ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Enterprise

Operations and Maintenance Guide

Product Version: 2109, Internal: V3.15.0

Document Version: 20211210

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.Apsara Uni-manager Operations Console Operations

1.1. User Guide

1.1.1. Overview

This topic describes the management framework of the Apsara Uni-manager Operations Console.

Management framework of the Apsara Uni-manager Operations Console

Alibaba Cloud Apsara Stack adopts the ISO 20000 standard and references the methods regulated by the Information Technology Service Standards (ITSS) and ITIL frameworks to build the management framework of the Apsara Uni-manager Operations Console. The following figure shows the management framework of the Apsara Uni-manager Operations Console.

Apsara Uni-manager Operations Console

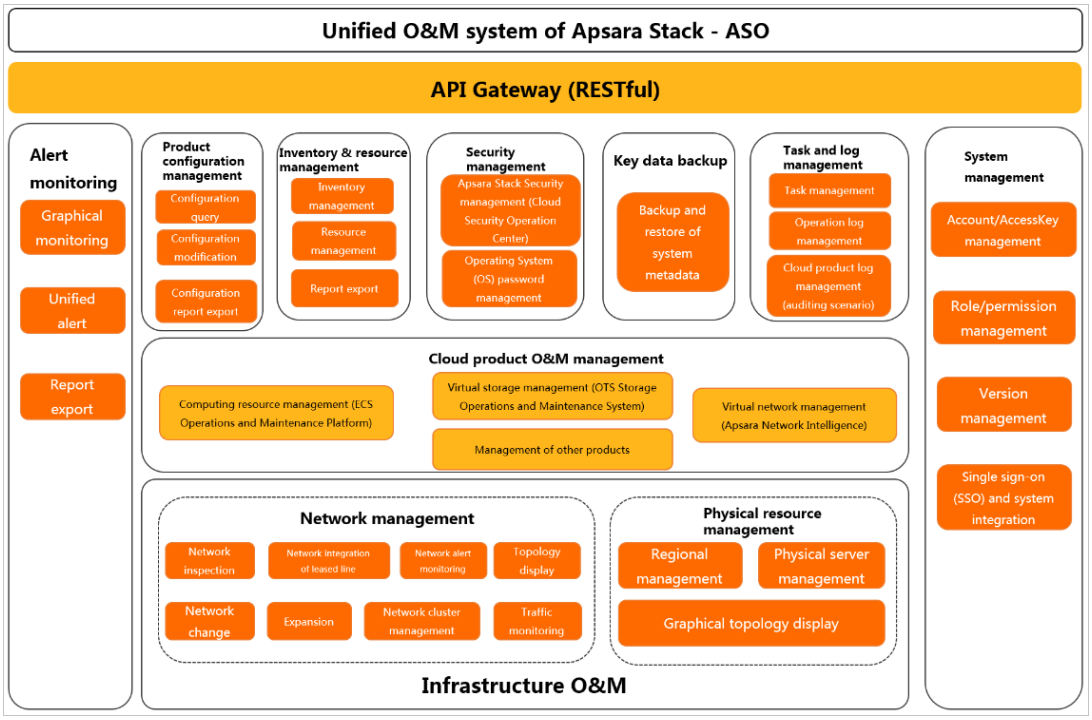

Based on ITIL and ISO 20000, the management framework of the Apsara Uni-manager Operations Console uses management support tools to adapt to various management modes in a process-oriented, normalized, and standardized manner. This has implemented the systematic management of the overall process of operations services The management framework of the Apsara Uni-manager Operations Console provides the full lifecycle management methods, management standards, management modes, management supporting tools, management objects, and process-based management methods of IT operations services.

The Apsara Uni-manager Operations Console defines various entities involved in operations activities and relationships between these entities. Relevant entities are well organized and coordinated based on the Apsara Uni-manager Operations Console and can provide different levels of operations services based on the service agreements.

Apsara Uni-manager Operations Console

The Apsara Uni-manager Operations Console is a unified and intelligent O&M platform. In accordance with the Information Technology Infrastructure Library (ITIL) and IT Service Management (ITSM) standards, the operations processes and requirements must be abstract, and automation is implemented by using intelligent operations tools. For customized operations, interfaces and multi-level approval must be used to reduce risks.

In the Apsara Uni-manager Operations Console, cloud operations is classified into the following layers: infrastructure, cloud service, and business operations.

Based on the operations experience and data accumulated and collected from three layers, Alibaba Cloud Apsara Stack aggregates data collected by the operations platform to the Configuration Management Database (CMDB) of the platform. The Apsara Uni-manager Operations Console consolidates, analyzes, and comprehensively processes the data and integrates rich practical experience and operations capabilities to the platform operations tools. The Apsara Uni-manager Operations Console is designed to be desired state-oriented and uses unified operations tools for the fault discovery and tracking, link display, ITIL process, and self-repaired faults of the platform to realize the ultimate goal of artificial intelligence for IT operations (AIOps).

The Apsara Uni-manager Operations Console provides a centralized operations portal that allows you to have a consistent operations experience. The Apsara Uni-manager Operations Console supports interconnections with third-party platforms and provides centralized API operations capabilities to deliver data to third-party systems by using APIs.

The Apsara Uni-manager Operations Console performs centralized operations management, such as automated deployments, upgrades, changes, and configurations, on physical devices, operating systems, computing resources, network, storage, databases, middleware, and business applications in the cloud computing environment. The Apsara Uni-manager Operations Console also provides the features 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, the Apsara Uni-manager Operations Console guarantees the continuous and stable running of cloud computing business applications and provides services and support for O&M processes to build an improved operations service management platform.

O&M support services

In addition to tools, process assurance and personnel management are essential to ensure the integrity of operations. 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 provides upgrade channels to support urgent problems of customers. As an autonomous and controllable platform, the Apsara Uni-manager Operations Console ensures that technical problems can be effectively solved in a timely manner.

1.1.2. Get started

1.1.2.1. Prepare an operations account

Before you perform O&M operations in the Apsara Uni-manager Operations Console, make sure that you have obtained an operations account that have corresponding permissions from an administrator.

Procedure

1. Log on to the Apsara Uni-manager Operations Console as a system administrator.
2. Create a role. For more information, see Role management.
3. Create an operations account and assign the created 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 **OAM**, grant permissions to the role, and then grant the role to the corresponding operations account as an administrator.

1.1.2.2. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager Operations Console.

