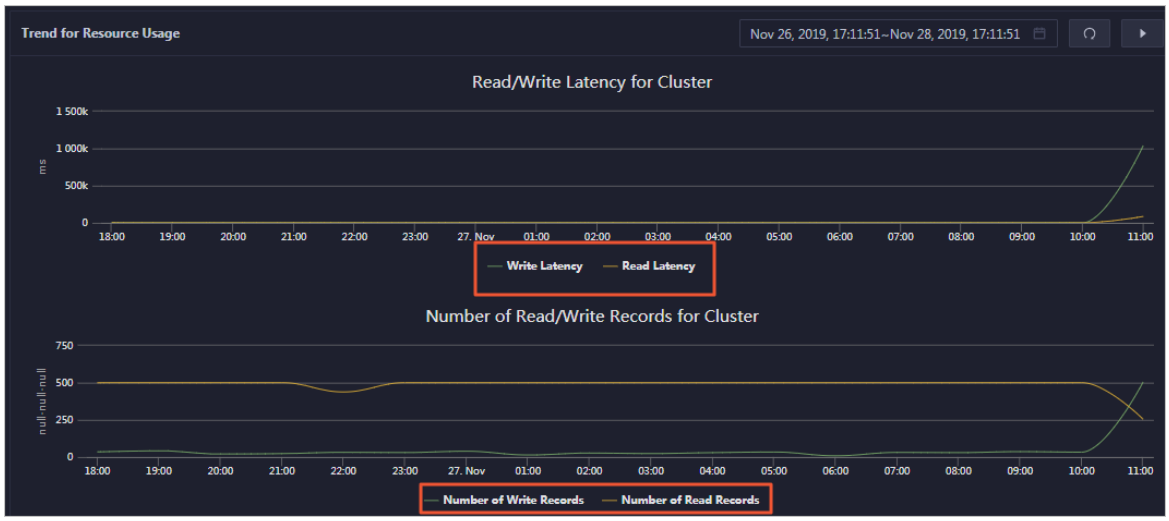

The trend charts, such as the trend charts of the read/write latency and the number of read/write records, appear in the Trend for Resource Usage section. Each chart displays the trend lines of the metrics over time in different colors. You can customize the metrics to display. You can click the name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

11.9.4.3. DataHub O&M overview

This topic describes the features of DataHub O&M and how to go to the DataHub O&M page.

Modules and features

DataHub O&M includes the business O&M, service O&M, cluster O&M, and host O&M modules. The following table describes the submodules and features contained in each module.

Module	Submodule or feature		Description
Business O&M	Projects		Displays the name, owner, the number of topics, read traffic, write traffic, storage usage of each project, and the time when a project was created.
	Topics		Displays the name, number of shards, storage usage, read traffic, and write traffic of each topic, the name of the project to which a topic belongs, and the time when a topic was created.
	Hotspot Analysis		Displays the distribution of shards on the hosts of a cluster for you to perform hotspot analysis.
Service O&M	Fuxi	Overview	Displays the key operation metrics of Job Scheduler, including the service overview, service status, health check result, health check history, resource usage, and overview of compute nodes. You can also view the trend charts of CPU and memory usage on this page.
		Instances	Displays the information about the Job Scheduler service roles, including the name, host, IP address, and status of a service role, and host status.
		Health Status	Displays the information about the checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts. In addition, you can log on to a host and perform manual checks on the host.
		Compute Nodes	Displays the information about compute nodes of a cluster, 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.
	Pangu	Overview	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 overview of storage nodes. You can also view the trend charts of storage usage and file count on this page.
		Instances	Displays the information about the Apsara Distributed File System service roles, including the name, host, IP address, and status of a service role, and host status.

Module	Submodule or feature		Description
		Health Status	Displays the information about the checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts. In addition, you can log on to a host and perform manual checks on the host.
		Storage Nodes	Displays the information about the storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, TTL, and send buffer size. You can also set the status of storage nodes and data disks on this page.
Clusters	Overview		Displays the overall running information about a cluster, including 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.
	Health Status		Displays the information about the checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts. In addition, you can log on to a host and perform manual checks on the host.
	Hosts		Displays the information about all hosts in a cluster, including the CPU usage, memory usage, root disk usage, packet loss rate, and packet error rate.
	Scale in Cluster and Scale out Cluster operations		Allow you to scale in or out a DataHub cluster by removing or adding physical hosts.
	Delete Topic from Smoke Testing operation		Allows you to delete topics from a DataHub test project and view the execution history.
	Reverse Parse Request ID operation		Allows you to reverse parse RequestId to obtain the time when a job was run and the IP address of the host. You can use the obtained information to query logs for troubleshooting.
Hosts	Overview		Displays the overall running information about a host in a DataHub 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.
	Charts		Displays the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission of a host.

Module	Submodule or feature	Description
	Health Status	Displays the information about the checkers of a host, including the checker details, check results, and schemes to clear alerts. In addition, you can log on to the host and perform manual checks on the host.
	Services	Displays the information about service instances and service roles of a host.

DataHub O&M entry

1. Log on to the ABM console.
2. Click in the upper-left corner and select **DataHub**.
3. On the DataHub page, click **O&M** in the upper-right corner. The **Business** tab appears.

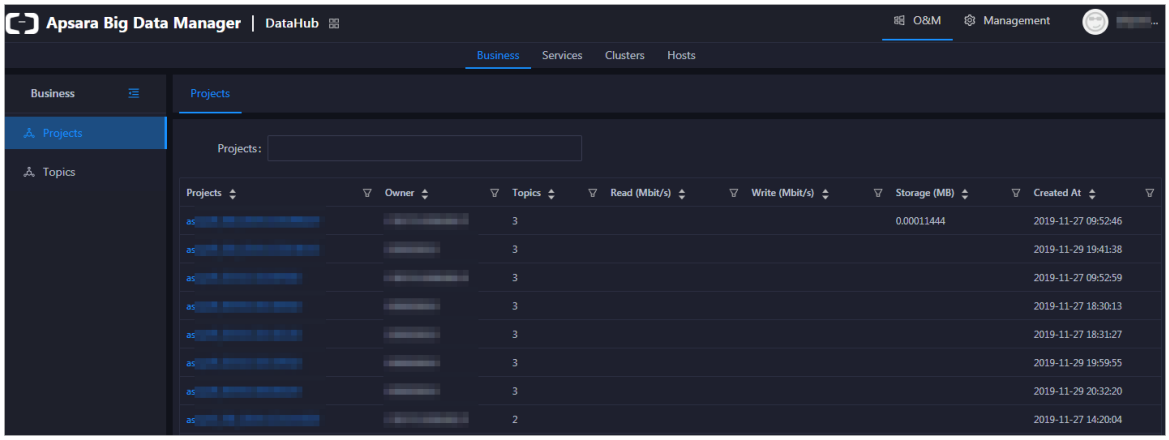

The **O&M** page includes four modules, namely, **Business**, **Services**, **Clusters**, and **Hosts**.

11.9.4.4. Business O&M

11.9.4.4.1. Business O&M entry

This topic describes how to go to the business O&M page for DataHub in the ABM console.

1. Log on to the ABM console.
2. Click in the upper-left corner and select **DataHub**.
3. On the DataHub page, click **O&M** in the upper-right corner, and then click the **Business** tab. The **Projects** page appears.

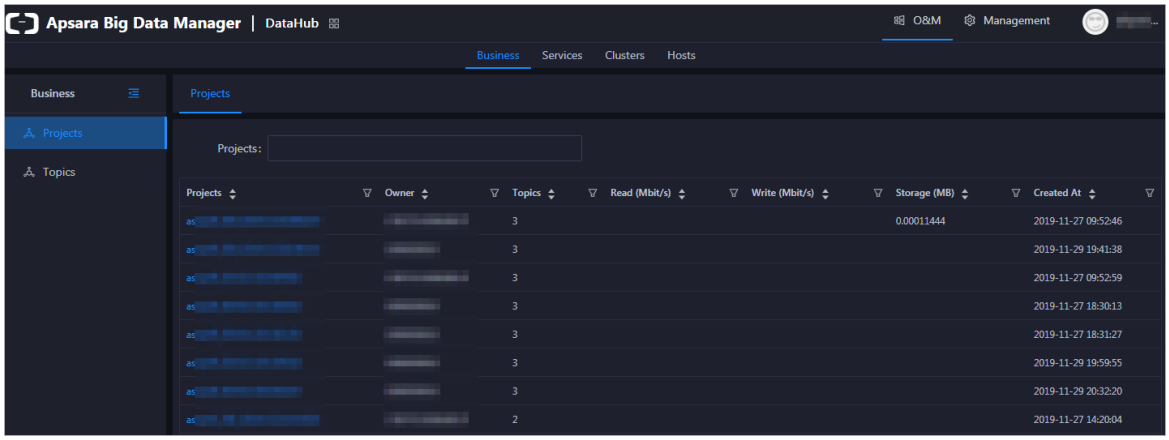

11.9.4.4.2. Projects

The Projects page displays the name, owner, the number of topics, read traffic, write traffic, storage of each project, and the time when a project is created.

Entry

On the **Business** page, click **Projects** in the left-side navigation pane. The Projects page appears on the right.

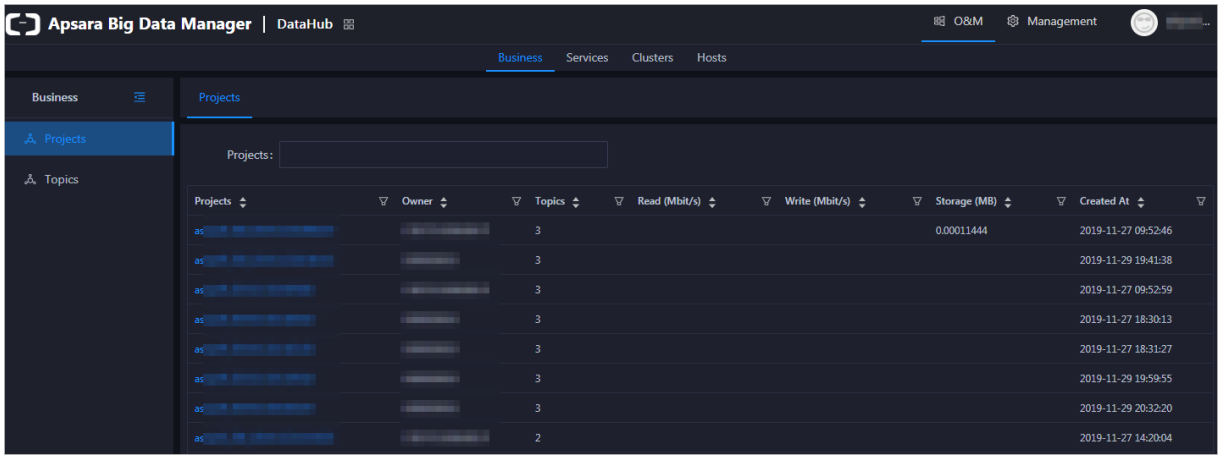

View project overview

On the **Projects** page, click the name of a project that you want to view. The Overview page for the project appears.

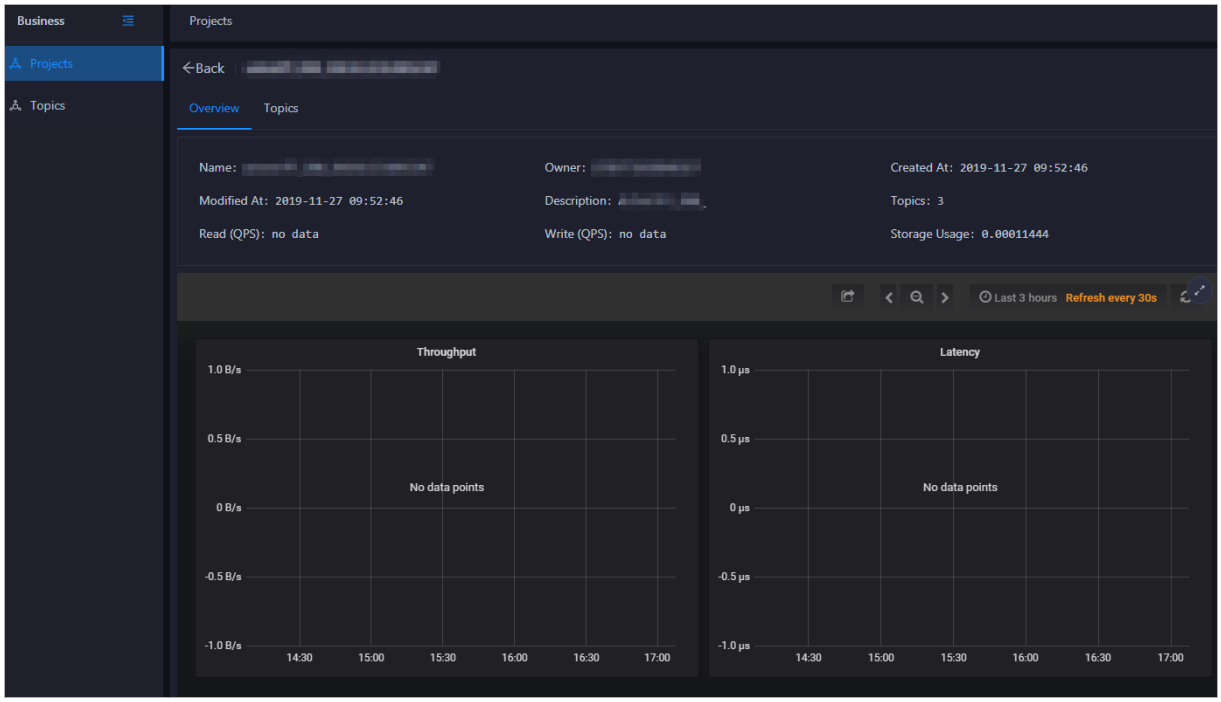

View topics of a project

On the **Projects** page, click the name of a project that you want to view. On the page that appears, click the **Topics** tab. All topics in the project appear.

11.9.4.4.3. Topics

The Log Sources page displays the name of a topic, the name of the project to which the topic belongs, the number of shards, storage usage, read traffic, and write traffic of the topic, and the time when the topic was created.

View a topic

1. On the **Business** tab, click **Topics** in the left-side navigation pane. The Log Sources page appears.

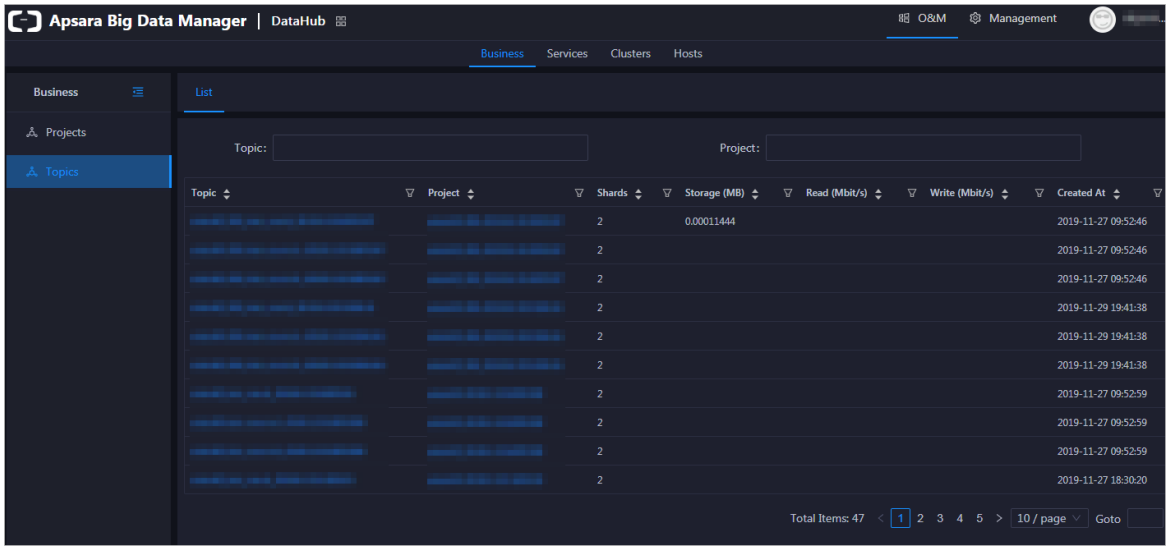

2. On the **Log Sources** page, click the name of the topic that you want to view. On the page that appears, you can view the number of shards, the time when the topic was created and modified, and the current storage usage, lifecycle, type, and description of the topic. You can also view more details about monitoring metrics, shards, subscriptions, DataConnectors, and schema.

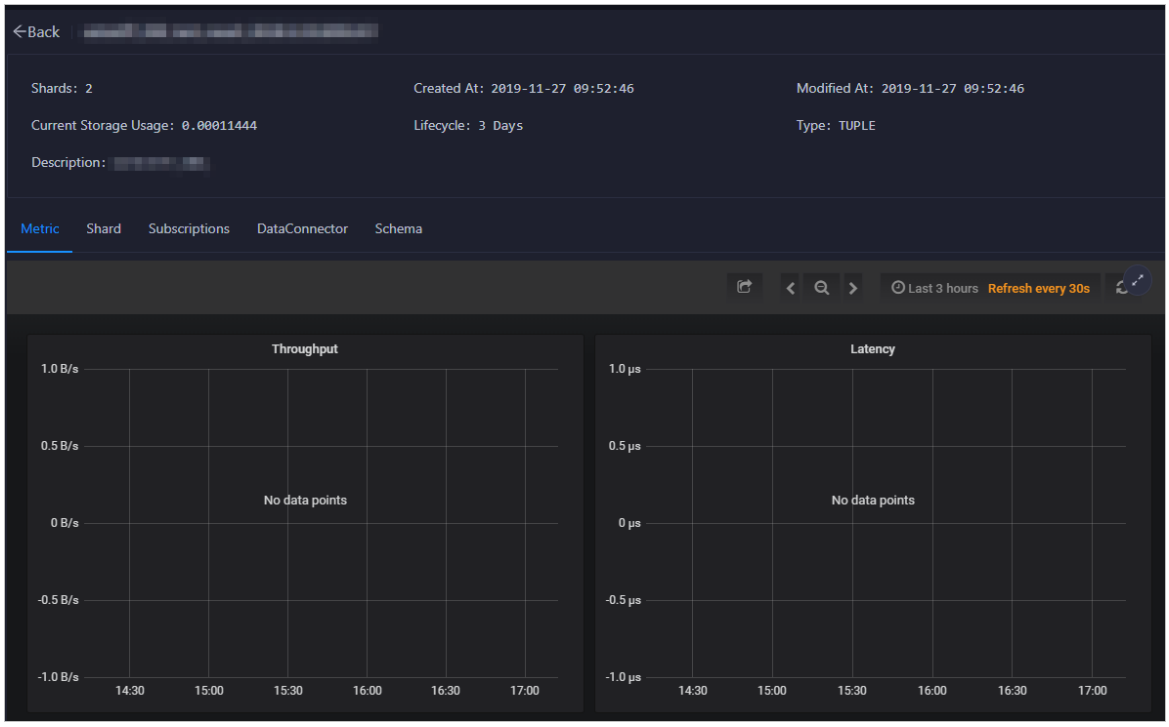

The details page of a topic contains the following tabs, where you can view different information:

- **Metric**: On the Metric tab, you can view information about the throughput and latency of a topic in quasi-real time.
- **Shard**: Shards are concurrent tunnels used for data transmission in a topic.

  On the Shard tab, you can view the ID, status, and active time of each shard.
- **Subscriptions**: The subscription feature of DataHub supports saving consumption offsets to the server and resuming data consumption from a saved consumption offset.

  On the Subscriptions tab, you can view the ID, status, owner, and description of each subscription and the time when the subscription was modified.
- **DataConnector**: DataConnectors synchronize the streaming data from DataHub to other Apsara Stack services. You can configure a DataConnector so that the data you write to DataHub can be used in other Apsara Stack services.

  On the DataConnector tab, you can view the name, ID, owner, and status of each DataConnector, and the time when the DataConnector was created and modified.
- **Schema**: The schema is used to define the data types of fields.

  On the Schema tab, you can view the data type and name of each field.

11.9.4.4.4. Hotspot analysis

The Hotspot Analysis page displays the distribution of shards on the hosts of a cluster for you to perform hotspot analysis.

Go to the Hotspot Analysis page

On the **Business** tab, click **Hotspot Analysis** in the left-side navigation pane. On the Hotspot Analysis page, you can view the distribution of shards on the hosts of a specific cluster in the column chart.

Refresh the column chart and filter the data

On the **Hotspot Analysis** page, you can click **Shards** to refresh the column chart. You can also set conditions in the list below the chart to filter the data.

11.9.4.5. Service O&M

11.9.4.5.1. Control Service O&M

The Overview page for the control service displays the overall running information about the service, including the service overview, service status, health check result, health check history, and trends of resource usage.

Go to the O&M page for the control service

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

11.9.4.5.2. Service O&M for Job Scheduler

11.9.4.5.2.1. Job Scheduler O&M entry

This topic describes how to go to the service O&M page for Job Scheduler in DataHub in the ABM console.

1. Log on to the ABM console.
2. Click  in the upper-left corner and select **DataHub**.
3. On the DataHub page, click **O&M** in the upper-right corner, and then click the **Services** tab.
4. On the **Services** tab, click **Fuxi** in the left-side navigation pane and select a cluster from the drop-down list. The **Overview** page for Job Scheduler appears.

11.9.4.5.2.2. Service overview

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

Go to the Overview page

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

2. Select a cluster from the drop-down list and click the **Overview** tab. The **Overview** page for Job Scheduler appears.

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

Services

This section shows the numbers of available services, unavailable services, and services that are being updated.

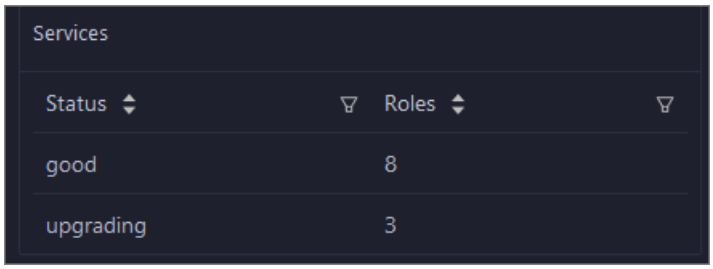

Roles

This section shows all Job Scheduler server roles and their states. You can also view the expected and actual numbers of machines for each server role.

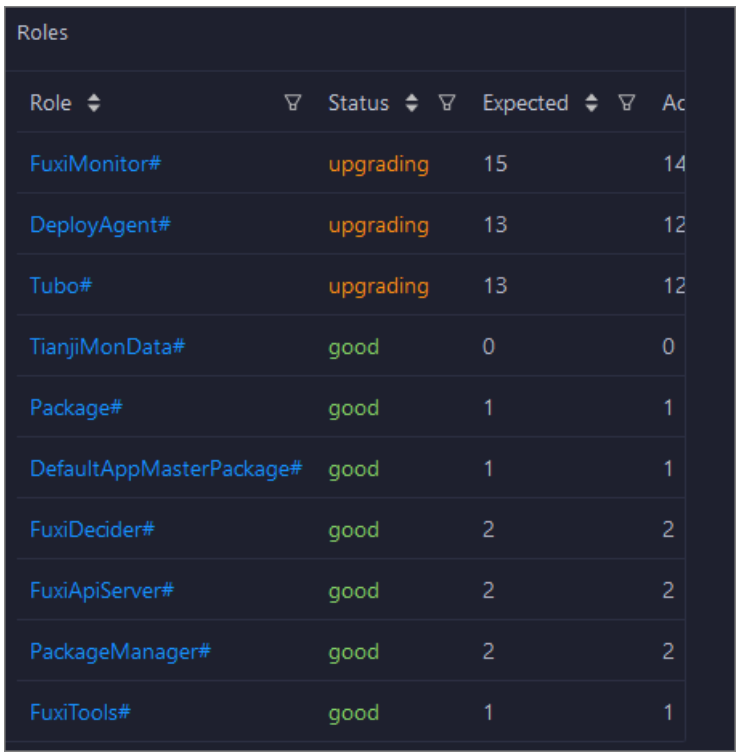

Click the name of a server role to go to the Apsara Infrastructure Management Framework console and view its details.

Saturability - Resource Usage

This section shows the allocation of CPU and memory resources.

- CPU (Core): shows the CPU utilization, the total number of CPU cores, the number of available CPU

cores, and the CPU cores for SQL acceleration.

- Memory (Bytes): shows the memory usage, the total memory size, the available memory size, and the memory size for SQL acceleration.

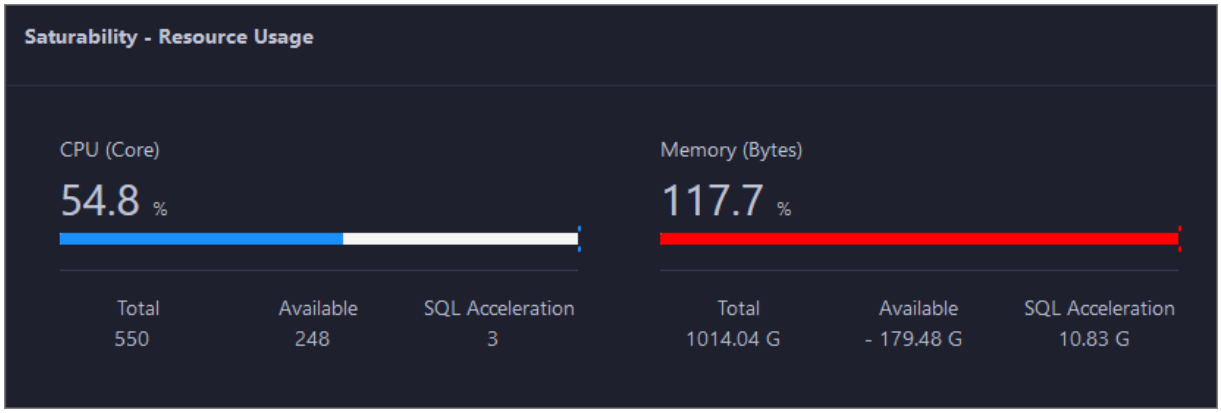

View the trend charts of CPU and memory usage

In the CPU Usage (1/100 Core) and Memory Usage (MB) sections, you can view the trend charts of CPU and memory usage of the selected cluster. Each trend chart displays the trend lines of the used quota, idle quota, and total quota of the relevant resource 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.

Compute Nodes

This section shows the details of compute nodes in Job Scheduler. The details include the percentage of online compute nodes, the total number of compute nodes, the number of online compute nodes, and the number of compute nodes in a blacklist.

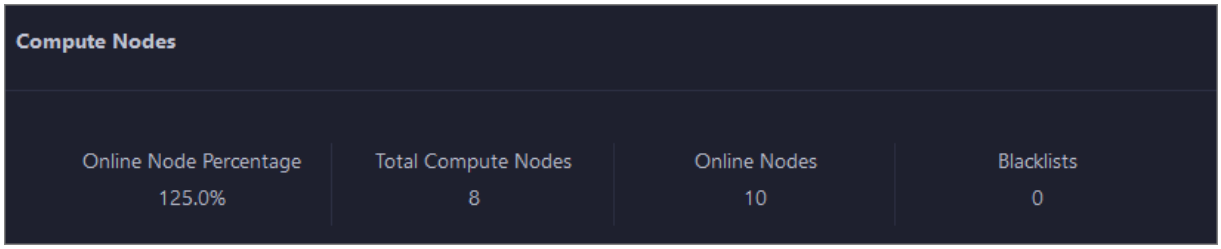

11.9.4.5.2.3. Service instances

The Instances page displays information about the Job Scheduler service roles, including the name, host, IP address, and status of a service role, and host status.

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.

   On the **Instances** page, you can view information about the Job Scheduler service roles, including the name, host, IP address, and status of a service role, and host status.

11.9.4.5.2.4. Service 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.

Go to the Health Status page

1. On the **Services** tab, click **Fuxi** in the left-side navigation pane.
2. Select a cluster from the drop-down list and 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 the information about the checkers of a cluster, including the checker details, hosts with alerts and alert causes, and schemes to clear alerts. In addition, you can log on to a host and perform manual checks on the host. For more information, see Cluster health.

11.9.4.5.2.5. 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.

Go to the Compute Nodes page

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

   On this page, you can view the details of compute nodes, 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.

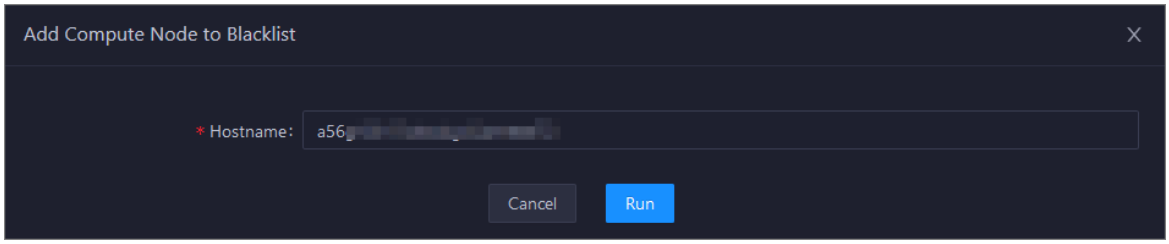

The value of the **Hostname** parameter is automatically filled. You do not need to specify a value for this parameter.

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

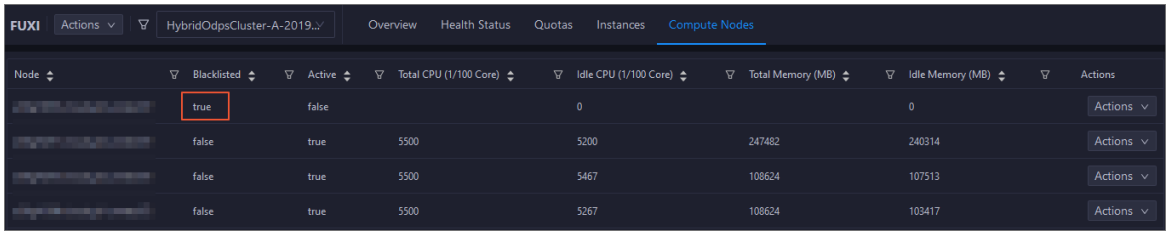

11.9.4.5.3. Service O&M for Apsara Distributed File System

11.9.4.5.3.1. Apsara Distributed File System O&M entry

This topic describes how to go to the service O&M page for Apsara Distributed File System in DataHub in the ABM console.

1. Log on to the ABM console.
2. Click in the upper-left corner and select **DataHub**.
3. On the DataHub page, click **O&M** in the upper-right corner, and then click the **Services** tab.
4. On the **Services** tab, click **Pangu** in the left-side navigation pane and select a cluster from the drop-down list. The **Overview** page for Apsara Distributed File System appears.

11.9.4.5.3.2. Service overview

The Overview tab shows the key operating information about Apsara Distributed File System. The information includes the service overview, service status, storage usage, storage node overview, and the trend charts of storage usage and file count.

Go to the Overview page

1. On the **Services** tab, click **Pangu** in the left-side navigation pane.
2. Select a cluster from the drop-down list and 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 overview of storage nodes. You can also view the trend charts of storage usage and file count on this page.

Services

This section shows the status of Apsara Distributed File System and the number of server roles.

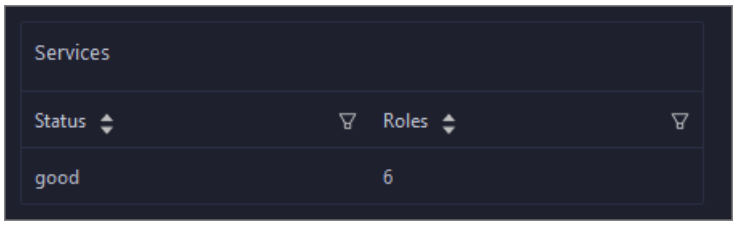

Roles

This section shows all server roles of Apsara Distributed File System and their states. You can also view the expected and actual numbers of hosts for each server role.

Saturability - Storage

This section shows the storage usage and file count.

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

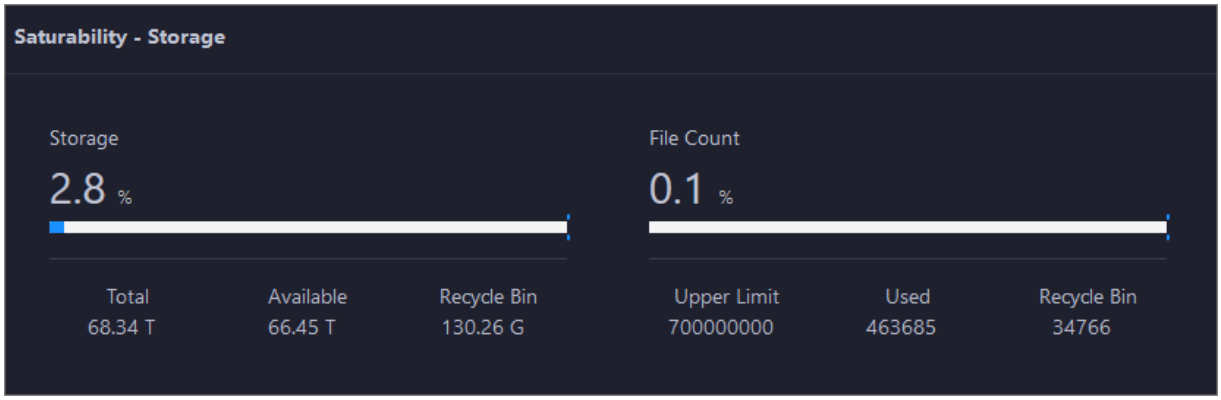

Storage Trend and File Count Trend

This section shows the trend charts of the storage usage and file count. The storage usage chart shows the trend lines of the total storage space, used storage space, and storage usage in different colors. The file count chart shows the trend line of the file count.

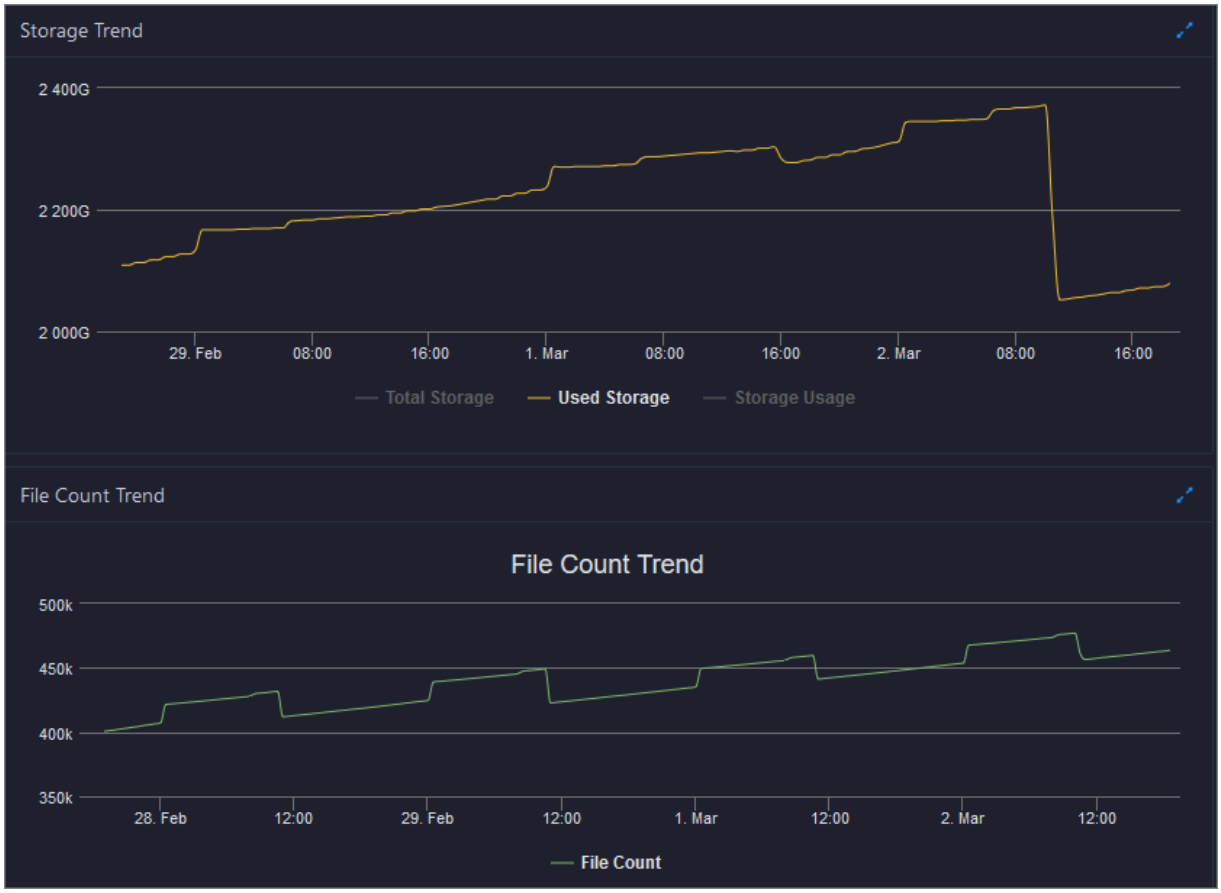

In the upper-right corner of the chart, click the icon to zoom in the chart. The following figure shows an enlarged chart of storage usage.

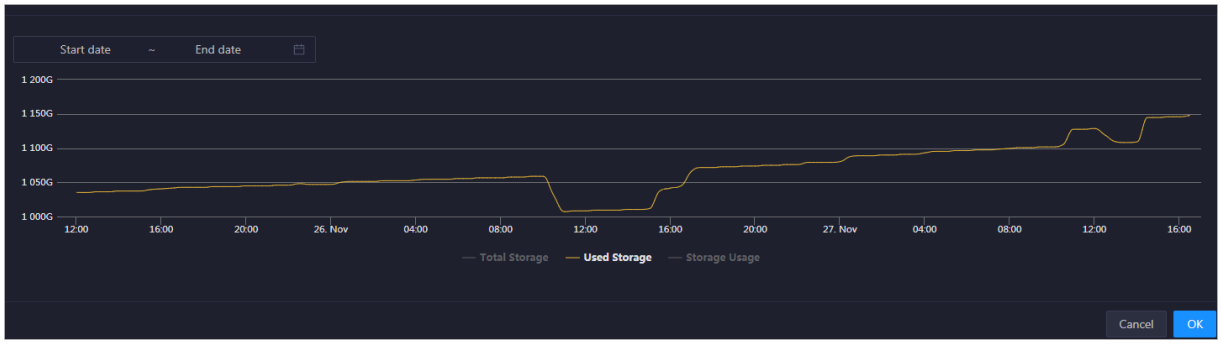

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

Storage Nodes

This section shows information about the storage nodes of Apsara Distributed File System. The information includes the numbers of data nodes, normal nodes, disks, and normal disks. You can also view the faulty node percentage and faulty disk percentage.

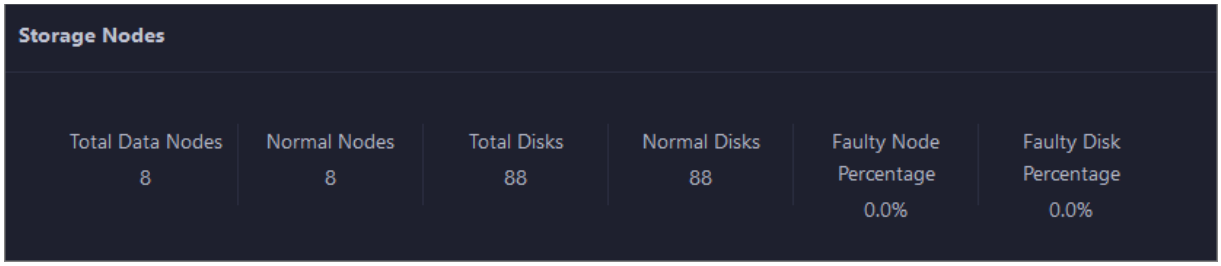

11.9.4.5.3.3. Service roles

The Instances page displays information about the Apsara Distributed File System service roles, including the name, host, IP address, and status of a service role, and host status.

Go to the Instances page

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

   On the **Instances** page, you can view information about the Apsara Distributed File System service roles, including the name, host, IP address, and status of a service role, 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 for Apsara Distributed File System and Run a checkpoint on the master nodes of Apsara Distributed File System.

11.9.4.5.3.4. Service 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.

Go to the Health Status page

1. On the **Services** tab, click **Pangu** in the left-side navigation pane.
2. Select a cluster from the drop-down list and 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.

11.9.4.5.3.5. Storage nodes

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.

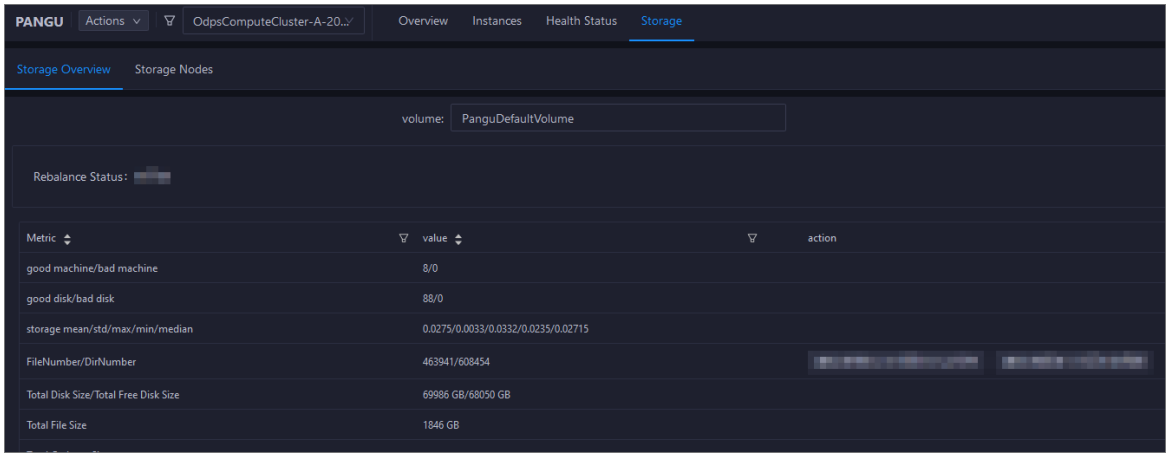

The **Storage Overview** page displays whether data rebalancing is enabled, key metrics and their values, suggestions to handle exceptions, and rack specifications of Apsara Distributed File System. The **Storage Nodes** page displays the information about all storage nodes of Apsara Distributed File 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.

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.

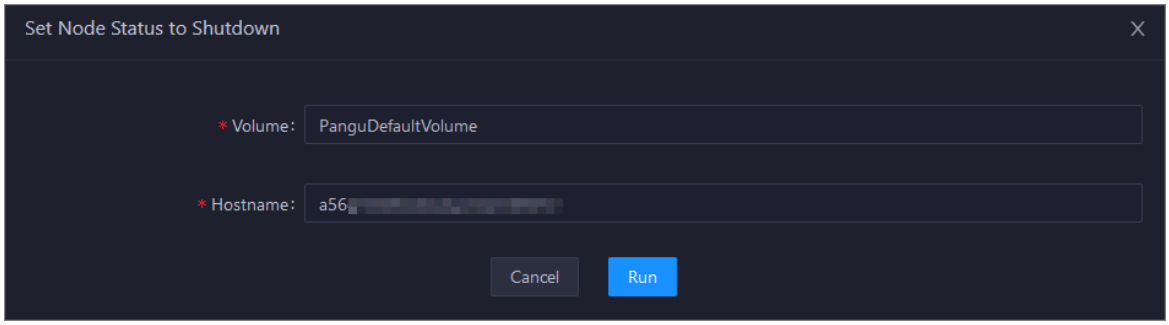

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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.

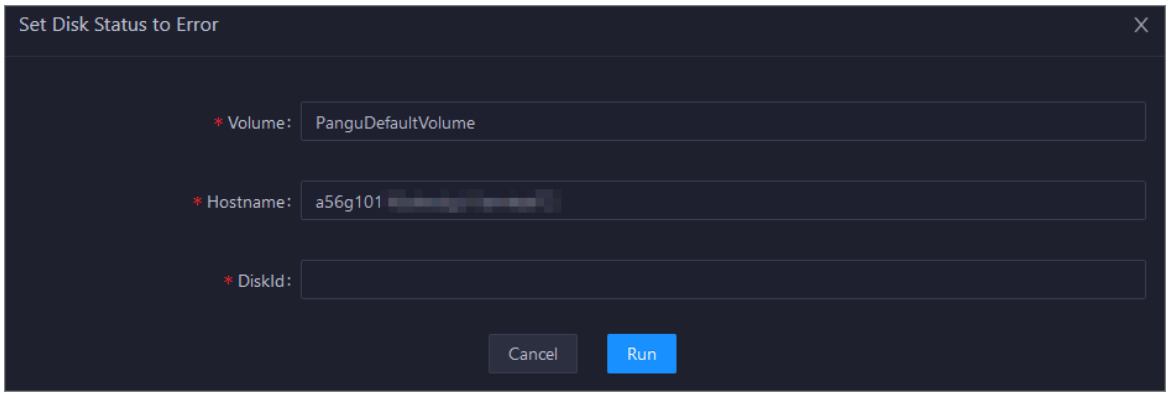

   The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.
3. Click **Run**. A message appears, indicating that the action has been submitted.

11.9.4.5.3.6. Empty the recycle bin of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to clear the recycle bin of Apsara Distributed File System to release storage space.

Prerequisites

Your ABM account has the permission to manage DataHub.

Procedure

1. On the **Services** tab, click **Pangu** in the left-side navigation pane and select a cluster from the drop-down list. The **Overview** page for Apsara Distributed File System appears.
2. Choose **Actions** > **Empty Recycle Bin** in the upper-right corner.
3. In the right-side pane that appears, set the **volume** parameter. The default value is **PanguDefaultVolume**.
4. Click **Run**. A message appears, indicating that the request has been submitted.
5. View the execution status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Empty Recycle Bin**. In the right-side pane that appears, 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 identify the cause of the failure.

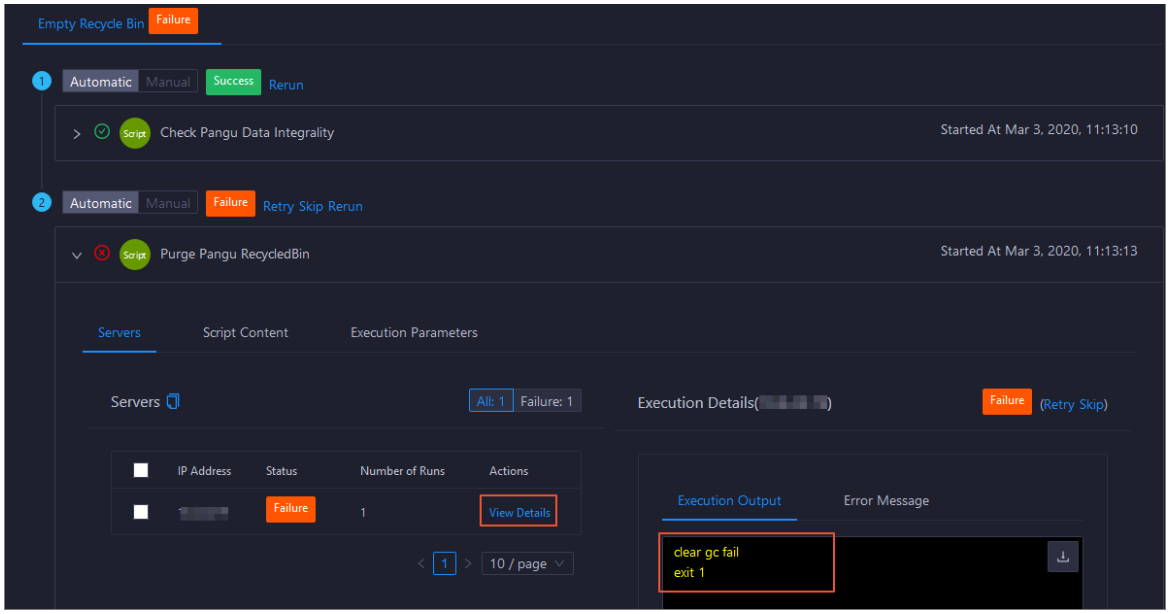

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

11.9.4.5.3.7. Enable or disable data rebalancing for Apsara Distributed File System

ABM allows you to enable or disable data rebalancing for Apsara Distributed File System.

Prerequisites

Your ABM account has the permission to manage DataHub.

Disable data rebalancing

1. On the **Services** tab, click **Pangu** in the left-side navigation pane and select a cluster from the drop-down list. The Overview page for Apsara Distributed File System appears.
2. Choose **Actions** > **Disable Data Rebalancing** in the upper-right corner.
3. In the right-side pane that appears, set the **volume** parameter. The default value is **PanguDefaultVolume**.

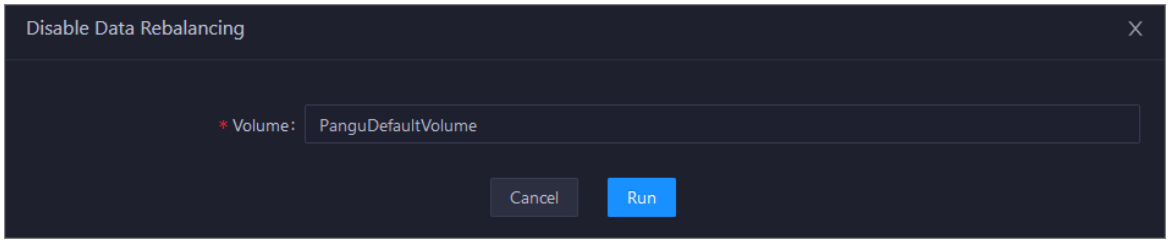

4. Click **Run**. A message appears, indicating that the request has been submitted.
5. View the execution status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Disable Data Rebalancing**. In the right-side pane that appears, 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** tab, click **Pangu** in the left-side navigation pane and select a cluster from the drop-down list. The Overview page for Apsara Distributed File System appears.
2. Choose **Actions** > **Enable Data Rebalancing** in the upper-right corner.
3. In the right-side pane that appears, set **volume**. The default value is **PanguDefaultVolume**.

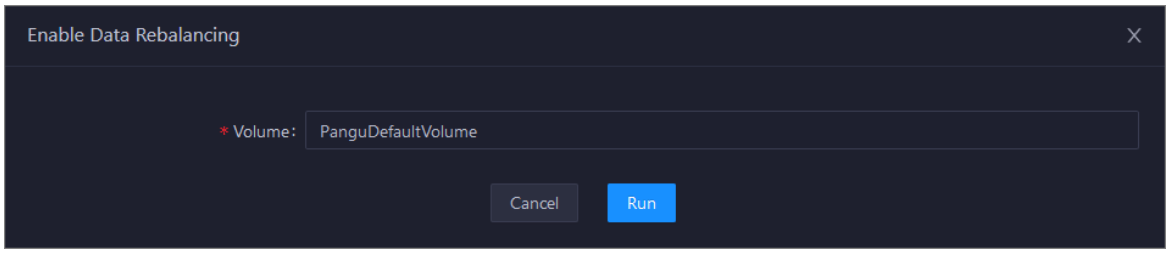

4. Click **Run**. A message appears, indicating that the request has been submitted.
5. View the execution status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Enable Data Rebalancing**. In the right-side pane that appears, 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. Find the target failed execution and click **Details** in the Details column.
2. In the right-side pane that appears, click **View Details** for a failed step to locate the failure cause.

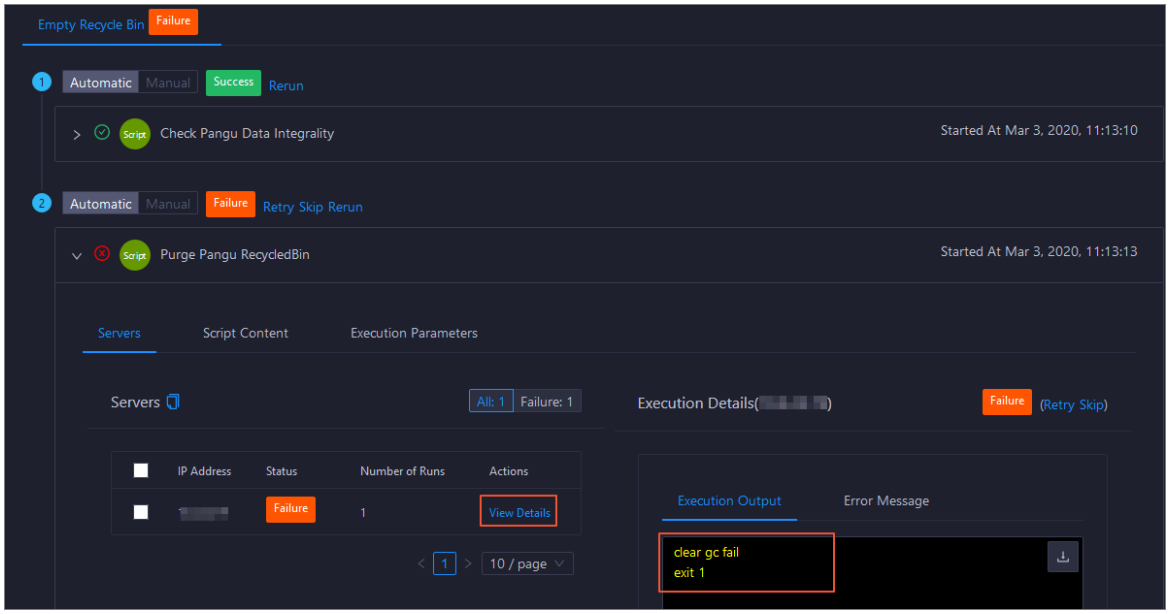

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

11.9.4.5.3.8. Run a checkpoint on master nodes of Apsara Distributed File System

ABM allows you to run checkpoints on master nodes of Apsara Distributed File System. This operation writes memory data to disks. When a failure occurs in 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 has the permission to manage DataHub.

Procedure

1. On the **Services** tab, click **Pangu** in the left-side navigation pane and select a cluster from the drop-down list. The Overview page for Apsara Distributed File System appears.
2. Choose **Actions** > **Run Checkpoint on Master Node** in the upper-right corner.
3. In the right-side pane that appears, set the **volume** parameter. The default value is **PanguDefaultVolume**.
4. Click **Run**. A message appears, indicating that the request has been submitted.
5. View the execution status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Run Checkpoint on Master Node**. In the right-side pane that appears, view the execution history.

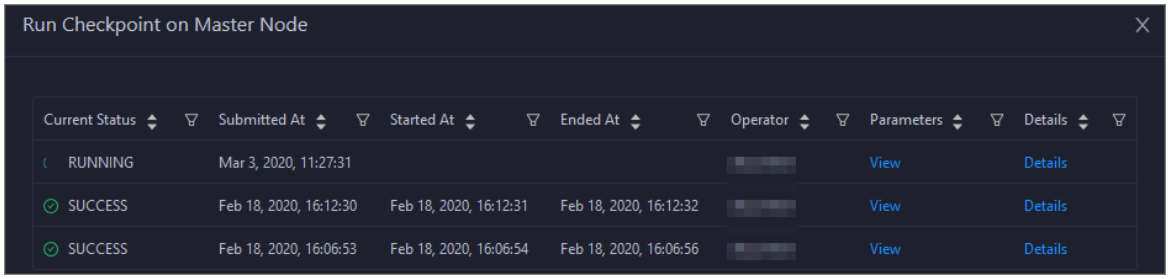

In the Current Status column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

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

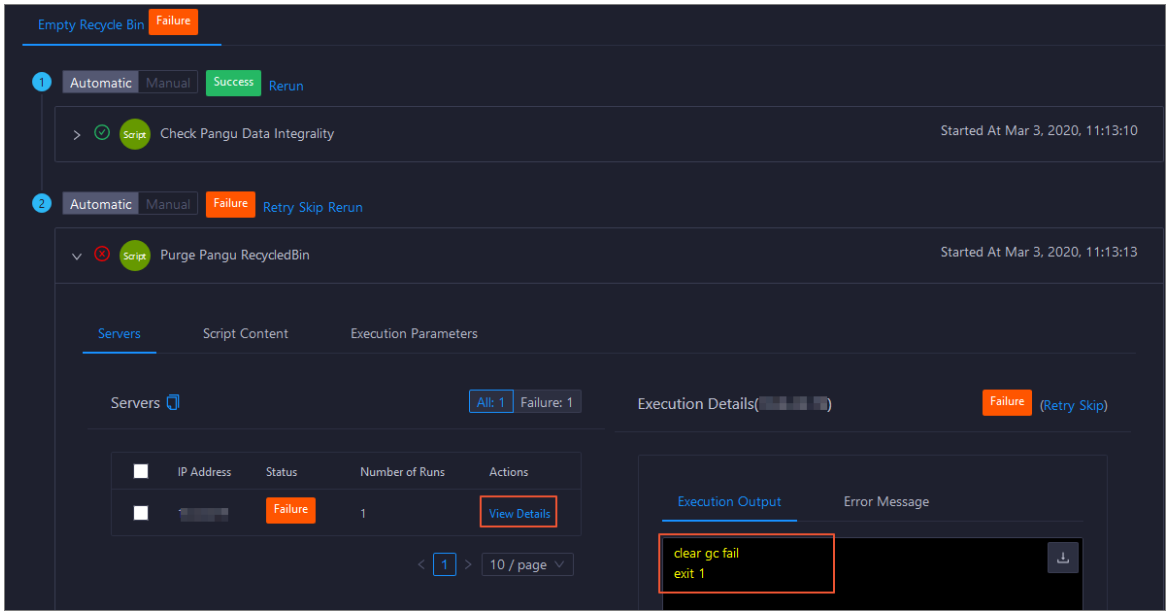

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

11.9.4.5.3.9. Change the primary master node of Apsara Distributed File System

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 has the permission to manage DataHub.
- You have obtained the roles of the primary and secondary master nodes in a volume. To view the role of a master node, log on to the Apsara Infrastructure Management Framework console and access the **PanguTools#** host in the DataHub cluster. Then, run the **puadmin gems** command on the host.
- You have obtained the hostname of the secondary master node that is to be changed to the new primary master node. To view the hostname, perform the following steps: Log on to the ABM console, go to the O&M page for DataHub, and then click **Services**. On the page that appears, click

Pangu in the left-side navigation pane and click the **Instances** tab. On the **Instances** page, view the hostnames of **PanguMaster#** hosts.

Background information

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. On the **Services** tab, click **Pangu** in the left-side navigation pane and select a cluster from the drop-down list. The **Overview** page for Apsara Distributed File System appears.
2. Choose **Actions** > **Change Primary Master Node** in the upper-right corner. In the right-side pane that appears, set the parameters.

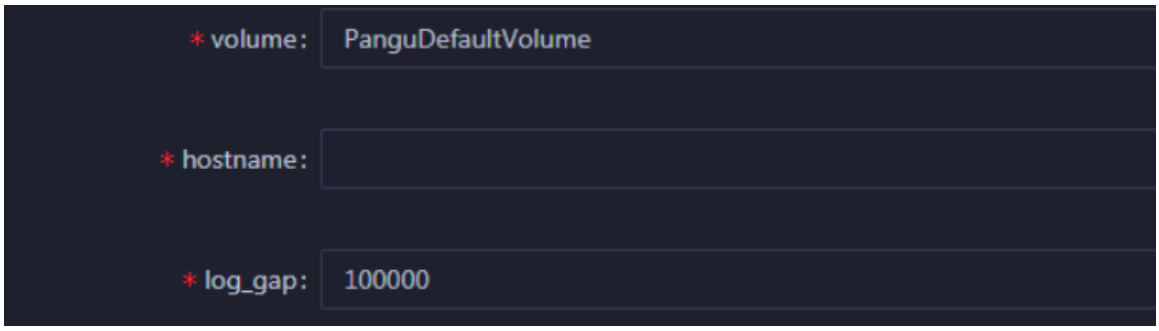

   You must set the following parameters in this step:

   - **volume**: the volume whose primary master node is 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 is to be changed.
   - **hostname**: the hostname of the secondary master node that is changed to be the new primary master node.
   - **log_gap**: the maximum log number gap between the original primary and secondary master nodes. 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**.
3. Click **Run**. A message appears, indicating that the request has been submitted.
4. View the execution status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Change Primary Master Node**. In the right-side pane that appears, view the execution history.

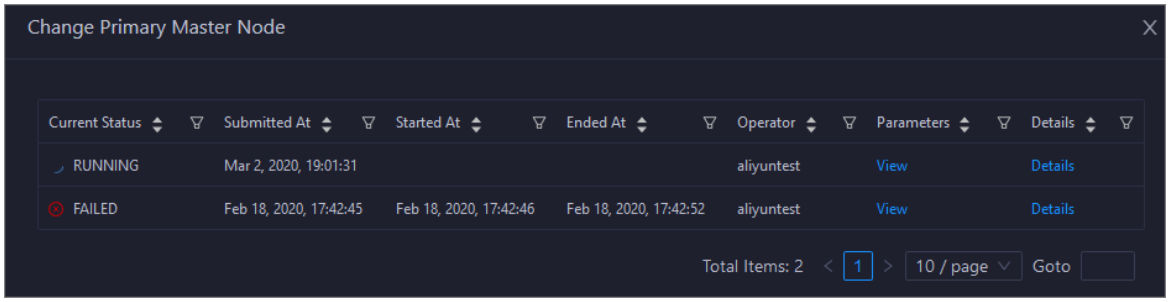

In the Current Status column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

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

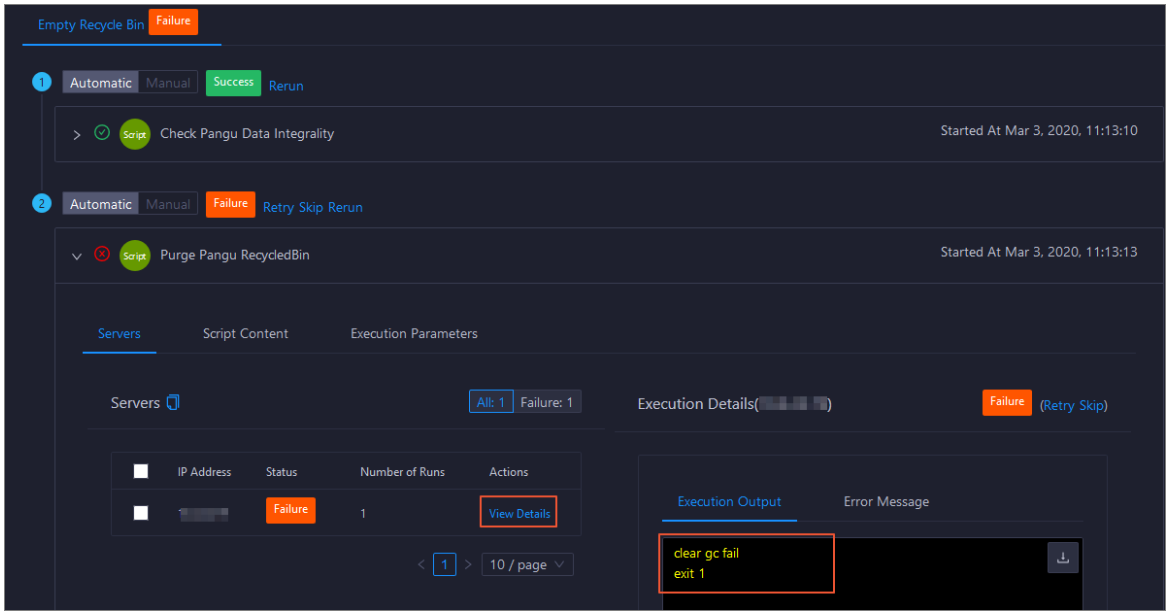

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

11.9.4.6. Cluster O&M

11.9.4.6.1. Cluster O&M entry

This topic describes how to go to the cluster O&M page for DataHub in the ABM console.

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

11.9.4.6.2. Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the host status, service status, 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.

Go to the Overview page for a cluster

1. On the **O&M** page, click the **Clusters** tab.
2. On the **Clusters** tab, select a cluster in the left-side navigation pane and click the **Overview** tab. The Overview page for the cluster appears.

Hosts

This section displays the respective number of hosts in different states in the cluster. A host may be in one of the following states: good, bad, and upgrading.

Services

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

Health Check

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

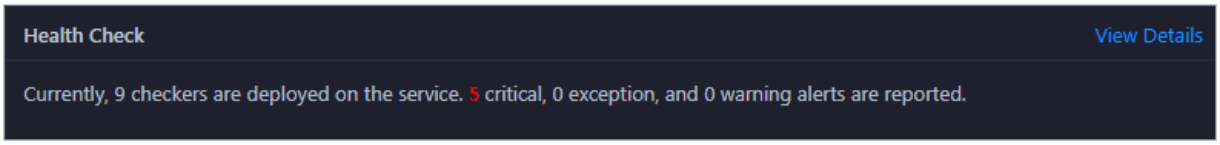

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Cluster health.

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.

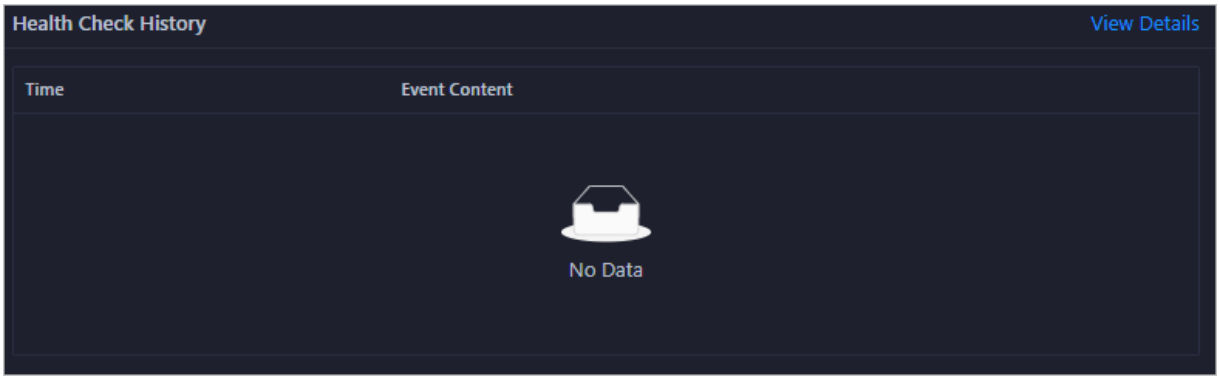

You can click the event content of a check to view the anomalous items.

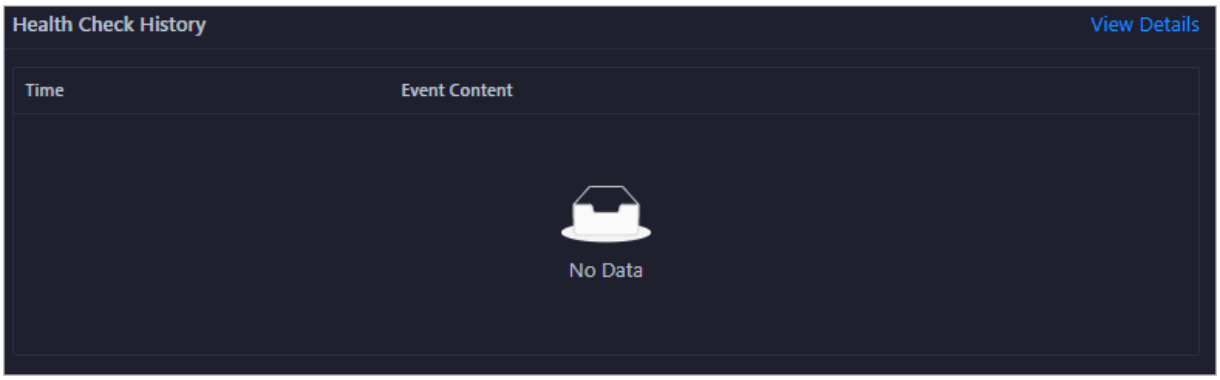

CPU

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

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

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

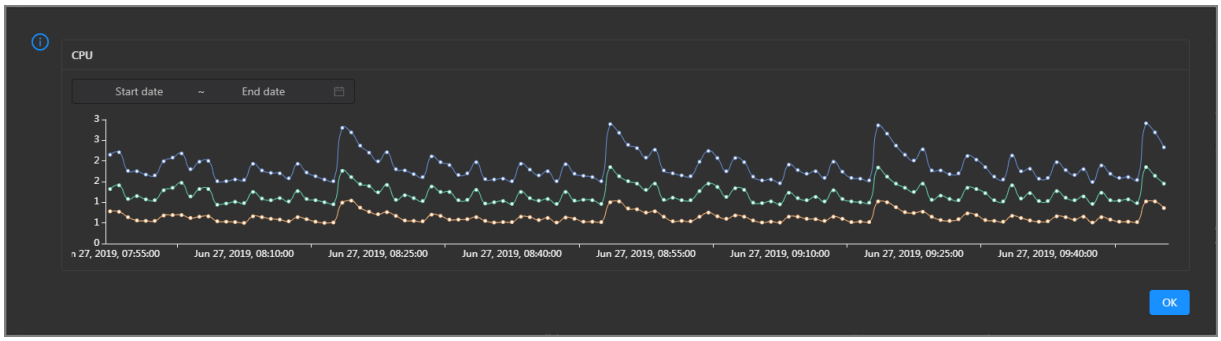

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

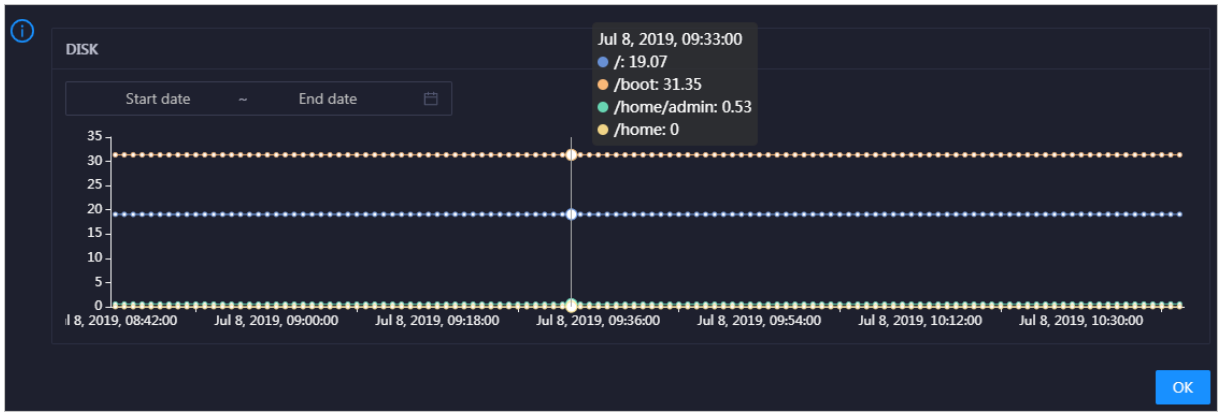

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

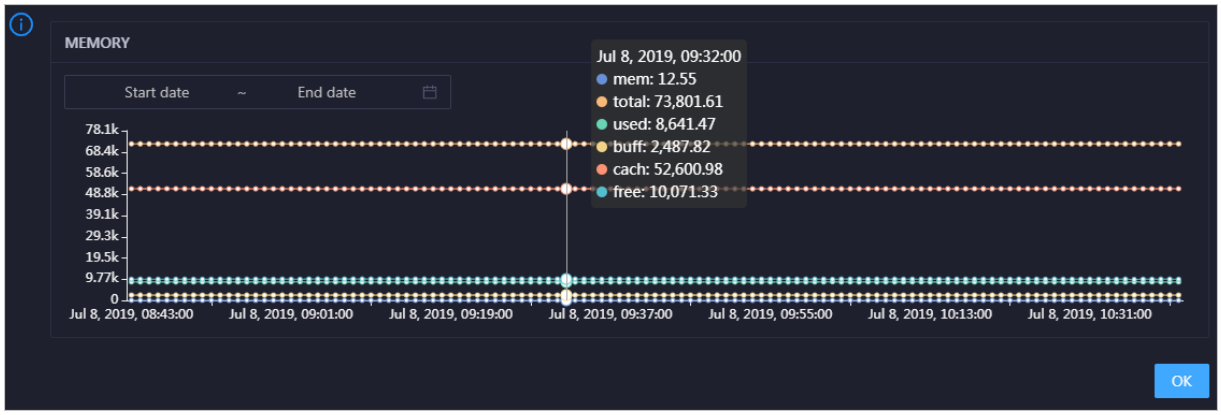

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

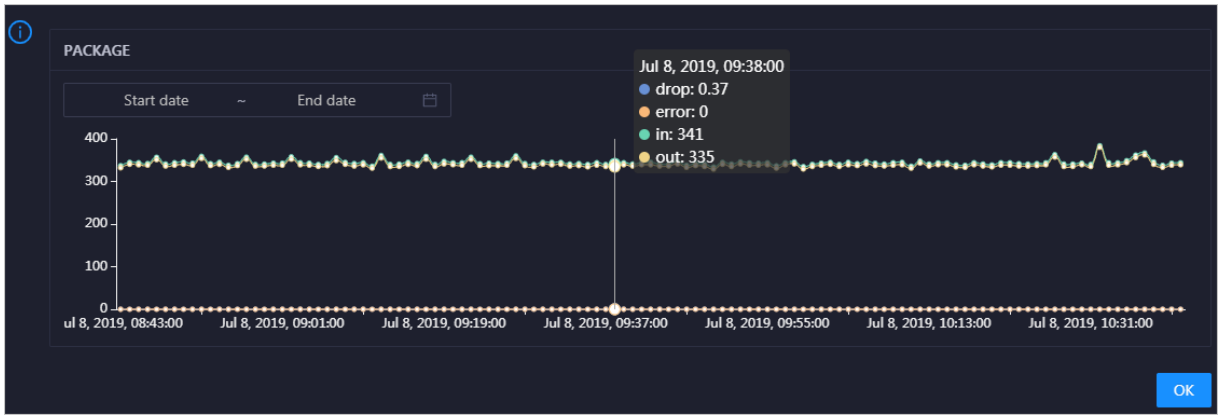

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

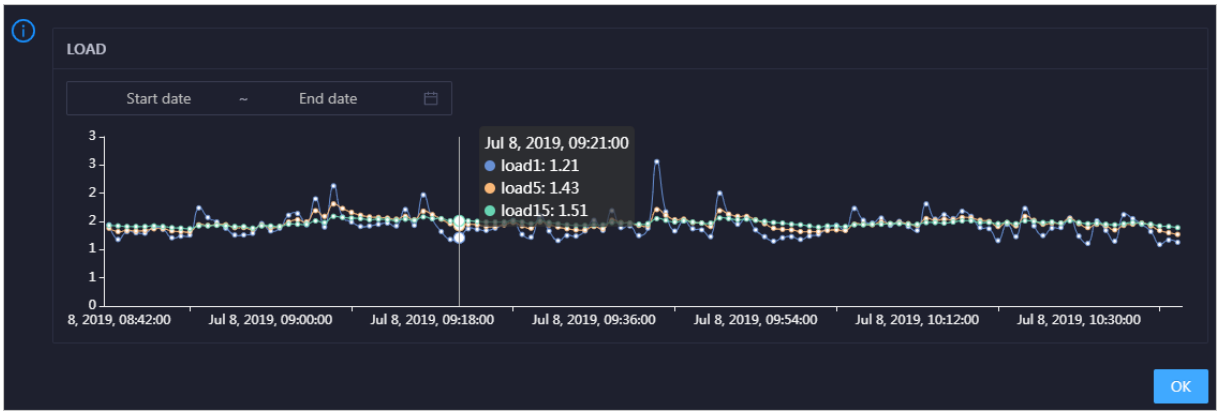

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.9.4.6.3. Cluster health

The Health Status page displays the information about the checkers of the selected cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts. In addition, you can log on to a host and perform manual checks on the host.

Go to the Health Status page

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

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

View checker details

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

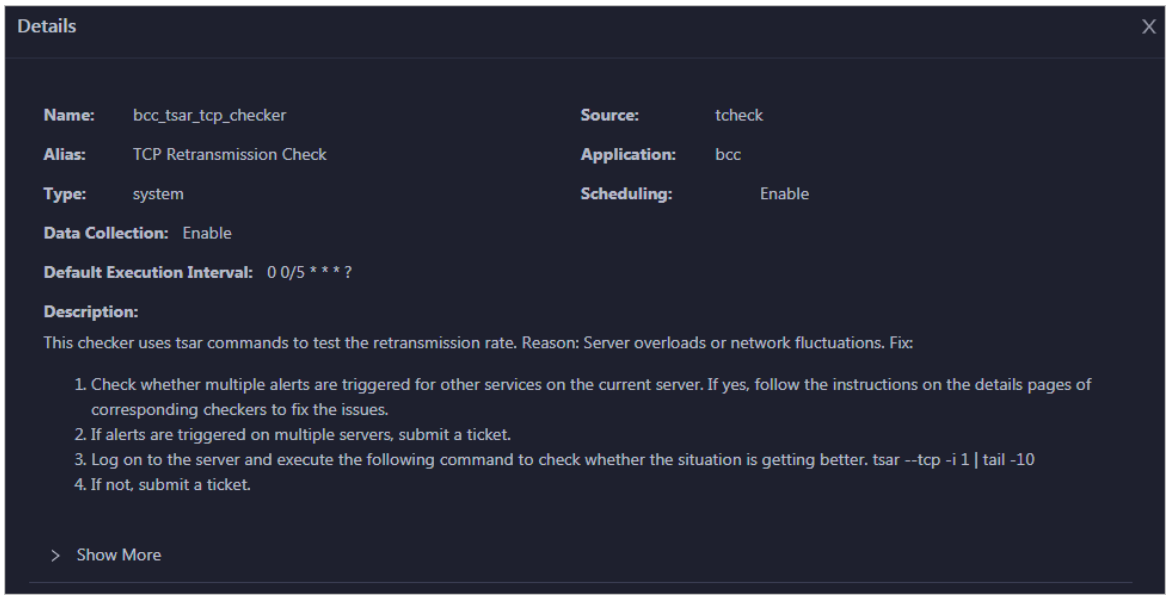

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

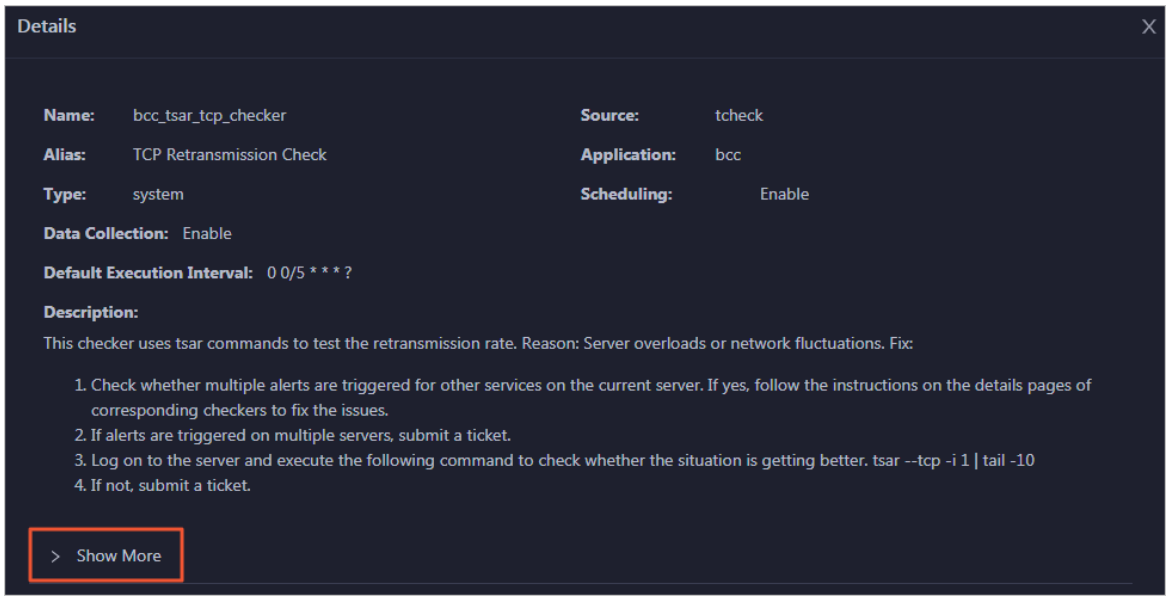

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

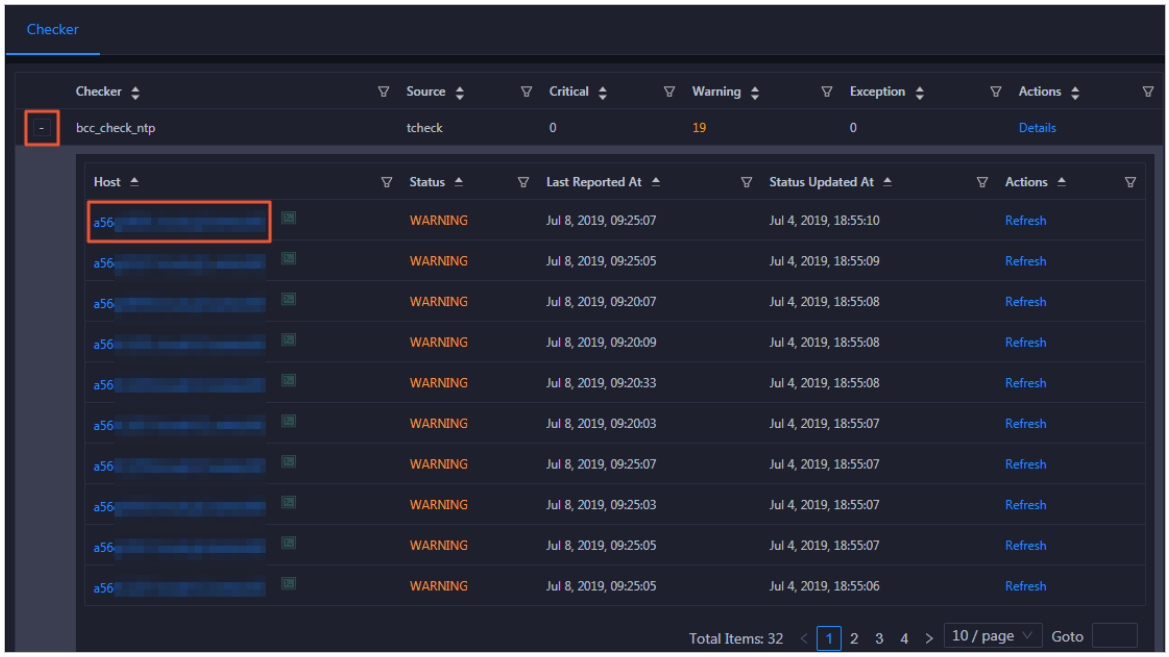

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

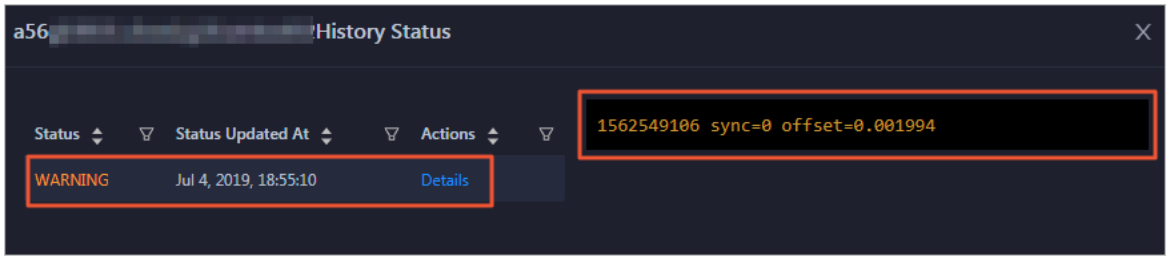

Clear alerts

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

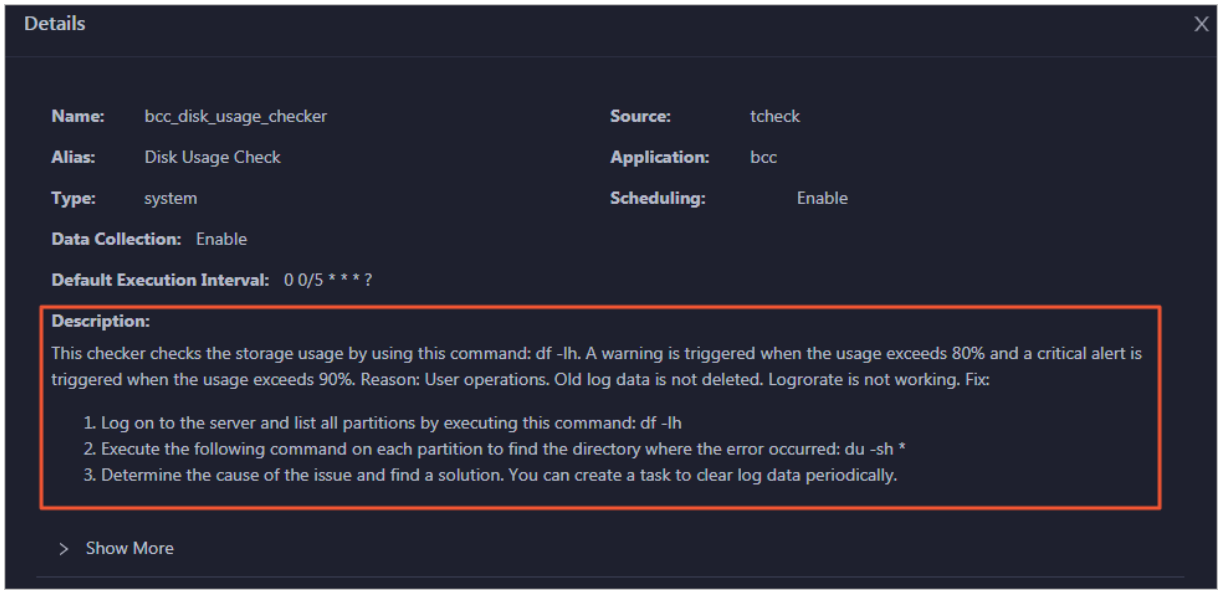

Log on to a host

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

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

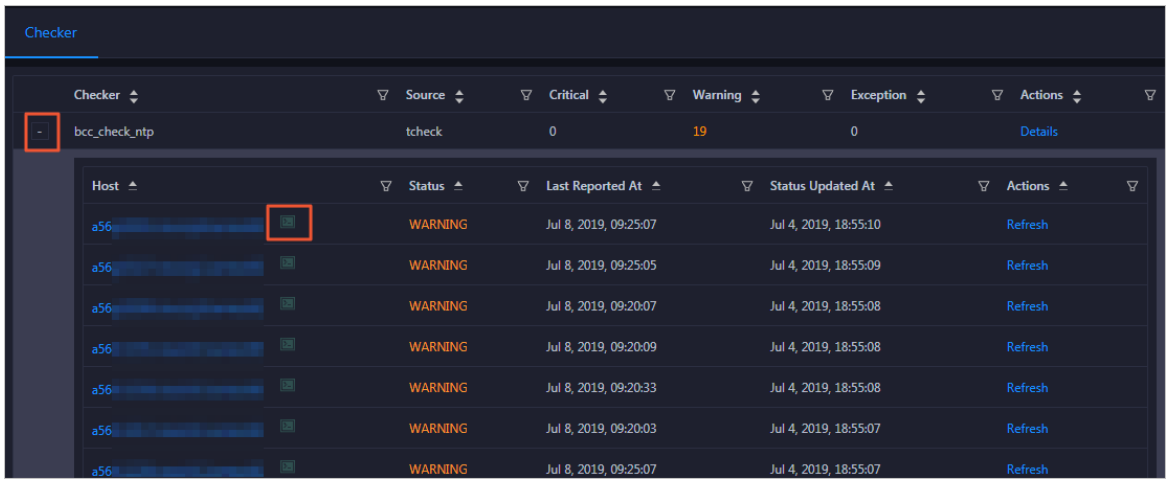

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

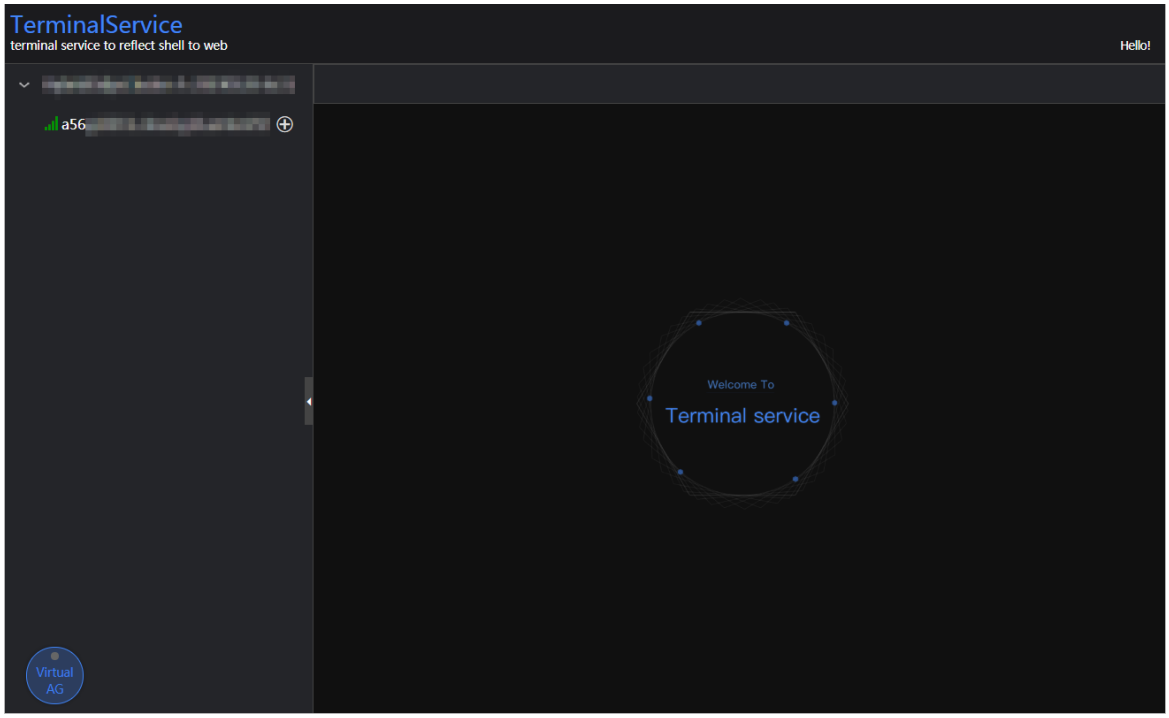

3. On the **TerminalService** page, click the hostname to log on to the host.

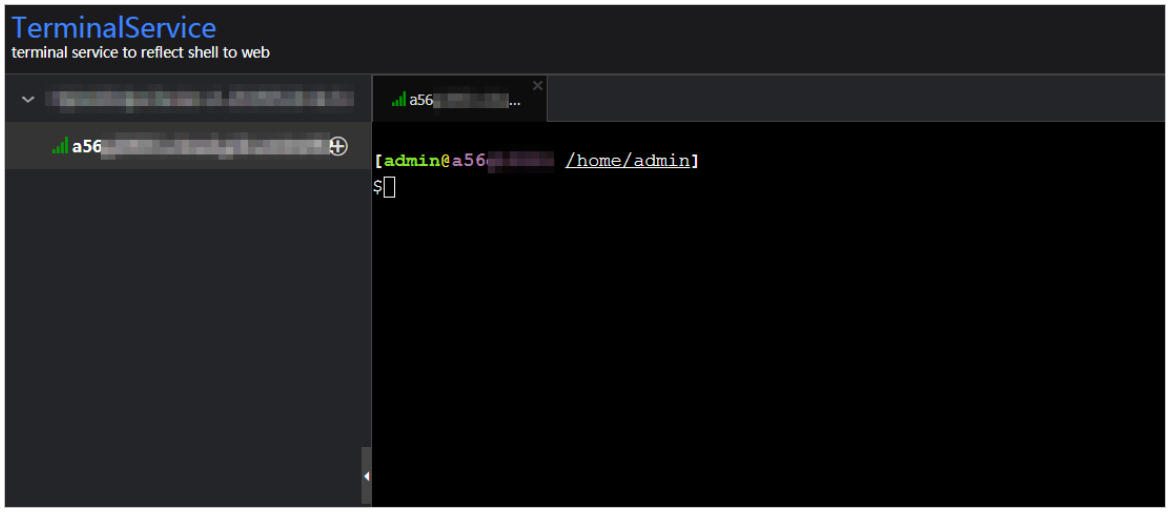

Run a checker again

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

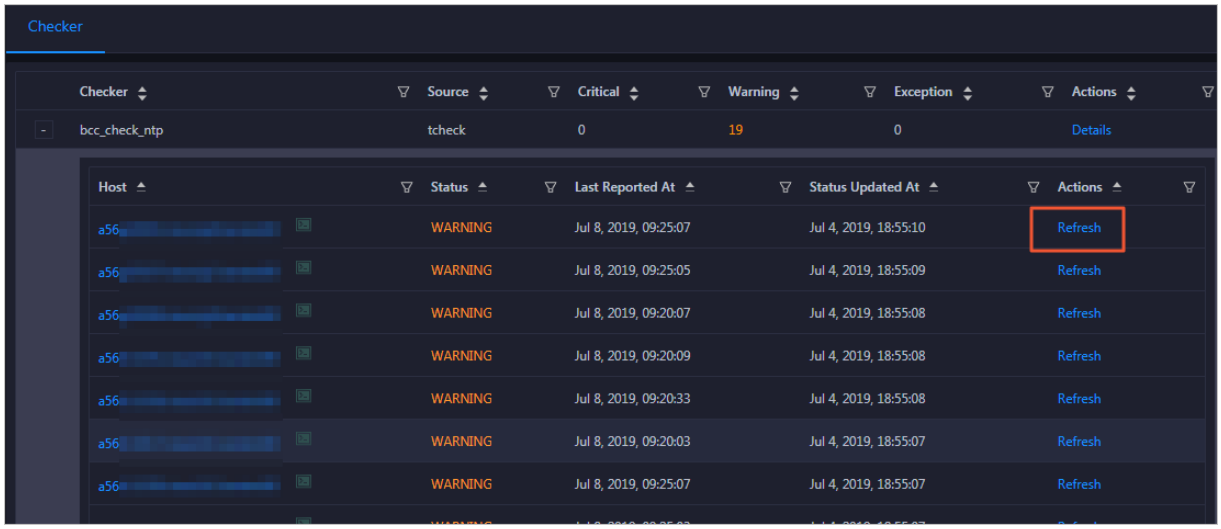

11.9.4.6.4. Cluster hosts

The cluster 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.