Prerequisites

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

  The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id*.ops.console.*intranet-domain-id*.
- A browser is available. We recommend that you use Google Chrome.

Procedure

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

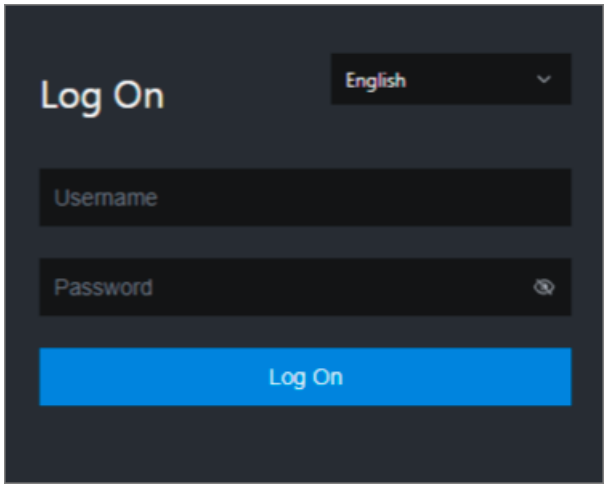

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains the following special characters: ! @ # $ %
- The password must be 10 to 20 characters in length.

4. Click **Log On**.

1.1.2.3. Apsara Uni-manager Operations Console homepage

This topic describes the basic operations on and features of the Apsara Uni-manager Operations Console.

The following table describes the sections on the homepage of the console.

No.	Section	Description

No.	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	View the alert knowledge base and upload other relevant HTML documents.
⑤	Current user	Show the name of the current logon user.
⑥	Top navigation bar	Select an O&M operation.
⑦	Language	Move the pointer over this section and select a language.
⑧	Current user information	Move the pointer over 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 console.
⑨	Operation	View information and perform operations.

1.1.2.4. Instructions for the homepage

The homepage allows you to view the statistics and summary data of Apsara Stack alerts, physical devices, and cloud service inventory.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Homepage**.
3. View the homepage.

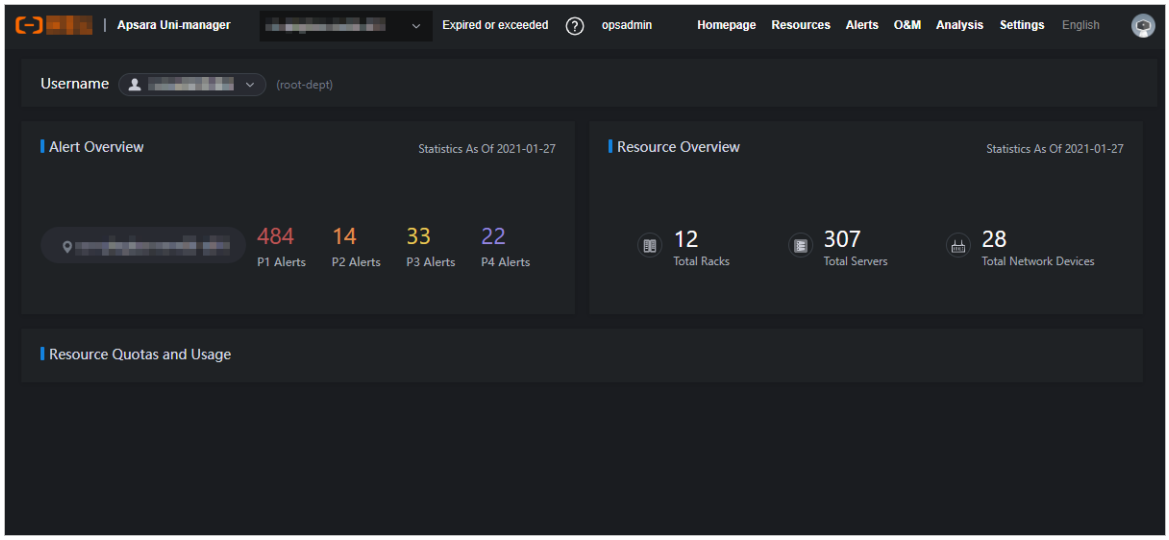

The homepage consists of the **opsadmin**, **Alert Overview**, **Resource Overview**, and **Resource Quotas and Usage** sections.

- In the **opsadmin** section, select a user. The department to which the user belongs is displayed on the right.
- In the **Alert Overview** section, view the total number of alerts at the P1, P2, P3, and P4 levels.
- In the **Resource Overview** section, view the total number of racks, servers, and network devices.
- In the **Resource Quotas and Usage** section, view the resource quotas and usage related to cloud services.

  Cloud service-related metrics are displayed in the following dimensions: total, used and available resources, and their usage.

Alibaba Cloud service	Statistical metric
ECS	CPU (Core)
	EBS (GB)
	Memory (GB)
RDS	CPU (Core)
	Disk (GB)
	Memory (GB)
SLB	Internal IP Address
	Public IP Address
OSS	Storage Capacity (GB)
DFS	Memory (GB)
	SLS-INNER

Alibaba Cloud service	Statistical metric
SLS	SLS-PUBLIC
OTS	Memory (GB)
NAS	Memory (GB)

1.1.2.5. View the version of the Apsara Uni-manager Operations Console

You can view the version and build information of the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click the icon. A dialog box displays the version and build information.

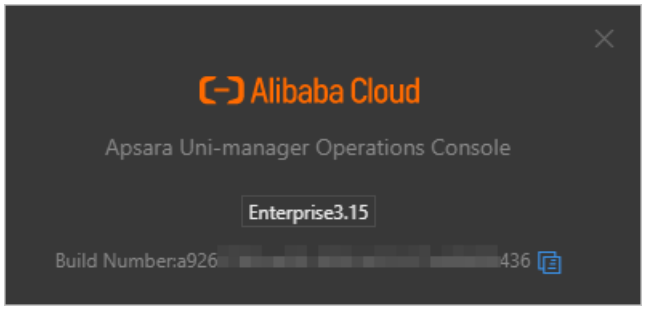

> **Note** Click the icon in the dialog box to copy the build information.

1.1.3. Settings

1.1.3.1. Default operations roles

This topic describes the default roles for the Apsara Uni-manager Operations Console and their responsibilities.

For quick reference, the following roles are preset in the Apsara Uni-manager Operations Console: 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. By default, this role is not displayed in the role list.

Role	Responsibility
System administrator	Manages platform nodes, physical devices, and virtual resources, backs up, restores, and migrates product data, as well as 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.