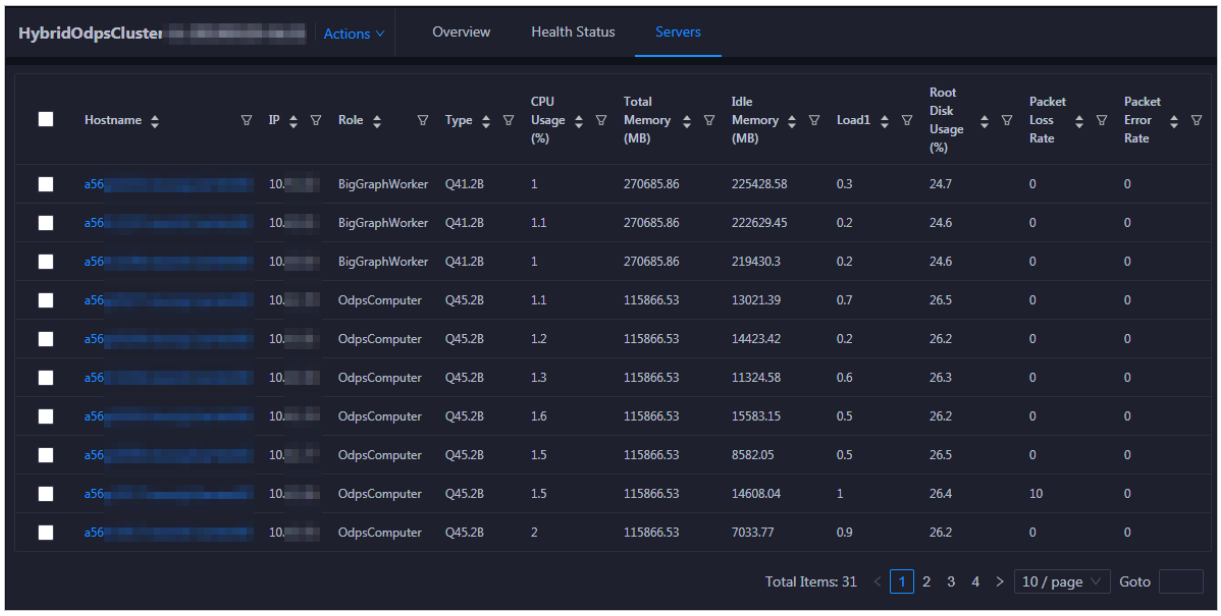

To view more information about a host, click the name of the host. The Overview tab of the Hosts page appears. For more information, see Host overview.

11.9.4.6.5. Cluster scale-out

This topic describes how to scale out a DataHub cluster in the ABM console. Cluster scale-out refers to the process of adding physical hosts in the default cluster of Apsara Infrastructure Management Framework to a DataHub cluster. The physical hosts of a DataHub cluster include chunkserver and frontend hosts.

Prerequisites

- The physical hosts to be added to a DataHub cluster are available in the default cluster of Apsara Infrastructure Management Framework.
- The default cluster of Apsara Infrastructure Management Framework has hosts whose **project** is **datahub**.

> **Note** Scale-out is only available for **chunkserver** and **frontend** hosts in a DataHub cluster.

Background information

In Apsara Stack, scaling out a cluster involves complex operations. You must 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 can be considered as an idle resource pool that provides resources for scaling out clusters 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.

Step 1: Obtain the name of the host to be added to a DataHub cluster

1. Log on to the ABM console.
2. Click  in the upper-left corner and select **TIANJI** to log on to the Apsara Infrastructure Management Framework console.
3. In the left-side navigation pane, choose **Operations > Machine Operations**.
4. On the **Machine Operations** page, search for a host whose project is **datahub** in the **default** cluster. Then, copy the name of the host.

Step 2: Add the host to the target DataHub cluster

You can add multiple hosts to a DataHub cluster at a time to scale out the cluster. To scale out a cluster, you must first specify an existing host as the template host. When you scale out the DataHub 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 Clusters tab, select the target cluster in the left-side navigation pane, click the **Hosts** tab, and then select a host whose role is **chunkserver** or **frontend** as the template host.
2. Choose **Actions > Scale out Cluster** in the upper-right corner. In the **Scale out Cluster** right-side pane, set relevant parameters.

   You must set the following parameters in this step:

   - **Refer Hostname**: the name of the template host. The name of the selected host is used by default.
   - **hostname**: the name of the host to be added to the DataHub 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 request has been submitted.
4. View the scale-out status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Scale out Cluster**. In the right-side pane that appears, view the execution status.

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

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

Locate the failure cause

1. On the **Clusters** tab, move the pointer over **Actions** in the upper-left corner and select **Execution History** next to **Scale out Cluster**. In the right-side pane that appears, view the execution history.
2. If the status of a record is FAILED, click **Details** to locate the failure cause.

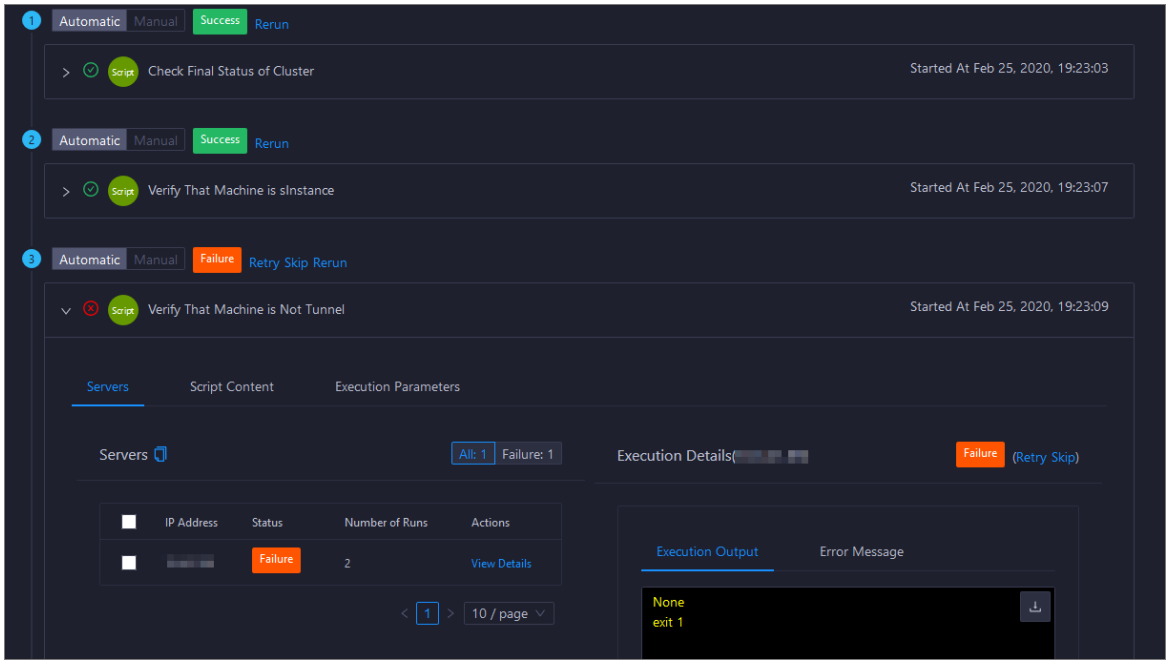

You can also view the parameter settings, host details, script, and execution parameters to locate the failure cause.

11.9.4.6.6. Cluster scale-in

This topic describes how to scale in a DataHub cluster in the ABM console. Cluster scale-in refers to the process of removing physical hosts from a DataHub cluster to the default cluster of Apsara Infrastructure Management Framework. The physical hosts of a DataHub cluster include chunkserver and frontend hosts.

Prerequisites

- Scale-in is only available for **chunkserver** and **frontend** hosts in a DataHub cluster.
- The following operations are performed before you remove one or more **chunkserver** hosts:
  - Run the df command to check the disk usage on each host. Calculate whether the disk will be full after a specific number of hosts are removed. If so, we recommend that you do not perform the scale-in.
  - Shards on the removed hosts will be migrated to other hosts. Therefore, you must log on to the webconsole host to calculate the shard load on each host after the scale-in. If the number of shards on a host exceeds 1,000, performance may be affected. In this case, we recommend that you do not perform the scale-in.
- The following operations are performed before you remove one or more **frontend** hosts:
  - Run the df command to check the disk usage on each host. Calculate whether the disk will be full after a specific number of hosts are removed. If so, we recommend that you do not perform the scale-in.
  - Shards on the removed hosts will be migrated to other hosts. Therefore, you must log on to the webconsole host to calculate the shard load on each host after the scale-in. If the number of shards on a host exceeds 1,000, performance may be affected. In this case, we recommend that you do not perform the scale-in.

- Check the traffic and queries per second (QPS). If the traffic exceeds 400 MBit/s or the QPS exceeds 15,000, we recommend that you do not perform the scale-in.

Background information

In Apsara Stack, scaling out a cluster involves complex operations. You must 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 can be considered as an idle resource pool that provides resources for scaling out clusters 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.

Procedure

1. On the O&M page of the ABM console, click the **Clusters** tab. On the Clusters tab, select the target cluster in the left-side navigation pane, click the **Hosts** tab, and then select one or more hosts whose role is **chunkserver** or **frontend**.
2. On the Clusters tab, choose **Actions > Scale in Cluster** in the upper-right corner. In the **Scale in Cluster** right-side pane, set the following parameter:

   Hostname: the name of the host to be removed from the DataHub cluster. The name of the selected host is used by default.
3. Click **Run**. A message appears, indicating that the request has been submitted.
4. View the scale-in status.

   Move the pointer over **Actions** in the upper-right corner and select **Execution History** next to **Scale in Cluster**. In the right-side pane that appears, view the execution status.

   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.

   > **Note** If the status is **FAILED**, click Details in the Details column to locate the failure cause. For more information, see Locate the failure cause.
5. (Optional)View the scale-in progress.

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

Locate the failure cause

1. On the **Clusters** tab, move the pointer over **Actions** in the upper-left corner and select **Execution History** next to **Scale in Cluster**. In the right-side pane that appears, view the execution history.
2. If the status of a record is FAILED, click **Details** to locate the failure cause.

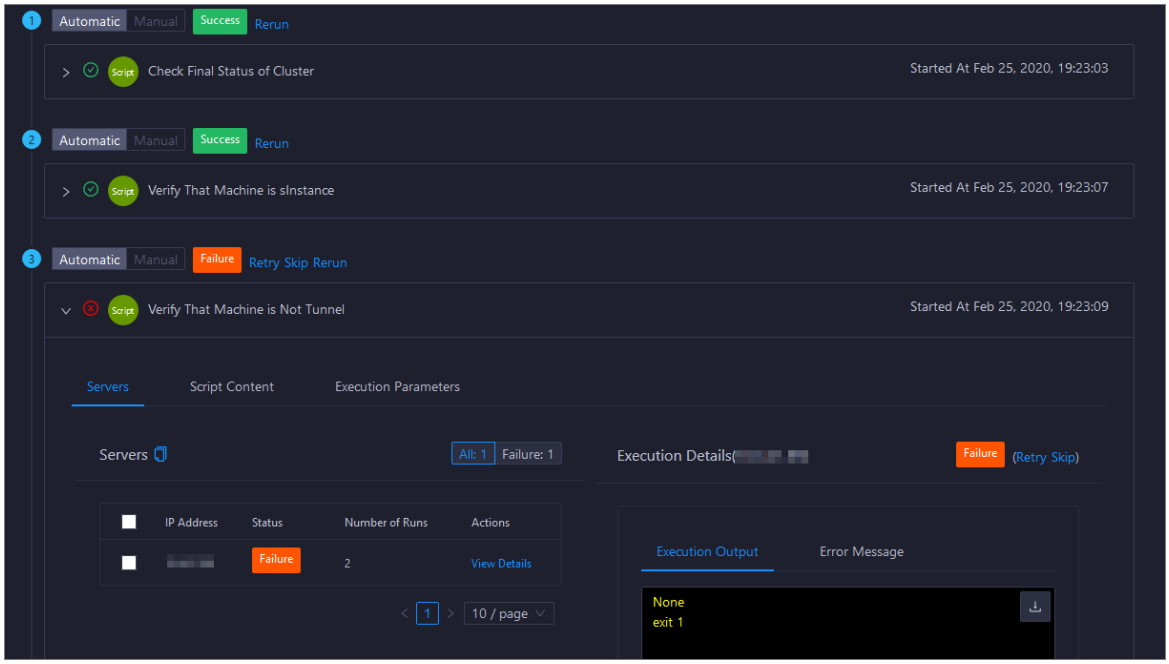

You can also view the parameter settings, host details, script, and execution parameters to locate the failure cause.

11.9.4.6.7. Delete topics from a smoke testing project

Apsara Bigdata Manager (ABM) allows you to delete topics from a DataHub test project and view the execution history.

1. On the **Clusters** page, select a cluster in the left-side navigation pane. Click the **Hosts** tab. The **Hosts** page for the cluster appears.
2. On the Clusters page, choose **Actions > Delete Topic from Smoke Testing**. The **Delete Topic from Smoke Testing** dialog box appears.
3. Click **Run**. A message appears, indicating that the action has been submitted.
4. View the history of deleting topics.

   Click **Actions** in the upper-left corner, and then click **Execution History** next to **Delete Topic from Smoke Testing** 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.

   - If the status is FAILED, click **Details** to locate the failure cause.

     You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.
   - If the status is SUCCESS, click **Details** to view the execution result. On the page that appears, click **View Details** in the **Actions** column. The execution result including the time when the job was run and the IP address of the host appears in the Execution Details section in the lower-right corner.

11.9.4.6.8. Reverse parse RequestId

Apsara Bigdata Manager (ABM) allows you to reverse parse RequestId in DataHub to obtain the time when a job was run and the IP address of the host. You can use the obtained information to query logs for troubleshooting.

1. On the **Clusters** page, select a cluster in the left-side navigation pane. Click the **Hosts** tab. The **Hosts** page for the cluster appears.
2. On the Clusters page, choose **Actions > Reverse Parse Request ID**. In the **Reverse Parse Request ID** dialog box that appears, set **Request Id**.
3. Click **Run**. A message appears, indicating that the action has been submitted.
4. View the reverse parsing status.

   Click **Actions** in the upper-left corner, and then click **Execution History** next to **Reverse Parse Request ID** 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.

   - If the status is FAILED, click **Details** to locate the failure cause.

     You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.
   - If the status is SUCCESS, click **Details** to view the execution result. On the page that appears, click **View Details** in the **Actions** column. The execution result including the time when the job was run and the IP address of the host appears in the Execution Details section in the lower-right corner.

11.9.4.7. Host O&M

11.9.4.7.1. Host O&M entry

This topic describes how to go to the host O&M page for DataHub in the ABM console.

1. Log on to the ABM console.
2. Click in the upper-left corner and select **DataHub**.
3. On the DataHub page, 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.

11.9.4.7.2. Host overview

The host overview page displays the overall running information about a host in a DataHub 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

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.

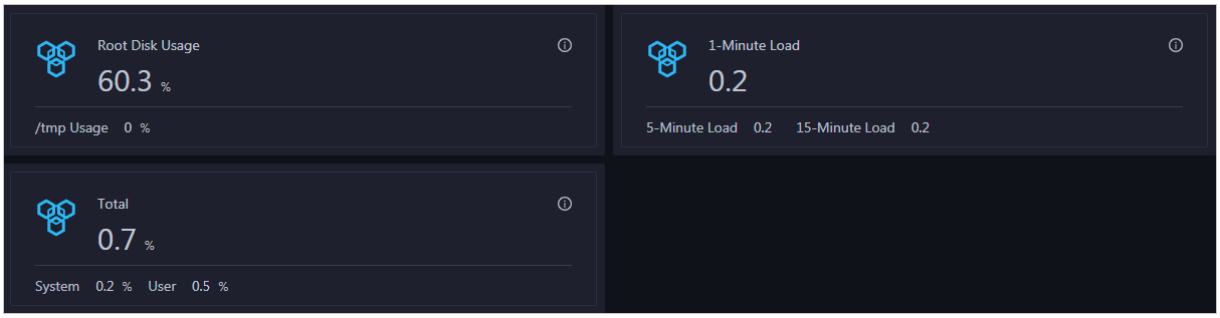

CPU

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.

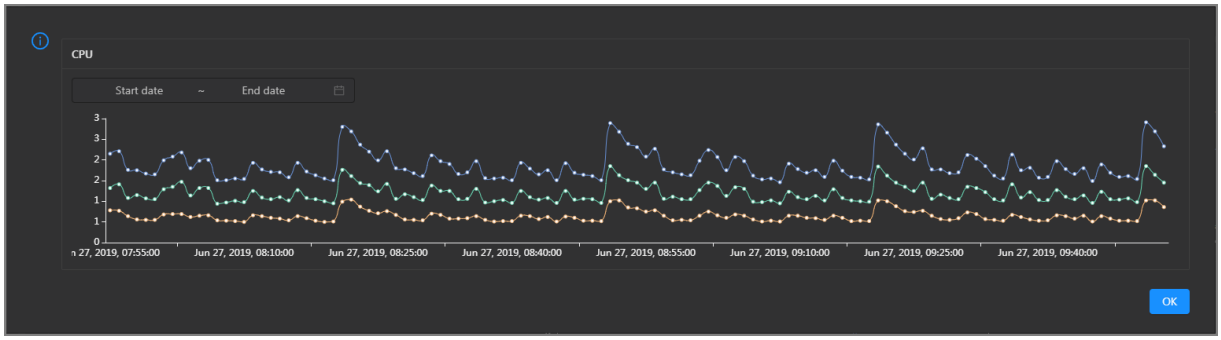

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

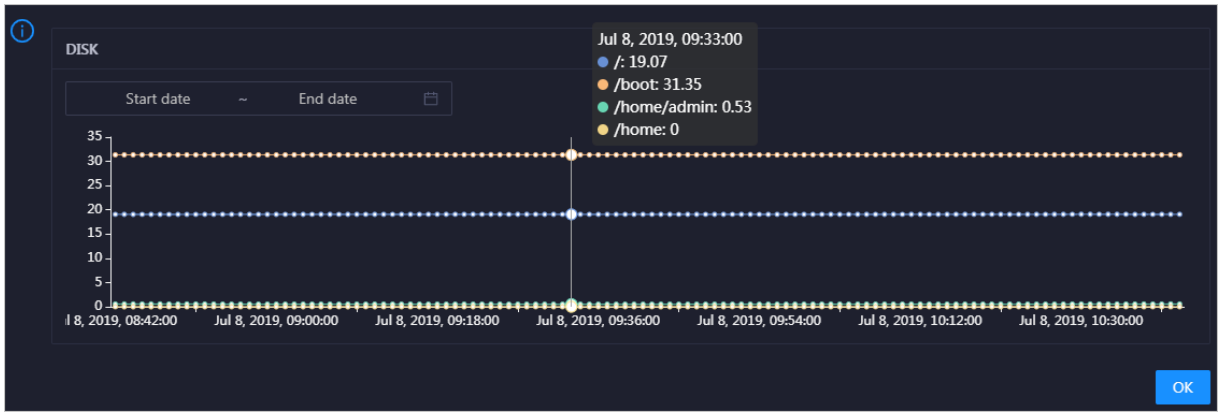

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

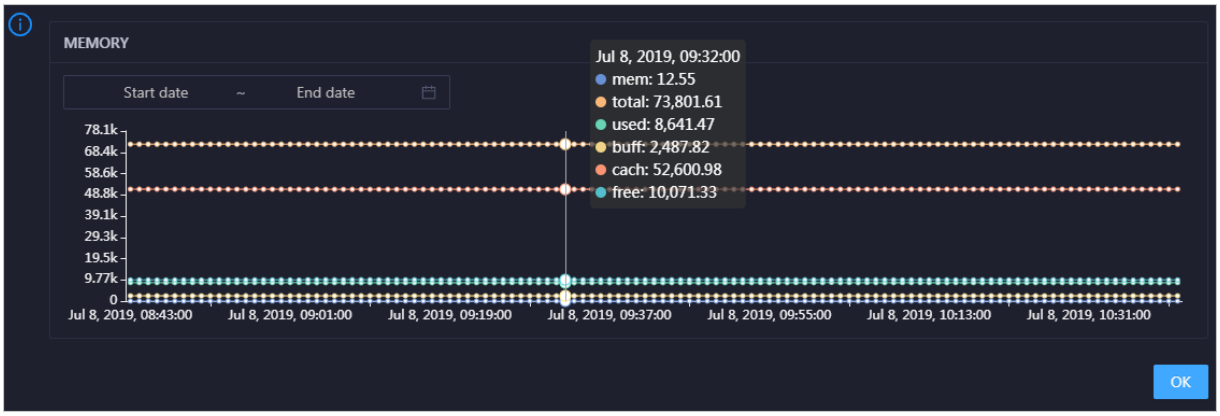

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

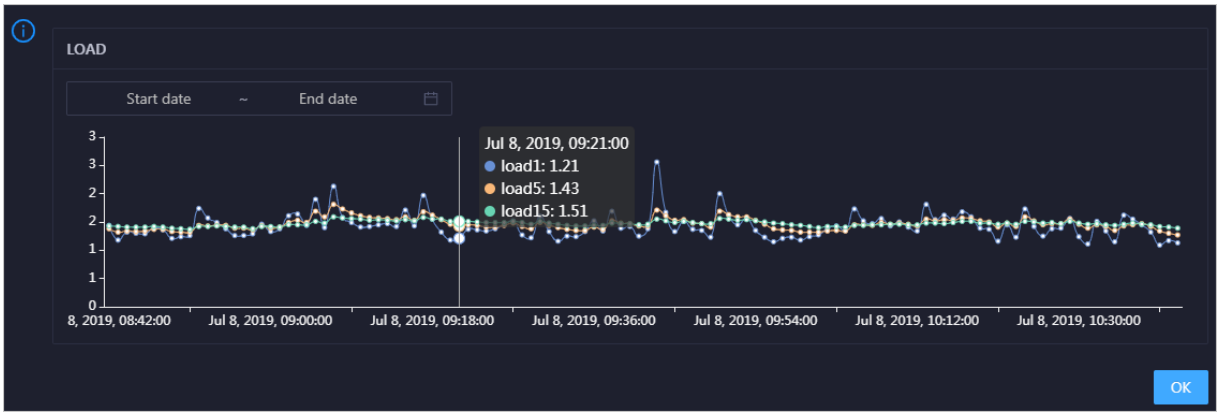

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

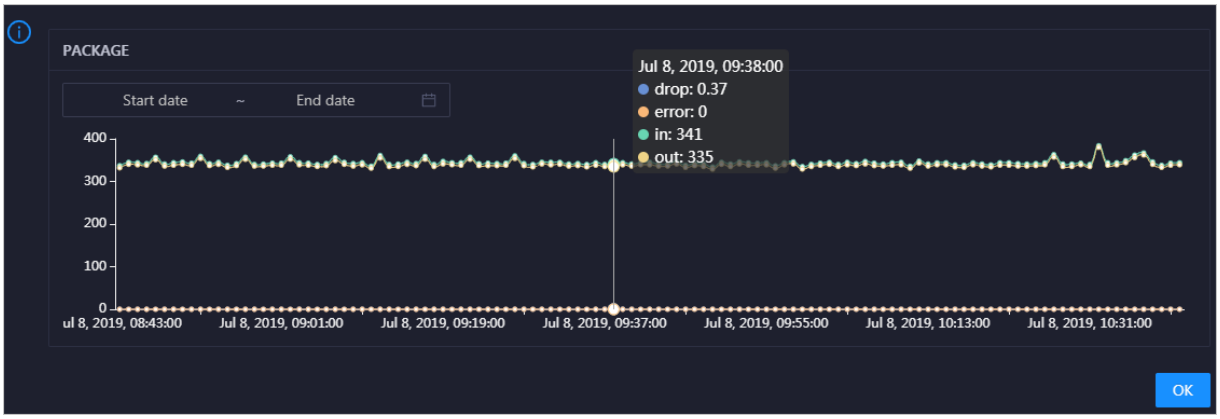

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

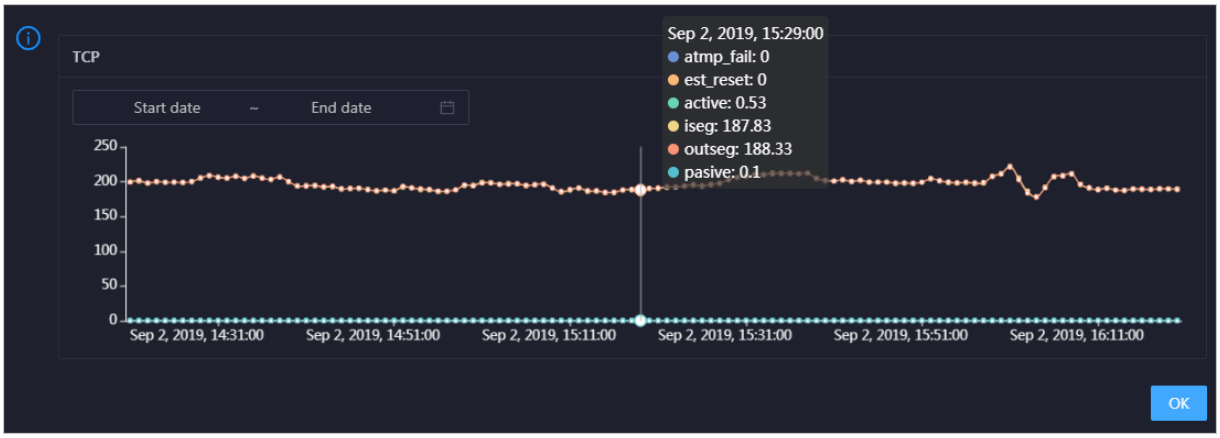

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

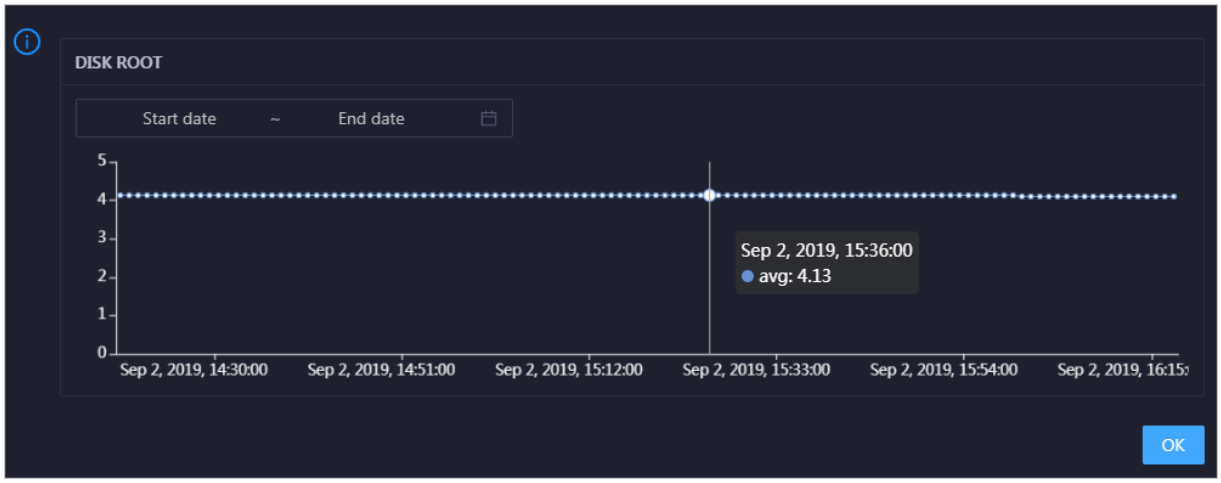

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

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.

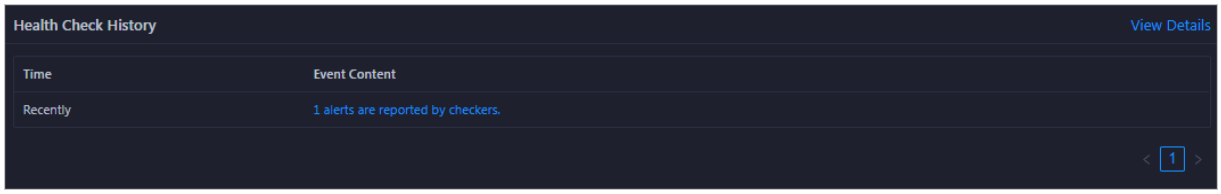

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

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

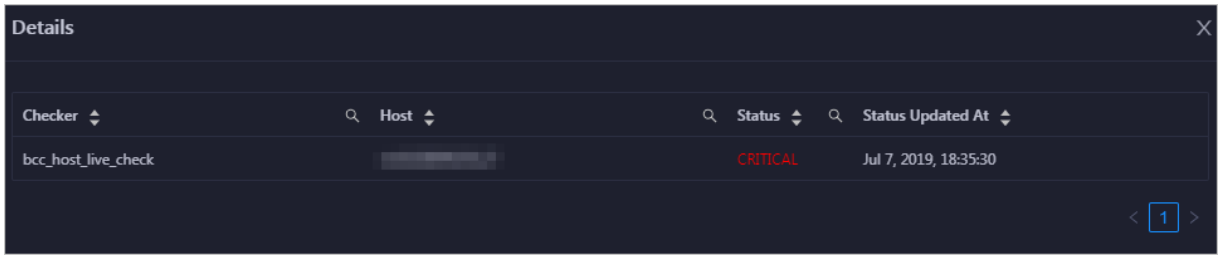

11.9.4.7.3. Host charts

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.

11.9.4.7.4. Host health

The Health Status page displays the information about the checkers of the selected host, including the checker details, check results, and schemes to clear alerts. In addition, you can log on to the host and perform manual checks on the host.

Go to the Health Status page

On the **Hosts** tab, select a host in the left-side navigation pane and 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.

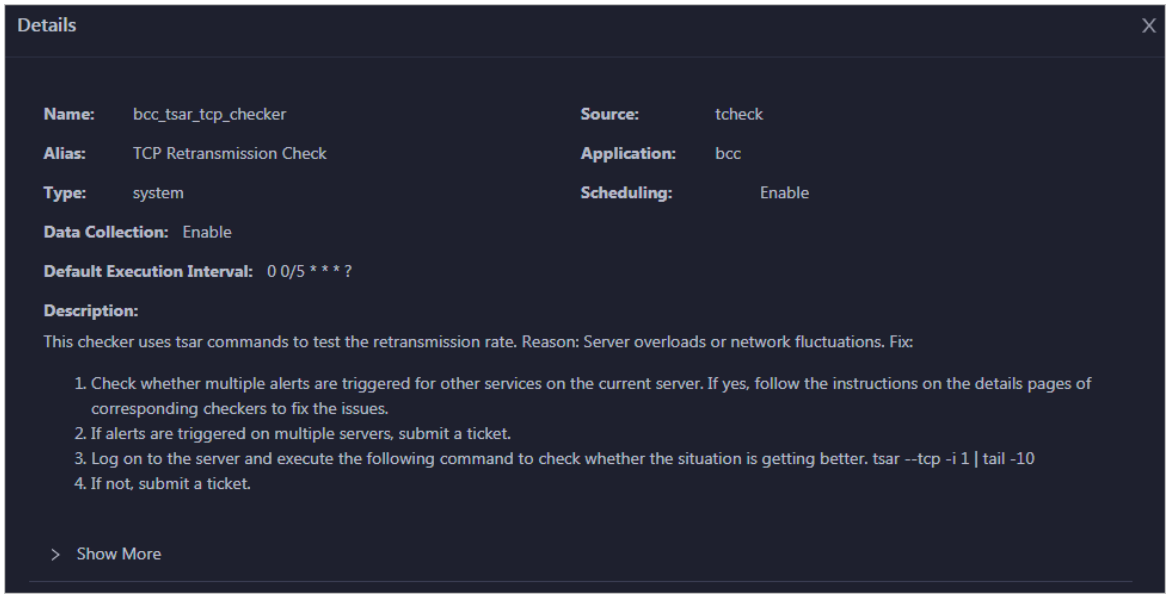

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

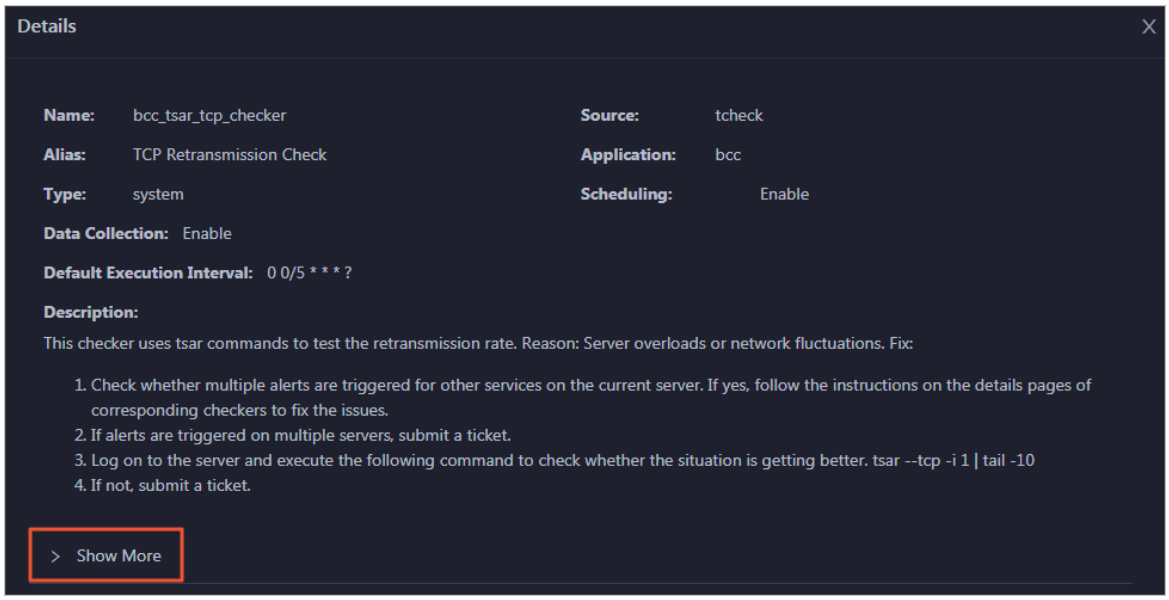

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

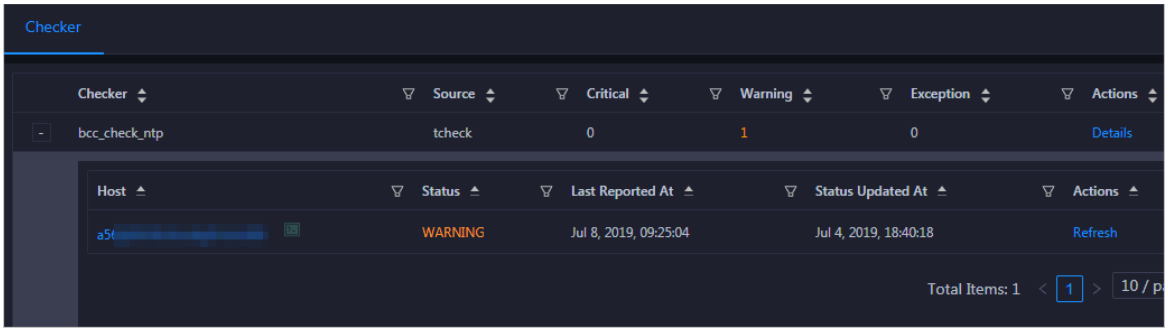

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

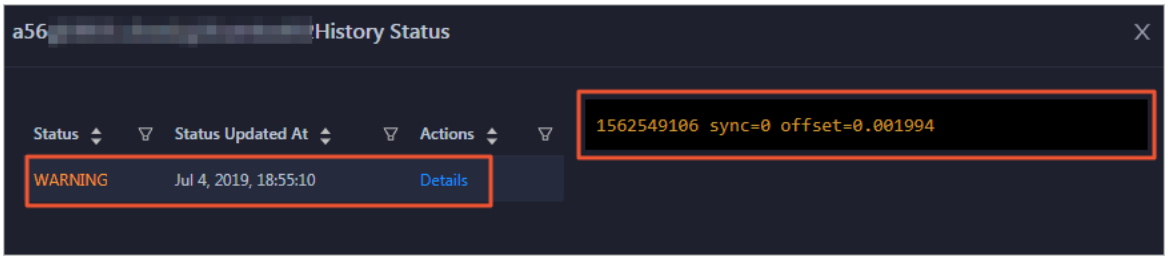

Clear alerts

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.

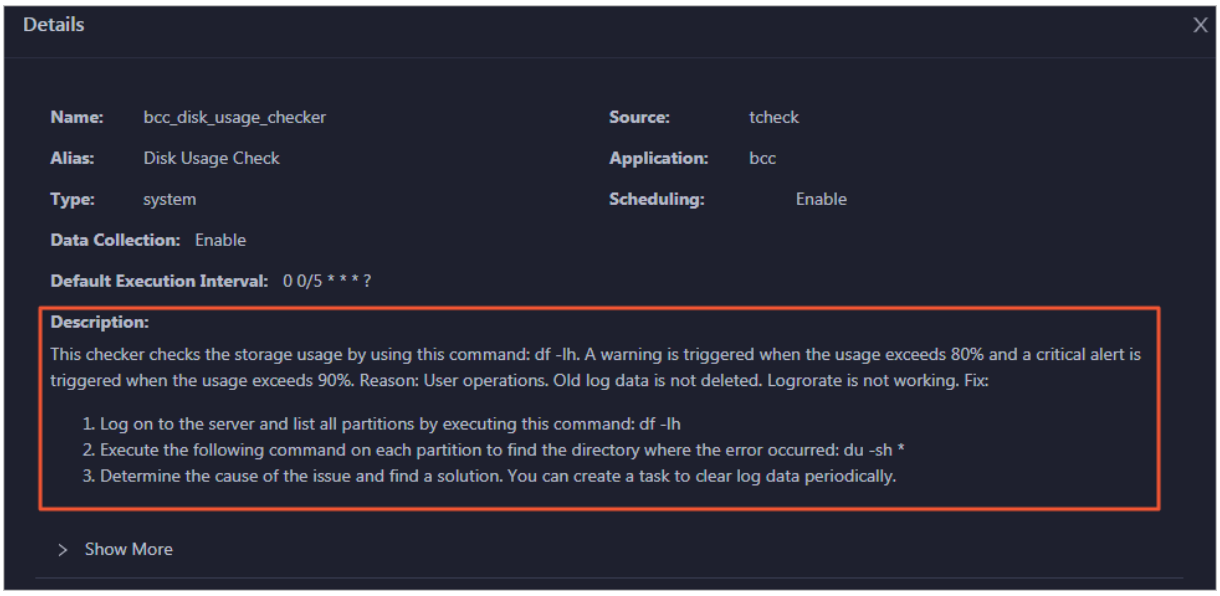

Log on to a host

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.

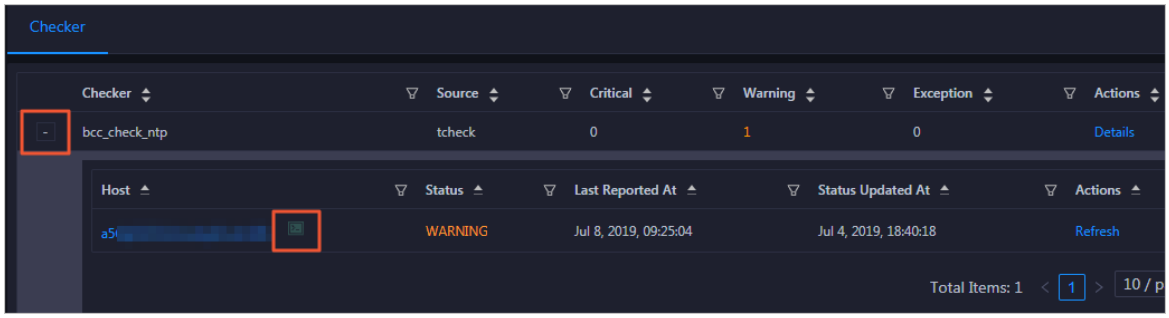

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

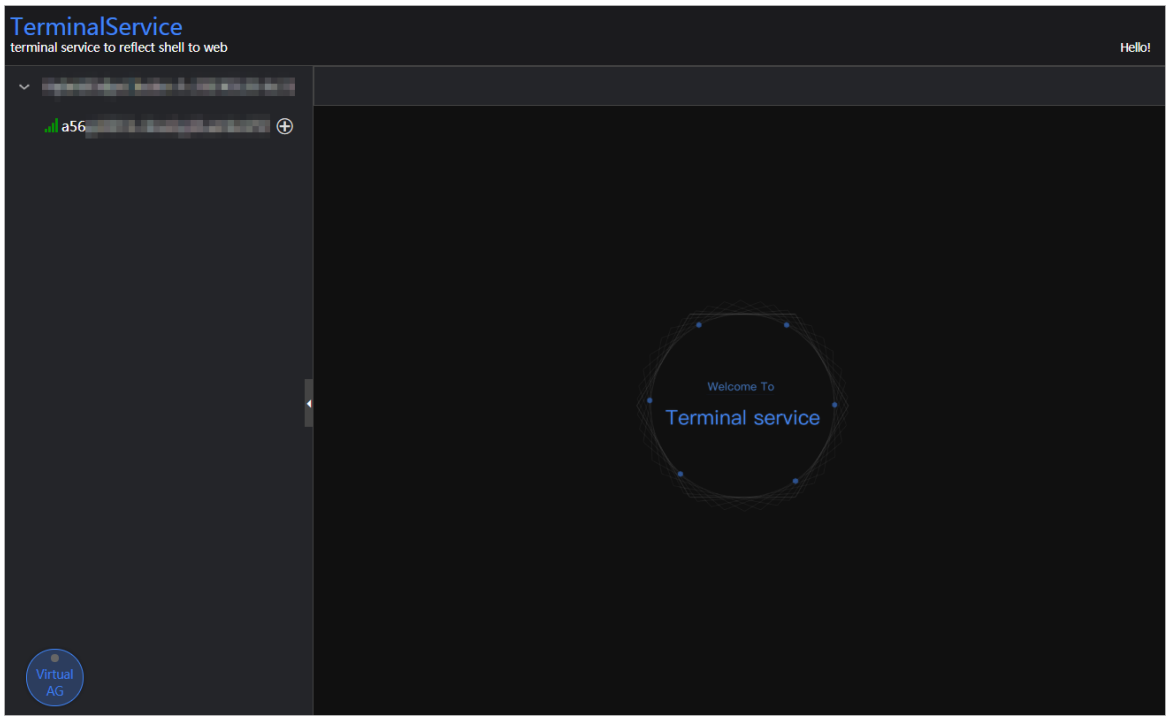

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

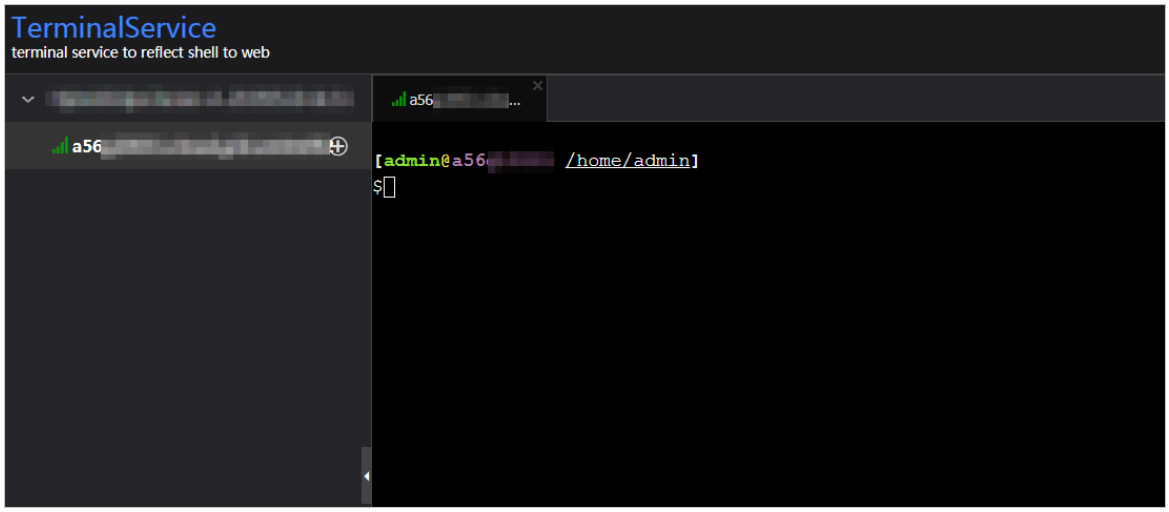

Run a checker again

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.

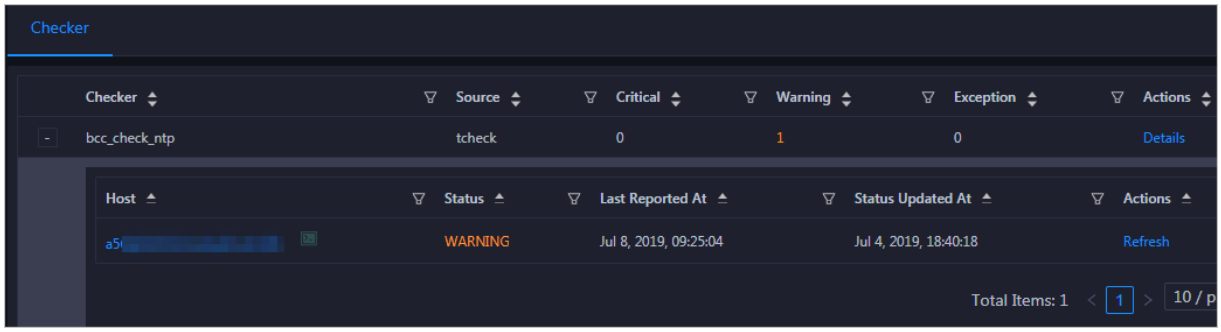

11.9.4.7.5. Host services

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.

11.9.5. Exceptions and solutions

This section describes some of the common error codes in the current version and corresponding solutions.

Error Code: LimitEceeded

Cause: The error code is returned because you can create up to 5 projects and 20 topics in a project in the previous version of DataHub.

Solution: In the latest version, you can create up to 10 projects and 1,000 topics in a project. Perform the following operations to change the project or topic limits:

1. Obtain the hostname of the ApsaraDB RDS for MySQL database from the following path: */home/admin/datahub/service/deploy/env.cfg*.
2. Access the corresponding ApsaraDB RDS for MySQL database. In the config_meta table, check the values of ProjectLimit4User and TopicLimit4Project.
3. Run the following commands to update the configurations. The new configurations take about 1 minute to take effect. You do not need to restart the database.

```
update config_meta set config_value = 10 where config_type = ‘ProjectLimit4User’;
```

```
update config_meta set config_value = 1000 where config_type = ‘TopicLimit4Project’;
```

Error code: IanlnvalidParameter

Cause: The error code is returned when StreamCompute attempts to capture records from DataHub by using an invalid timestamp. The timestamp you submit to the StreamCompute task is later than the current time, which may be caused by inaccurate local system time.

Solution: Correct your local system time by using the Network Time Protocol (NTP) or specify a timestamp that is for example 10 minutes earlier than the local system time.

Error code: InvalidCursor

Cause: The error code is returned when StreamCompute attempts to capture records from DataHub by using an invalid or expired cursor. An error may have occurred while StreamCompute is processing records from several days ago. When the time-to-live of the records expires and the records are deleted from DataHub, the cursor of these records is invalid.

Solution: Contact technical support for StreamCompute to learn about the cause of the task.

Error code: Parse response failed

Cause: This is probably caused by an invalid endpoint. For example, you may enter the console address as endpoint.

Solution: Perform a smoke test to check whether the system is running properly. If yes, check whether the endpoint is incorrect in the Apsara Infrastructure Management Framework console. Find the endpoint from the following path in the console: **DataHubCluster > Cluster Dashboard > Cluster Resource > Service: datahub-frontend > dns in the Parameters and Result columns**.

Error code: InternalServerError

Cause: Retry the smoke test or StreamCompute task. If the error code is still returned, an internal server error may occur. If the galaxy logs record this type of errors that occurred a long time ago, ignore these errors.

Solution: Use the following methods to search for corresponding logs to diagnose the issue. If you have any problems, screenshot the logs and contact technical support.

- In the logs directory of DataHubServer, search for the log files based on the specific time that the error occurred. The specific time can be found in the RequestId. RequestId is the unique ID of the request generated by DataHubServer.
- If more than one error occur, find the logs that are marked as **ERROR** in the logs directory of DataHubServer.

11.9.6. Appendix

11.9.6.1. Installation environment

Operation system: AliOS5U7-x86-64

Template: Bigdata

11.9.6.2. Deployment directories and services

Services

Name	Type	Description
service-datahub-service	Controller	The service that is used to deploy DataHub backend services and used as the admin gateway of Apsara system.
service-datahub-webconsole	Controller	The service that is used to deploy the DataHub console and configured on the same container as service-datahub-service.

Name	Type	Description
service-datahub-frontend	Worker	The service that is used to deploy frontend servers and used as chunk servers.
Chunkserver	Worker	The service that is used to deploy chunk servers in Apsara Distributed File System.
PanguMaster	Controller	The service that is used to deploy three masters in Apsara Distributed File System.
NuwaMaster	Controller	The service that is used to deploy three masters of Apsara Name Service and Distributed Lock Synchronization System.
FuxiMaster	Controller	The service that is used to deploy two masters of Job Scheduler.

Deployment directories and corresponding services

Module	Directory	Service
Datahub/XStreamServicex	/home/admin/datahub_service	service-datahub-service
Datahub/ShipperServicex	/home/admin/datahub_service	service-datahub-service
Datahub/CoordinatorServicex	/home/admin/datahub_service	service-datahub-service
WebConsole	/home/admin/datahub_webconsole	service-datahub-webconsole
Smoke	/home/admin/datahub_smoke	service-datahub-frontend
Frontend	/home/admin/datahub_frontend_server	service-datahub-frontend

11.9.6.3. Error codes

Error codes

Error code	HTTP status code	Description
InvalidUriSpec	400	The error code is returned when the request URI is invalid. This is probably caused by invalid topic or project names.
InvalidParameter	400	The error code is returned when a parameter is invalid. For more information about the cause of the error, see the error message.

Error code	HTTP status code	Description
Unauthorized	401	The error code is returned when the signature is incorrect. This is usually caused by an incorrect AccessKey or a time difference of more than 15 minutes between the client and the server.
NoPermission	403	The error code is returned when you do not have the permission to perform the operation.
InvalidSchema	400	The error code is returned when the schema format is invalid.
InvalidCursor	400	The error code is returned when the cursor is invalid or has expired.
NoSuchProject	404	The error code is returned when the specified project does not exist.
NoSuchTopic	404	The error code is returned when the specified topic does not exist.
NoSuchShard	404	The error code is returned when the specified shard ID does not exist.
ProjectAlreadyExist	400	The error code is returned when the project name already exists.
TopicAlreadyExist	400	The error code is returned when the topic name already exists.
InvalidShardOperation	405	The error code is returned when the operation on the shard is not allowed. For example, you are not allowed to write data into a shard when it is in Deactivated status.
LimitExceeded	400	The error code is returned when a specified threshold is exceeded. For example, you create no more than 512 shards in a topic and 20 topics in a project.
InternalServerError	500	The error code is returned when an unknown or internal error occurs or when the system is being upgraded. For more information about the cause of the error, obtain the request ID or search DataHub server logs for **InternalServerError**.

11.10. E-MapReduce (EMR)

11.10.1. Methods for logging on to O&M platforms

11.10.1.1. Log on to the Apsara Infrastructure Management Framework console

This topic describes how to log on to the Apsara Infrastructure Management Framework console.

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Context

Apsara Infrastructure Management Framework is a cluster monitoring and management tool. It displays basic cluster information as well as machine and server monitoring information. It monitors system load, CPU, memory, disk, and transmission metrics and the status of service instances. This helps you detect exceptions and take actions to fix the exceptions. After you log on to the Apsara Infrastructure Management Framework console, you can perform command-line operations on cluster machines.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   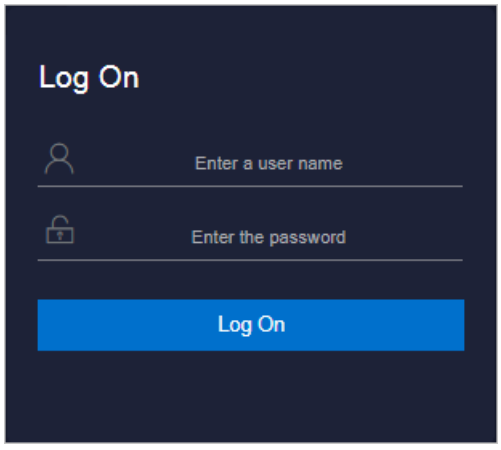

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, click **Products** and then **Product List**.
6. Click **Apsara Infrastructure Management Framework**.

11.10.2. Routine maintenance

11.10.2.1. O&M in the Apsara Infrastructure Management Framework console

This topic describes how to perform O&M in the Apsara Infrastructure Management Framework console.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click the C tab and select emr from the Project drop-down list.

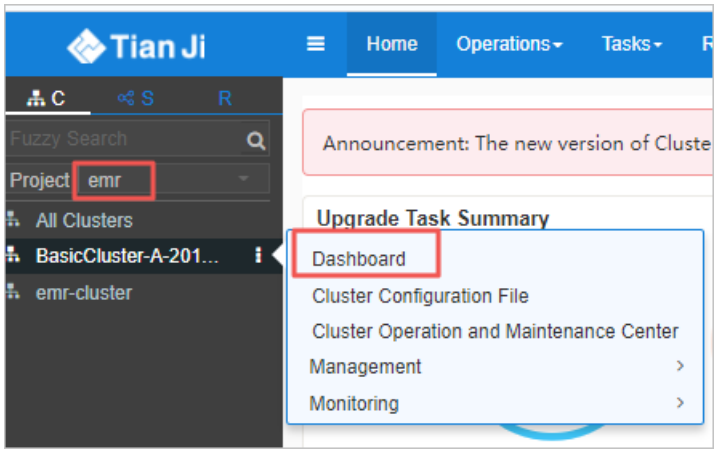

3. Move the pointer over the icon next to an EMR cluster and select **Dashboard** to go to the Cluster Dashboard page.
4. In the **Service Instances** section, click **Details** in the Actions column that corresponds to emr-service to go to the Service Instance Information Dashboard page.

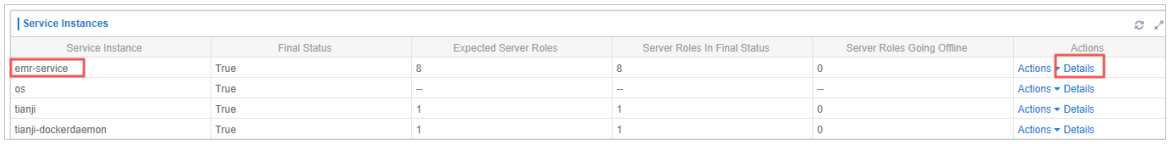

5. On the **Services** tab, find the **emr-service** service and click **Details** in the Action column.
6. On the **Service Details** page, click a server role.
7. In the **Machines** section, click **Restart Server Role** in the **Actions** column for a machine. Perform

this operation for the other machines.The service is restarted.

11.10.3. Troubleshooting

11.10.3.1. Troubleshooting methods

If you detect a system fault during routine maintenance, read the Routine Maintenance part of this documentation for reference.

If you fail to rectify the fault, collect related information such as system information and fault symptoms and contact Alibaba Cloud technical support for help.

After you rectify a fault, analyze its causes, review the troubleshooting process, and make improvements.

11.11. Dataphin

11.11.1. What is Apsara Bigdata Manager?

Apsara Bigdata Manager (ABM) is an operations and maintenance (O&M) platform tailored for big data products. As an O&M tool for Dataphin, ABM supports O&M on the business, services, clusters, and hosts of Dataphin. In addition, ABM allows you to deploy upgrade patches for Dataphin, customize alert configurations, and view the O&M history.

On-site Apsara Stack engineers can manage Dataphin by using ABM. For example, they can view operation metrics, modify configurations, and check and handle alerts.

11.11.2. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

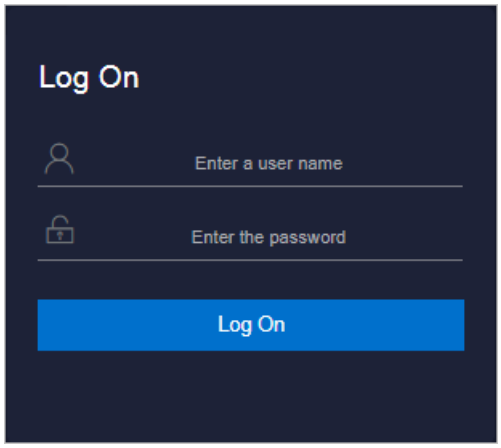

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.11.3. O&M overview

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

Modules

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

Module	Feature	Description
	Service 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.

Service O&M Module	Feature	Description
	Service hosts	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Cluster O&M	Cluster overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	Displays the check results for a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
Host O&M	Host 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, 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.
	Host health	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.

Entry

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

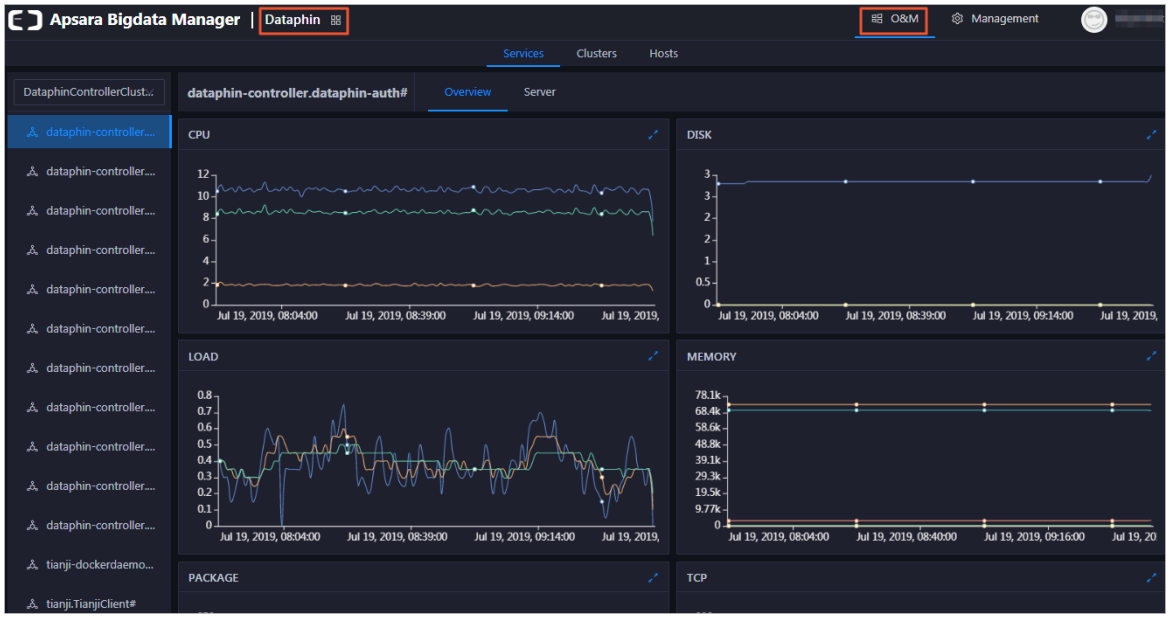

The **O&M** page includes three modules, namely, **Services**, **Clusters**, and **Hosts**.

11.11.4. Service O&M

11.11.4.1. Service overview

The service overview page lists all Dataphin 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.

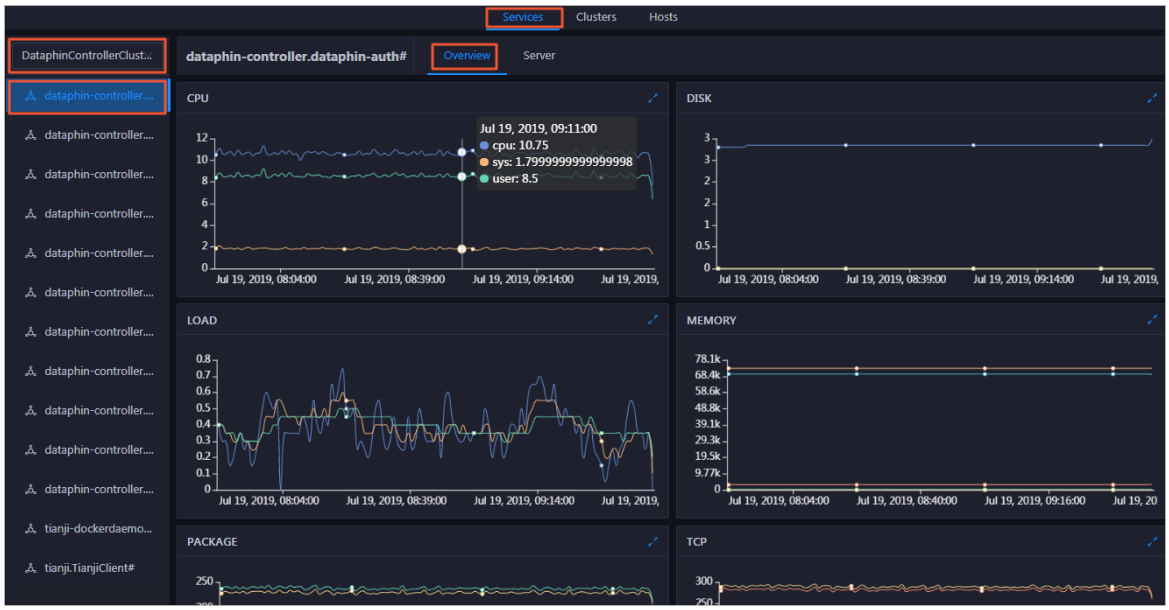

On the **Overview** page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

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.

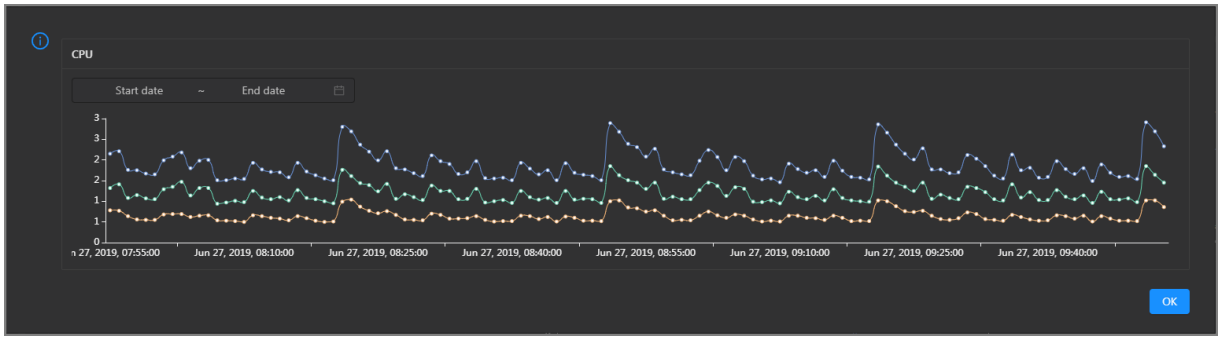

DISK

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.

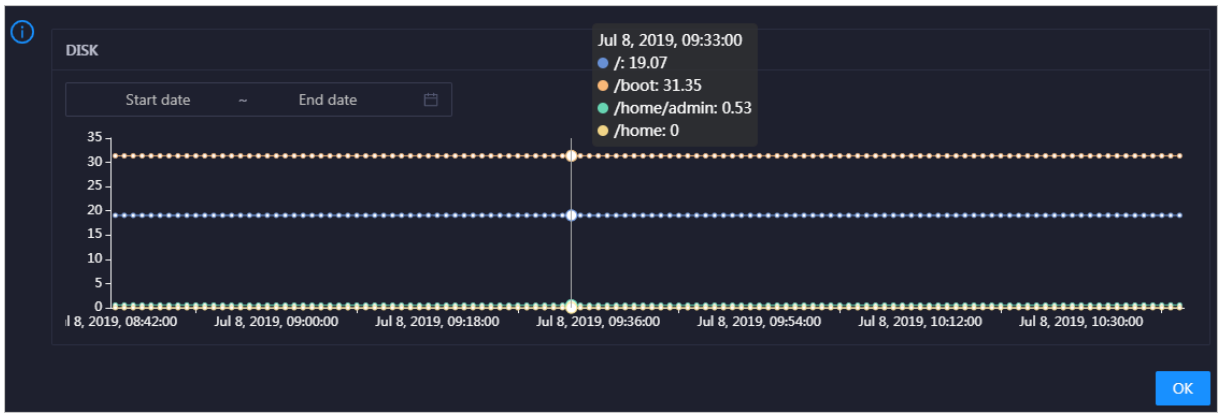

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

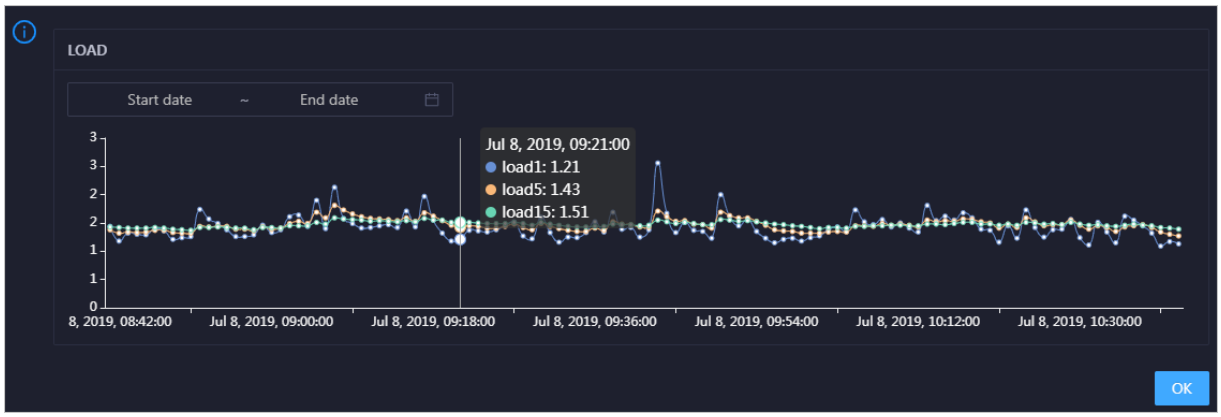

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

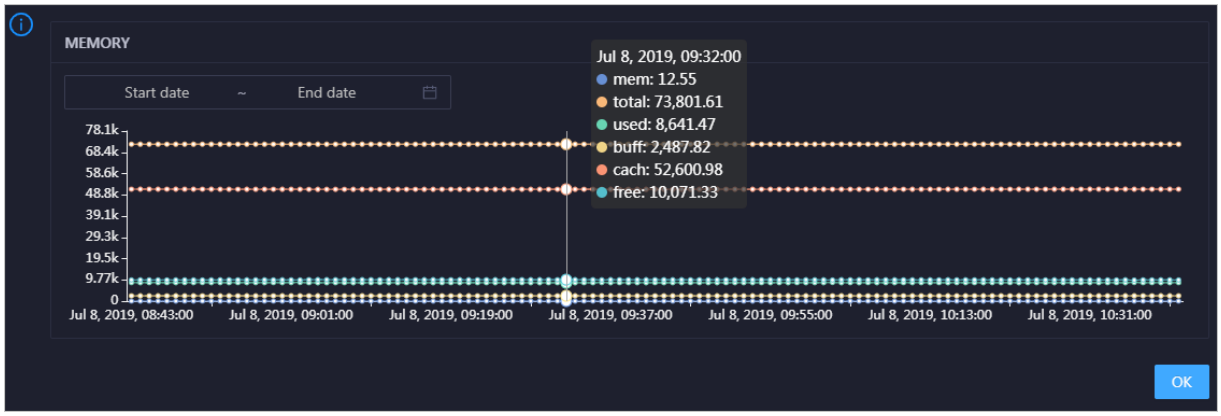

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

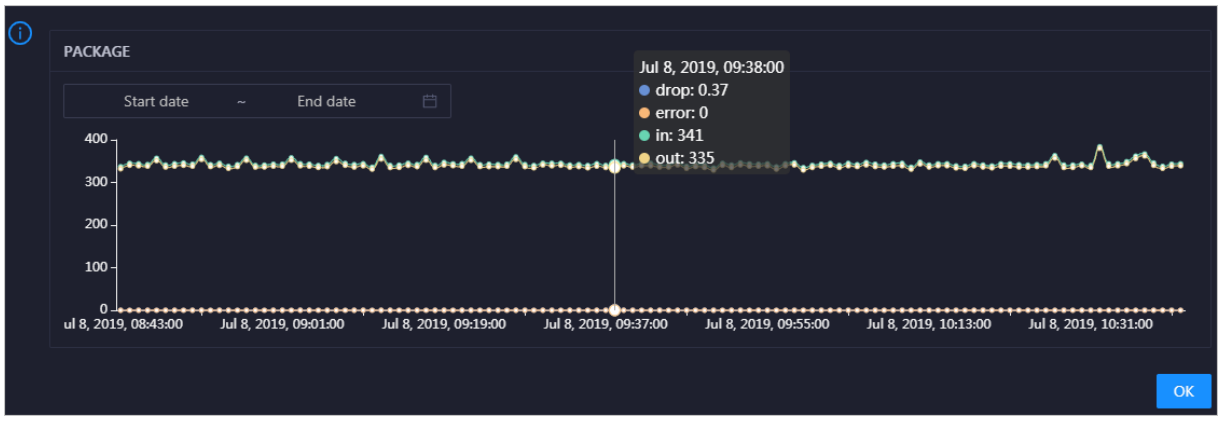

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

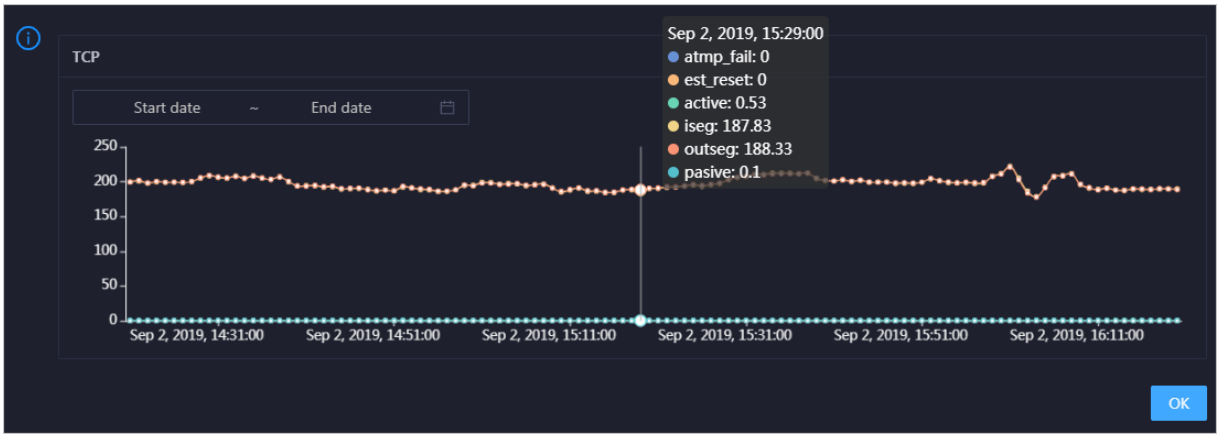

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

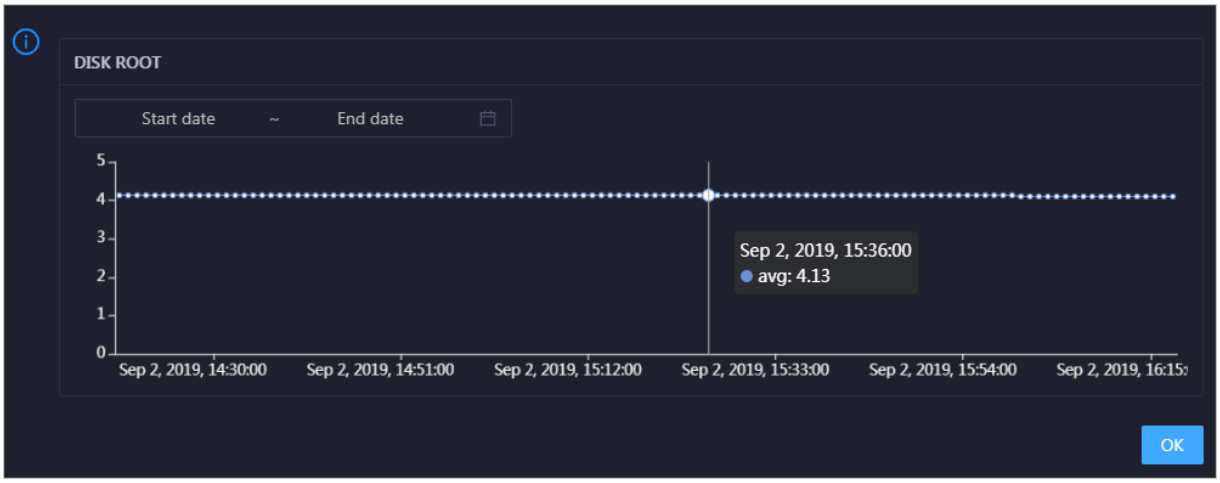

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

11.11.4.2. Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each Dataphin 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.

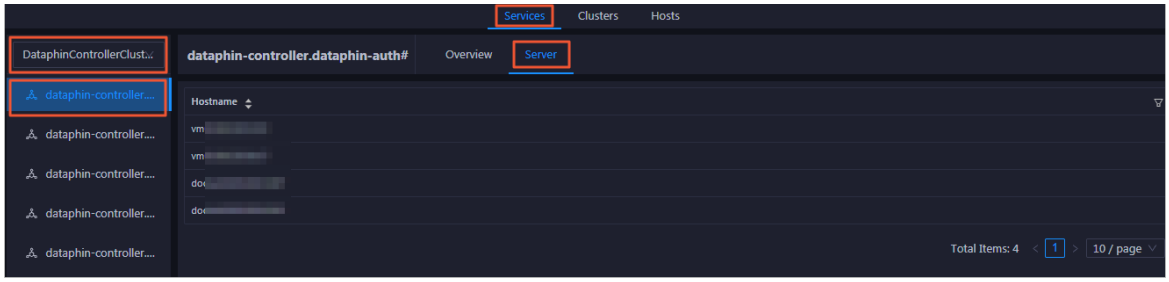

On the **Server** page, you can view the hosts where the selected service is run.

11.11.5. Cluster O&M

11.11.5.1. 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

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.

The **Overview** page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. To view information about a cluster, select a **region** in the left-side navigation pane, and then select a **cluster** in the region.

CPU

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

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

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

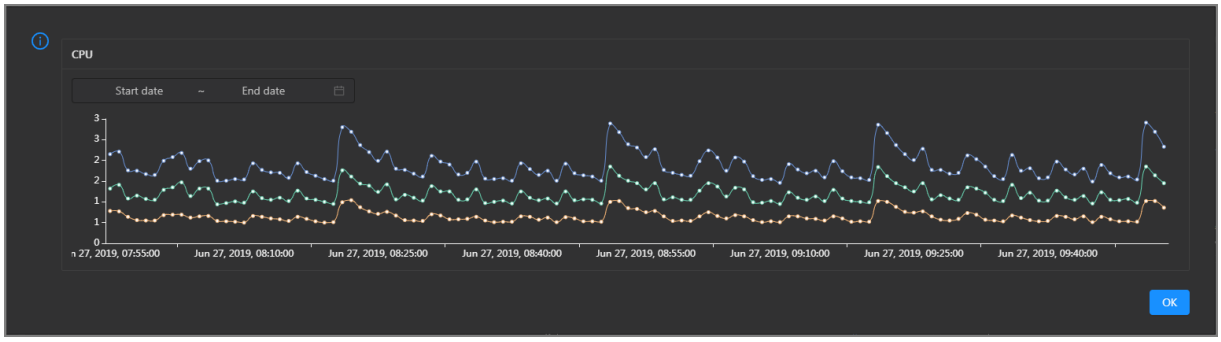

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

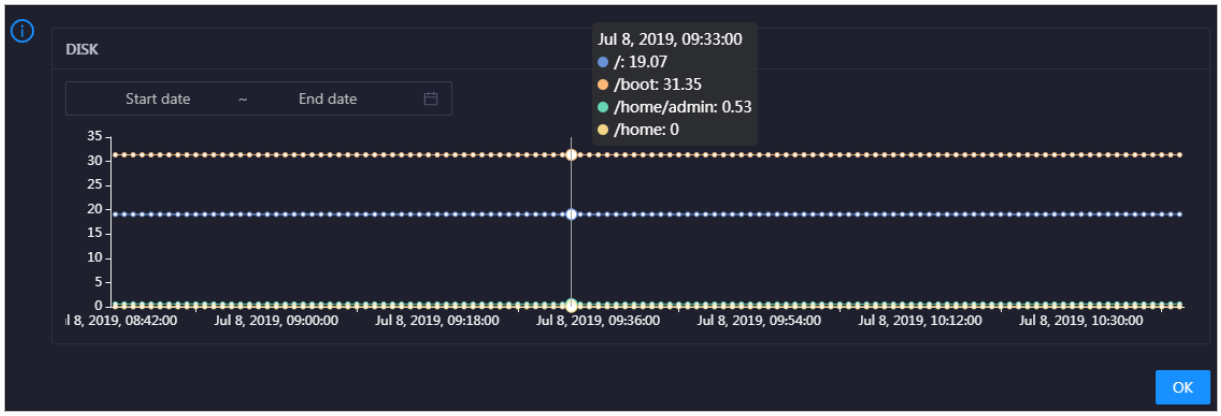

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

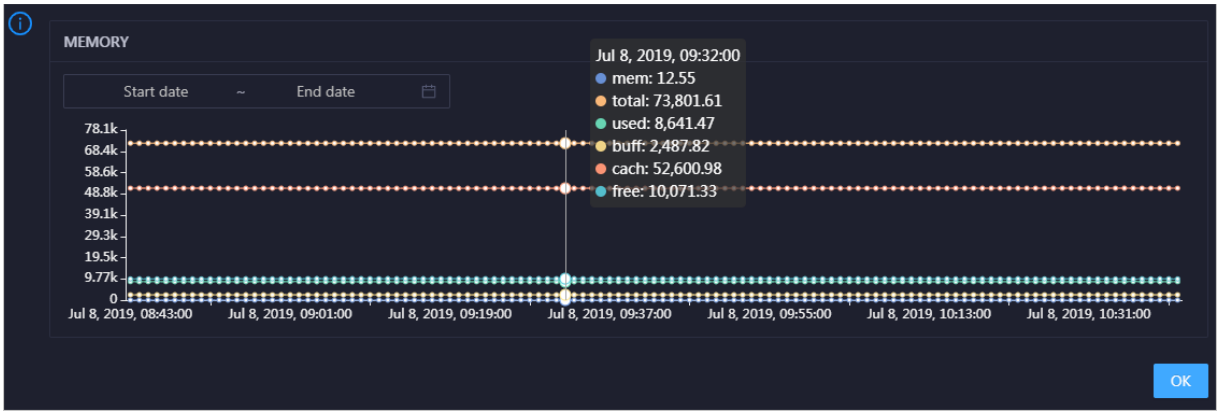

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

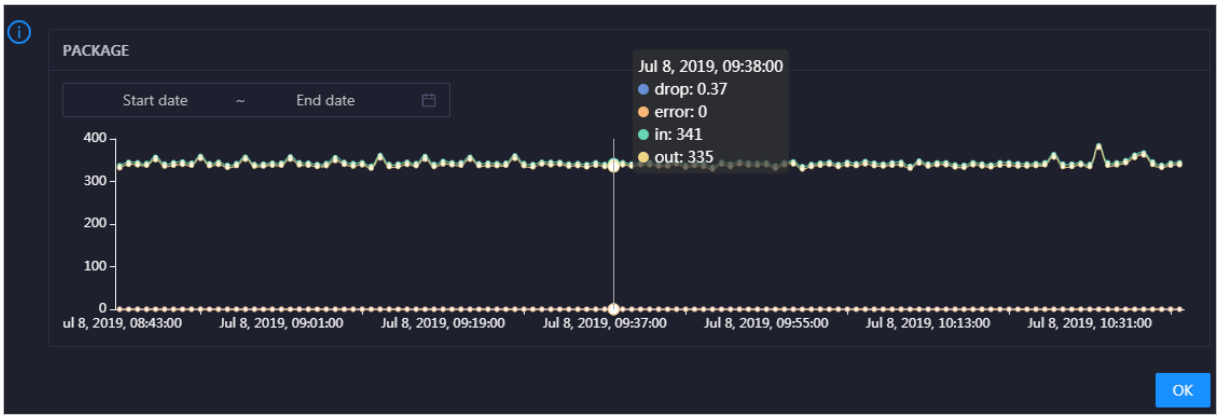

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

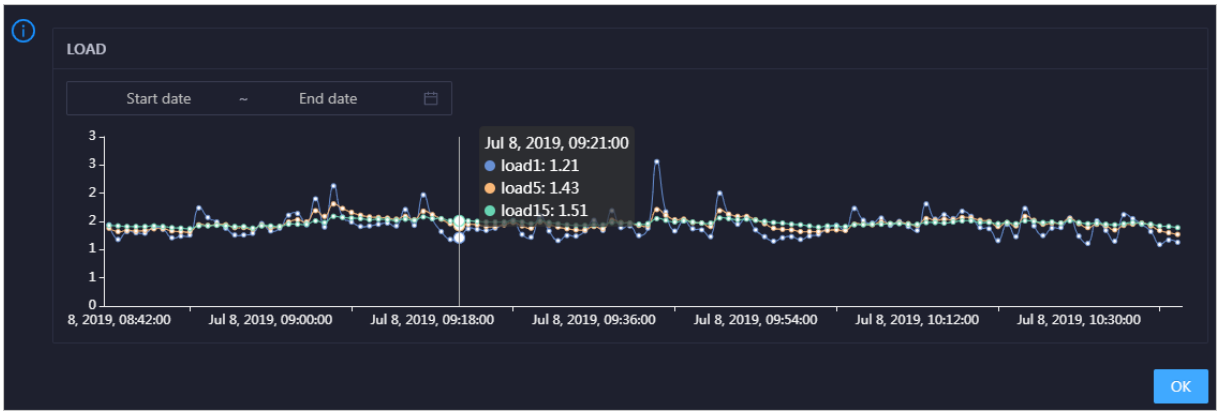

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.11.5.2. 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

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.

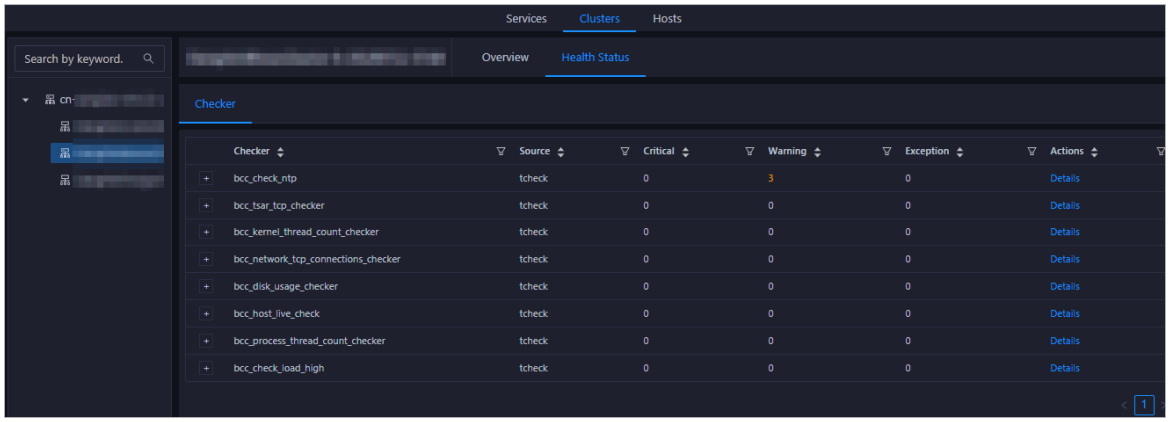

On the **Health Status** page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, **Critical**, **Warning**, and **Exception** results are alerts. You need to pay special attention to them, especially the **Critical** and **Warning** results.