1.1.3.2. Security policies

1.1.3.2.1. Logon policies

As an administrator, you can configure logon policies to control what times and IP addresses are valid for logon.

Context

The system provides a default policy. You can configure logon policies based on your needs to better control the read and write permissions of users and improve the system security.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Security Policies > Logon Policy**.
4. On the **Logon Policies** page, perform the following operations:
   - Query policies

     In the upper-left corner of the page, enter a policy name in the **Policy Name** search box and click **Search** to view the policy information in the list. You can also click **Reset** to clear the previous search conditions.
   - Add a policy

     Click **Add Policy**. In the Add Policy dialog box, set the policy name, start time, end time, and allowed or prohibited logon IP addresses. Click **OK**.

> **Note** On the **Logon Settings** page:
>
> - If you select **Blacklist** for the **Logon Policies** field, the **Prohibited Logon IP Addresses** field is displayed in the Add Policy dialog box.
> - If you select **Whitelist** for the **Logon Policies** field, the **Allowed Logon IP Addresses** field is displayed in the Add Policy dialog box.
>
> You can specify the value accordingly. For more information, see Modify logon settings.

- Modify a policy

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

- Delete a policy

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

> **Notice** A logon policy that is bound to a user cannot be deleted. You must unbind the policy before you can delete it.

1.1.3.2.2. Physical server passwords

The Server Password module allows you to configure and manage all physical servers in the Apsara Stack environment, and query their information.

Context

The Server Password module provides the following features:

- The system automatically collects information of all the servers in the Apsara Stack environment.
- The server password is periodically updated.
- You can configure the expiration period and length of passwords.
- You can manually update the passwords 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Security > Server Password**.

   The **Manage Passwords** tab shows the passwords of all servers in the current Apsara Stack environment.
4. Perform the following operations:
   - Query servers

     On the **Manage Passwords** tab, select a product or hostname, or enter an IP address, and then click **Search**. You can also click **Reset** to clear the previous search conditions.

- Query a password
  a. On the **Manage Passwords** tab, find the server whose password you want to query.
  b. Click the icon in the **Password** column. The server password in plaintext is displayed and is converted into cipher text after 10 seconds. Alternatively, click the icon to show the password in cipher text.
- Update a password
  a. On the **Manage Passwords** tab, find the server for which you want to update the password.
  b. Click **Update Password** in the **Actions** column.
  c. In the dialog box that appears, specify **Password** and **Confirm Password** and click **OK**.
  The password of the server is updated.
- Batch update passwords
  a. On the **Manage Password** tab, select multiple servers.
  b. Click **Batch Update** in the upper part of the tab.
  c. In the dialog box that appears, enter **New Password** and **Confirm Password** and click **OK**.
  The passwords of the selected servers are updated.
- Configure the password expiration period
  a. On the **Manage Passwords** tab, select one or more servers.
  b. Click **Configure** in the upper part of the tab.
  c. In the **Configuration Items** dialog box, specify **Password Expiration Period** and **Unit** and click **OK**.
  Server passwords are immediately updated and are updated again after the expiration period.
- Query the update history of server passwords
  Click the **History Password** tab. Select a product, hostname, or IP address, and then click Search to view the update history of server passwords in the search results.
- Query historical passwords of a server
  a. On the **History Password** tab, find the server whose historical passwords you want to query.
  b. Click the icon in the **Password** column. The server password in plaintext is displayed and is converted into cipher text after 10 seconds. Alternatively, click the icon to show the password in cipher text.
- Query and modify the password configuration policy

a. Click the **Configuration** tab and view the metadata of server password management, including the initial password, password length, and retry times. The following parameters in the Input Parameters section must be specified:
   - **Initial Password** indicates the password assigned when server password management is deployed in the Apsara Stack environment. This parameter is required to modify the password of a server in the Apsara Stack environment.
   - **Password Length** indicates the length of passwords updated by the system.
   - **Retry Times** indicates a limit of how many times a password can fail to be updated before the system stops trying.
   - **Status** indicates whether the configuration takes effect. By default, the switch is turned off. To show the status, turn on  .

b. Click **Save**.

1.1.3.3. Offline backup

Offline Backup is used to back up the key metadata of Apsara Stack. Only the metadata of Apsara Distributed File System can be backed up. The backup metadata information is used for quick recovery from Apsara Stack failures.

1.1.3.3.1. Add a backup product

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Product Management**.
4. Click **Add Product**.
5. In the **Add Product** dialog box, add information of a product as described in the following table and click **OK**.

Parameter	Description	Example
Product	The product name. You can select a service from the drop-down list.	pangu
Backup Item	The name of the backup item. Set this parameter based on the information of the cloud service to be backed up.	ecs_pangu
Script	The name of the backup script.	metadata_backup.py
Retry Times	The number of retry attempts after an error occurs. Typically, set this parameter to 3.	3

6. To add more product backup items, repeatedly perform the preceding steps.

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

Result

You can view the added product on the **Backup Configurations** page.

1.1.3.3.2. Configure backup

After you add a product backup item, you must configure the backup in the Apsara Uni-manager Operations 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 services, such as ecs pangu, ots pangu, oss pangu, and ads pangu.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Backup Configurations**.

   In the left side of the Backup Configurations page, the current backup configurations is displayed in a hierarchical tree-like structure. The root node is a product list and shows the products whose data can be backed up in the current backup system. Only the metadata of Apsara Distributed File System and OPS DNS can be backed up.
4. Click a product backup item on the left and then configure the parameters on the right.

   > **Note** If the #FTPMaster server is deployed in the cluster of the service, the system will automatically enter the backup information of ecs pangu.

Parameter	Description
Product Cluster IP Address	The IP address of the actual transfer server.
Backup File Folder	A folder on the transfer server. You are required only to enter a folder path in the field without manually creating a folder to store backup files. Examples: ○ pangu: /apsarapangu/disk8/pangu_master_bak/*product name*_pangu/bak ○ opsdns: /apsarapangu/disk8/opsdns_bak/opsdns/bak

Parameter	Description
Script Execution Folder	A folder on the transfer server. You are required only to enter a folder path in the field without manually creating a folder to store script executions. Examples: ○ pangu: /apsarapangu/disk8/pangu_master_bak/*product name*_pangu/bin ○ opsdns: /apsarapangu/disk8/opsdns_bak/opsdns/bak
Script Parameters	Required. The execution parameters for the script. You must enter the value in the `--ip=xxx.xxx.xxx.xxx` format. ○ pangu: The IP address is one of the IP addresses of the pangu master. ○ opsdns: We recommend that you enter `--ip=127.0.0.1` .
Backup Schedule	The execution period of recurring executions. In this example, a value of 1 is entered to specify that the backup is performed only once.
Backup Schedule Unit	The unit of the execution period. Valid values: **Day**, **Hour**, and **Minute**. In this example, **Hour** is selected to specify that the backup is performed by hour.
Timeout Period	The timeout period, in seconds. In this example, set the value to **3600**.