View checker details

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

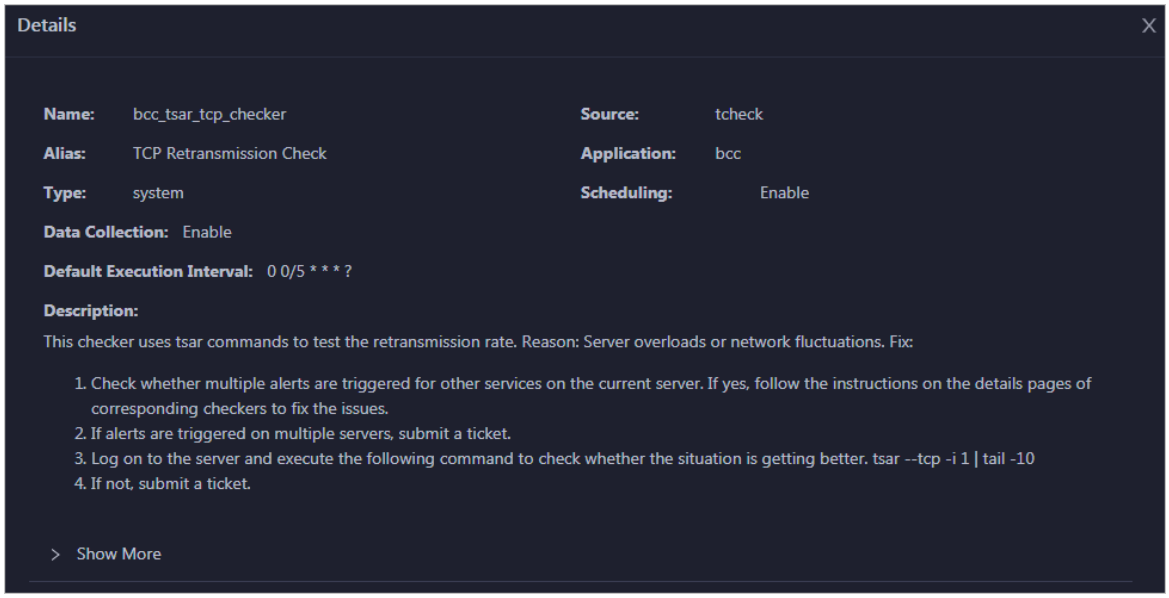

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

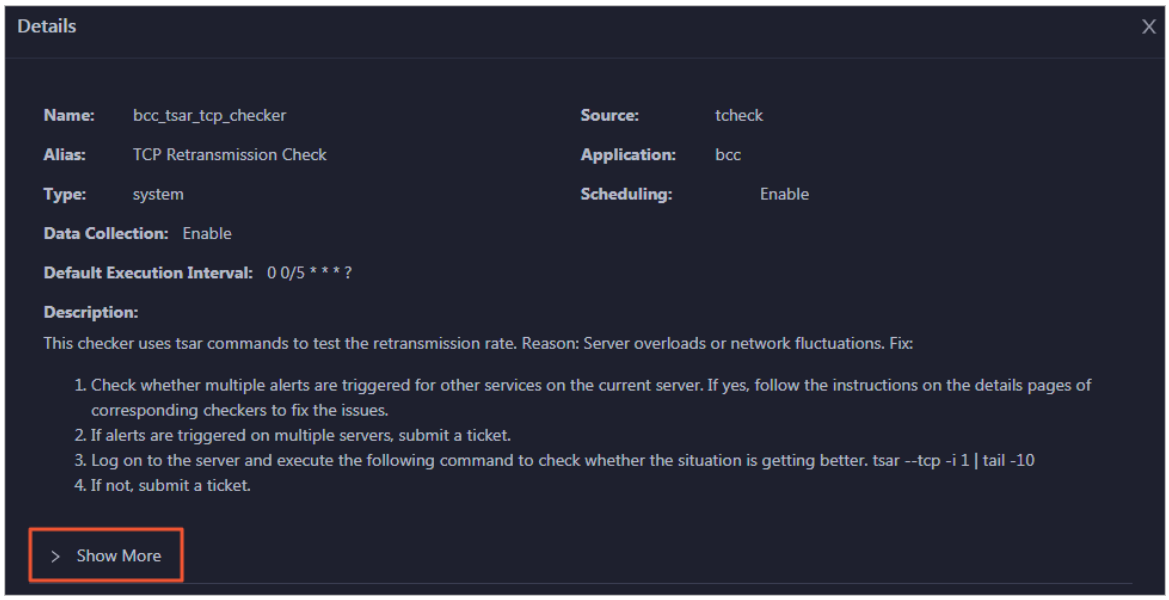

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

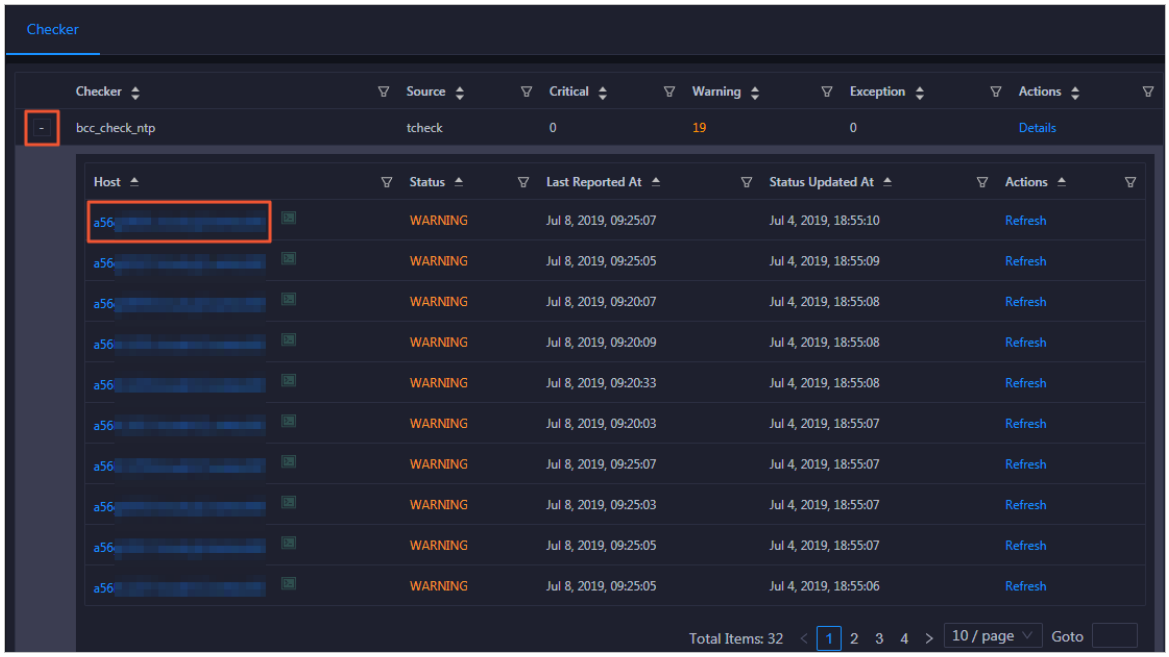

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

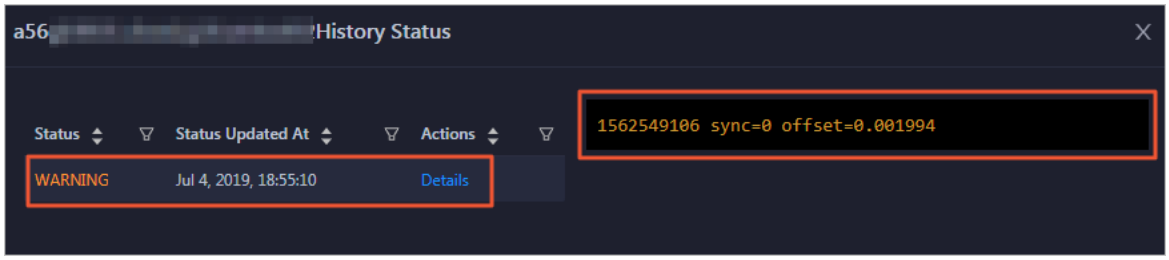

Clear alerts

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

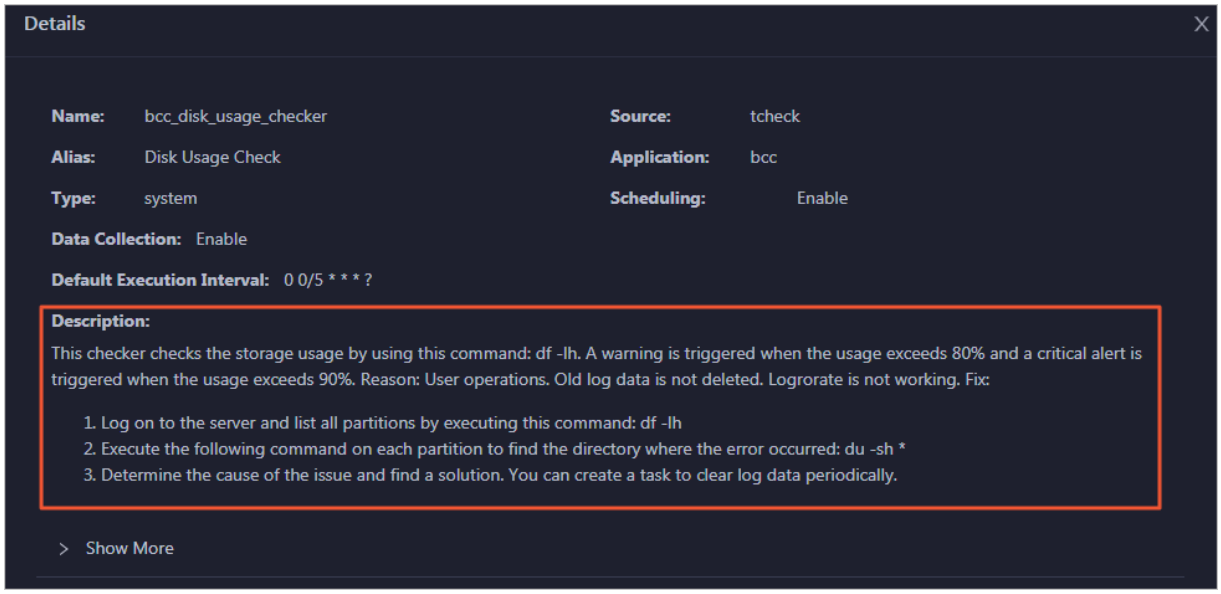

Log on to a host

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

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

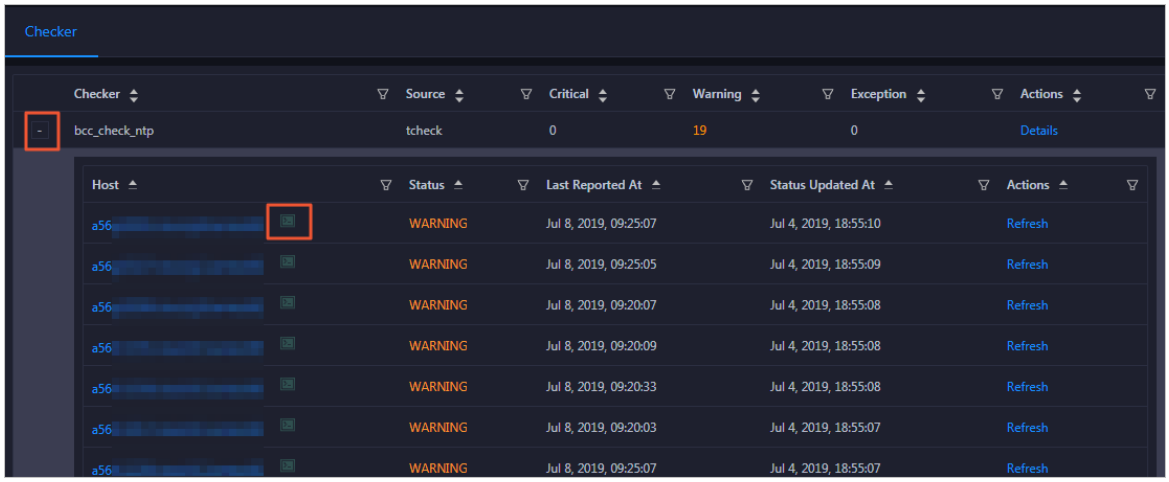

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

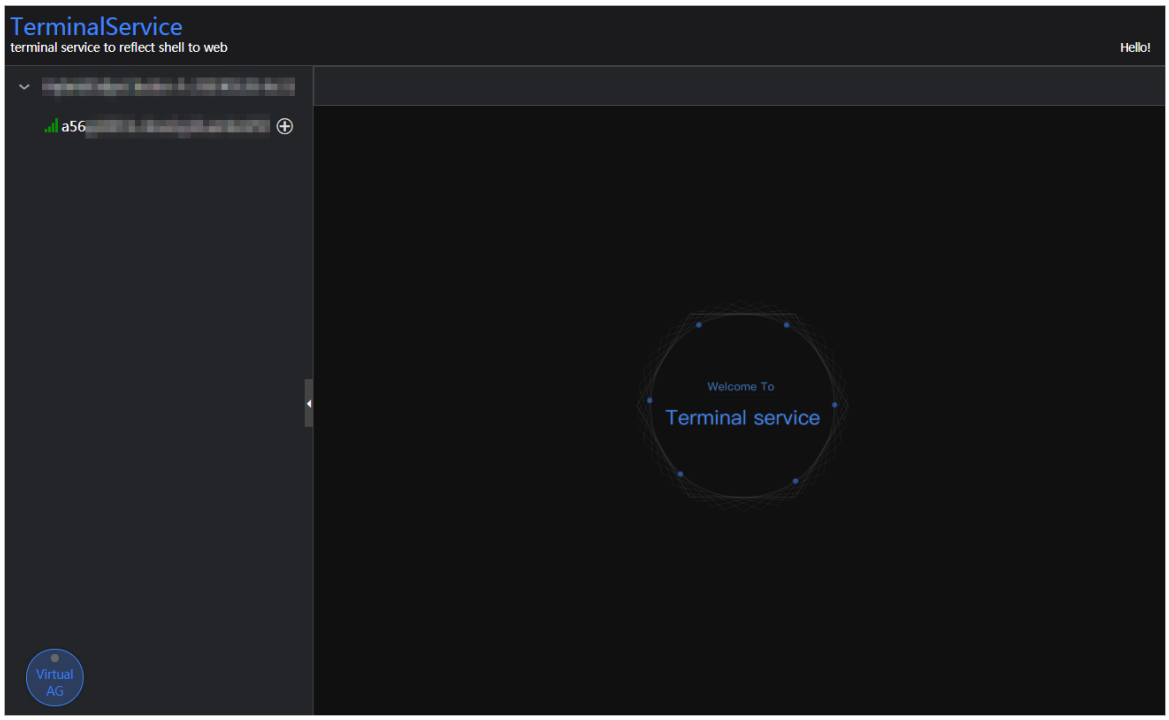

3. On the **TerminalService** page, click the hostname to log on to the host.

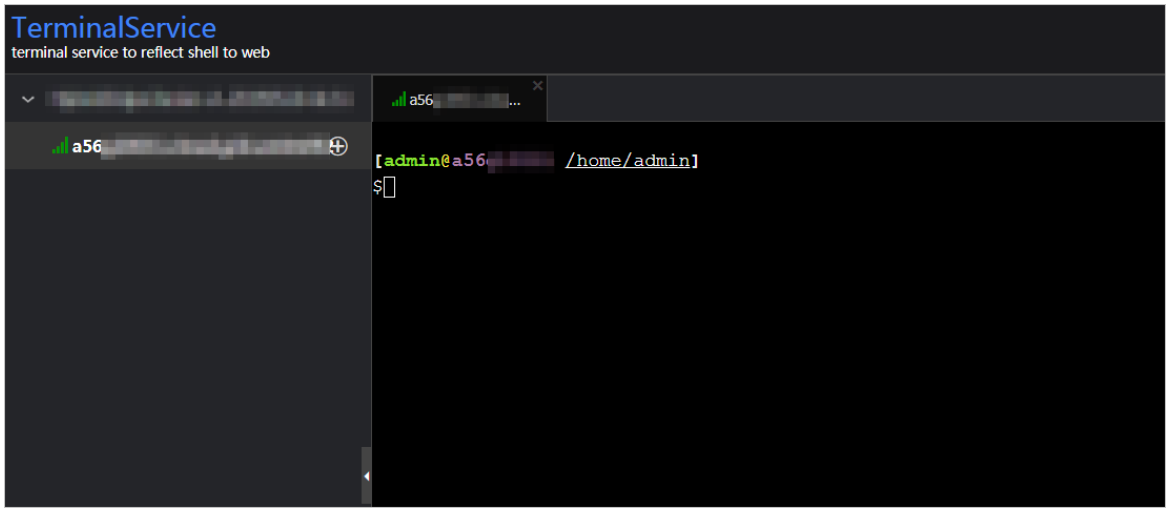

Run a checker again

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

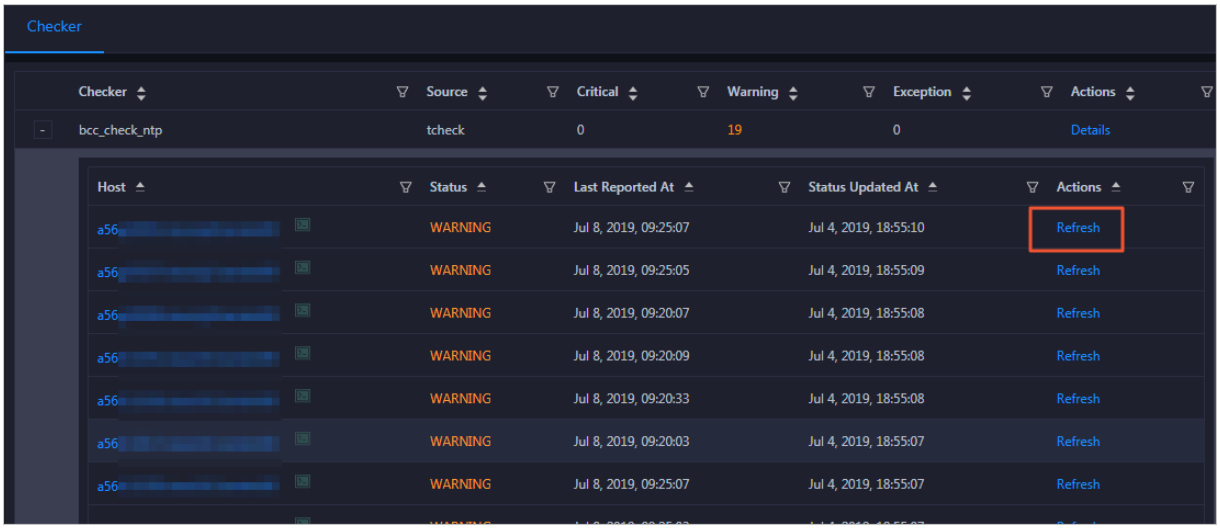

11.11.6. Host O&M

11.11.6.1. Host overview

The Overview tab displays the overall running and health check information about a host in a Dataphin 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. 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.

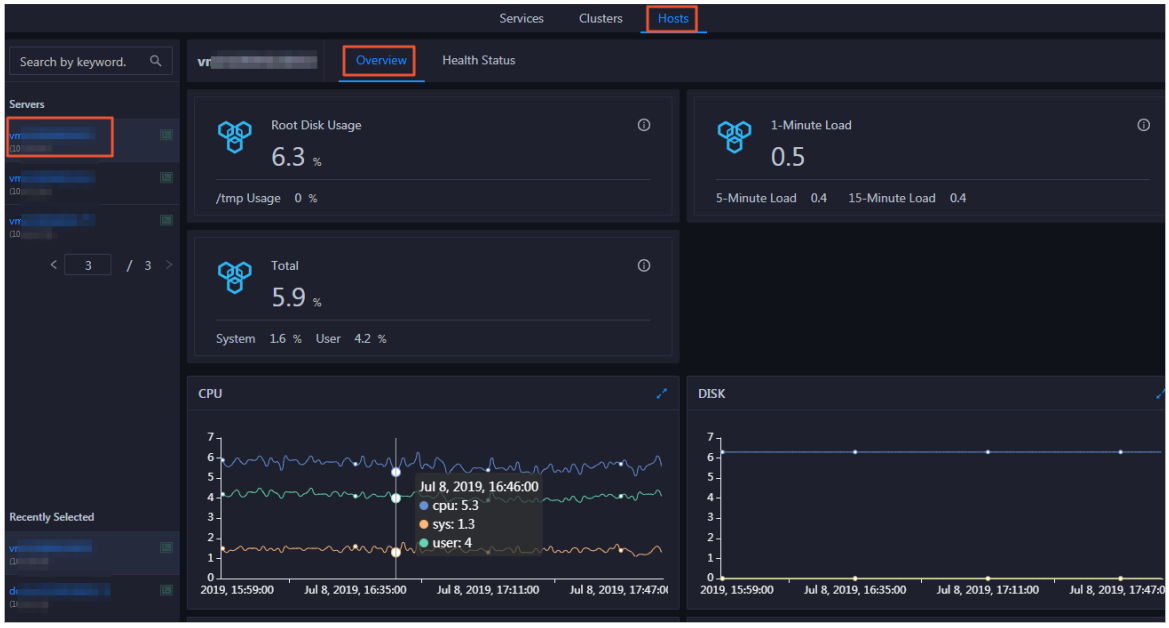

On the **Overview** tab, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root 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.

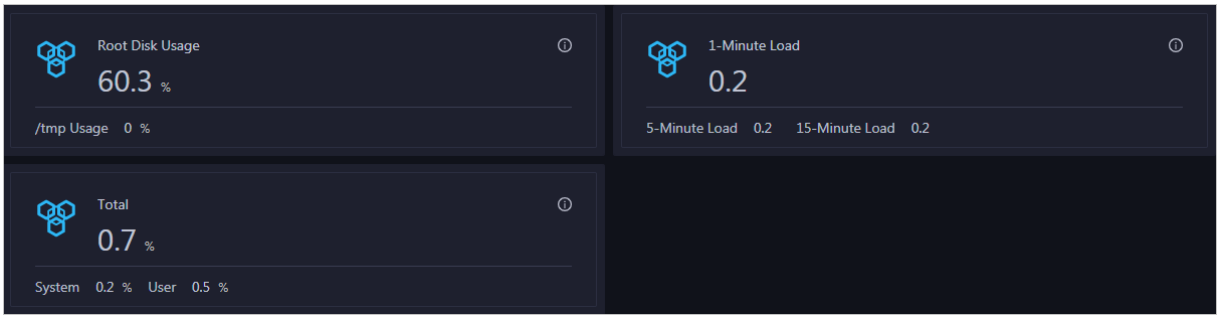

CPU

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.

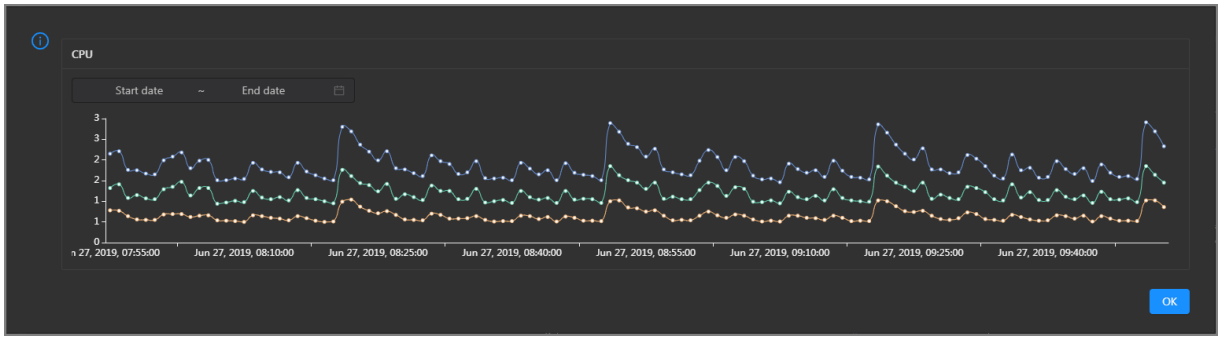

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

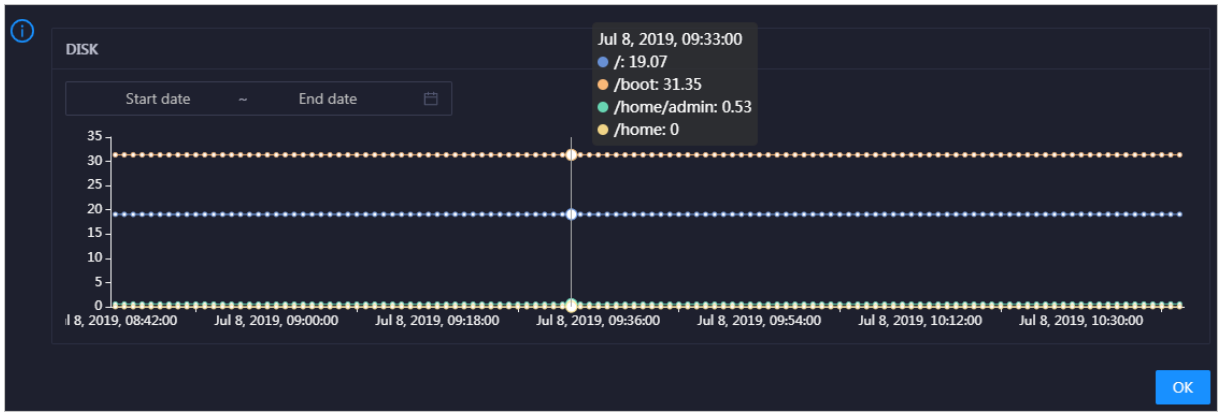

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

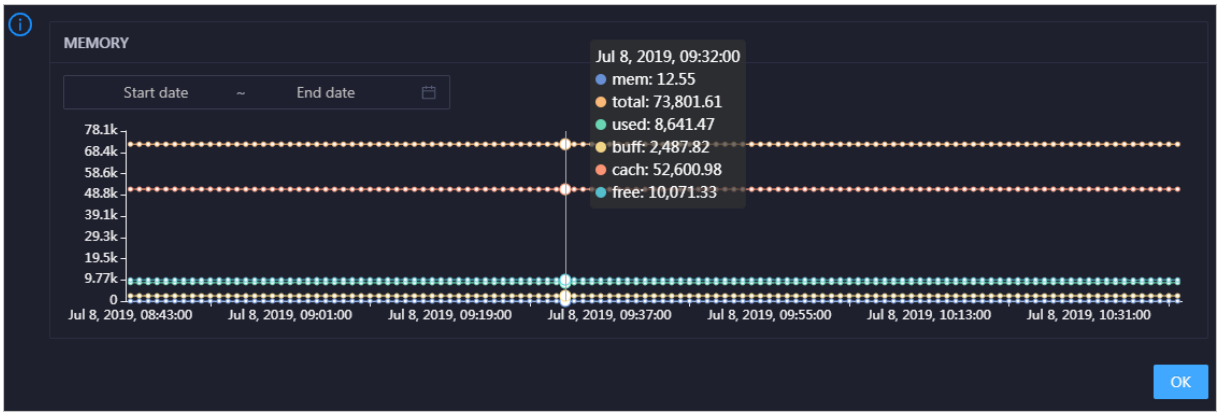

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

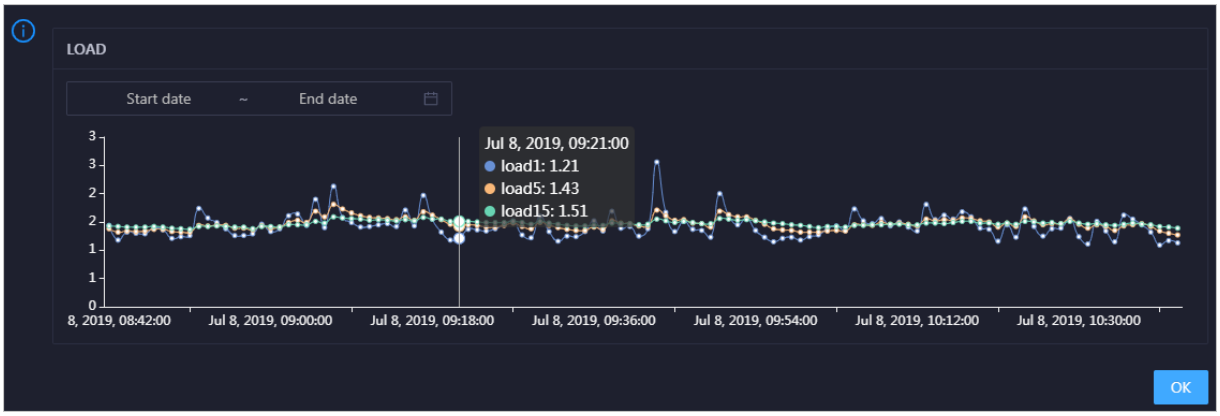

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

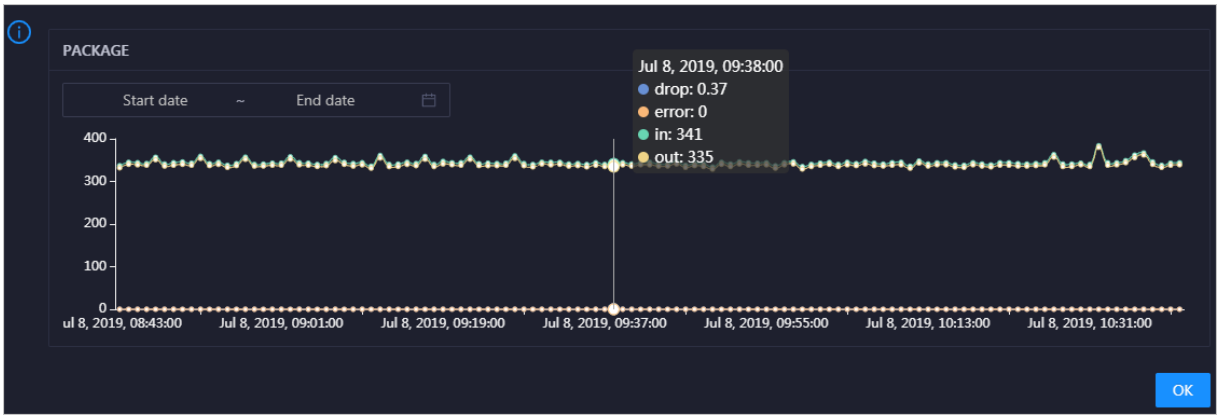

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

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.

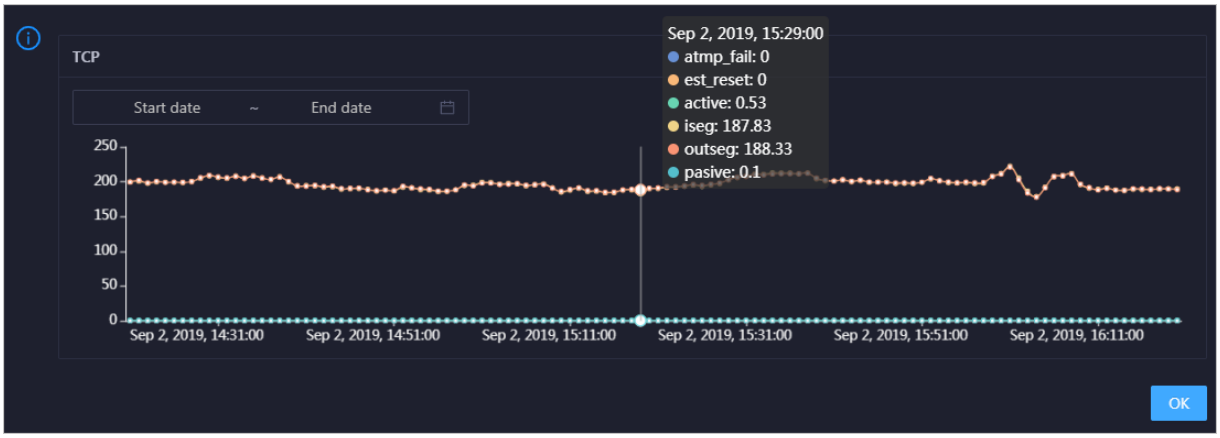

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

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.

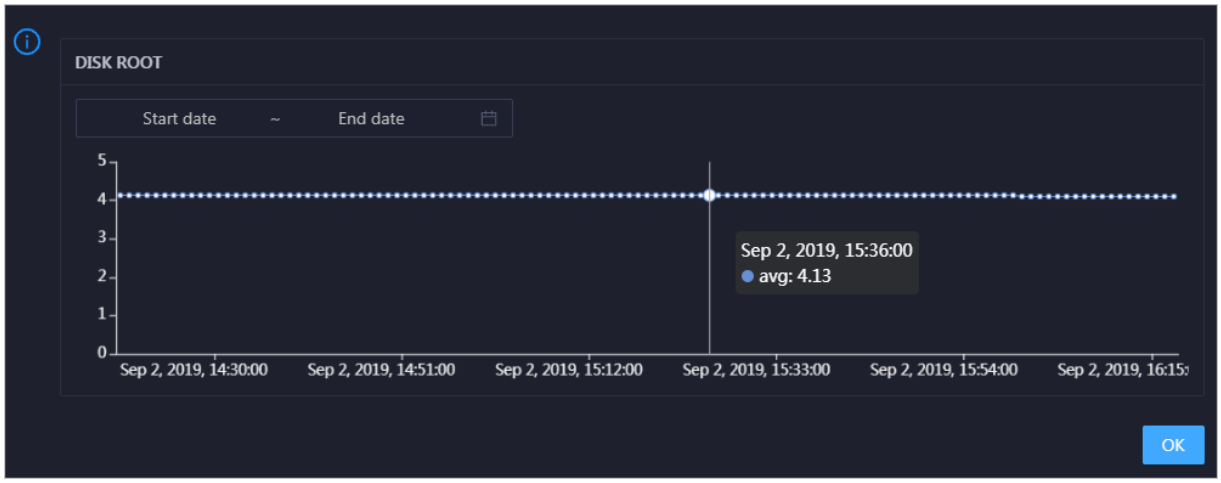

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

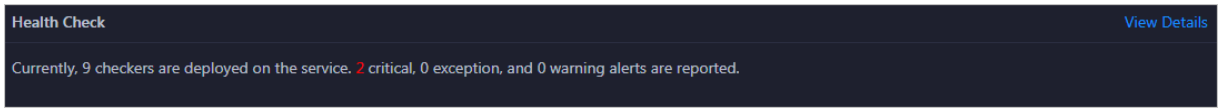

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

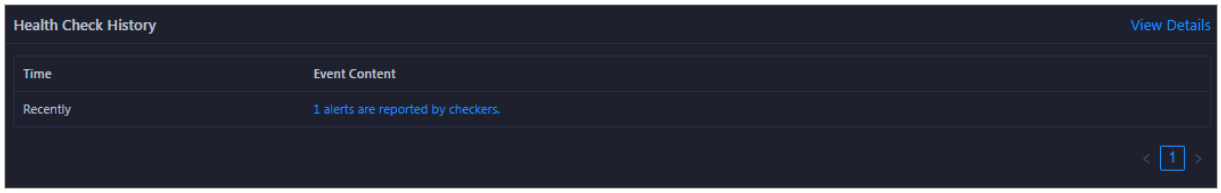

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

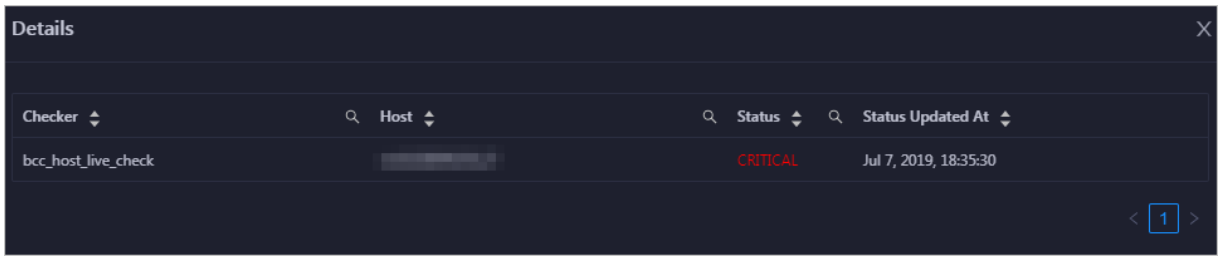

11.11.6.2. Host health

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 a host and perform manual checks on the host.

Entry

1. At the top of the **O&M** page, click the **Hosts** tab.
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.

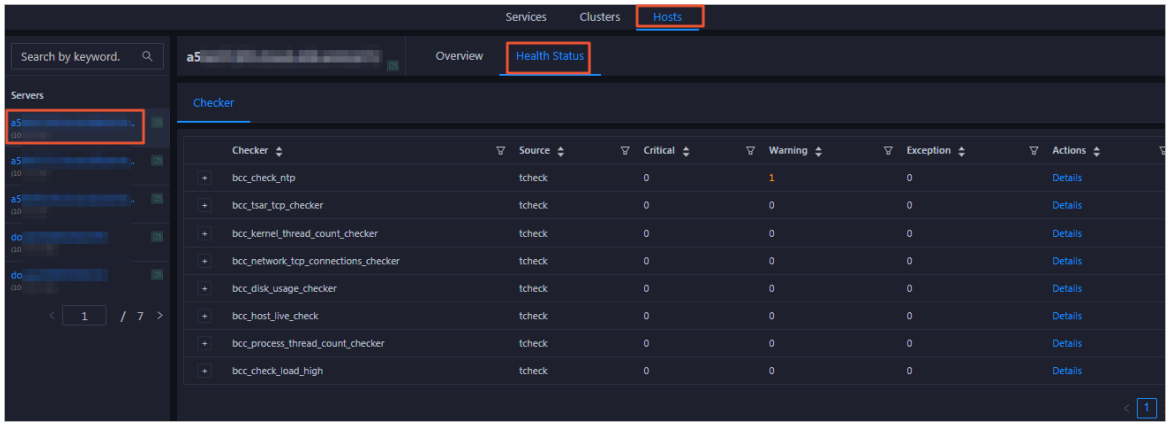

On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into the **Critical**, **Warning**, **Exception**, and **OK** types. They are displayed in different colors. Among them, the **Critical**, **Warning**, and **Exception** results are alerts. You need to pay attention to them, especially the **Critical** and **Warning** results.

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.

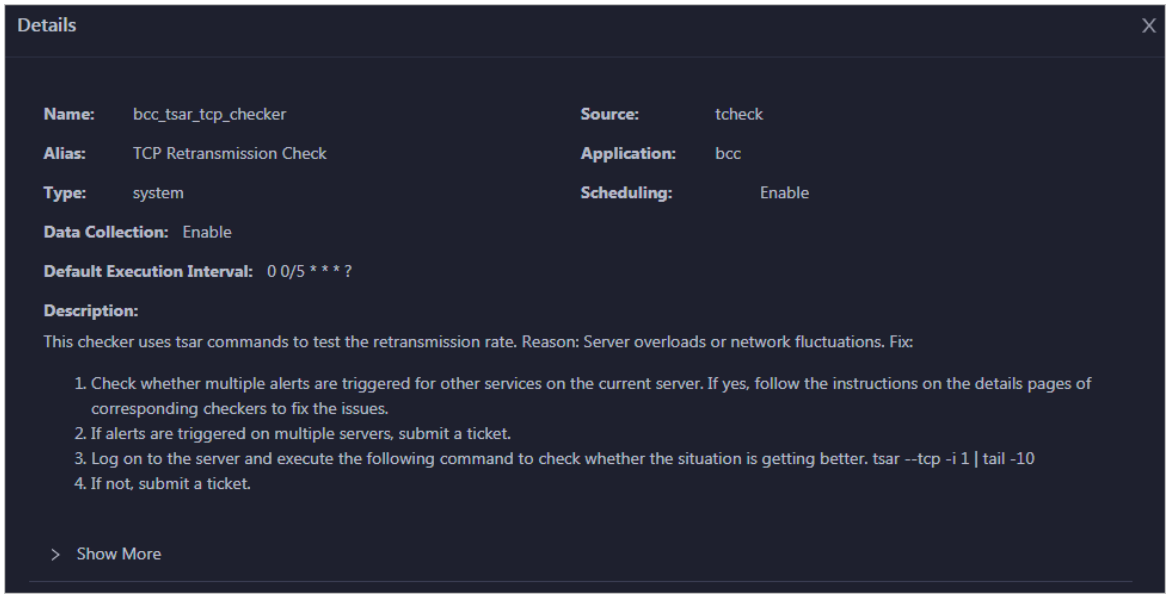

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

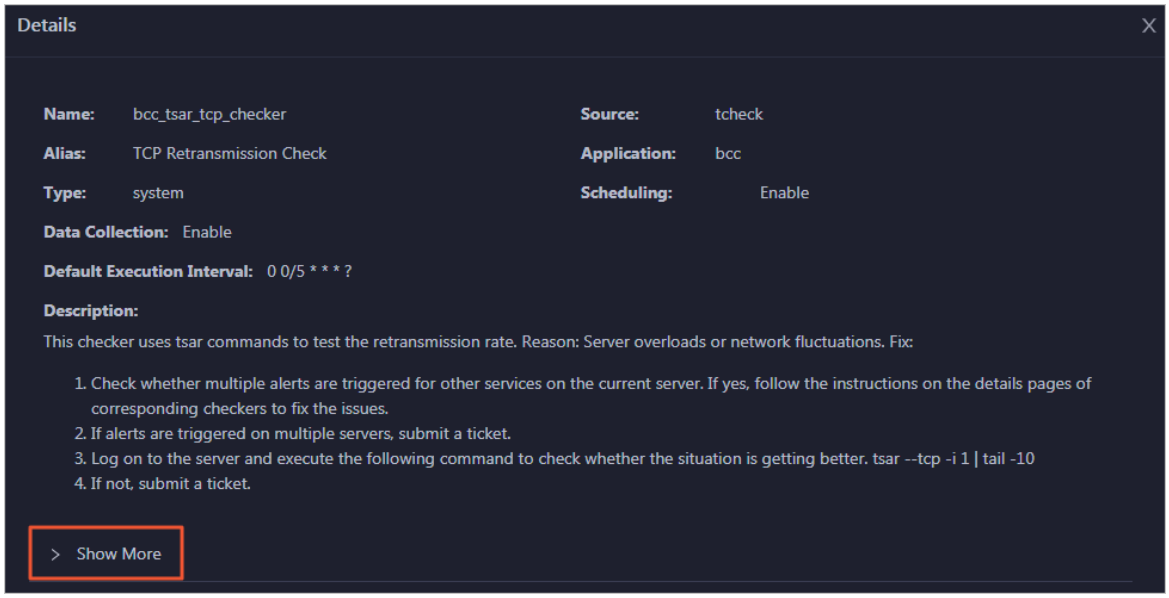

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

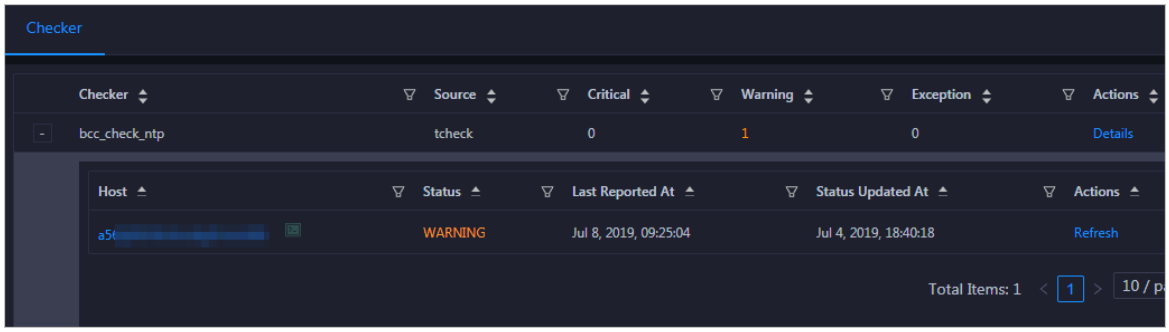

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

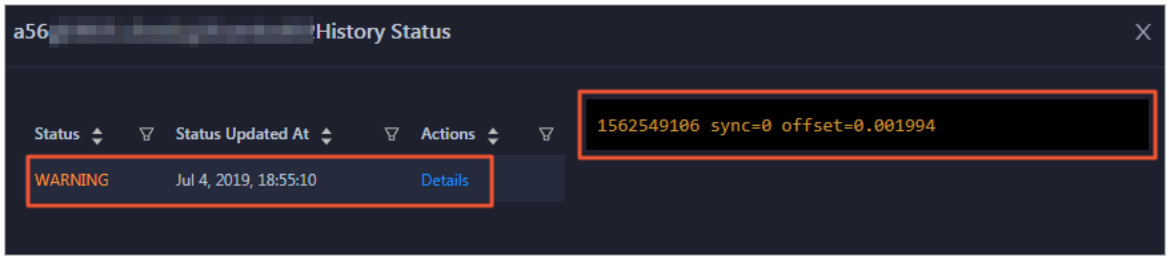

Clear alerts

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.

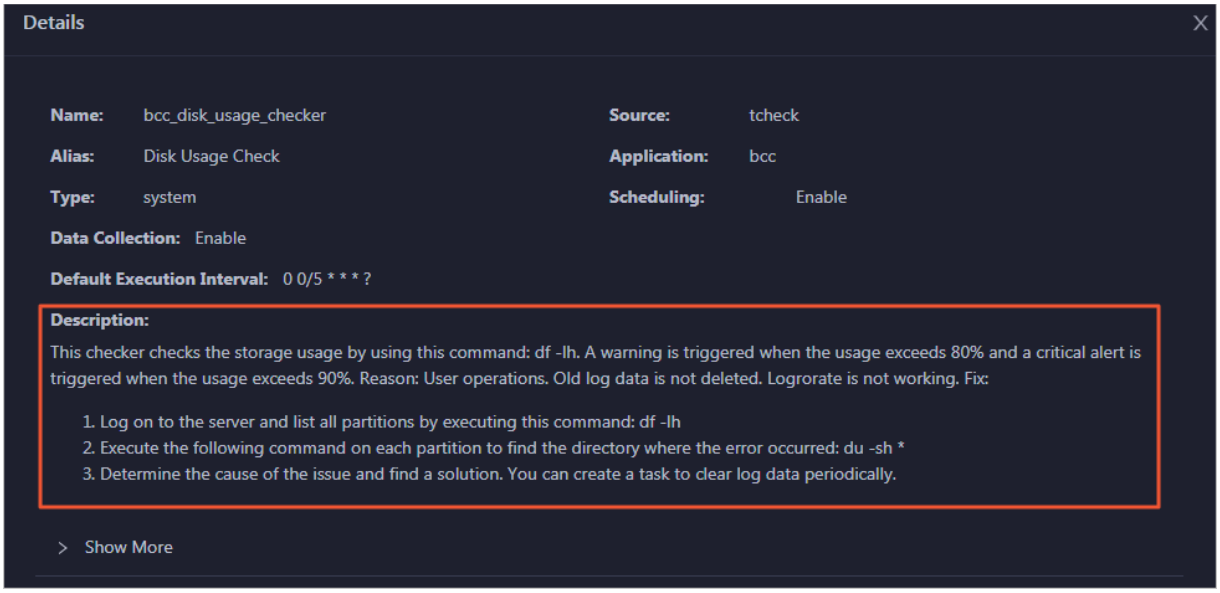

Log on to a host

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.

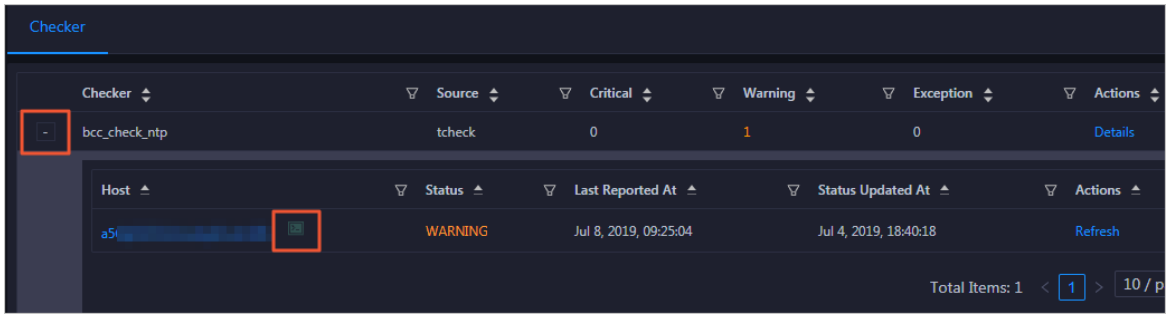

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

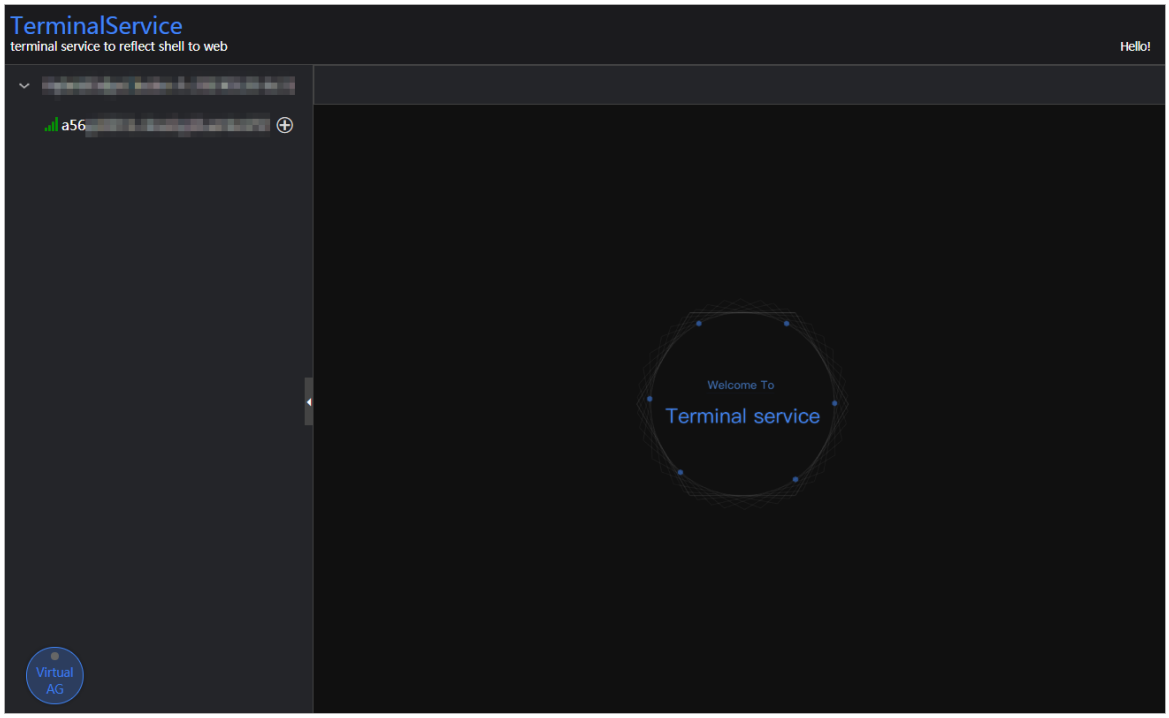

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

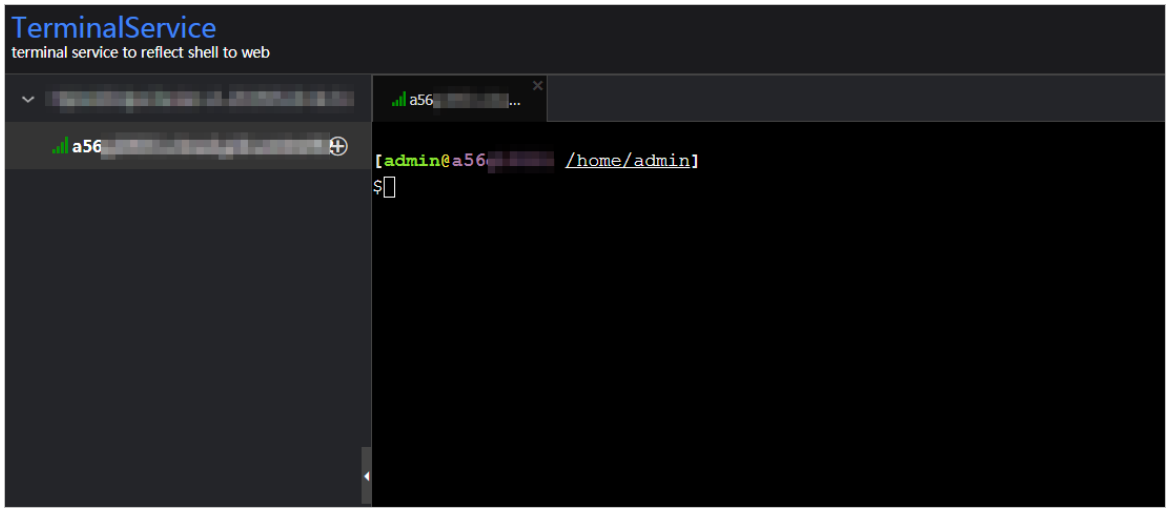

Run a checker again

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.

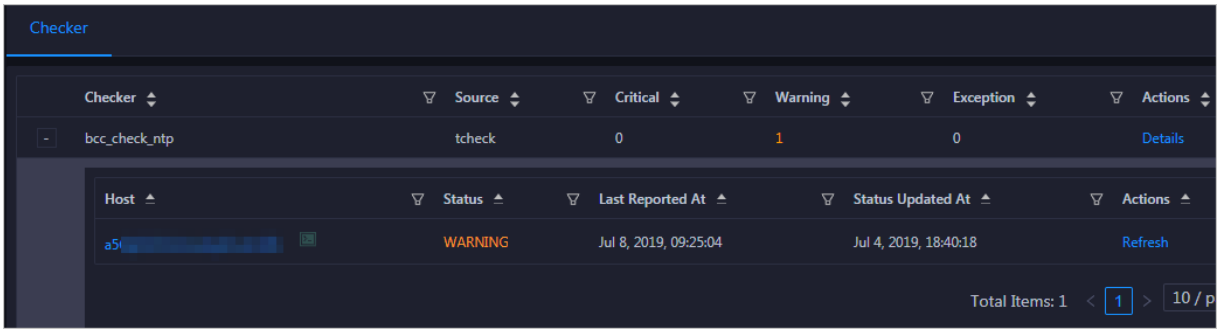

11.12. Elasticsearch (on ECS)

11.12.1. What is Apsara Bigdata Manager?

Apsara Bigdata Manager (ABM) is an operations and maintenance (O&M) platform designed for big data products. You can use ABM to perform O&M operations on Apsara Stack Elasticsearch from the perspectives of businesses, services, clusters, and hosts. You can also perform patch updates for Elasticsearch, customize alerting configurations, and view O&M history in ABM.

ABM helps on-site Apsara Stack engineers manage Elasticsearch. For example, the engineers can view resource usage, check and handle alerts, and modify configurations.

11.12.2. Log on to the ABM console

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

Context

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

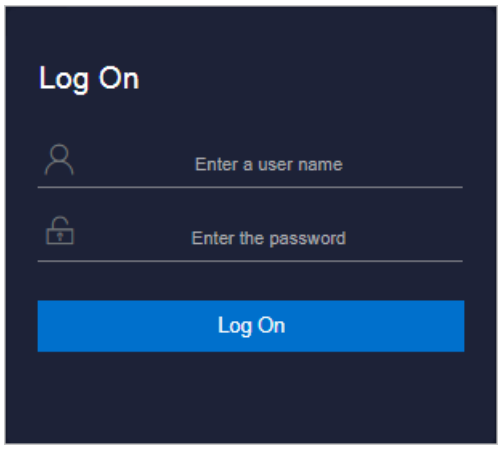

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

When you log on to the ASO console for the first time, you must change the password of your username as prompted.

To enhance security, a password must meet the following requirements:

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- It must be 10 to 20 characters in length.

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**. In the **Big Data Services** section of the page that appears, click **Apsara Bigdata Manager** to go to the homepage of ABM.

11.12.3. Elasticsearch O&M overview

This topic describes the features of Apsara Stack Elasticsearch O&M. It also provides information about how to navigate to the Elasticsearch O&M page.

Modules

Elasticsearch O&M contains four modules: Business, Services, Clusters, and Hosts. The following table describes these modules.

Module	Feature	Description
Business	Cluster Configuration	Allows you to view and modify the cluster configuration files in the **worker** and **kibana** lists for Elasticsearch.
	System Configuration	Allows you to view and modify the system configuration files for Elasticsearch.
Services	Overview	Displays all Elasticsearch services in a cluster. You can view the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.
	Server	Displays all hosts where each Elasticsearch service is running. This allows you to quickly understand service deployment on hosts.
Clusters	Overview	Displays the overall running and health check information about a cluster. On this tab, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Health Status	Displays all checkers for 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 tab, 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 utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Charts	Displays the enlarged trend charts of CPU utilization, memory usage, disk usage, load, and packet transmission.
	Health Status	Displays the health check results of a host. A host has the following health states: CRITICAL, WARNING, EXCEPTION, and OK.
	Services	Displays information about service instances and service instance roles of a host.

Entry

1. Log on to Apsara Bigdata Manager.
2. Click the icon in the upper-left corner and then **Elasticsearch**.
3. In the top navigation bar of the page that appears, click **O&M**. The **Business** page appears.

   O&M contains four modules: **Business**, **Services**, **Clusters**, and **Hosts**.

11.12.4. Business O&M

11.12.4.1. Cluster configuration

This topic describes how to view and modify the cluster configuration files in the work and kibana lists for Elasticsearch in the Apsara Bigdata Manager (ABM) console.

Entry

1. In the upper part of the page, click the **Business** tab.
2. In the left-side navigation pane of the **Business** tab, click **Cluster Configuration**.
3. In the **worker** or **kibana** list, click the cluster configuration file that you want to view. The details of the file appear on the right part of the page.

Modify a cluster configuration file

1. Click the cluster configuration file you want to modify and click **Edit**. Then, modify the file as needed.
2. Click **Save**.
3. Click **Preview**.
   i. In the **Preview** dialog box, compare the differences before and after the file modification.
   ii. If the modification is correct, click **OK**.

4. Click **Submit** in the lower part of the page. The modification is complete.

   If you want to cancel the modification, click **Undo**.

Upload a plug-in

> **Notice** Custom plug-ins may affect the stability of your cluster. Make sure that the custom plug-in you want to upload is reliable, secure, and ready to use. Plug-ins are not automatically updated with Elasticsearch. To update a plug-in, you must manually upload a new version of the plug-in.

1. Select a cluster to which you want to upload a plug-in from the drop-down list. Click **Upload Plug-in**.
2. In the **Upload Plug-in** dialog box, click **Click here to select files for upload** to upload one or more files.

   To delete a file that is not required for the upload, click the icon next to the file.

3. Select the check box in the dialog box and click **OK**.

11.12.4.2. System configuration

This topic describes how to view and modify the system configuration files for Elasticsearch in Apsara Bigdata Manager (ABM).

Entry

1. At the top of the **O&M** page, click the **Business** tab.
2. On the **Business** page that appears, click **System Configuration** in the left-side navigation pane.
3. Click a configuration file that you want to view. The details of the file appear on the right.

Modify a system configuration file

1. Click a system configuration file to be modified and click **Edit** to modify the configuration file.
2. Click **Save**.
3. Click **Preview**.
   i. In the **Preview** dialog box that appears, you can compare the differences before and after the file modification.
   ii. If the modification is correct, click **OK**.
4. Click **Submit** at the bottom of the page. The modification is completed.

   If you want to undo the modification, click **Undo**.

11.12.5. Service O&M

11.12.5.1. Service overview

The service overview page lists all Elasticsearch 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 the **Services** tab.
2. On the **Services** page that appears, select a service in the left-side navigation pane. Click the **Overview** tab.

   On the Overview page that appears, 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.

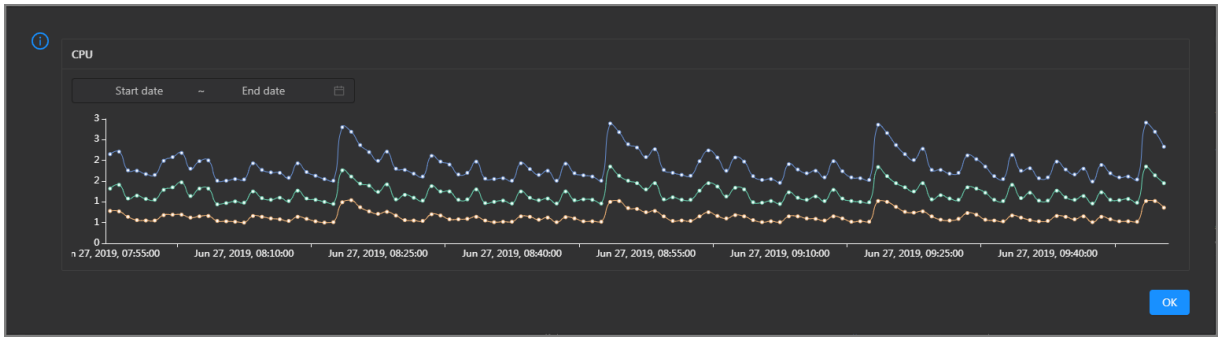

DISK

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.

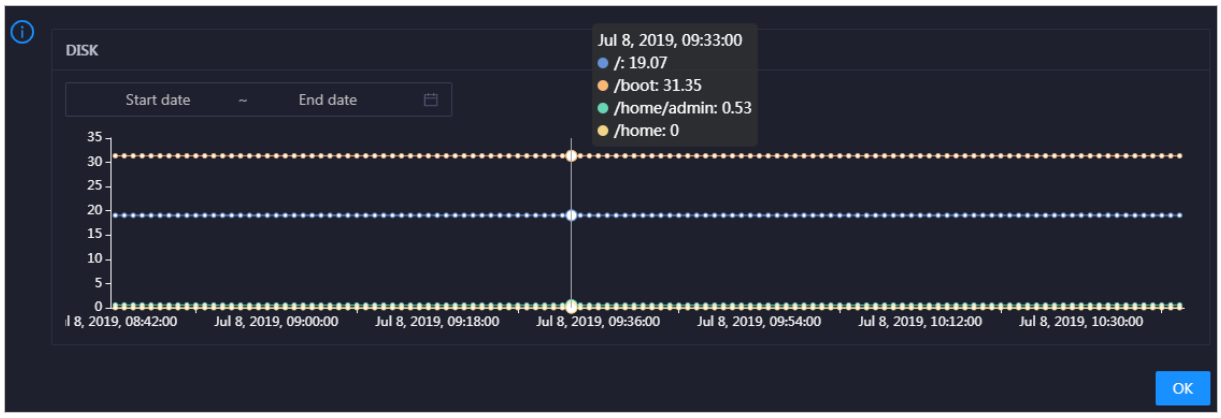

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

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.

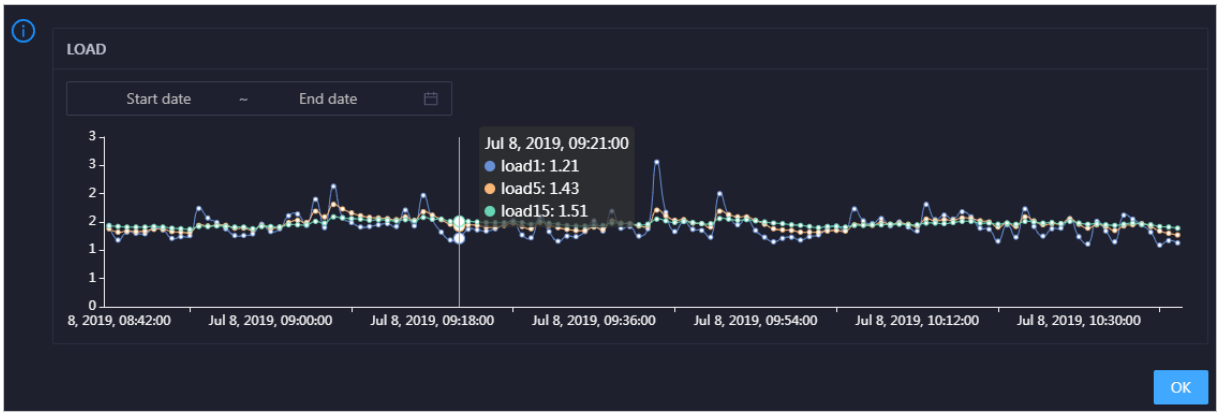

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

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.

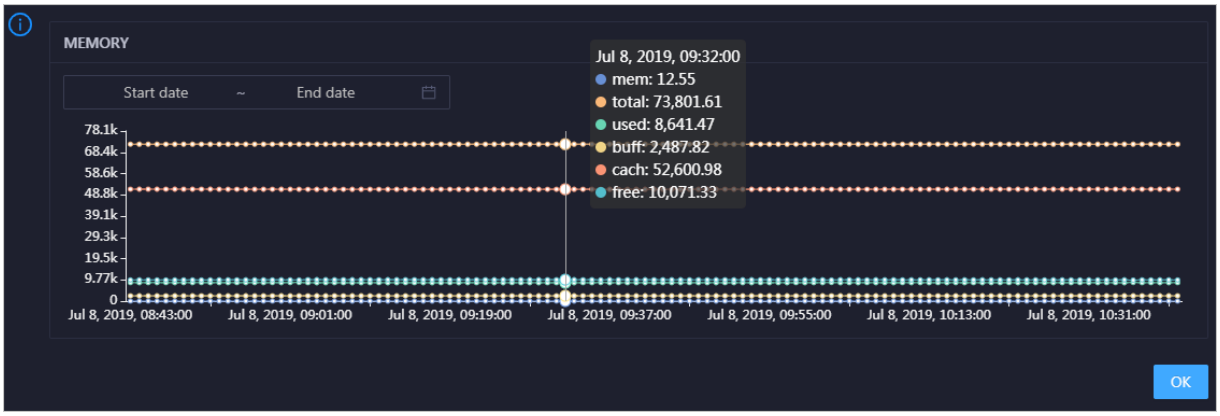

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

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.

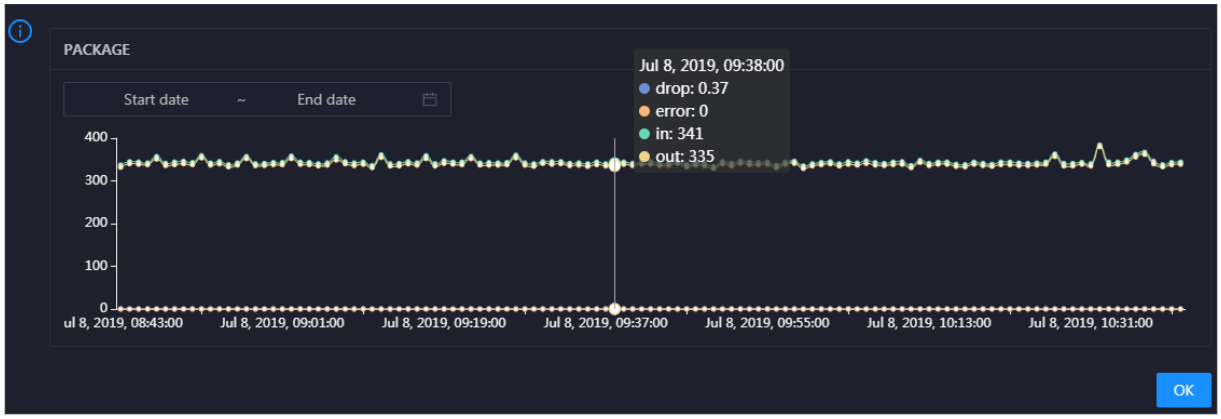

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

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.

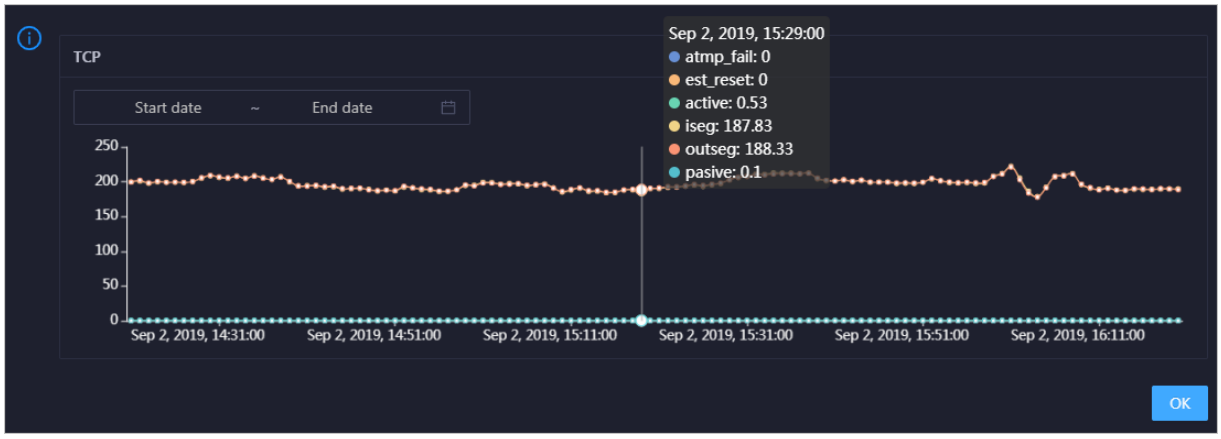

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

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.

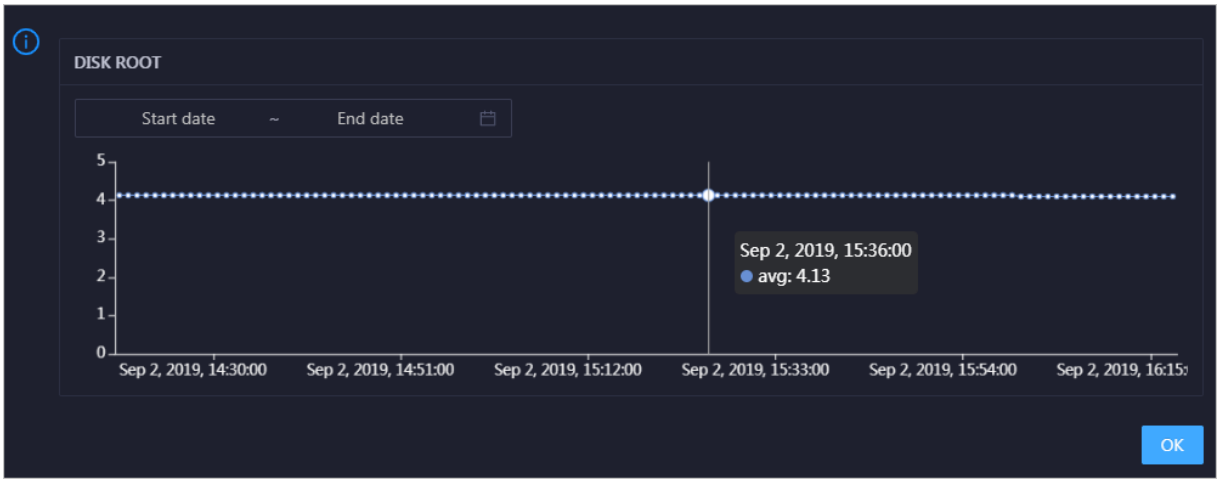

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

11.12.5.2. Service hosts

This topic describes how to view all hosts where each Elasticsearch service is run.

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

1. At the top of the **O&M** page, click the **Services** tab.
2. On the **Services** page that appears, select a service in the left-side navigation pane.
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.

11.12.6. Cluster O&M

11.12.6.1. Cluster overview

The Overview tab of a cluster displays the overall running and health check information about the cluster. On this tab, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

1. In the upper part of the page, click the **Clusters** tab.
2. On the **Clusters** tab, select a cluster in the left-side navigation pane and click the **Overview** tab.

Hosts

This section displays host states and the number of hosts in each state. The host states include **good** and **bad**.

Service Status

This section displays all the services that are deployed in the cluster. It also provides information about the numbers of available and unavailable services.

Health Check

This section displays the number of checkers for the cluster and the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

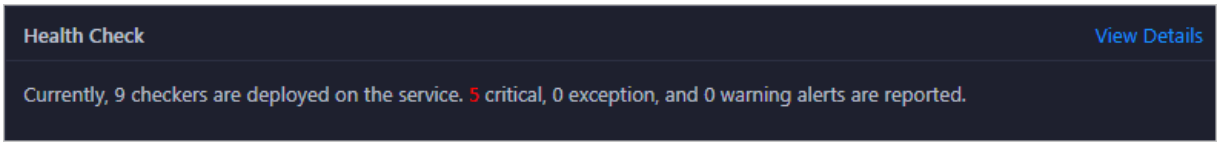

You can click **View Details** to go to the Health Status tab. On this tab, you can view the health check details.

Health Check History

This section displays the historical health checks that are performed on the cluster.

You can click **View Details** to go to the Health Status tab. On this tab, you can view the health check details.

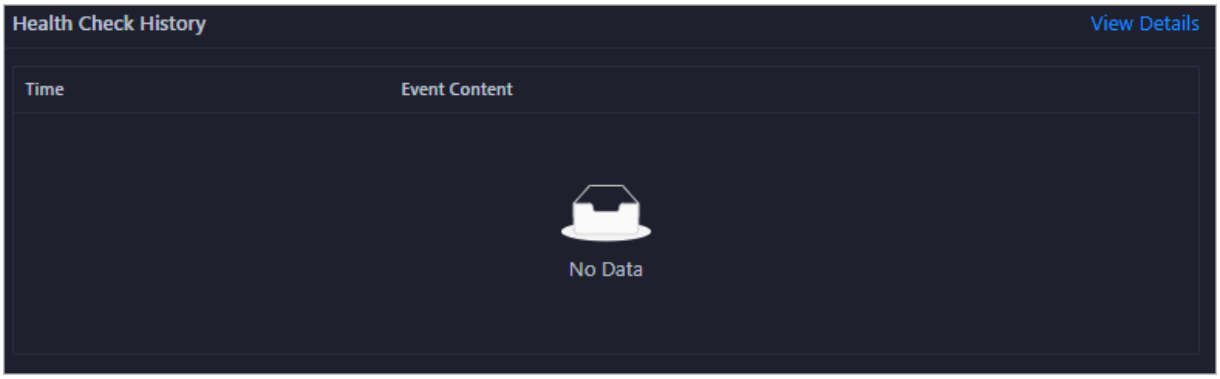

You can click the event content of a check to view exception items.

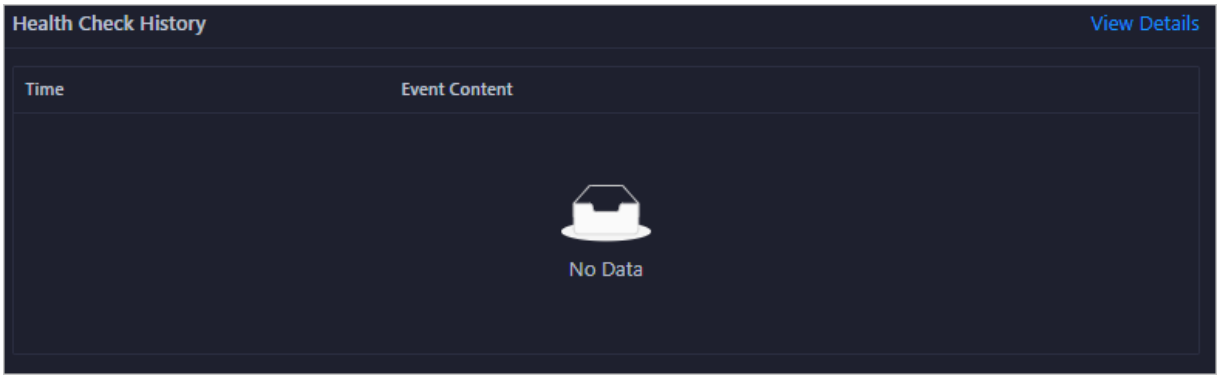

CPU

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

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

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

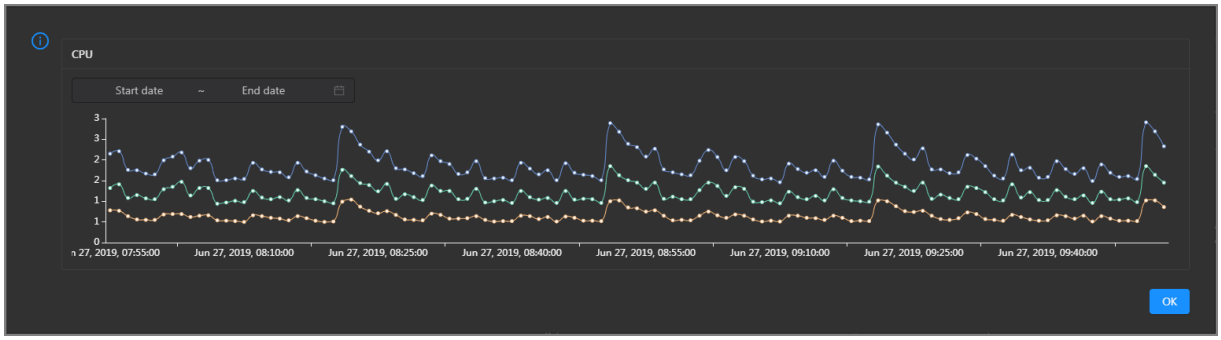

DISK

This chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the cluster in different colors.

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

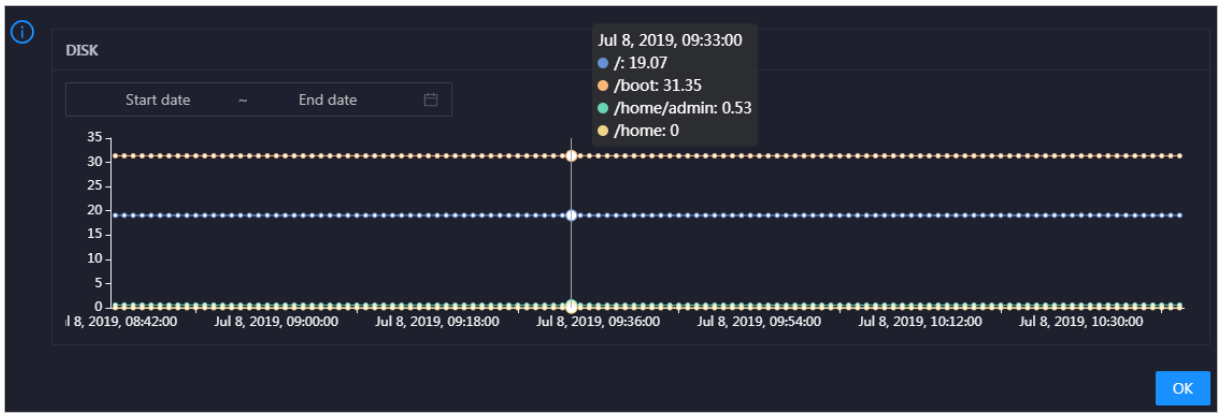

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

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

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

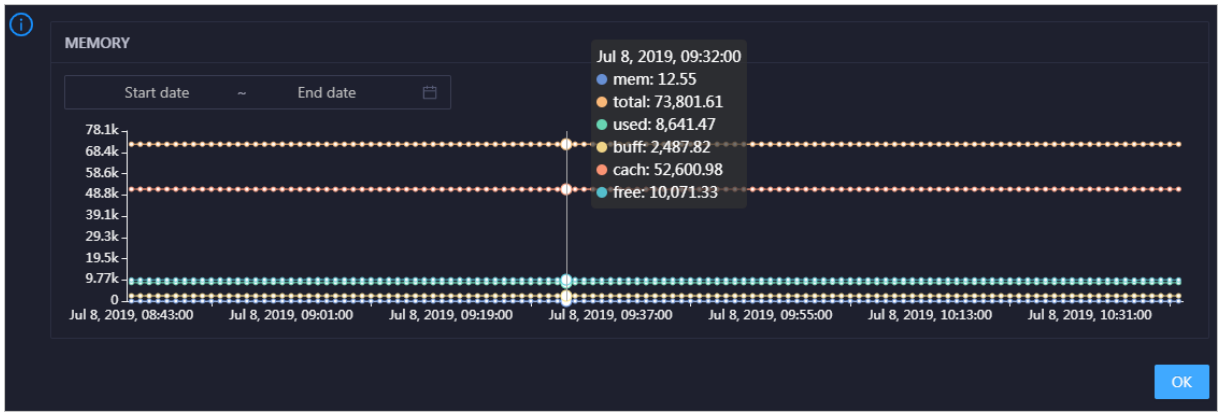

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart shows the trend lines of the numbers of dropped packets (drop), error packets (error), received packets (in), and sent packets (out) for the cluster in different colors. These trend lines reflect the data transmission status of the cluster.

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

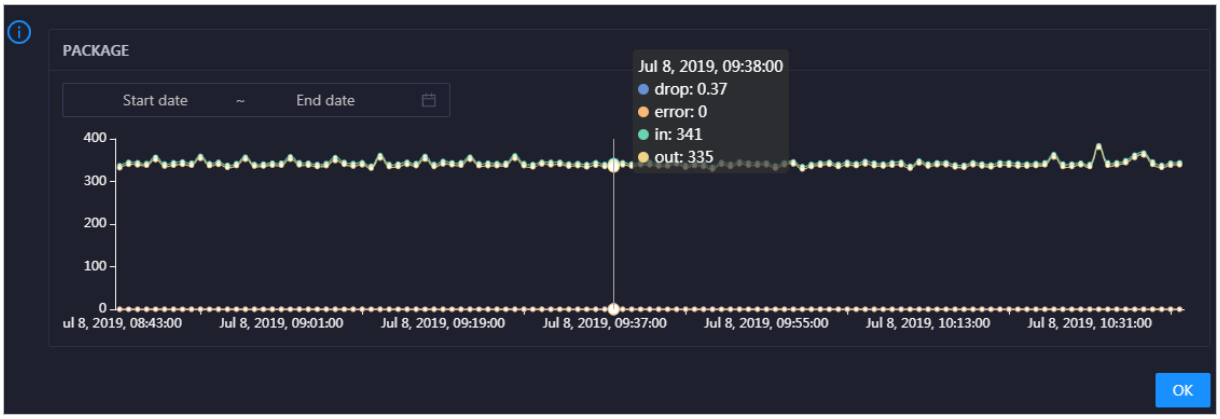

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

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

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

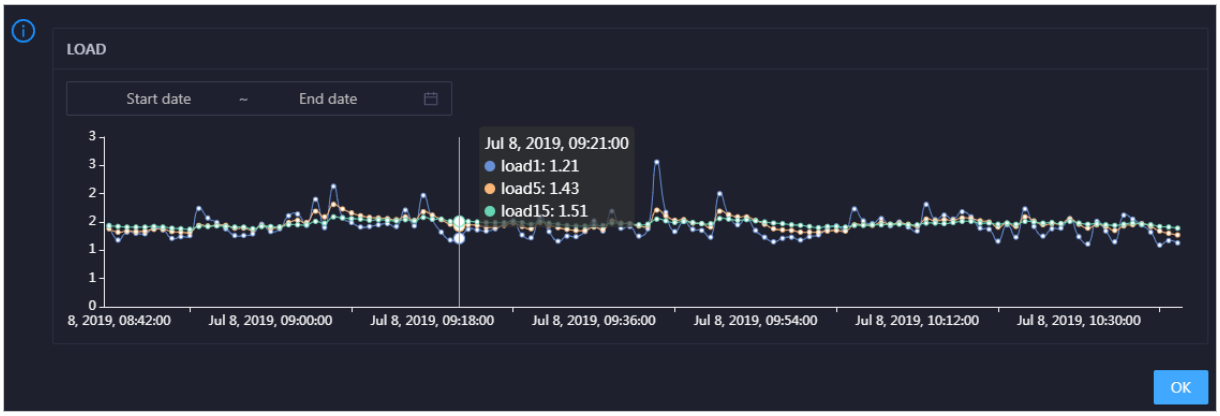

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

11.12.6.2. Cluster health

On the Health Status tab of a cluster, you can view all checkers for the 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

In the upper part of the page, click the **Clusters** tab. Then, select a cluster in the left-side navigation pane and click the **Health Status** tab.

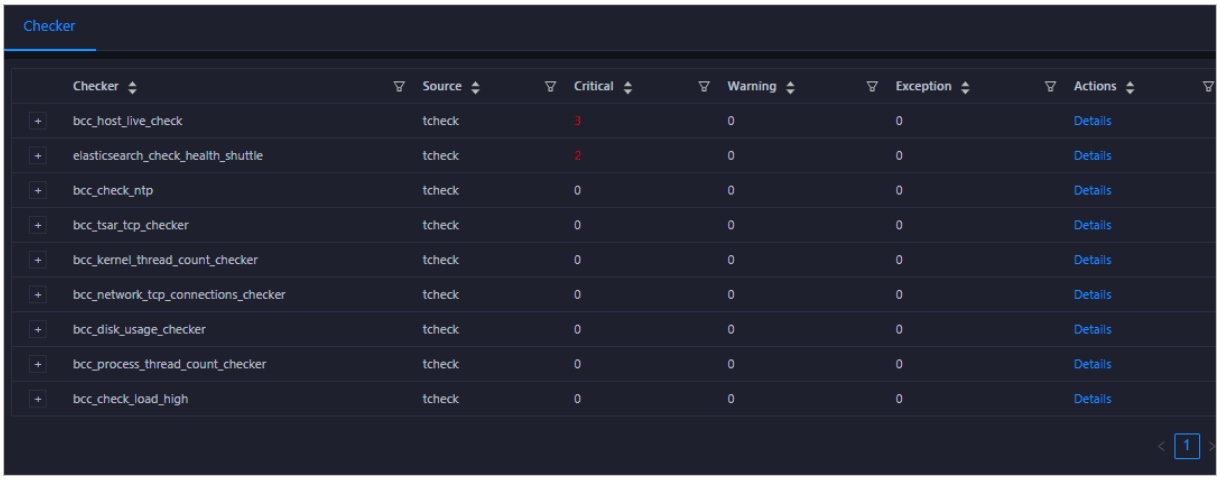

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are repesented in different colors. Handle the alerts in a timely manner,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

View checker details

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

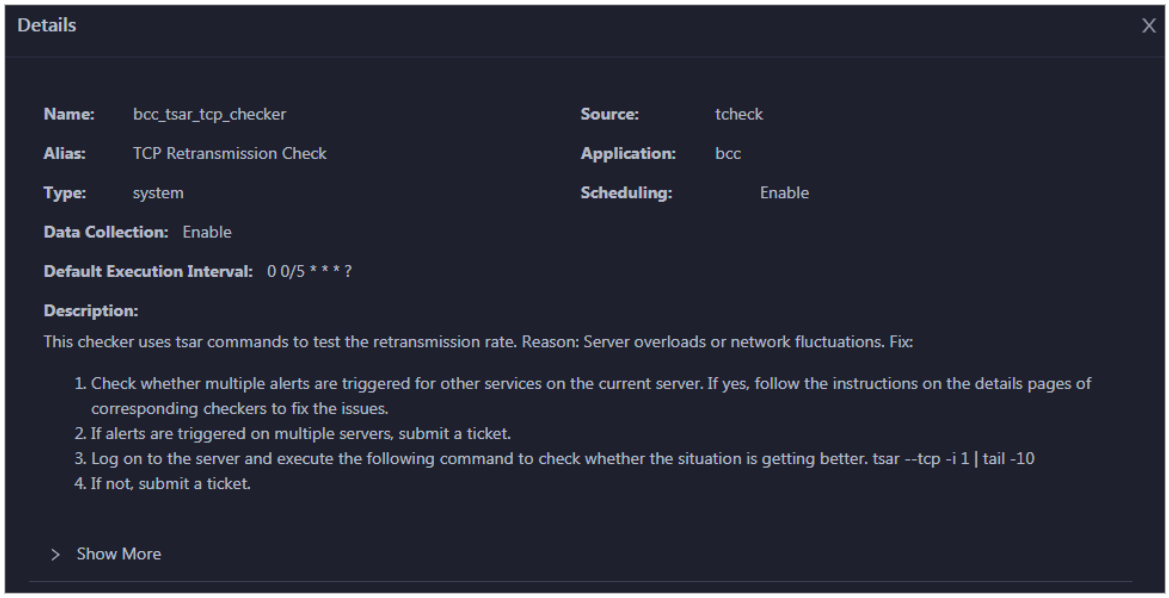

   The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

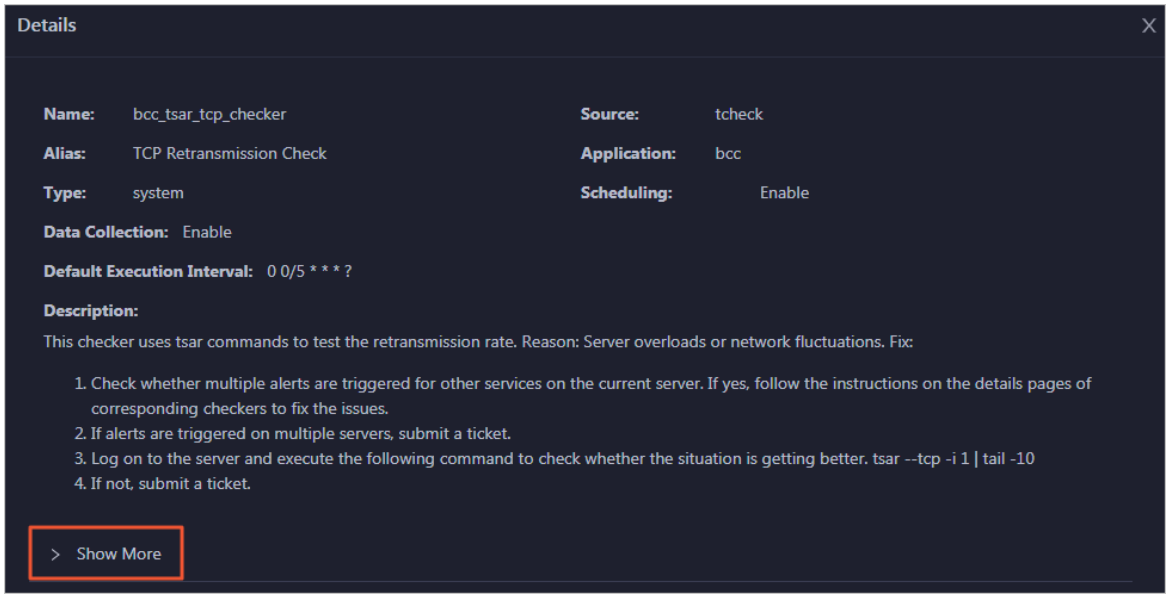

You can view information about **Script**, **Target**, **Default Threshold**, and **Mount Point**.

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

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

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

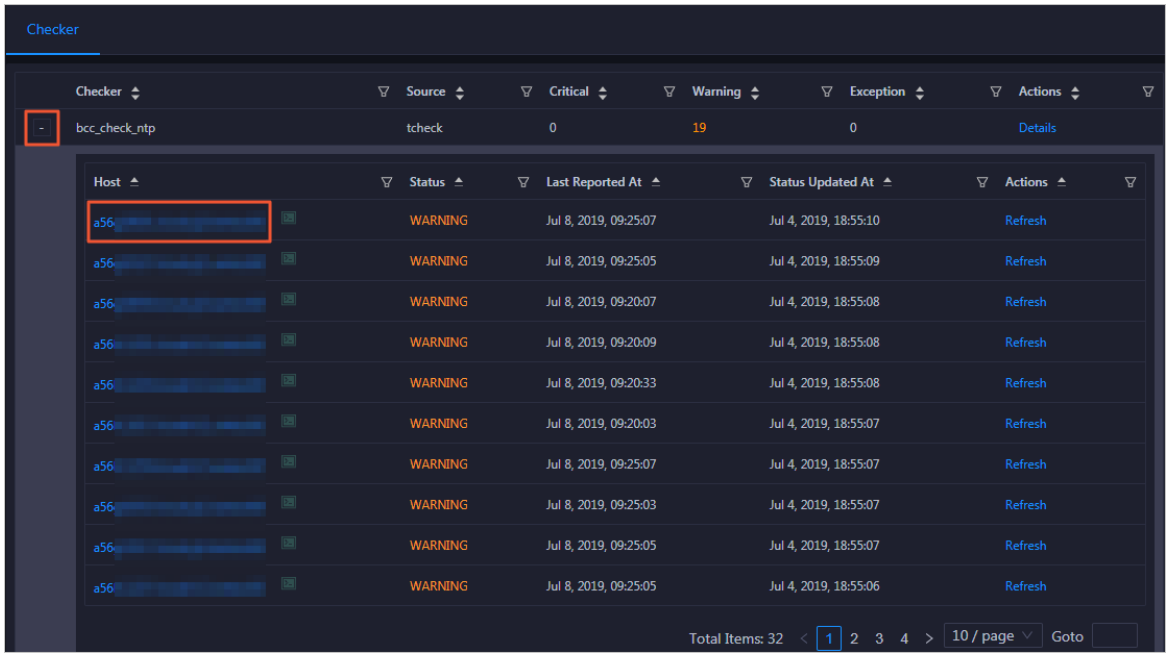

2. Click a hostname. In the pane that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

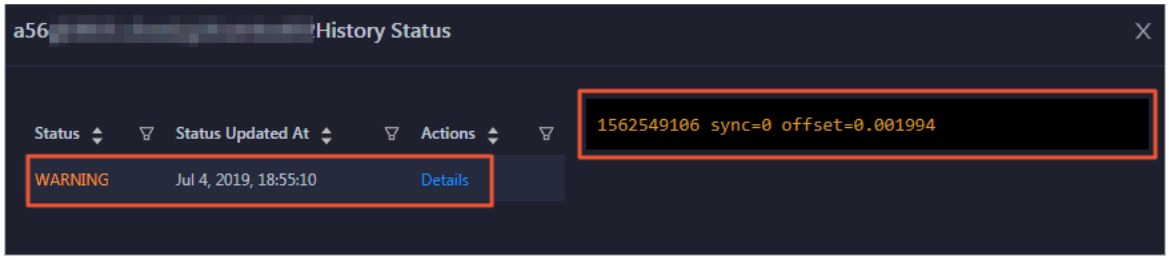

Clear alerts

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

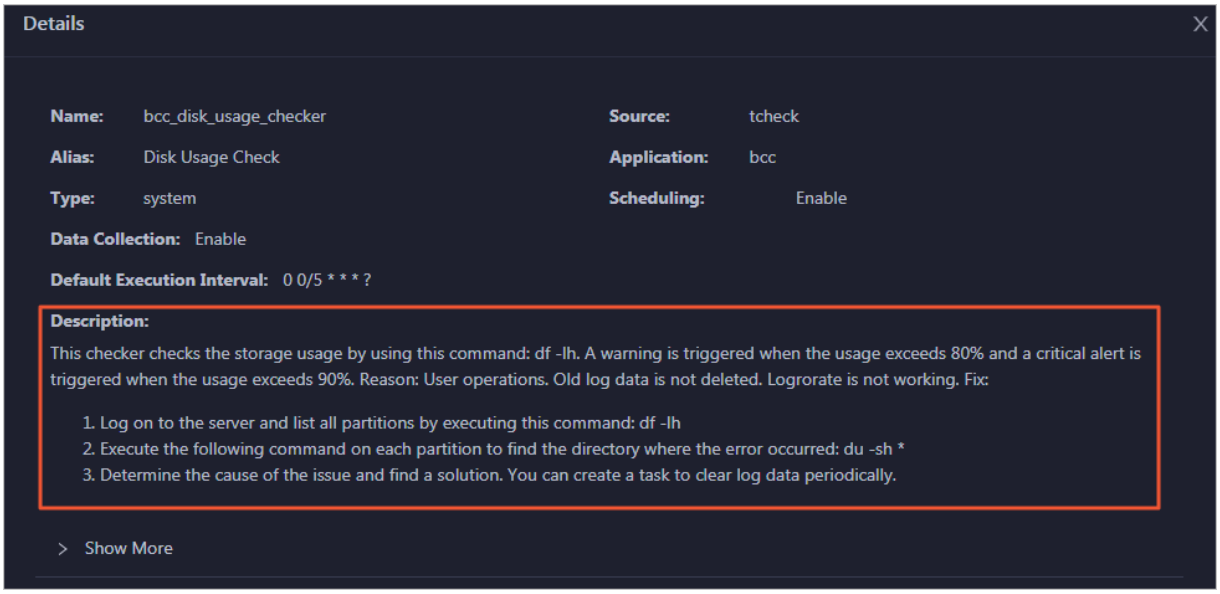

Log on to a host

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

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

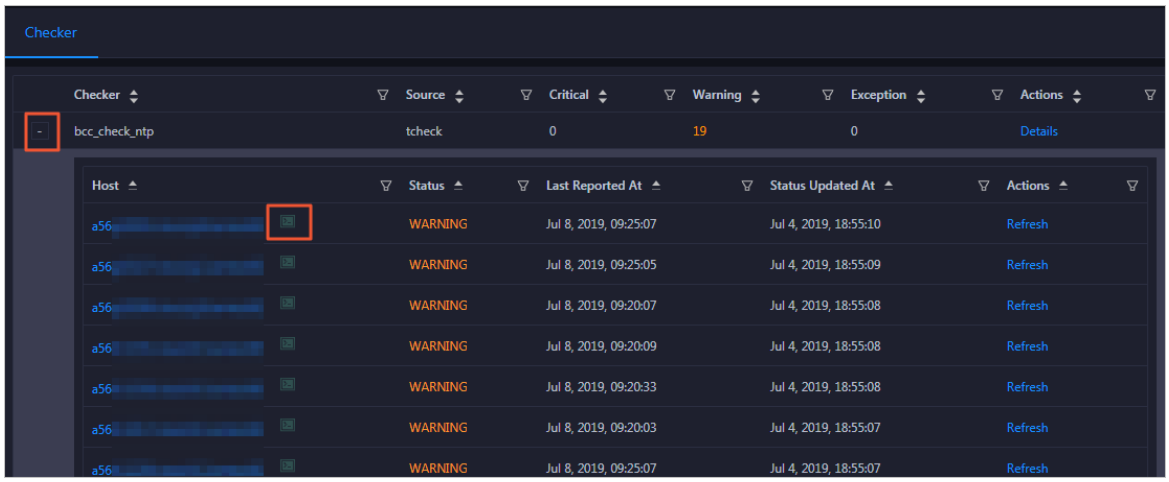

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

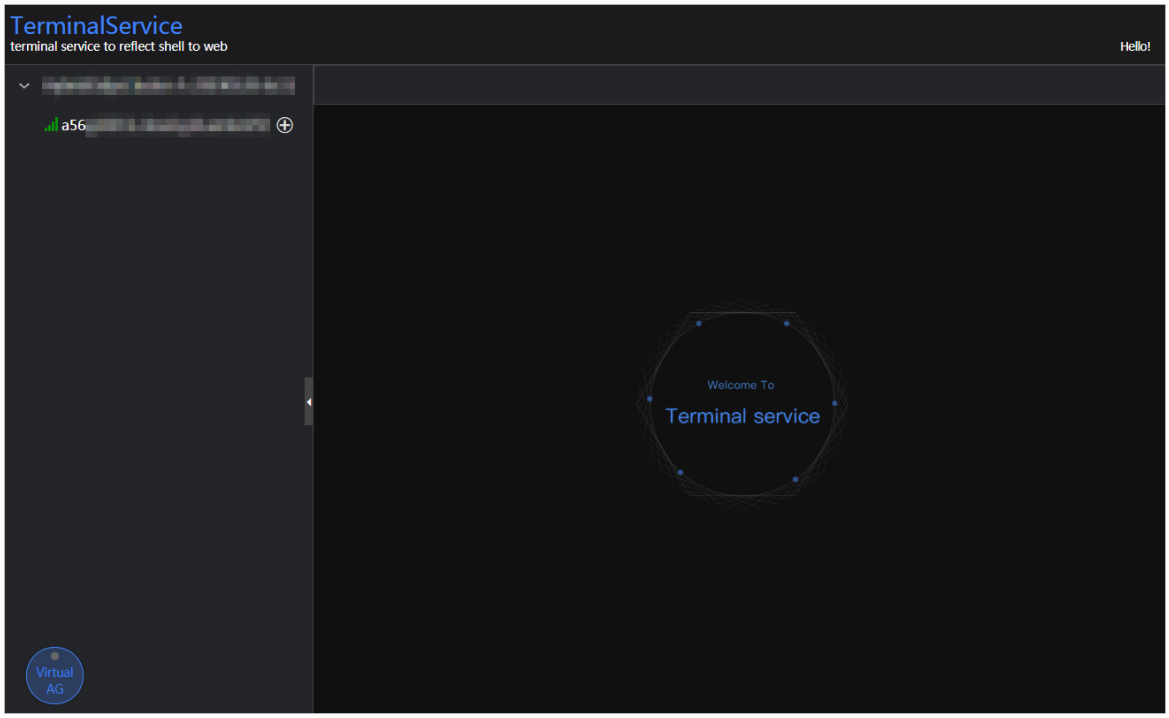

3. On the **TerminalService** page, click the hostname to log on to the host.

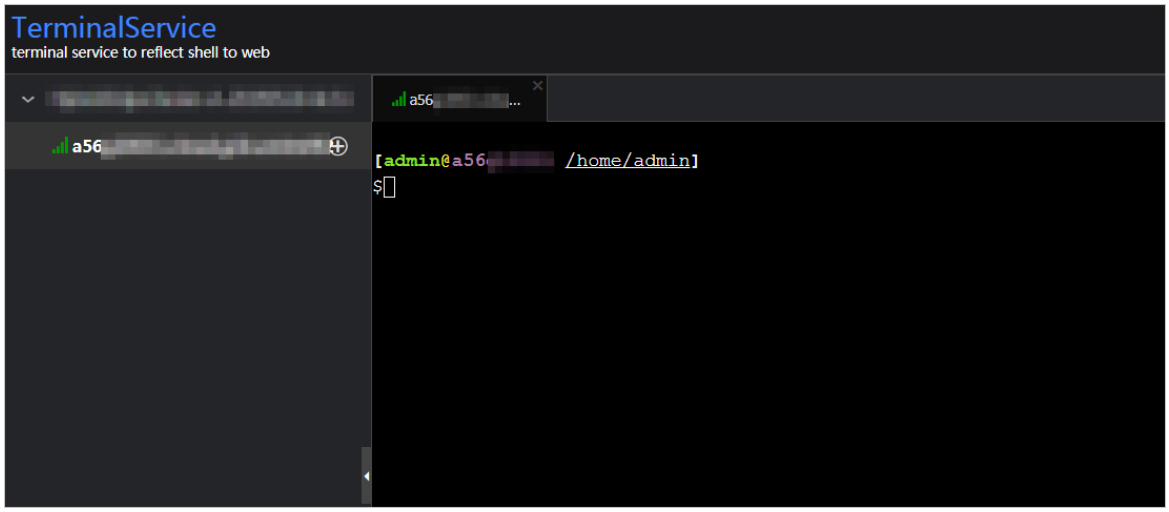

Run a checker again

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

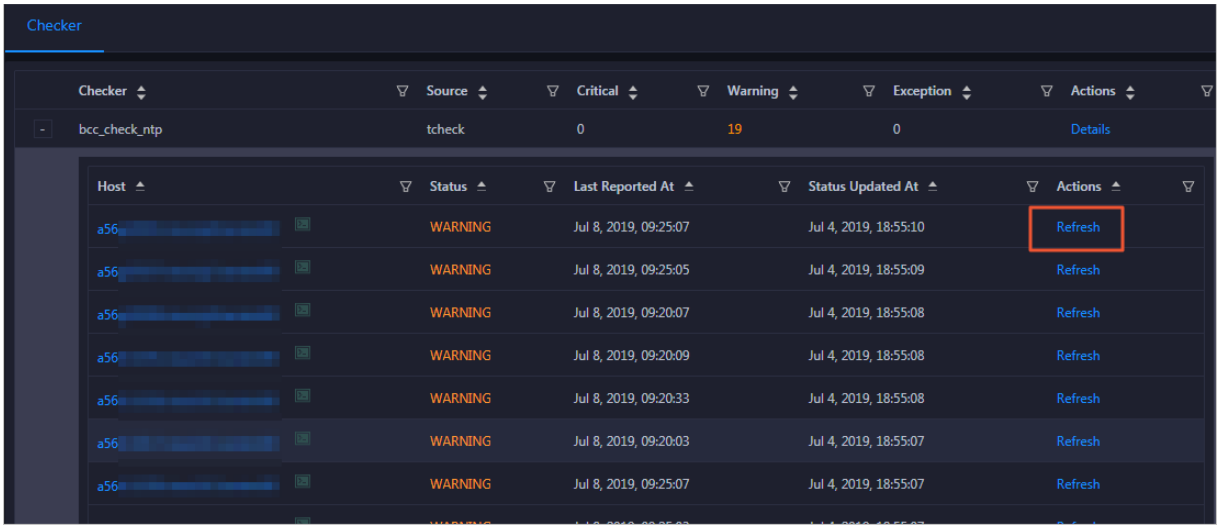

11.12.7. Host O&M

11.12.7.1. 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 information, service role status, 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.