5. Click **Modify** to complete the configurations and trigger the backup.
6. Perform the preceding steps to configure all the product backup items.

1.1.3.3.3. View the backup details

During the backup, you can view the backup details of each backup item on the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Backup Details**.
4. On the **Backup Details** page, enter the service and backup item, select the start date, and then click **Search**.
5. View the backup details of a backup item, including the product, backup item, the name of the file to be backed up, start time, and status.

   The backup status includes **Not started**, **In transmission**, **Complete**, **Time-out**, and **Failed**.

6. (Optional)You can also click **Reset** to clear the previous search conditions.

1.1.3.3.4. Configure a backup server

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Backup Server Configurations**.

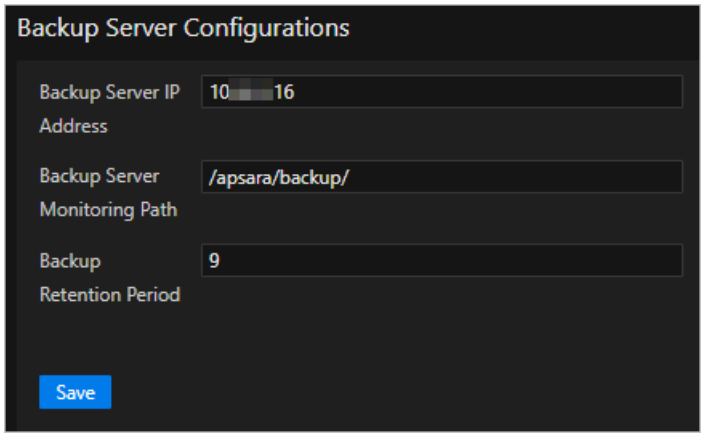

4. Configure the backup server parameters.

> **Note** If the #FTPMaster server is deployed in the cluster of the service, the system will automatically enter the backup server information of ecs pangu.

The following table describes 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 network of the backup server is connected to other servers in Apsara Stack. ○ Apsara Distributed File System cannot be deployed on the server or on the disk where backup metadata is stored.

Parameter	Description
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 Period	The maximum time period in days that backup files are stored. Backup files whose retention periods exceed the specified time period are deleted.

5. Click **Save**.

1.1.3.3.5. Use cases

To ensure the availability of services, you can back up the data of different services stored on Apsara Distributed File System.

1.1.3.3.5.1. Preparations

This topic describes preparations to make before you perform backup operations.

Before the backup, take note of the following items:

- A buffer server is required as the backup server.

  If no buffer servers are available, select a physical server that has a large disk capacity and good network performance. Otherwise, the security of the backup data cannot be ensured.

  > **Note** Offline backup files cannot be stored on objects to be backed up. If no more physical servers are available and if disk capacity is insufficient in the on-site environment, the system is unable to perform offline backup. In this case, you must add physical servers or increase disk capacity before the offline backup.

- A transfer server is required to store one-time backup data and backup scripts of each product.

  No other requirements are needed for transfer servers.

- The network of the backup server must be connected with that of the Docker container where the offline backup service is located. This ensures that backup containers in the clusters of the Apsara Uni-manager Operations Console can log on to the transfer server and backup server by using SSH key pairs, without the need to provide the username and password.

1.1.3.3.5.2. Collect the Apsara Distributed File System information of each product

The Products module allows you to collect the Apsara Distributed File System information of products to be backed up, which helps add the backup product information to the Apsara Uni-manager Operations Console.

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.
   i. Log on to the Apsara Uni-manager Operations Console.
   ii. In the top navigation bar, click **O&M**.
   iii. In the left-side navigation pane, choose **Product Management > Products**.
   iv. In the **Apsara Stack O&M** section, choose **Basic O&M > Apsara Infrastructure Management Framework**.

   > **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 **Service** search box to search for the Apsara Distributed File System service.
4. Click **Operations** in the **Actions** column corresponding to pangu to go to the service details page.
5. Click the **Clusters** tab.
6. Click the name of a cluster.
7. On the **Services** tab, click **pangu.PanguMaster#**.

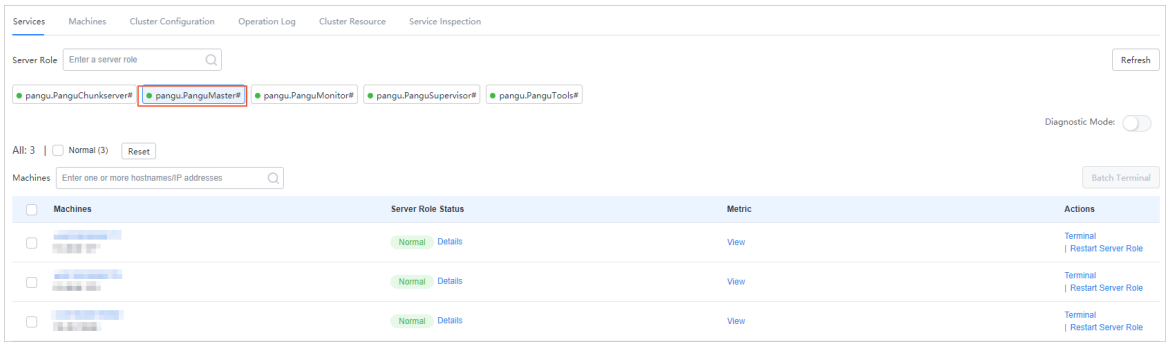

8. View and record the IP addresses of Apsara Distributed File System master in the server list.

   Record one of the three IP addresses of **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
ECS-IO7-A-xx	10.10.10.2	ecs
ads-A-xx	10.10.10.3	ads

Cluster name	pangumaster IP	Product name
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.

1.1.3.3.5.3. Configure a backup server

The Backup Server Configurations module allows you to configure parameters related to a backup server.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Backup Server Configurations**.

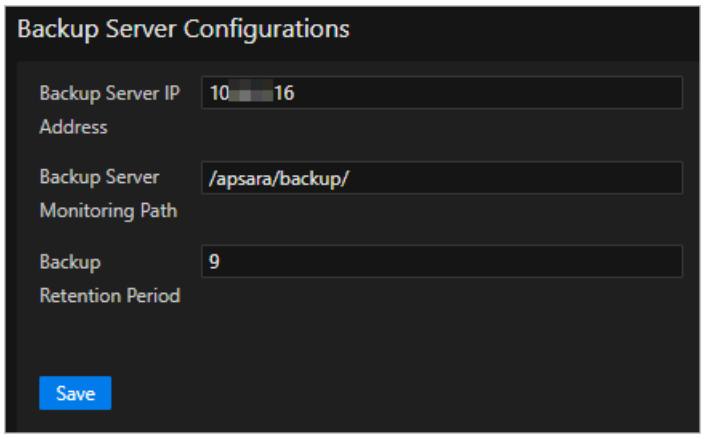

4. Configure the backup server parameters.

Note If the #FTPMaster server is deployed in the cluster of the service, the system will automatically enter the backup server information of ecs pangu.

The following table describes the parameters.

Parameter	Description

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 network of the backup server is connected to 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 Period	The maximum time period in days that backup files are stored. Backup files whose retention periods exceed the specified time period are deleted.

5. Click **Save**.

1.1.3.3.5.4. Add a backup product

The Product Management module allows you to add backup product information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Product Management**.
4. Click **Add Product**.
5. In the **Add Product** dialog box, add information of a product as described in the following table and click **OK**.

Parameter	Description	Example
Product	The product name. You can select a service from the drop-down list.	pangu

Parameter	Description	Example
Backup Item	The name of the backup item. Set this parameter based on the product information described in the Collect the Apsara Distributed File System information of each product topic.	ecs_pangu
Script	The name of the backup script.	metadata_backup.py
Retry Times	The number of retry attempts after an error occurs. Typically, set this parameter to 3.	3

6. Repeat the preceding steps to add all the backup items.

Result

You can view the added product on the **Backup Configurations** page.

1.1.3.3.5.5. Configure backup parameters

After you add backup items, you must configure the backup in the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Backup Configurations**.
4. Click a product backup item on the left and then configure the parameters on the right.

Parameter	Description
Product Cluster IP Address	The IP address of the actual transfer server.
Backup File Folder	A folder on the transfer server. You are required only to enter a folder path in the field without manually creating a folder to store backup files. Example: /apsarapangu/disk8/pangu_master_bak/*product name*_pangu/bak
Script Execution Folder	A folder on the transfer server. You are required only to enter a folder path in the field without manually creating a folder to store backup files. Example: /apsarapangu/disk8/pangu_master_bak/*product name*_pangu/bin

Parameter	Description
Script Parameters	The execution parameters for the script. You must enter the value in the `--ip=xxx.xxx.xxx.xxx` format, in which the IP address is one of the IP addresses of the pangu master described in the Collect the Apsara Distributed File System information of each product topic.
Backup Schedule	The execution period of recurring executions. In this example, a value of 1 is entered to specify that the backup is performed only once.
Backup Schedule Unit	The unit of the execution period. Valid values: **Day**, **Hour**, and **Minute**. In this example, **Hour** is selected to specify that the backup is performed by hour.
Timeout Period	The timeout period. Unit: seconds. In this example, set the value to **3600**.

5. Click **Modify** to complete the configurations and trigger the backup.
6. Perform the preceding steps to configure all the product backup items.

1.1.3.3.5.6. View the backup details

After you configure backup items, you can check whether the backup items function normally in the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Offline Backup > Backup Details**.
4. 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 functions normally.

   > **Note** When a backup task is complete, you must check whether the MD5 values of the offline backup service and the backup server are consistent with each other. If yes, the backup was successful.

1.1.3.4. Log cleanup

The Log Cleanup module allows you to clean up logs from specified log files in the specified containers (Docker) or physical machines (virtual machines or bare metal machines) in the system.

1.1.3.4.1. Import the log cleanup rules of containers or physical servers

If you have configured log cleanup rules on your computer, you can batch import multiple cleanup rules of containers or physical servers.

Context

Before you import a cleanup rule, take note of the following items:

- Imported rules are incrementally added.
- You must check the values of Product, Service, ServerRole, SrcPath, MatchFile, Threshold, and Method to determine whether a cleanup rule already exists. If all values in the environment are the same as the values specified in the rule to be imported, the rule already exists. If a rule already exists, it is not imported.
- Before you import a rule, you must contact technical support to obtain the encryption sequence.
- After you have imported a rule, special characters such as spaces, carriage returns, line feeds, and tabs in the rule are automatically deleted.
- The maximum disk usage range specified by a rule is [0%,100%], and the value before the percent sign (%) must be an integer. Otherwise, the rule is automatically filtered out when you import it. We recommend that you set the maximum disk usage to 75%.
- Make sure that the cleanup methods specified by rules are tested and can be normally executed. Otherwise, exceptions may occur when you use these methods to clean up logs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Container** or **Physical Machines** tab.
5. Click **Import**.
6. Select the XLS or XLSX files that you want to import and click **Open**. You can import multiple log cleanup rules.

   After you import rules, corresponding execution plans are asynchronously generated.

1.1.3.4.2. Export the log cleanup rules of containers or physical servers

You can batch export multiple log cleanup rules of containers or physical servers.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Containers** or **Servers** tab.
5. Perform the following operations to export the log cleanup rules of containers or physical servers:
   - Click **Export** to export all cleanup rules.
   - In the upper part of the page, select a product, service, and server role, and click **Search**. In the

search result, select the cleanup rules that you want to export and click **Export**.

> **Note** By default, the **Product**, **Service**, and **Service Role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

1.1.3.4.3. Modify a log cleanup rule

You can modify log cleanup rules to suit your business needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Containers** or **Physical Machines** tab.
5. (Optional)In the upper part of the tab, select the product, service, and server role, and click **Search** to query cleanup rules that meet the filter conditions.

   > **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

6. Find the cleanup rule that you want to modify and click **Modify** in the **Actions** column.
7. In the panel that appears, modify the maximum disk usage and specify whether to automatically clean up logs that match the cleanup rule.