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.

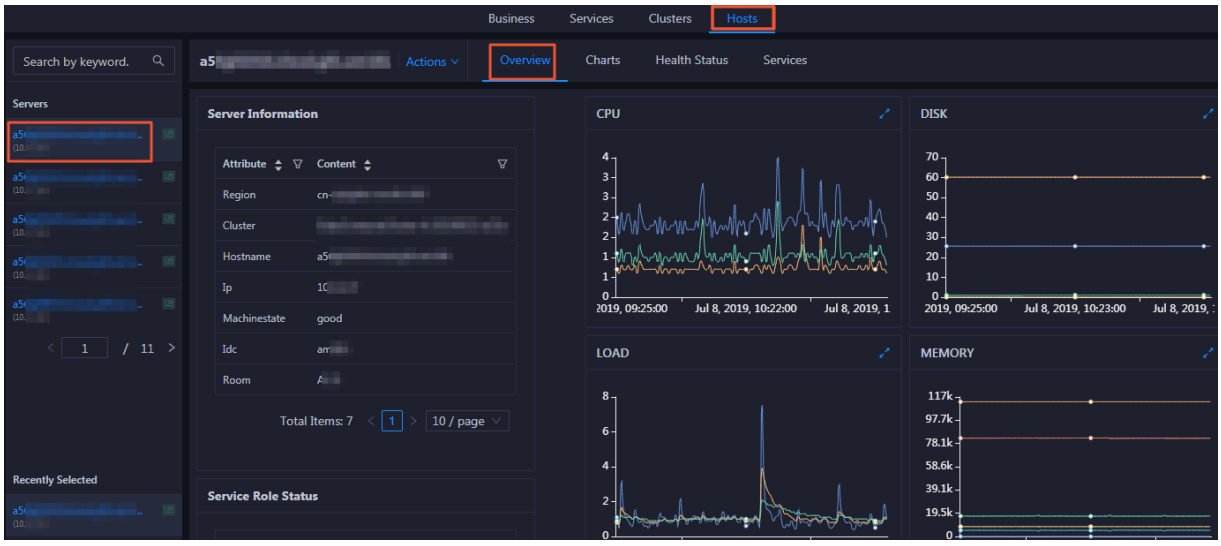

On the **Overview** page, you can view the attributes, services, service roles, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host.

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.

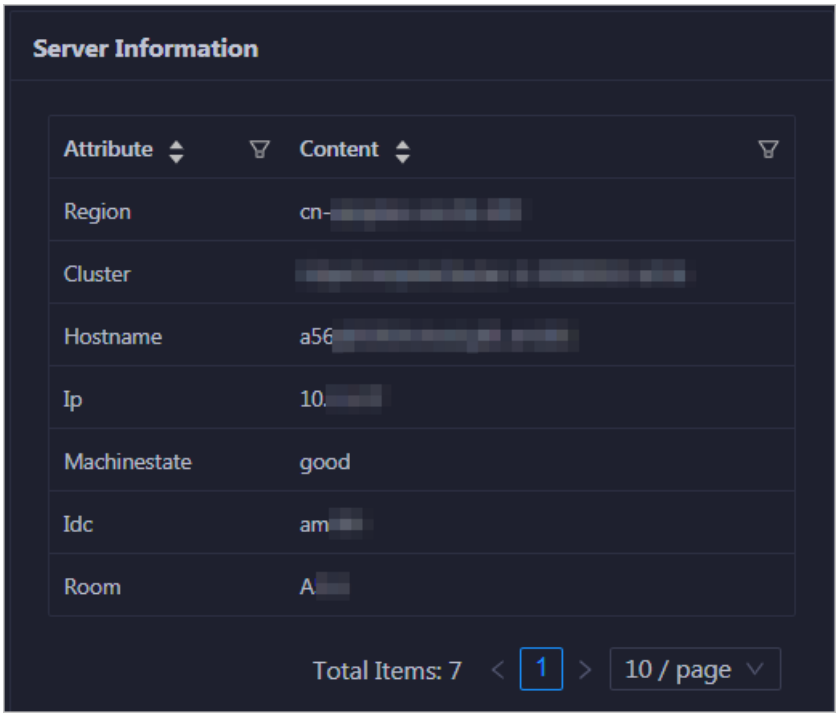

Service Role Status

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.

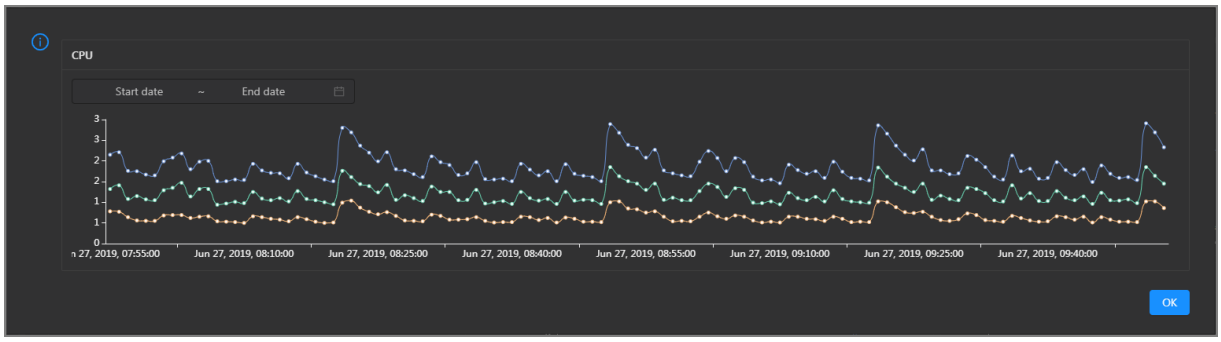

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

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.

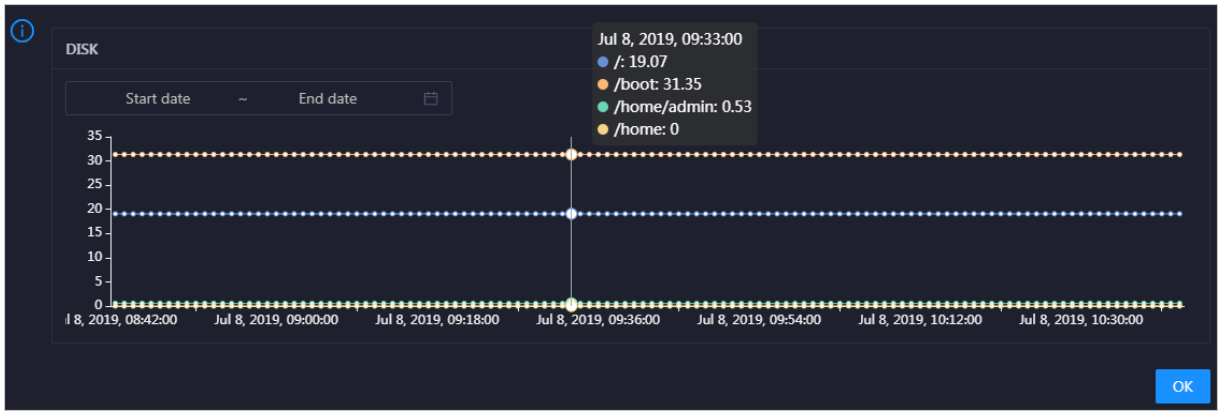

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

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.

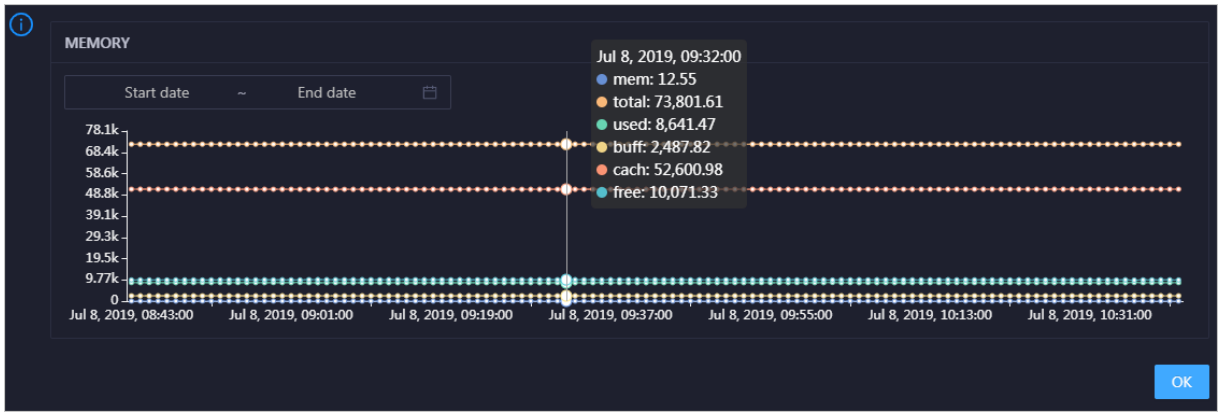

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

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.

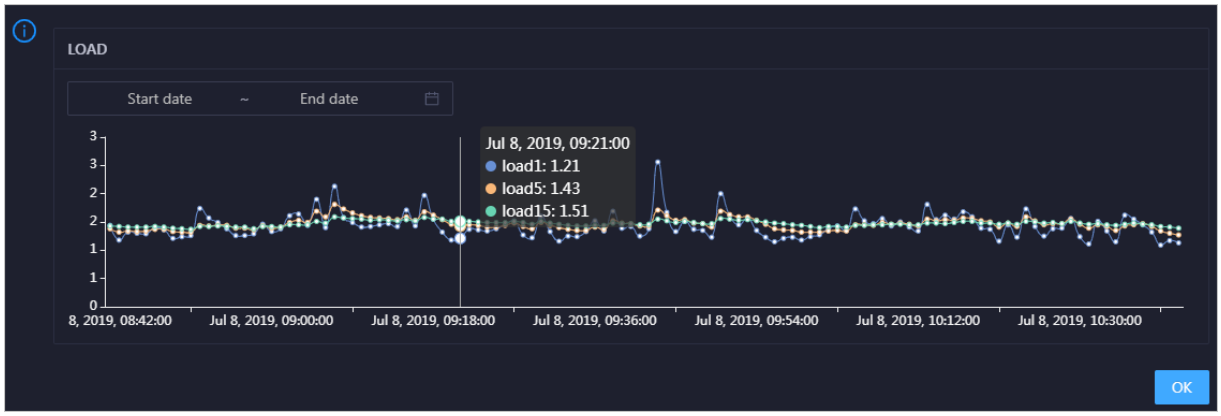

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

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 transmission status of the host.

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

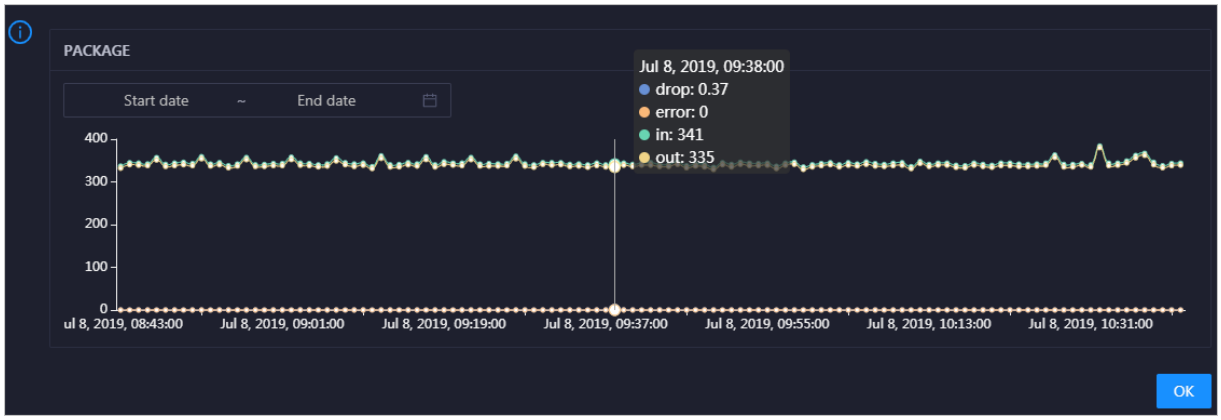

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

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

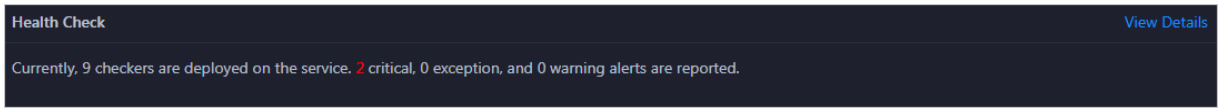

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

This section displays a record of the health checks performed on the host.

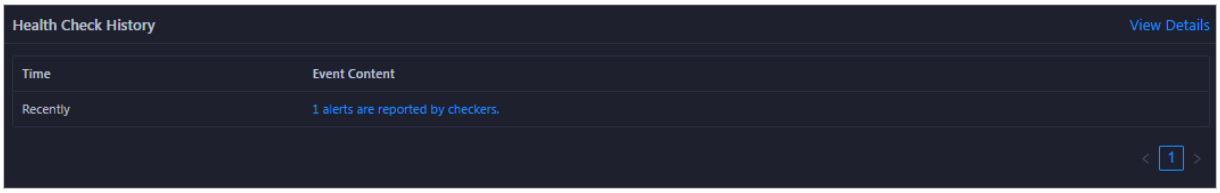

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

You can click the event content of a check to view the exception items.

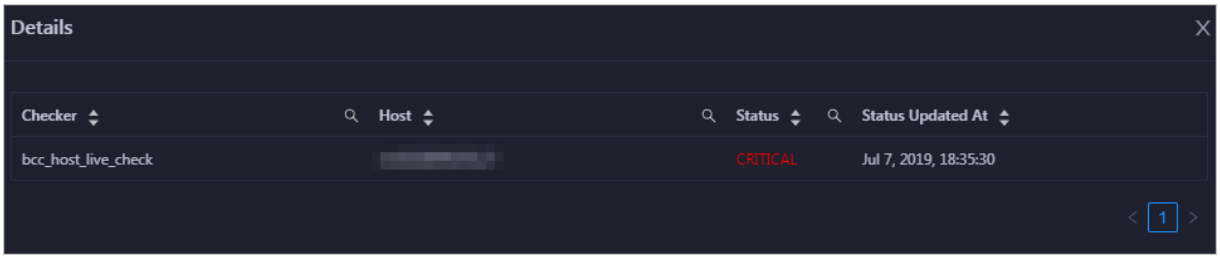

11.12.7.2. Host charts

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** tab, select a host in the left-side navigation pane and click the **Charts** tab. The **Charts** tab for the host appears.

The **Charts** tab displays the trend charts of CPU utilization, disk usage, memory usage, load, and packet transmission for the host. For more information, see Host overview.

11.12.7.3. Host health

On the Health Status tab of a host, you can view the checkers for 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 a host and perform manual checks on the host.

Entry

In the upper part of the page, click the **Hosts** tab. Then, select a cluster in the left-side navigation pane and click the **Health Status** tab.

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.

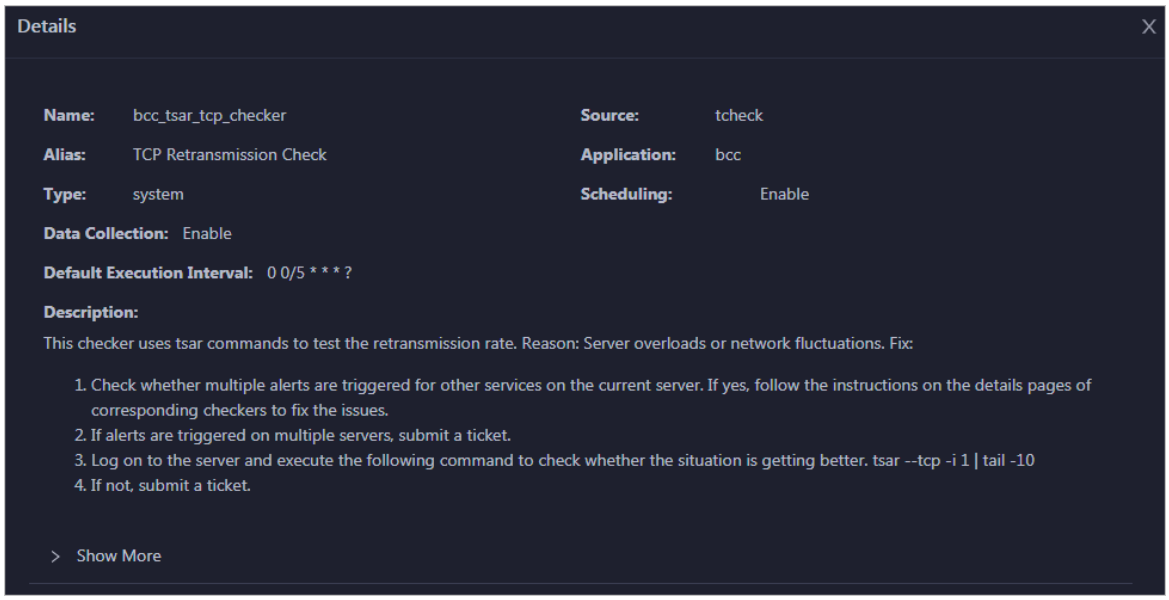

   The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click **Show More** at the bottom to view more information about the checker.

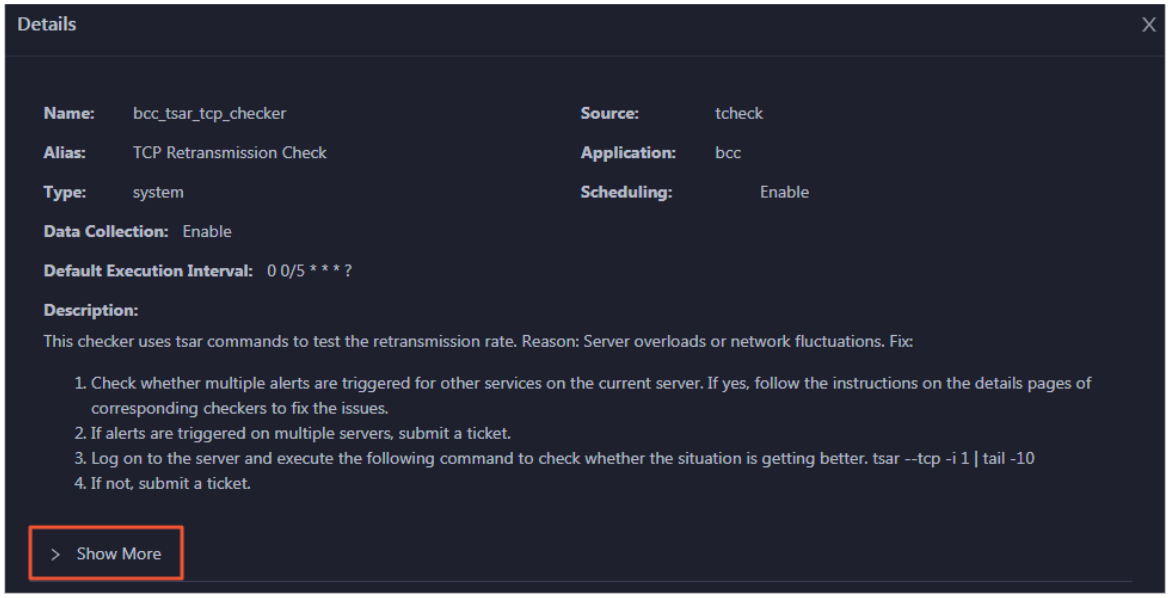

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

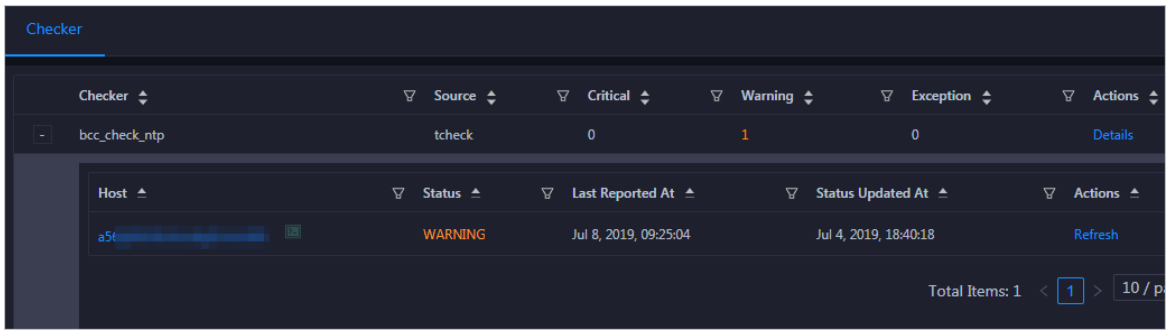

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

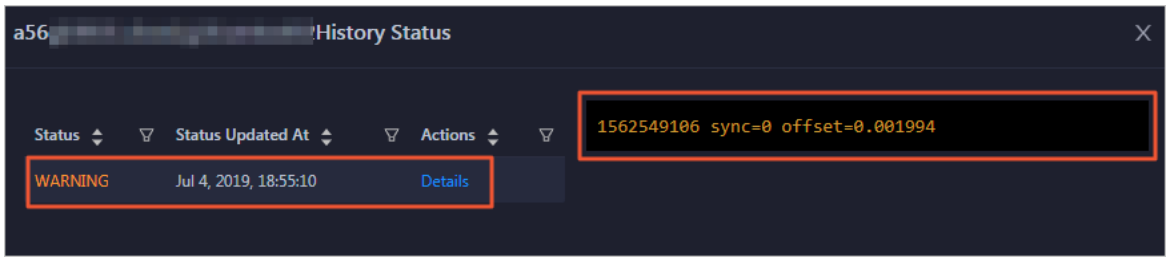

Clear alerts

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.

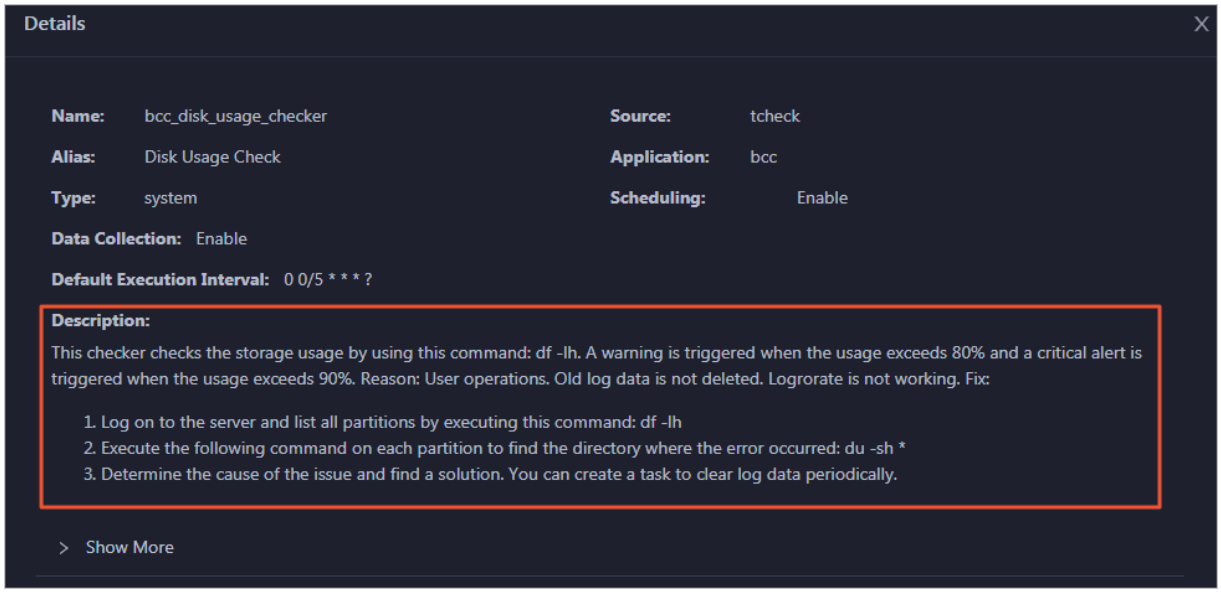

Log on to a host

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.

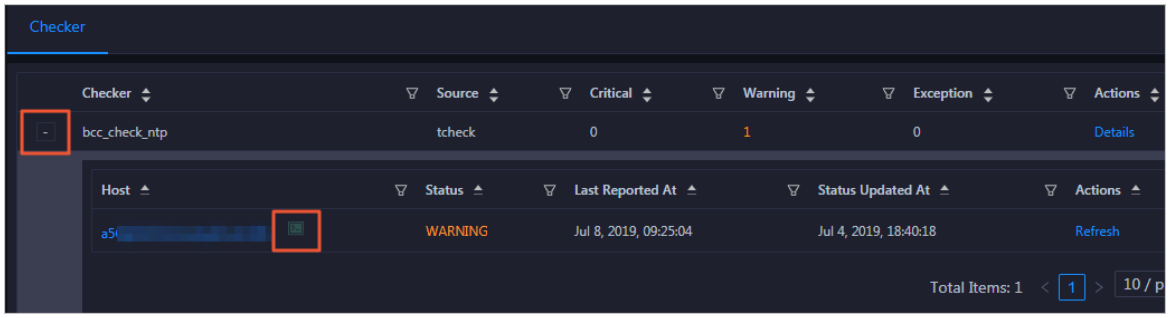

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

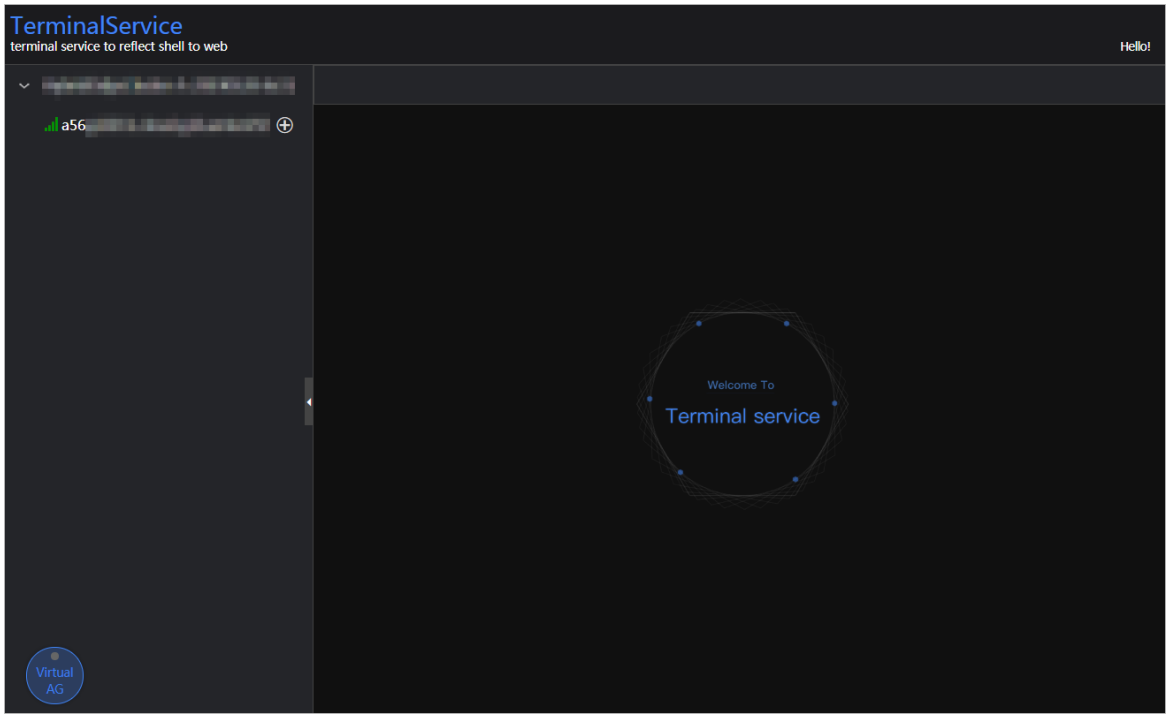

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

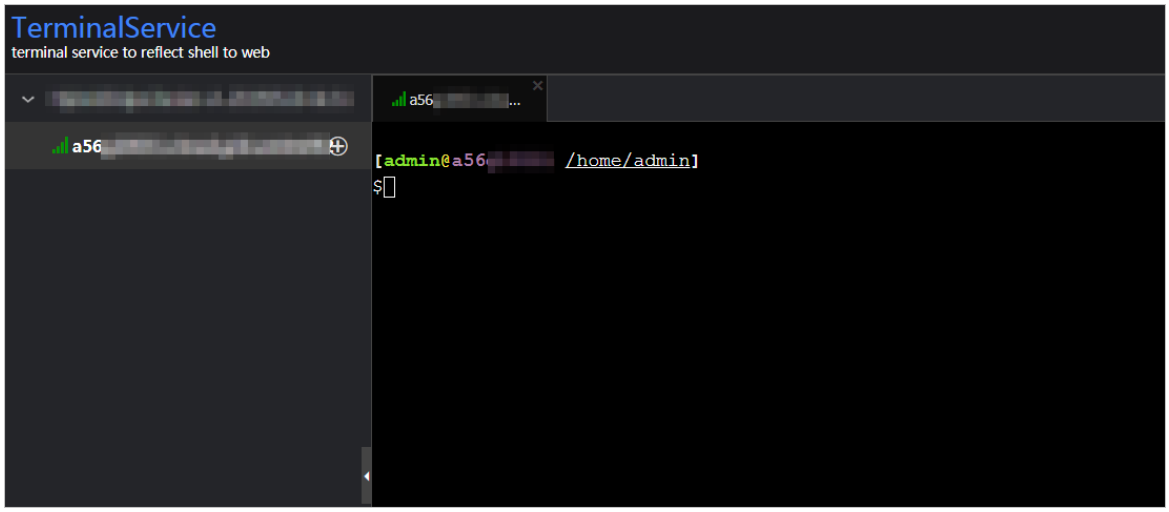

Run a checker again

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.

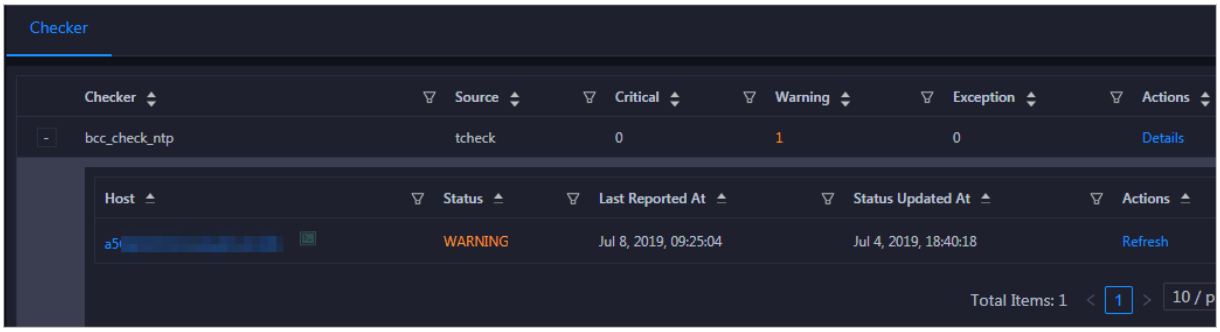

11.12.7.4. Host services

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.

11.12.8. Online O&M

11.12.8.1. Cluster health

You can view the statistical information of Elasticsearch clusters. The cluster health information is the most important. An Elasticsearch cluster has three health states: red, yellow, and green. This topic describes how to view the health status of an Elasticsearch cluster. It also provides more information about the preceding states.

You can run the following command to view the health status of a cluster:

```
curl -u Username:Password http://domain:9200/_cluster/health
```

State	Description	Remarks
red	Not all of the shards are available.	One or more indexes have unassigned shards.
yellow	All shards are available, but not all of the replicas are available.	One or more indexes have unassigned replicas.
green	All shards and replicas are available.	All indexes in the cluster are healthy and do not have unassigned shards or replicas.

Notice To ensure that the health state of your Elasticsearch cluster is green, all shards and replicas must be available at all times. We recommend that the number of replicas be less than or equal to amount_Node minus one. amount_Node represents the number of nodes. This ensures that the health state of your Elasticsearch cluster is green after it is restarted when dedicated master nodes are used.

11.12.9. Common failure troubleshooting

11.12.9.1. Resolve the issue that the health state of an Elasticsearch cluster is yellow

If the health state of your Elasticsearch is yellow, operations such as password resets and cluster upgrades are time-consuming. We recommend that you perform these operations when the health state of the cluster is green.

Cause: Some replicas are unassigned. You must check which indexes in the cluster have unassigned replicas.

11.12.9.2. Query index status

This topic describes how to view the status of an index in an Elasticsearch cluster.

You can run the following command to check which indexes have unassigned replicas:

```
curl -u Username:Password http://domain:9200/_cat/indices
```

If the cause of the issue is that the number of replicas is greater than amount_Node minus one, you must change the number of replicas for those indexes.

11.12.9.3. Recover an index

If your Elasticsearch cluster has three nodes and one or more indexes have three replicas, the health state of the cluster is yellow. This topic describes how to recover the indexes.

You can run the following command to set the number of replicas to 2:

```
curl -XPUT -u Username:Password http://domain:9200/Name of the index with unassigned replicas/_settings
-H 'Content-Type:
application/json' -d '{"index":{"number_of_replicas":(amount_Node - 1)}'
```

Note After you perform operations such as restart, scale-out, and configuration modification, set an appropriate number of replicas based on the number of nodes. This improves the reliability and stability of your Elasticsearch cluster.

12.Appendix

12.1. Operation Access Manager (OAM)

12.1.1. OAM introduction

Overview

Operation Access Manager (OAM) is a centralized permission management platform of Apsara Stack Operations (ASO). OAM uses a simplified role-based access control (RBAC) model. Administrators can use OAM to assign roles to operations personnel, granting them corresponding operation permissions to operations systems.

OAM permission model

In RBAC, administrators do not directly grant system operation permissions to users. Instead, they create a collection of roles between a collection of users and a collection of permissions. Each role corresponds to a group of permissions. If a role is assigned to a user, the user is granted all the operation permissions of that role. Therefore, when creating a user, administrators are only required to assign a role to the user, saving the trouble to grant specific permissions to the user. In addition, the frequency of role permission changes is less than that of user permission changes, simplifying the user permission management and reducing the system overhead.

See the OAM permission model as follows.

Permission model

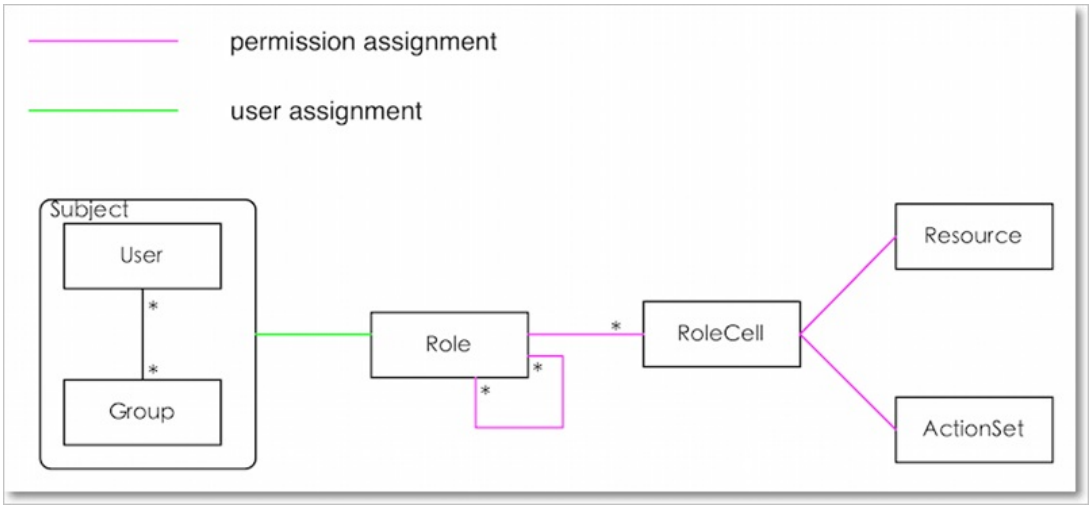

12.1.2. Instructions

Before using Operation Access Manager (OAM), you must know the following basic concepts about permission management.

subject

Operators of the access control system. OAM has two types of subjects: users and groups.

user

Administrators and operators of operations systems.

group

A collection of users.

role

The core of the role-based access control (RBAC) system.

Generally, a role can be regarded as a collection of permissions. A role can contain multiple RoleCells or roles.

RoleHierarchy

In the OAM system, a role can contain other roles to form RoleHierarchy.

RoleCell

The specific description of a permission. A RoleCell consists of resources, ActionSets, and available authorizations.

resource

The description of an authorized object. For more information about resources of operations platforms, see Permission lists of operations platforms.

ActionSet

The description of authorized actions. An ActionSet can contain multiple actions. For more information about actions of operations platforms, see Permission lists of operations platforms.

available authorizations

The maximum number of authorizations in the cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.

For example, if administrator A sets **Available Authorizations** to 5 when granting a permission to administrator B, the permission can be granted for another five times at most. When administrator B grants the permission to administrator C, the value of **Available Authorizations** cannot be greater than 4. If **Available Authorizations** is set to 0 when administrator B grants the permission to operator D, operator D can only use the permission but cannot grant it to others.

> **Note** Currently, OAM does not support the cascaded revocation for cascaded authorization. Therefore, administrator C and operator D still have the permission even if the permission is revoked for administrator B.

12.1.3. Quick Start

By completing the steps in this guide, you will learn how to create and assign roles for O&M.

12.1.3.1. Log on to OAM

This topic describes how to log on to Operation Administrator Manager (OAM).

Prerequisites

- The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

  The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL *region-id*.aso.*intranet-domain-id*.com and press the Enter key.

   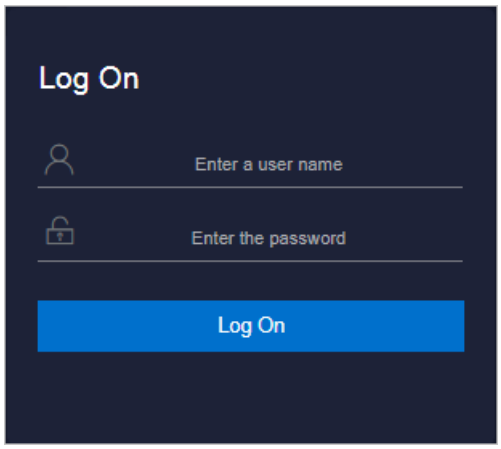

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

   > **Note** Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

   When you log on to the ASO console for the first time, you must change the password of your username as prompted.

   To enhance security, a password must meet the following requirements:

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**.
6. In the **Apsara Stack O&M > Basic O&M** section, click **OAM**.

12.1.3.2. Create groups

You can create user groups for centralized management.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. On the **Owned Groups** page, click **Create Group** in the upper-right corner. In the **Create Group** dialog box that appears, set **Group Name** and **Description**.

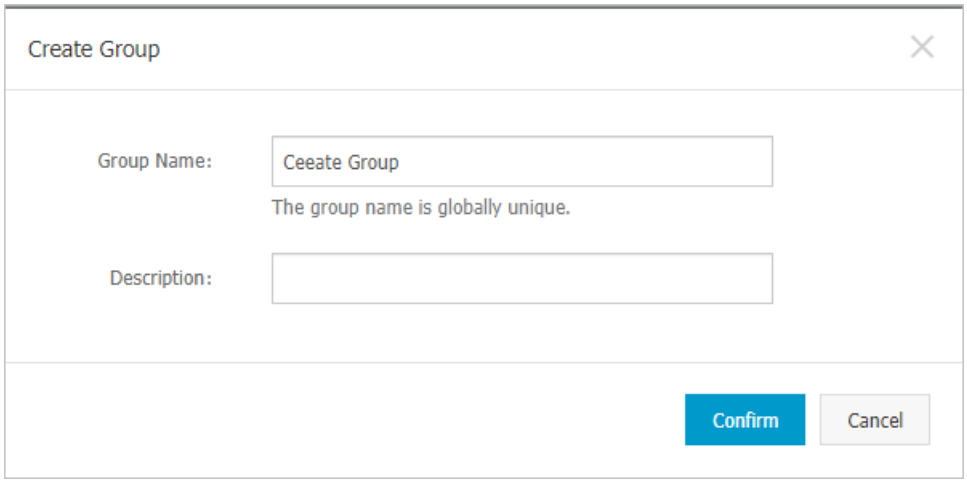

4. Click **Confirm**.After the group is created, it is displayed on the **Owned Groups** page.

12.1.3.3. Add group members

You can add members to an existing group to grant permissions to the group members in a centralized manner.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.
4. In the upper-right corner of the **Group Member** section, click **Add Member**.

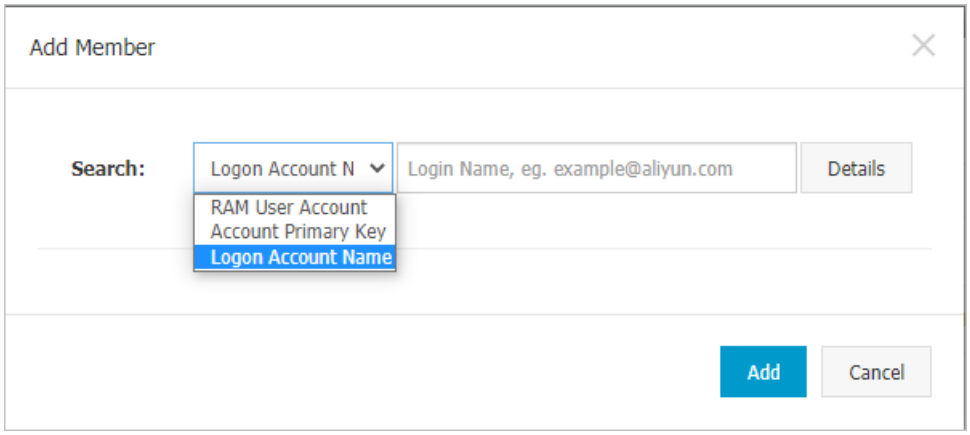

5. Select a search mode, enter the corresponding information, and click **Details**. Details of the specified user are displayed.

   Three search modes are available:

   - **RAM User Account**: Search in the format of RAM user@Apsara Stack tenant account ID.
   - **Account Primary Key**: Search by using the unique ID of the Apsara Stack tenant account.
   - **Logon Account Name**: Search by using the logon name of the Apsara Stack tenant account.
6. Click **Add**.
7. You can repeat the preceding steps to add multiple group members.To remove a member from a group, click **Remove** in the Actions column corresponding to the member.

12.1.3.4. Add group roles

You can add roles to an existing group.

Prerequisites

- The role to be added is created. For more information, see Create roles.
- You are the owner of the group and the role.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.
4. In the upper-right corner of the **Role List** section, click **Add Role**.

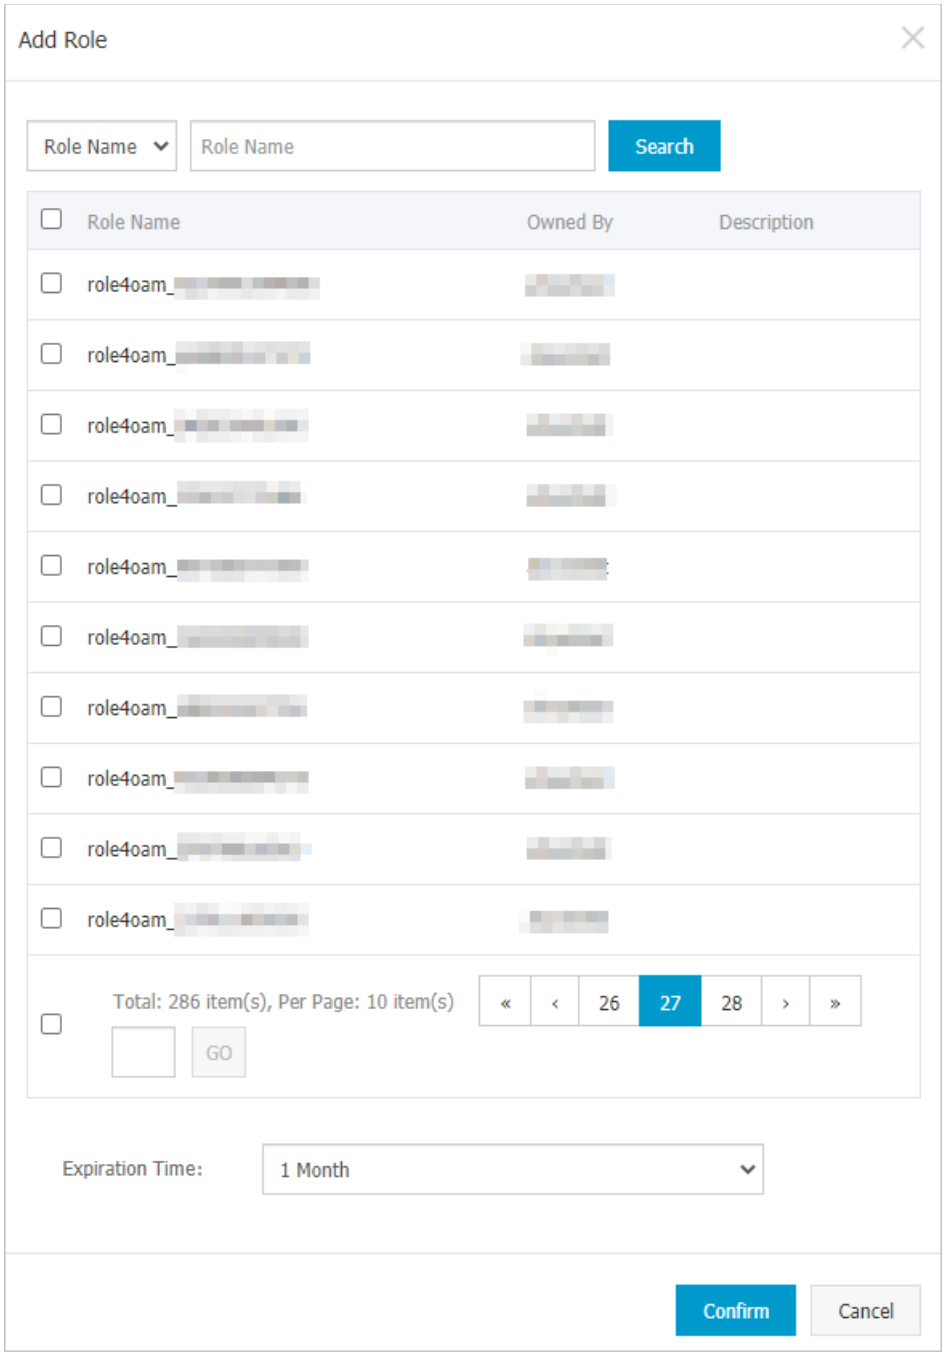

5. Search for roles by **Role Name**. Select one or more roles and set Expiration Time.
6. Click **Confirm**.

   To remove a role from a group, find the role in **Role List**, and click **Remove** in the **Actions** column.

12.1.3.5. Create roles

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. On the **Owned Roles** page, click **Create Role** in the upper-right corner.

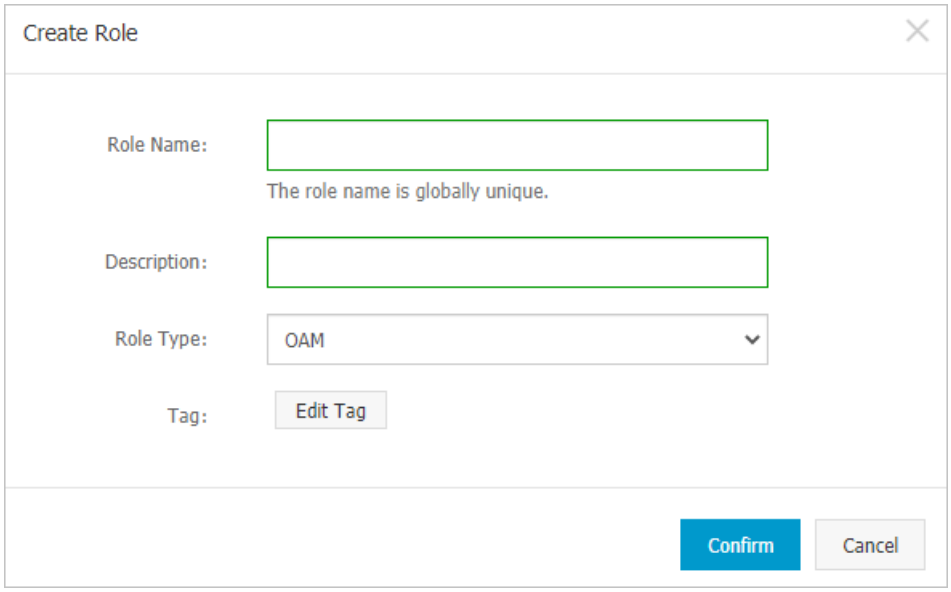

4. In the Create Role dialog box that appears, set **Role Name**, **Description**, and **Role Type**.
5. (Optional)Configure the role tags, which can be used to filter roles.
   i. Click **Edit Tags**.

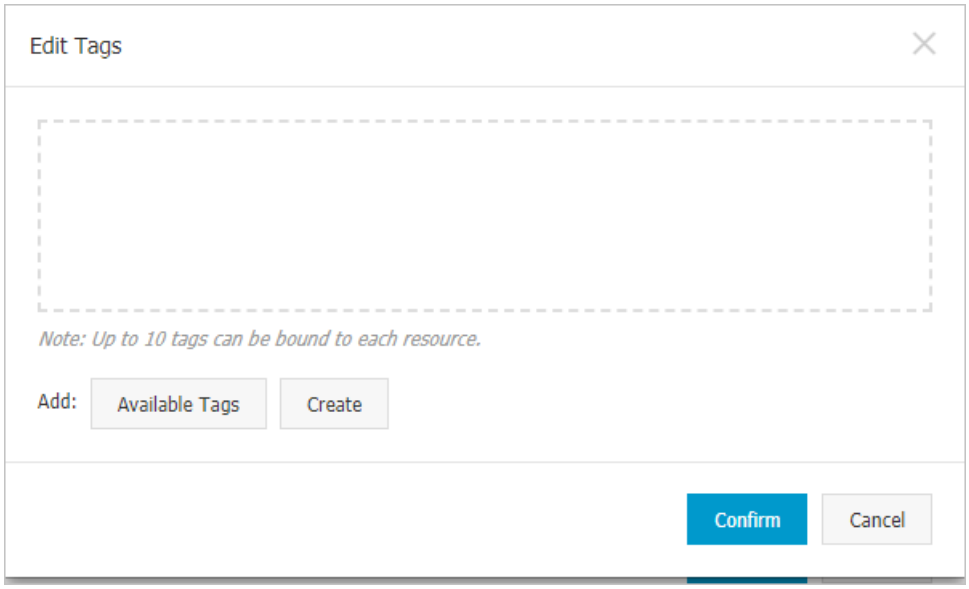

   ii. In the **Edit Tags** dialog box that appears, click **Create**.

iii. Set **Key** and **Value** for the tag and click **Confirm**.

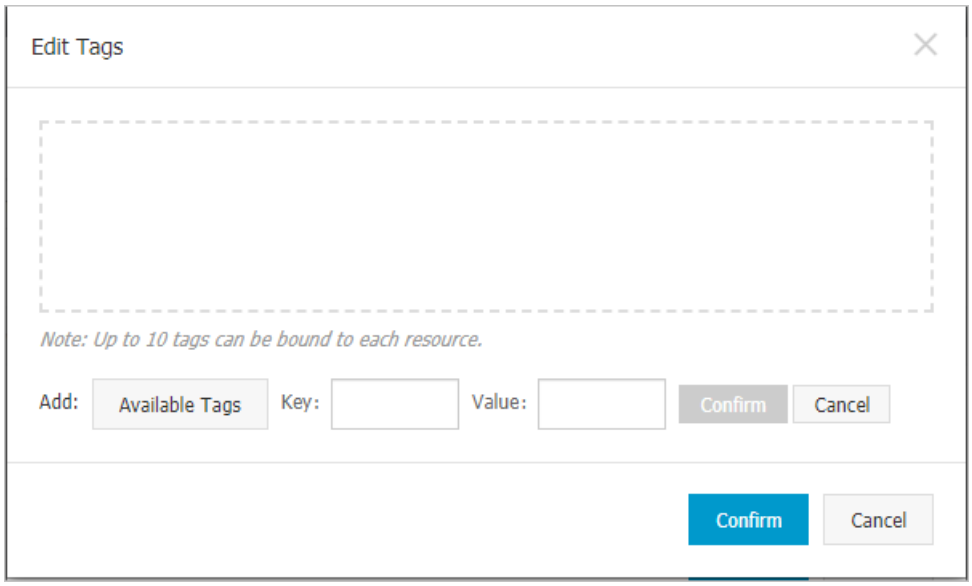

iv. Repeat the preceding step to create more tags.

The created tags are displayed inside the dotted box.

v. Click **Confirm** to create the tags and exit the **Edit Tags** dialog box.

6. Click **Confirm** to create the role.

12.1.3.6. Add inherited roles to a role

You can add inherited roles to a role to grant the permissions of the inherited roles to the role.

Prerequisites

You are the owner of the current role and the inherited role to be added.

For more information about how to query your owned roles, see Query roles.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add an inherited role and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Inherited Role** tab.
5. Click **Add Role**. In the **Add Role** dialog box that appears, search for roles by **Role Name**. Select one or more roles.

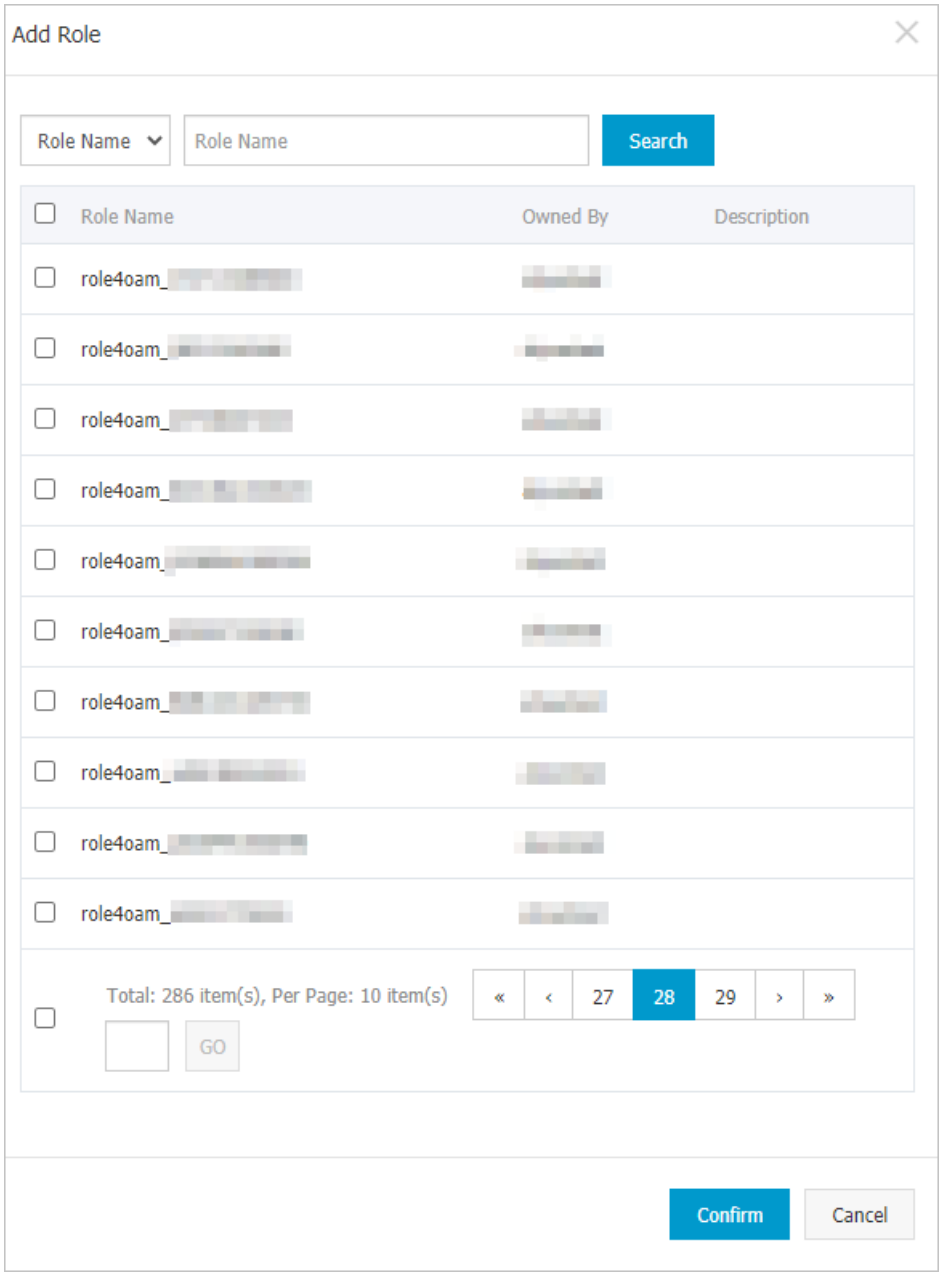

6. Click **Confirm**.

12.1.3.7. Add resources to a role

You must add resources to a created role.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add a resource and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Resource List** tab.
5. Click **Add Resource**.

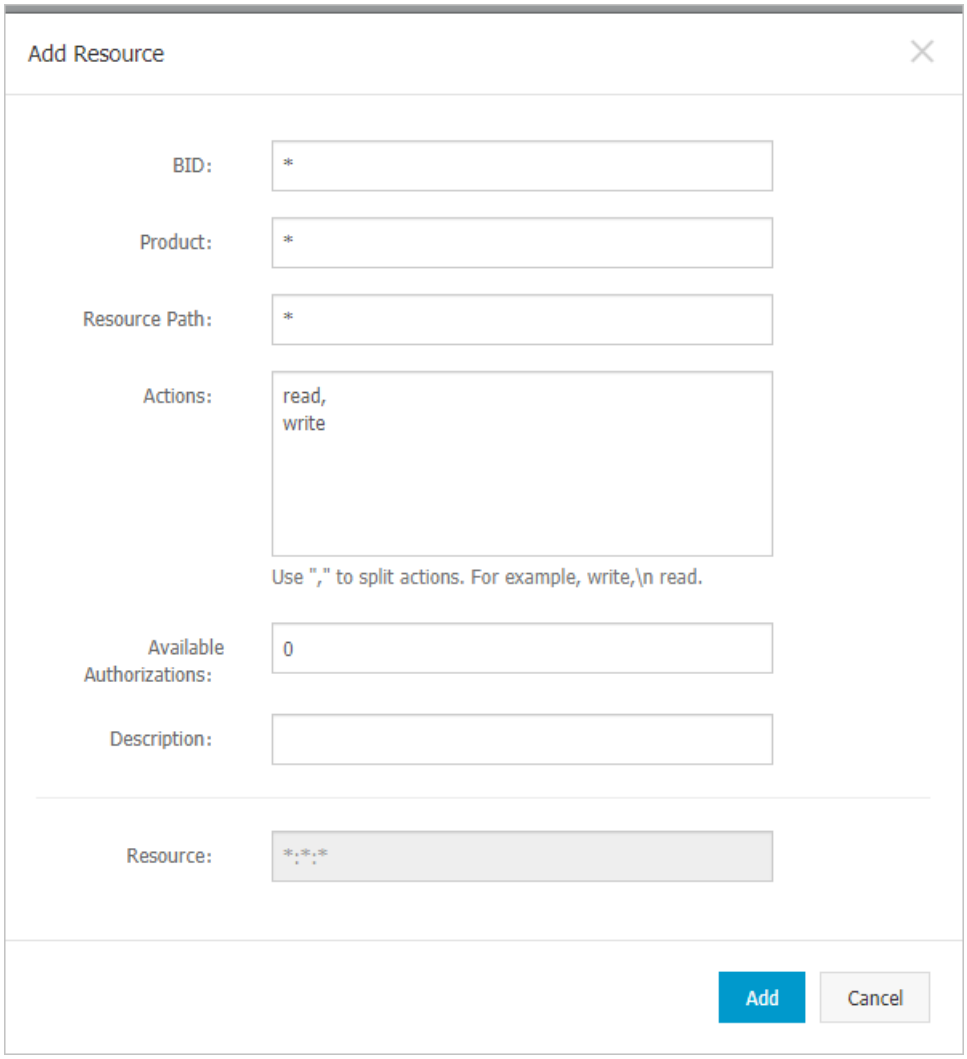

6. In the **Add Resource** dialog box, complete the configurations. For more information, see Parameters.

Parameters

Parameter	Description
BID	The deployment region ID.
Product	The cloud product to be added, such as rds. **Note** The cloud product name must be lowercase. For example, enter **rds** instead of **RDS**.
Resource Path	The resources of the cloud product. For more information about resources of the O&M platforms, see Permission lists of operations platforms.

Parameter	Description
Actions	An action set, which can contain multiple actions. For more information about actions on the O&M platforms, see Permission lists of operations platforms.
Available Authorizations	The maximum number of authorizations in cascaded authorization, which must be an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.
Description	The description of the resource.

7. Click **Add**.

12.1.3.8. Add authorized users to a role

You can assign an existing role to users or user groups.

Prerequisites

The corresponding users or user groups are created. Users are created in the Apsara Stack Cloud Management (ASCM) console. For more information about how to create a user group, see Create groups.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add an authorized user and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Authorized Users** tab.
5. Click **Add User** in the upper-right corner.

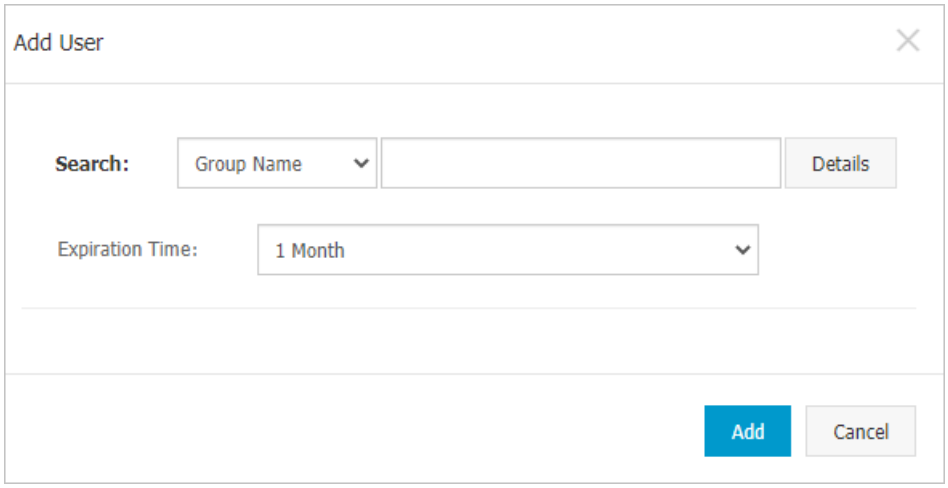

6. Select a search mode and enter the corresponding information.

   Four search modes are available:

- **RAM User Account**: Search in the format of *RAM user@Apsara Stack tenant account ID*.
- **Account Primary Key**: Search by using the unique ID of the Apsara Stack tenant account.
- **Logon Account Name**: Search by using the logon name of the Apsara Stack tenant account.
- **Group Name**: Search by group name.

> **Note** You can search for a single user or user group. For more information about how to create a user group, see Create groups.

7. Set Expiration Time.When the specified expiration time is due, the user no longer has the permissions of the role. To authorize the user again, the role creator must click **Renew** in the Actions column corresponding to the authorized user on the **Authorized Users** tab to modify the expiration time.
8. Click **Add** to assign the role to the user.To cancel the authorization, click **Remove** in the Actions column corresponding to the authorized user on the **Authorized Users** tab.

12.1.4. Manage groups

Group management allows you to view, modify, and delete groups.

12.1.4.1. Modify group information

After you create a group, you can modify the group name and description on the Group Information page.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.
4. On the **Group Information** page, click **Modify** in the upper-right corner.
5. In the **Modify Group** dialog box that appears, modify the group name and description.
6. Click **Confirm**.

12.1.4.2. View group role details

You can view information about the inherited roles, resource list, and inheritance tree of a group role.

Prerequisites

A role is added to the group.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.

4. In **Role List** section, click **Details** in the Actions column corresponding to a role.
5. On the **Role Information** page, perform the following operations:
   - Click the **Inherited Role** tab to view the information about the inherited roles of the role.

     To view the detailed information of an inherited role, click **Details** in the **Actions** column corresponding to the inherited role.
   - Click the **Resource List** tab to view the resource information of the role.

     For information about how to add other resources to this role, see Add resources to a role.
   - Click the **Inheritance Tree** tab to view the basic information and resource information of the role and its inherited roles by using the inheritance tree on the left.

12.1.4.3. Delete groups

You can delete groups that are no longer needed.

Prerequisites

The group to be deleted does not contain members.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group to be deleted and click **Delete** in the **Actions** column.

12.1.4.4. View authorized groups

You can view the groups to which you are added on the Authorized Groups page.

Context

You can view only the groups to which you belong, but cannot view groups of other users.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Authorized Groups**.
3. On the **Authorized Groups** page, view the name, owner, description, and modification time of the group to which you belong.

12.1.5. Manage roles

Role management allows you to view, modify, transfer, and delete roles.

12.1.5.1. Query roles

You can view your owned roles on the Owned Roles page.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Enter a role name in the Role Name field and click **Search** to search for roles that meet the search criteria.

> **Note** If the role you want to search for has a tag, you can click **Tag** and select the tag key to search for the role based on the tag.

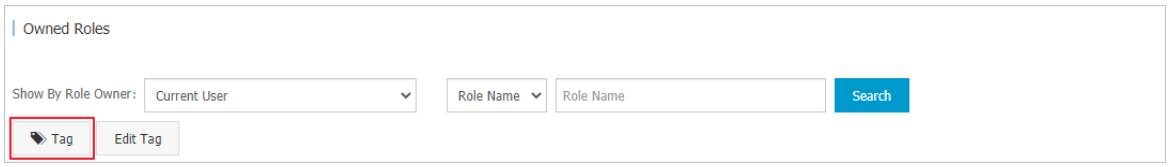

12.1.5.2. Modify role information

After you create a role, you can modify the role information.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role whose information you want to modify and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click **Modify** in the upper-right corner.
5. In the **Modify Role** dialog box that appears, set **Role Name**, **Description**, **Role Type** and **Tag**.
6. Click **Confirm**.

12.1.5.3. View the role inheritance tree

You can view the role inheritance tree to learn about the basic information and resource information of a role and its inherited roles.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role whose information you want to modify and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Inheritance Tree** tab.

   View the basic information and resource information of the role and its inherited roles by using the inheritance tree on the left.

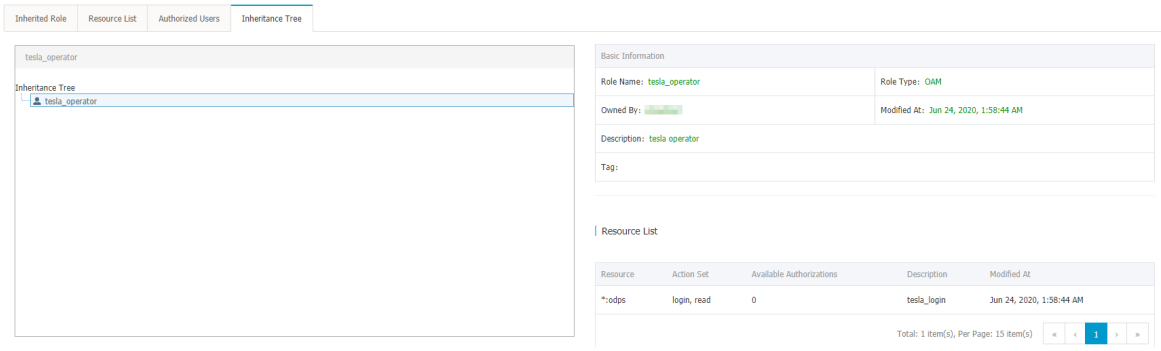

12.1.5.4. Transfer roles

You can transfer roles to other groups or users based on business requirements.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. On the **Owned Roles** page, configure the search condition and search for the roles to be transferred.
4. Select one or more roles in the search results and click **Transfer** in the lower-left corner.
5. In the **Transfer** dialog box that appears, select a search mode, enter the corresponding information, and then click **Details**. Details of the user or group are displayed.

   Four search modes are available:

   - **RAM User Account**: Search in the format of RAM user@Apsara Stack tenant account ID.
   - **Account Primary Key**: Search by using the unique ID of the Apsara Stack tenant account.
   - **Logon Account Name**: Search by using the logon name of the Apsara Stack tenant account.
   - **Group Name**: Search by group name.

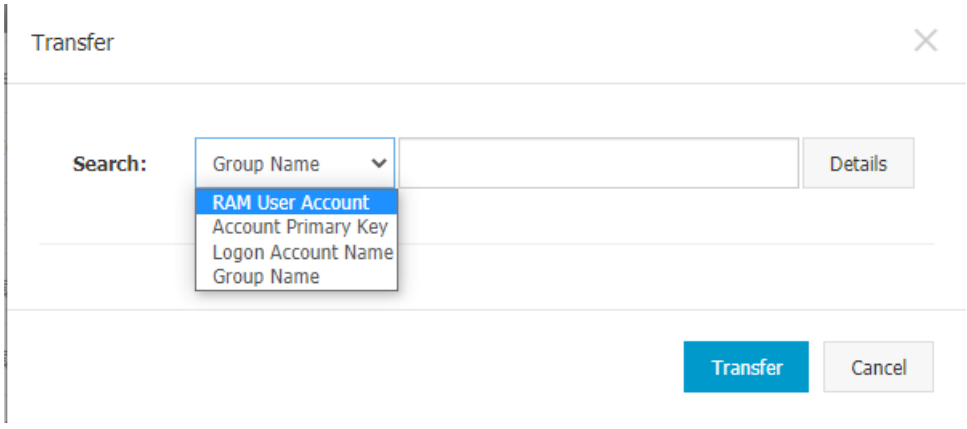

6. Click **Transfer**.

12.1.5.5. Delete a role

You can delete a role that is no longer in use according to business requirements.

Prerequisites

The role to be deleted does not contain inherited roles, resources, or authorized users.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. At the right of the role to be deleted and then click **Delete**.

12.1.5.6. View assigned roles

You can view the roles assigned to you and permissions granted to the roles.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Authorized Roles**.
3. On the **Authorized Roles** page, you can view the name, owner, description, modification time, and expiration time of each role assigned to you.You can also click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

12.1.5.7. View all roles

You can view all roles in Operation Administrator Manager (OAM) on the All Roles page.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > All Roles**.
3. On the **All Roles** page, view all the roles in the system.You can search for roles by **Role Name**.
4. Click **Details** in the Actions column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

12.1.6. Search for resources

You can search for resources to view the roles to which the resources are assigned.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, click **Search Resource**.
3. Set **Resource** and **Action**, and click **Search** to search for the roles that meet the specified conditions.

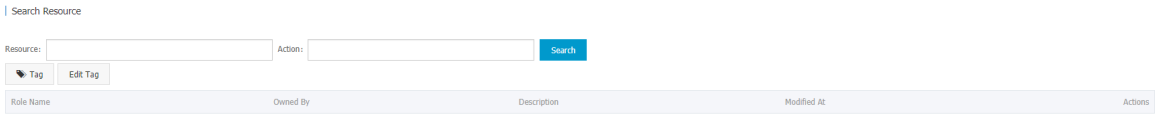

4. Click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources,

and inheritance tree of the role.

12.1.7. View personal information

You can view the personal information of the current user on the Personal Information page and test the user permissions.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, click **Personal Information**.
3. In the **Basic Information** section, you can view the username, type, creation time, AccessKey ID, and AccessKey secret of the current user.

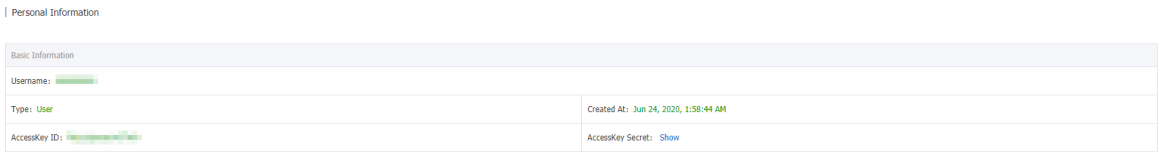

> **Note** You can click **Show** or **Hide** to show or hide the AccessKey secret.

4. In the **Test Permission** section, you can check whether the current user has a specific permission.
   i. Enter the resource information in the **Resource** field.

   > **Note** Use the English input method when you enter values in the **Resource** and **Action** fields.

   ii. Enter the permissions such as create, read, and write in the **Action** field. Separate multiple permissions with commas (,).

12.1.8. Default roles and permissions

12.1.8.1. Default roles and their functions

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

12.1.8.1.1. Default role of OAM

This topic describes the default role of Operation Access Manager (OAM) and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Super administrator	An administrator with root permissions	*:*	*	10

12.1.8.1.2. Default roles of Apsara Infrastructure Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations.

Apsara Infrastructure Management Framework is a distributed data center management system, used to manage applications on clusters containing multiple machines, and provides basic functions such as deployment, upgrade, expansion, contraction, and configuration change.

For more information about the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations, see the following table.

Role name	Role description	Resource	Actions	Available authorizations
Tianji_Project read-only	Has the read-only permission to Apsara Infrastructure Management Framework projects, which allows you to view the configurations and statuses of all projects and clusters	*:tianji:projects	["read"]	0
Tianji_Project administrator	Has all the permissions to Apsara Infrastructure Management Framework projects, which allows you to view and modify the configurations and statuses of all projects and clusters	*:tianji:projects	["*"]	0

Role name	Role description	Resource	Actions	Available authorizations
Tianji_Service read-only	Has the read-only permission to Apsara Infrastructure Management Framework services, which allows you to view the configurations and templates of all services	*:tianji:services	["read"]	0
Tianji_Service administrator	Has all the permissions to Apsara Infrastructure Management Framework services, which allows you to view and modify the configurations and templates of all services	*:tianji:services	["*"]	0
Tianji_IDC administrator	Has all the permissions to Apsara Infrastructure Management Framework data centers, which allows you to view and modify the data center information	*:tianji:idcs	["*"]	0
Tianji administrator	Has all the permissions to Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastructure Management Framework configurations	*:tianji	["*"]	0

12.1.8.1.3. Default roles of Webapp-rule

This topic describes the default roles of Webapp-rule and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Webapp-rule operations administrator	Has all the permissions to Webapp-rule projects, which allows you to view, modify, add, and delete all the configurations and statuses	26842:webapp-rule:*	["read", "write"]	0
Webapp-rule read-only	Has the read-only permission to Webapp-rule projects, which allows you to view all the configurations and statuses	26842:webapp-rule:*	["read"]	0

12.1.8.1.4. Default roles of the workflow console

This topic describes the default roles of the workflow console and the corresponding available authorizations.

The workflow console Grandcanal is an internally distributed process development framework. Developers can assemble, retry, roll back, and manually intervene the process based on this framework. Operations engineers can use the workflow console to manually intervene the corresponding processes.

For more information about the default roles of the workflow console and the corresponding available authorizations, see the following table.

Role name	Role description	Resource	Actions	Available authorizations
grandcanal.ADMIN	The workflow console administrator, who can query the workflow and activity details, and retry, roll back, terminate, and restart a workflow	26842:grandcanal	[“write” ,"read"]	0

Role name	Role description	Resource	Actions	Available authorizations
grandcanal.Reader	Has the read-only permission to the workflow console and can only perform the read operation	26842:grandcanal	["read"]	0

12.1.8.1.5. Default role of Tianjimon

This topic describes the default role of Tianjimon and the corresponding available authorizations.

Tianjimon, as the monitoring module of Apsara Infrastructure Management Framework, is used for the basic monitoring function of physical machines and services deployed based on Apsara Infrastructure Management Framework.

For more information about the default role of Tianjimon and the corresponding available authorizations, see the following table.

Role name	Role description	Resource	Actions	Available authorizations
Tianjimon operations	Has all Tianjimon permissions, which allows you to perform basic monitoring and operations	26842:tianjimon:*	["*"]	0

12.1.8.1.6. Default roles of Rtools

This topic describes the default roles of Rtools and the corresponding grant options.

Rtools is an assistant O&M system of Distributed Relational Database Service (DRDS). It is used to query metadata in the Diamond configuration management system.

The following table describes the default roles of Rtools and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
Rtools administrator	Has all permissions in the Rtools console.	26842:drds:rtools:*	*	0

12.1.8.1.7. Default roles of Opsapi

This topic describes the default roles of the Apsara Opsapi Management (Opsapi) system and the corresponding grant options.

Opsapi is a platform that manages O&M APIs and SDKs in the Apsara Stack environment in a centralized manner. This system also manages API and SDK versions.

The following table describes the default roles of Opsapi and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
Opsapi platform administrator	Has all the permissions on Opsapi.	*:opsapi:*	["read","write"]	0
Opsapi platform developer	Has the read permissions on Opsapi and the permissions to call API operations.	*:opsapi:*	["read","invoke"]	0
Opsapi common user	Has the read permissions on Opsapi.	*:opsapi:*	["read"]	0

12.1.8.1.8. Default roles of ASO

This topic describes the default roles of the Apsara Stack Operations (ASO) system and the corresponding grant options.

ASO is a centralized O&M management system that is developed for the Apsara Stack O&M personnel to perform centralized O&M operations.

The following table describes the default roles of ASO and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
ASO system administrator	Has the permissions to manage platform nodes, physical devices, and virtual resources, back up, restore, and migrate product data, and query and back up system logs.	*:aso:api-adapter:*	["read","write"]	0
		:aso:auth:	["read"]	0
		:aso:backup:	["read","write"]	0
		:aso:cmdb:	["read","write"]	0
		:aso:doc:	["read","write"]	0
		:aso:fullview:	["read","write"]	0
		:aso:init:	["read","write"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read","write"]	0
		:aso:lock:	["read","write"]	0
		:aso:physical:	["read","write"]	0

Role	Role description	Resource	Action	Grant option
		:aso:psm:	["read","write"]	0
		:aso:scm:	["read","write"]	0
		:aso:serviceWhitelist:	["read","write"]	0
		:aso:slalink:	["read","write"]	0
		:aso:task:	["read","write"]	0
ASO security officer	Has the permissions to manage permissions, security polices, and network security, and review and analyze security logs and activities of security audit officers.	*:aso:auth:*	["read","write"]	0
		:aso:plat-access:	["read","write"]	0
		:aso:twoFactorAuth:	["read","write"]	0
ASO security auditor	Has the permissions to audit, track, and analyze the activities of the system administrator and security officer.	*:aso:audit:*	["read","write"]	0
		:aso:auth:	["read"]	0
		:aso:serviceWhitelist:	["read"]	0
ASO product O&M officer	Has the permissions to perform O&M operations such as data import and export, modification, configuration, upgrade, and troubleshooting	*:aso:api-adapter:*	["read"]	0
		:aso:backup:	["read"]	0
		:aso:cmdb:	["read"]	0
		:aso:doc:	["read"]	0
		:aso:fullview:	["read","write"]	0
		:aso:init:	["read"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read"]	0
		:aso:lock:	["read"]	0
		:aso:physical:	["read","write"]	0

Role	coordination. Role description	Resource	Action	Grant option
		:aso:psm:	["read"]	0
		:aso:scm:	["read"]	0
		:aso:slalink:	["read"]	0
		:aso:task:	["read"]	0
ASO common O&M officer	Has the permissions to perform daily health checks, query service status, query inventory information, and query product usage.	*:aso:api-adapter:*	["read"]	0
		:aso:backup:	["read"]	0
		:aso:cmdb:	["read"]	0
		:aso:doc:	["read"]	0
		:aso:fullview:	["read"]	0
		:aso:init:	["read"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read"]	0
		:aso:lock:	["read"]	0
		:aso:physical:	["read","write"]	0
		:aso:psm:	["read"]	0
		:aso:scm:	["read"]	0
		:aso:slalink:	["read"]	0
		:aso:task:	["read"]	0
ASO duty observer	Has the permissions to view and monitor the dashboard, and monitor system alerts.	*:aso:doc:*	["read"]	0
		:aso:fullview:	["read"]	0

12.1.8.1.9. Default roles of PaaS

This topic describes the default roles of the Platform as a Service (PaaS) console and the corresponding grant options.

The PaaS console is an O&M platform designed for the PaaS platform and products, and is used to view, manage, and upgrade the products deployed on the PaaS platform.

The following table describes the default roles of the PaaS console and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
PaaS_Operation_Manager	Has all the permissions in the PaaS console.	*:paas-ops:*	["*"]	0

12.1.8.1.10. Default roles of OCP

This topic describes the default roles of the OceanBase Cloud Platform (OCP) and the corresponding grant options.

OCP is an enterprise-level database management platform with ApsaraDB for OceanBase as the core. It provides full lifecycle management for ApsaraDB for OceanBase components related to clusters, tenants, and databases, and manages ecosystem tools of ApsaraDB for OceanBase.

The following table describes the default roles of OCP and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
ocp_readonly	Has the read-only permissions on OCP.	*:oceanbase:role:ocp_readonly	["access"]	0
ob_dev	Has permissions on the performance and monitoring modules.	*:oceanbase:role:ob_dev	["access"]	0
ocp_dev	Has all permissions on OCP, but does not have the grant permission.	*:oceanbase:role:ocp_dev	["access"]	0

12.1.8.1.11. Default roles of Apsara Stack Security

This topic describes the default roles of Apsara Stack Security and the corresponding grant options.

Apsara Stack Security is a solution that provides Apsara Stack with a full suite of security features, such as network security, server security, application security, data security, and security management.

The following table describes the default roles of Apsara Stack Security and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
Apsara Stack Security administrator	Has all the permissions on Apsara Stack Security and can manage data in all the Apsara Stack Security modules.	*:yundun-luban:*	["*"]	0
Apsara Stack Security viewer	Has the read permissions on Apsara Stack Security and can read data in all the Apsara Stack Security modules.	*:yundun-luban:*	["read"]	0

12.1.8.1.12. Default roles of Apsara Network Intelligence

This topic describes the default roles of Apsara Network Intelligence and the corresponding grant options.

Apsara Network Intelligence is a system designed for network traffic analysis. It provides data to facilitate resource planning, diagnosis, monitoring, system management, and user behavior analysis.

The following table describes the default roles of Apsara Network Intelligence and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
Apsara Network Intelligence instance querier	Has permissions to query various instance resources.	• *:qitian:instance:* • *:qitian:user:*	["read","create","delete","update"]	0
Apsara Network Intelligence product O&M personnel	Has permissions to use the functions under the "Products" menu of Apsara Network Intelligence.	*:qitian:product:*	["read","create","delete","update"]	0
Apsara Network Intelligence R&D and O&M personnel	Has permissions to use the functions under the "System" menu of Apsara Network Intelligence.	*:qitian:system:*	["read","create","delete","update"]	0

12.1.8.2. Operation permissions on O&M platforms

This topic describes the operation permissions on O&M platforms.

12.1.8.2.1. Permissions on Apsara Infrastructure Management Framework

This topic describes the operation permissions on Apsara Infrastructure Management Framework.

Resource	Action	Description
*:tianji:services:[sname]:tjmontemplates:[tmplname]	delete	DeleteServiceTjmonTmpl
*:tianji:services:[sname]:tjmontemplates:[tmplname]	write	PutServiceTjmonTmpl
*:tianji:services:[sname]:templates:[tmplname]	write	PutServiceConfTmpl
*:tianji:services:[sname]:templates:[tmplname]	delete	DeleteServiceConfTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTjmonTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:tssessions	terminal	CreateTsSessionByService
*:tianji:services:[sname]:serviceinstances:[siname]:template	write	SetServiceInstanceTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:template	delete	DeleteServiceInstanceTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:template	read	GetServiceInstanceTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanceProductTagInService
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	write	AddServiceInstanceProductTagInService

Resource	Action	Description
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:resources	read	GetServerroleResourceInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines:[machine]	write	OperateSRMachineInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines:[machine]	read	GetMachineSRInfoInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines:[machine]	delete	DeleteSRMachineActionInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	read	GetMachinesSRInfoInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	delete	DeleteSRMachinesActionInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	write	OperateSRMachinesInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:apps:[app]:resources	read	GetAppResourceInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:apps: [app]:machines: [machine]:tianjilogs	read	TianjiLogsInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles	read	GetServiceInstanceServerrolesInService

Resource	Action	Description
*:tianji:services:[sname]:serviceinstances:[siname]:schema	write	SetServiceInstanceSchema
*:tianji:services:[sname]:serviceinstances:[siname]:schema	delete	DeleteServiceInstanceSchema
*:tianji:services:[sname]:serviceinstances:[siname]:rollings:[version]	write	OperateRollingJobInService
*:tianji:services:[sname]:serviceinstances:[siname]:rollings	read	ListRollingJobInService
*:tianji:services:[sname]:serviceinstances:[siname]:resources	read	GetInstanceResourceInService
*:tianji:services:[sname]:serviceinstances:[siname]:machines:[machine]	read	GetMachineAllSRInfoInService
*:tianji:services:[sname]:serviceinstances:[siname]	write	DeployServiceInstanceInService
*:tianji:services:[sname]:serviceinstances:[siname]	read	GetServiceInstanceConf
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:name	read	GetMachineAppFileListInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:download	read	GetMachineAppFileDownloadInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:content	read	GetMachineAppFileContentInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:filelist	read	GetMachineFileListInService

Resource	Action	Description
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:dockerlogs	read	DockerLogsInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:debuglog	read	GetMachineDebugLogInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps	read	GetMachineAppListInService
*:tianji:services:[sname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	DockerInspect
*:tianji:services:[sname]:schemas:[schemaname]	write	PutServiceSchema
*:tianji:services:[sname]:schemas:[schemaname]	delete	DeleteServiceSchema
*:tianji:services:[sname]:resources	read	GetResourceInService
*:tianji:services:[sname]	delete	DeleteService
*:tianji:services:[sname]	write	CreateService
*:tianji:projects:[pname]:machinebuckets:[bname]:machines:[machine]	read	GetMachineBucketMachineInfo
*:tianji:projects:[pname]:machinebuckets:[bname]:machines	read	GetMachineBucketMachines
*:tianji:projects:[pname]:machinebuckets:[bname]	write	CreateMachineBucket
*:tianji:projects:[pname]:machinebuckets:[bname]	write	OperateMachineBucketMachines
*:tianji:projects:[pname]:machinebuckets:[bname]	delete	DeleteMachineBucket

Resource	Action	Description
*:tianji:projects:[pname]:machinebuckets:[bname]	read	GetMachineBucketMachinesLegacy
*:tianji:projects:[pname]:machinebuckets	read	GetMachineBucketList
*:tianji:projects:[pname]:projects:[pname]:clusters:[cname]:tssessions:[tssessionname]:tsses	terminal	UpdateTsSessionTssByCluster
*:tianji:projects:[pname]:projects:[pname]:clusters:[cname]:tssessions	terminal	CreateTsSessionByCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTjmonTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	delete	DeleteServiceInstanceTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	write	SetServiceInstanceTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	read	GetServiceInstanceTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	write	AddServiceInstanceProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanceProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	GetServerroleResourceInCluster

Resource	Action	Description
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:files:name	read	GetMachineAppFileList
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:download	read	GetMachineAppFileDownload
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:content	read	GetMachineAppFileContent
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:filelist	read	GetMachineFileList
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:dockerlogs	read	DockerLogsInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:debuglog	read	GetMachineDebugLog
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps	read	GetMachineAppList

Resource	Action	Description
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines:[machine]	read	GetMachineSRInfoInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines:[machine]	write	OperateSRMachineInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines:[machine]	delete	DeleteSRMachineActionInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	write	OperateSRMachinesInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	delete	DeleteSRMachinesActionInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	read	GetAllMachineSRInfoInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:apps:[app]:resources	read	GetAppResourceInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:apps: [app]:machines: [machine]:tianjilogs	read	TianjiLogsInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	DockerInspectInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles	read	GetServiceInstanceServerrolesInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	delete	DeleteServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	write	SetServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:resources	read	GetInstanceResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	delete	DeleteServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	write	CreateServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	read	GetServiceInstanceConfInCluster
*:tianji:projects:[pname]:clusters:[cname]:rollings:[version]	write	OperateRollingJob
*:tianji:projects:[pname]:clusters:[cname]:rollings	read	ListRollingJob
*:tianji:projects:[pname]:clusters:[cname]:resources	read	GetResourceInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:quota	write	SetClusterQuotas
*:tianji:projects:[pname]:clusters:[cname]:machinesinfo	read	GetClusterMachineInfo
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	read	GetMachineAllSRInfo
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	write	SetMachineAction
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	delete	DeleteMachineAction
*:tianji:projects:[pname]:clusters:[cname]:machines	write	OperateClusterMachines
*:tianji:projects:[pname]:clusters:[cname]:difflist	read	GetVersionDiffList
*:tianji:projects:[pname]:clusters:[cname]:diff	read	GetVersionDiff
*:tianji:projects:[pname]:clusters:[cname]:deploylogs:[version]	read	GetDeployLogInCluster
*:tianji:projects:[pname]:clusters:[cname]:deploylogs	read	GetDeployLogListInCluster
*:tianji:projects:[pname]:clusters:[cname]:builds:[version]	read	GetBuildJob
*:tianji:projects:[pname]:clusters:[cname]:builds	read	ListBuildJob
*:tianji:projects:[pname]:clusters:[cname]	write	OperateCluster
*:tianji:projects:[pname]:clusters:[cname]	delete	DeleteCluster
*:tianji:projects:[pname]:clusters:[cname]	read	GetClusterConf

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]	write	DeployCluster
*:tianji:projects:[pname]	write	CreateProject
*:tianji:projects:[pname]	delete	DeleteProject
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	CreateRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	SetRackunitAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	delete	DeleteRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	SetRackAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	CreateRack
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	delete	DeleteRack
*:tianji:idcs:[idc]:rooms:[room]	write	CreateRoom
*:tianji:idcs:[idc]:rooms:[room]	delete	DeleteRoom
*:tianji:idcs:[idc]:rooms:[room]	write	SetRoomAttr
*:tianji:idcs:[idc]	delete	DeleteIdc
*:tianji:idcs:[idc]	write	SetIdcAttr
*:tianji:idcs:[idc]	write	CreateIdc

12.1.8.2.2. Permission list of Webapp-rule

This topic describes the permissions of Webapp-rule.

Resource	Action	Description
26842:webapp-rule:*	write	Adds, deletes, and updates configuration resources
26842:webapp-rule:*	read	Queries configuration resources

12.1.8.2.3. Permission list of the workflow console

This topic describes the permissions of the workflow console.

Resource	Action	Description
26842:grandcanal	read	Queries the workflow activity details and summary
26842:grandcanal	write	Restarts, retries, rolls back, and terminates a workflow

12.1.8.2.4. Permissions on Monitoring System of Apsara Infrastructure Management Framework

This topic describes the operation permissions on Monitoring System of Apsara Infrastructure Management Framework.

Resource	Action	Description
26842:tianjimon:monitor-manage	manage	Monitoring and O&M

12.1.8.2.5. Permissions on Rtools

This topic describes the operation permissions on Rtools.

Resource	Action	Description
26842:drds:rtools:tddl	all	Publishes Taobao Distributed Data Layer (TDDL) configurations in the Rtools console.
26842:drds:rtools:jade	all	Queries and modifies configurations in the Rtools console.
26842:drds:rtools:gemini	all	Performs operations on gemini in the Rtools console.
26842:drds:rtools:system	all	Performs other operations in the Rtools console.