   > **Note** The maximum disk usage range specified by a rule is [0%,100%], and the value before the percent sign (%) must be an integer. We recommend that you set the maximum disk utilization to 75%.

8. Click **OK**.

1.1.3.4.4. Delete a log cleanup rule

You can delete log cleanup rules that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Containers** or **Physical Machines** tab.
5. (Optional)In the upper part of the tab, select the product, service, and server role, and click **Search**

to query cleanup rules that meet the filter conditions.

> **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

6. Find the cleanup rule and click **Delete** in the **Actions** column.
7. In the message that appears, click **OK**.

> **Note** The execution plan corresponding to a cleanup rule is not deleted when you delete the rule. At 02:00 every day, the system cleans up existing execution plans and generates new execution plans based on the current cleanup rules.

1.1.3.4.5. Obtain the usage information of containers or physical servers

This topic describes how to query the disk usage information of containers or physical servers.

Method 1

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Plans**.
4. Click the **Containers** or **Servers** tab.
5. Perform the following operations to obtain the disk usage information of a container or physical server:
   - In the upper part of the page, select a product, service, and server role, and click **Search**. In the search results, find the container or physical server for which you want to query disk usage information. Click **Query Usage** in the **Actions** column to obtain the disk usage information of the container or physical server.

     > **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

   - Select multiple containers or physical servers and click **Batch Query Usage** to obtain the disk usage information of multiple containers or physical servers.

     > **Note** The operation used to obtain the usage information is asynchronous. You must refresh the page to view the results. If the current usage of the disk is higher than the specified maximum disk usage, the value is displayed in red.

Solution 2

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Containers** or **Physical Machines** tab.
5. (Optional)In the upper part of the tab, select a product, service, and server role, and click **Search**.

   > **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

6. In the search results, select the cleanup rule of the container or physical server in which you want to obtain the disk usage information and click **Execution Plans** in the **Actions** column. The **Execution Plans** page appears.
7. Perform the following operations to obtain the disk usage information of a container or physical server:
   - Find the container or physical server for which you want to query disk usage information. Click **Query Usage** in the **Actions** column to obtain the disk usage information of the container or physical server.
   - Select multiple containers or physical servers and click **Batch Query Usage** to obtain the disk usage information of multiple containers or physical servers.

     > **Note** The operation used to obtain the usage information is asynchronous. You must refresh the page to view the results. If the current usage of the disk is higher than the specified maximum disk usage, the value is displayed in red.

1.1.3.4.6. Clean up the logs of containers or physical servers

You can clean up the logs of containers or physical servers in a timely manner based on disk usage information of the containers or physical servers.

Method 1

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Plans**.
4. Click the **Containers** or **Physical Machines** tab.
5. Perform the following operations to clean up logs of containers or physical servers:
   - In the upper part of the tab, select a product, service, and server role, and click **Search**. In the search results, find the container or physical server for which you want to clean up logs and click **Execute Clearance** in the **Actions** column.

> **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

- Select multiple containers or physical servers and click **Batch Clear** in the upper part of the tab to clean up the log information of multiple containers or physical servers at a time.

> **Note** The log cleanup operation is asynchronous, and you must view the log cleanup results on the **Records** page.

Method 2

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Containers** or **Physical Machines** tab.
5. (Optional)In the upper part of the tab, select a product, service, and server role, and click **Search**.

> **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

6. Select the cleanup rule of the container or physical server in which you want to obtain the disk usage information and click **Execution Plans** in the **Actions** column. The **Execution Plans** page appears.
7. Perform the following operations to clean up logs of containers or physical servers:
   - Find the container or physical server for which you want to clean up logs and click **Execute Clearance** in the **Actions** column to clean up the logs of a single container or physical server.
   - Select multiple containers or physical servers and click **Batch Clear** in the upper part of the tab to clean up the logs of multiple containers or physical servers at a time.

> **Note** The log cleanup operation is asynchronous, and you must view the log cleanup results on the **Records** page.

1.1.3.4.7. Configure automatic cleanups for container or physical server logs

You can configure automatic cleanups for container or physical server logs that meet the specified cleanup rules.

Context

At 02:00 every day, the system cleans up existing execution plans and generates new execution plans based on the current cleanup rules. If you turn on **Automatic Deletion** or enable automatic cleanup, the system cleans up the container or physical server logs that meet the cleanup rules based on execution plans at 02:30 every day.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Rules**.
4. Click the **Containers** or **Physical Machines** tab.
5. Perform the following operations to configure automatic cleanups for container or physical server logs that meet the specified cleanup rules:
   - In the upper part of the page, select a product, service, and server role, and click **Search**. In the search results, find the cleanup rule for which you want to set automatic cleanups and turn on **Automatic Deletion**. The system cleans up the container or physical server logs that meet the cleanup rule.

     > **Note** By default, the **Product**, **Service**, and **Service role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

     If you want to disable the automatic cleanup feature, you can turn off **Automatic Deletion**.
   - Select multiple cleanup rules and click **Enable Automatic Clearance**. The system cleans up the container or physical server logs that meet the selected cleanup rules.

     To disable the automatic cleanup feature, click **Disable Automatic Clearance**.

1.1.3.4.8. View cleanup records

After you clean up logs, you can view detailed cleanup records.

Context

When you perform operations on the **Records** page, take note of the following items:

- Each time you perform a log cleanup operation, the numbers of cleanup executions, SR, and machines are increased by one.
- Number of cleanup log files shows the number of log files that match all the available rules and can be cleaned up, rather than the number of log files that have been cleaned up.
- Clean up space shows the accumulated available space after you clean up logs.

Method 1

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Records**.
4. (Optional)In the upper part of the tab, select a product, service, and server role, and click **Search**.

> **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

5. Find the cleanup record that you want to view and click **View Details** in the **Details** column to view the detailed cleanup information.

Method 2

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Log Cleanup > Plans**.
4. Click the **Containers** or **Physical Machines** tab.
5. (Optional)In the upper part of the tab, select a product, service, and server role, and click **Search**.

> **Note** By default, the **product**, **service**, and **server role** parameters on the tab do not have options in their drop-down lists. When you specify these parameters for the first time, you must enter the product, service, and server role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

6. Find the cleanup record that you want to view and click **Cleanup Records** in the **Actions** column. The **Records** page appears.
7. Find the cleanup record that you want to view and click **View Details** in the **Details** column to view the detailed cleanup information.

1.1.3.5. System configurations

The System Settings module allows you to manage departments, roles, and users involved in the Apsara Uni-manager Operations Console in a centralized manner. This makes it easy to grant different resource access permissions to different users. The System Settings module is a core module in managing permissions. It integrates the features such as department management, role management, logon policy management, user management, and password management.

1.1.3.5.1. 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. Log on to the Apsara Uni-manager Operations Console.

2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Management**.

   The **Users** tab appears.
4. On the **Users** tab, perform the following operations:
   - Query users

     > **Note** To search for users in the Apsara Uni-manager Operations Console, you must have the security officer role or system administrator role.

     In the upper part of the page, set **Username**, **Role**, and **Department**, and click **Search** to view the information about the user in the list.

     (Optional) Click **Reset** to clear the filter conditions.
   - Add a user

     > **Note** To add a user in the Apsara Uni-manager Operations Console, you must have the security officer role.

     Click **Add** in the upper part of the tab. In the **Add User** dialog box, set **Username** and **Password** and click **OK**.

     The added user is displayed in the user list. The value in the **Primary Key Value** column corresponding to the added user is used to call API operations of applications. When you want to call applications in the Apsara Uni-manager Operations Console for other applications, you must use the primary key value for authentication.
   - Modify a user

     > **Note** To modify a user in the Apsara Uni-manager Operations Console, you must have the security officer role.

     In the user list, find the user that you want to modify and click **Modify** in the **Actions** column. In the **Modify User** dialog box, modify the parameters and click **OK**.
   - Delete a user

     In the user list, find the user that you want to modify and click **Delete** in the **Actions** column. In the message that appears, click **OK**.

     > **Note** Deleted users are displayed on the **Recycle Bin** tab. To restore a deleted user, click the **Recycle Bin** tab. Find the user that you want to restore and click **Recover** in the **Actions** column. In the message that appears, click **OK**.
   - Attach a logon policy

     In the user list, find the user to which you want to bind a logon policy and click **Bind Logon Policy** in the upper part of the page. In the **Bind Logon Policy** dialog box, select the logon policy to attach and click **OK**.
   - Query the personal information of the current user

Move the pointer over the profile picture in the upper-right corner of the page and click **Personal Information**. On the **User Profile** page, view the personal information of the current user, such as the **Username** and **Department**.

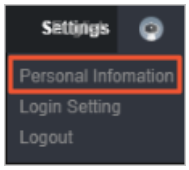

On the **User Profile** page, you can also change the password that the current user uses to log on to the Apsara Uni-manager Operations Console. For more information about how to change the logon password, see Change the logon password.

- Configure logon settings

  Move the pointer over the profile picture in the upper-right corner of the page and click **Login Setting**. On the **Logon Settings** page, you can modify the logon timeout period, maximum allowed password retries, logon policies, and validity period of the current account, and specify whether to allow multi-terminal logon. For more information about how to modify logon settings, see Modify logon settings.

1.1.3.5.2. 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Group Management**.
4. In the upper part of the page, click **Add**.
5. In the **Add User Group** dialog box, enter a user group name, select a department, and then click **OK**.

Modify the name of a user group

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Group Management**.
4. (Optional)Select a department name, enter a user group name and username, and then click **Search**.

   You can also click **Advanced**, select a department name and role name, enter a user group name and username, and then click **Search**. If you have specified filter conditions, you can click **Reset** to remove the conditions.
5. In the user group list, find the user group that you want to modify and click **Edit User Group** in the **Actions** column.
6. In the dialog box that appears, modify the user group name.
7. Click **OK**.

Manage users in a user group

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Group Management**.
4. (Optional)Select a department name, enter a user group name and username, and then click **Search**.

   You can also click **Advanced**, select a department name and role name, enter a user group name and username, and then click **Search**. If you have specified filter conditions, you can click **Reset** to remove the conditions.
5. In the user group list, find the user group for which you want to manage users and click **Manage Users** in the **Actions** column.
6. In the dialog box that appears, you can add or delete users in the user group.
   - Click **Add**. In the **Add** dialog box, select one or more users and click **OK**.
   - Click the icon to delete the user.
7. Click **OK**.

   Added users are displayed in the **Users** column corresponding to the user group.

   Deleted users are no longer displayed in the **Users** column corresponding to the user group.

Add a role to a user group

You can add only a single role to a user group.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Group Management**.
4. (Optional)Select a department name, enter a user group name and username, and then click **Search**.

   You can also click **Advanced**, select a department name and role name, enter a user group name and username, and then click **Search**. If you have specified filter conditions, you can click **Reset** to remove the conditions.
5. In the user group list, find the user group to which you want to add a role and click **Add Role** in the **Actions** column.
6. Select a role from the **Role** drop-down list.
7. Click **OK**.

   The 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Group Management**.
4. (Optional)Select a department name, enter a user group name and username, and then click

Search.

You can also click **Advanced**, select a department name and role name, enter a user group name and username, and then click **Search**. If you have specified filter conditions, you can click **Reset** to remove the conditions.

5. In the user group list, find the user group from which you want to delete a role and click **Delete Role** in the **Actions** column.
6. In the message that appears, click **OK**.

   The deleted role is no longer displayed in the **Role** column corresponding to the user group. The permissions of the role are rescinded from all users in the group.

Delete a user group

Notice Before you delete a user group, make sure that no users or roles are bound to the user group.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > User Group Management**.
4. (Optional)Select a department name, enter a user group name and username, and then click **Search**.

   You can also click **Advanced**, select a department name and role name, enter a user group name and username, and then click **Search**. If you have specified filter conditions, you can click **Reset** to remove the conditions.
5. In the user group list, find the user group that you want to delete and click **Delete User Group** in the **Actions** column.
6. In the message that appears, click **OK**.

1.1.3.5.3. Manage roles

You can customize roles in the Apsara Uni-manager Operations Console to implement 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 Open Application Model (OAM) system. You cannot modify or delete these roles. Custom roles can be modified and deleted.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Role Management**.
4. On the **Role Management** page, perform the following operations:
   - Query roles

> **Note** To query roles in the Apsara Uni-manager Operations Console, you must have the security officer role or system administrator role.

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

- Add a role

  > **Note** Only users that have security officer roles of the Apsara Uni-manager Operations Console can add a role in the console.

  Click **Add** in the upper part of the tab. In the **Add Role** dialog box, set **Role Name**, **Role Description**, and **Base Role**, and click **OK**.

- Modify a role

  > **Note** Only users that have security officer roles of the Apsara Uni-manager Operations Console can modify a role in the console.

  Find the role that you want to modify in the role list and click **Edit** in the **Actions** column. In the **Edit Role** dialog box, modify **Role Name** and **Role Description**, select a basic role, and then set menu permissions. Click **OK**.

- Delete a role

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

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

1.1.3.5.4. Menu management

The Menu Settings module allows you to add, hide, modify, or delete a menu based on your business needs.

1.1.3.5.4.1. Add a level-1 menu

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Menu Settings**.
4. In the upper part of the page, click **Add Menu Data**.
5. In the Add Menu panel, configure the parameters for the menu.

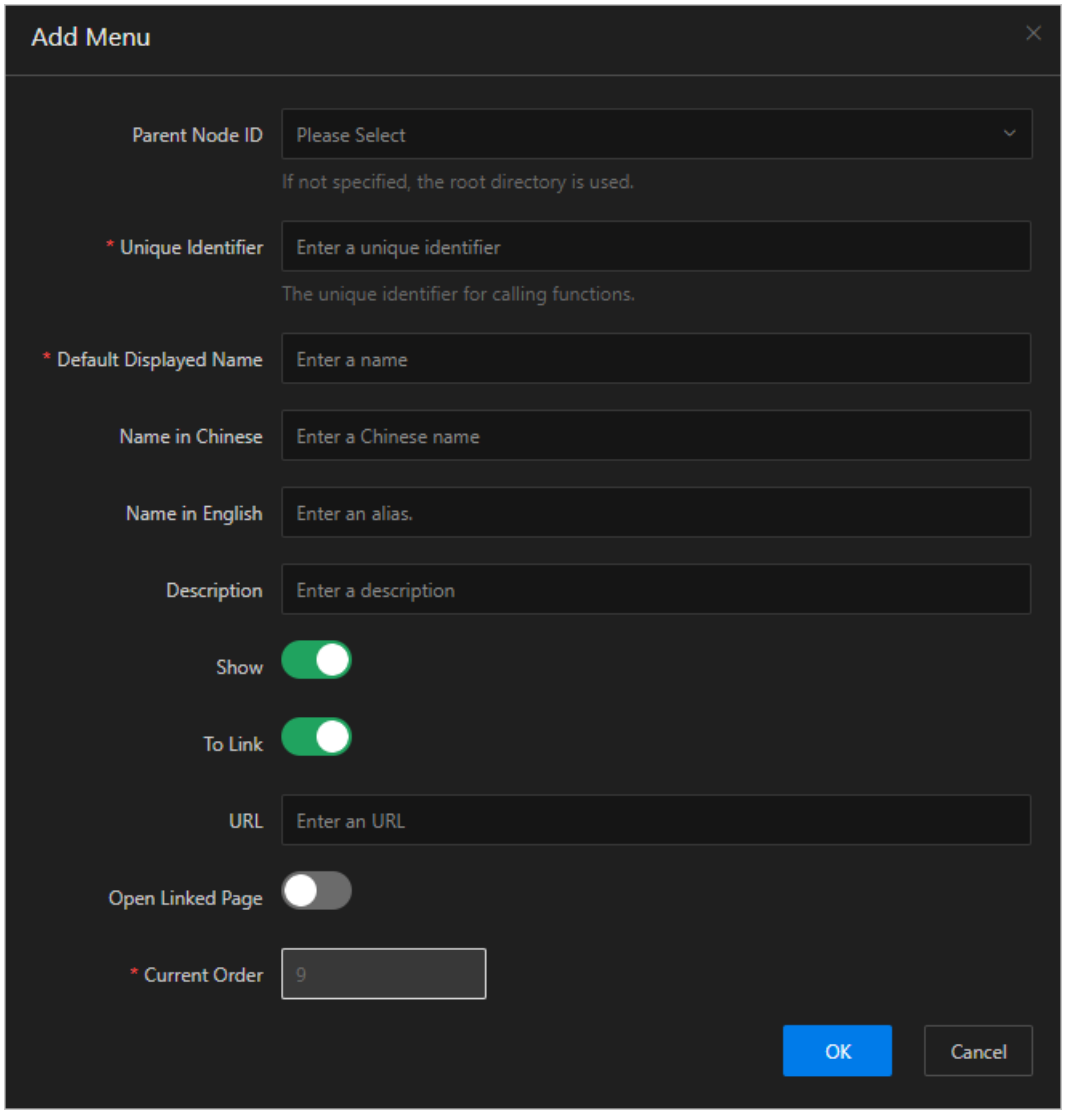

The following table describes the parameters.

Parameter	Description
Parent Node ID	The parent menu. This parameter does not need to specified when you add a level-1 menu.
Unique Identifier	The unique identifier used to call functions. It can consist only of letters and can be 5 to 20 characters in length.
Default Display Name	The default display name of the submenu.
Name in Chinese	The submenu name in Chinese. In the Chinese language environment, if the Chinese name of the submenu is specified, the default display name of the submenu is the specified Chinese name.
Name in English	The submenu name in English. In the English language environment, if the English name of the submenu is specified, the default display name of the submenu is the specified English name.

Parameter	Description
Description	The description of the submenu.
Show	Specifies whether to show the submenu after it is added. You can turn on or off **Show**. By default, Show is turned on.
To Link	Specifies whether to go to another page when you click the submenu. You can turn on or off **To Link**. By default, To Link is turned off.
URL	This parameter appears only when **To Link** is turned on. Set this parameter to the URL to go to when you click the submenu. ○ If the URL of a page within the current system is used, enter the absolute path or a relative path of the page. Example: /aso/aso-alarm/dashboard. ○ If the URL of a third-party system is used, enter the absolute path of the page. Example: http://example.com/TaskManageTool/#/taskView.
Open Linked Page	Specifies whether to open a new page for the URL to go to after you click the submenu. You can turn on or off **Open Linked Page**. By default, the switch is off.
Current Order	The order of the menu among all level-1 menus. You cannot configure the order in the panel. You can modify the configuration on the **Menu Settings** page after you create the menu.

6. Click **OK**.

Result

After you have added a level-1 menu, you can view the menu in the menu list and the top navigation bar.

1.1.3.5.4.2. Add a submenu

This topic describes how to add a submenu.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Menu Settings**.
4. Add a submenu.
   i. Find the menu to which you want to add a submenu and click **Add Submenu** in the **Actions** column.
   ii. In the Add Menu panel, configure the parameters for the submenu.

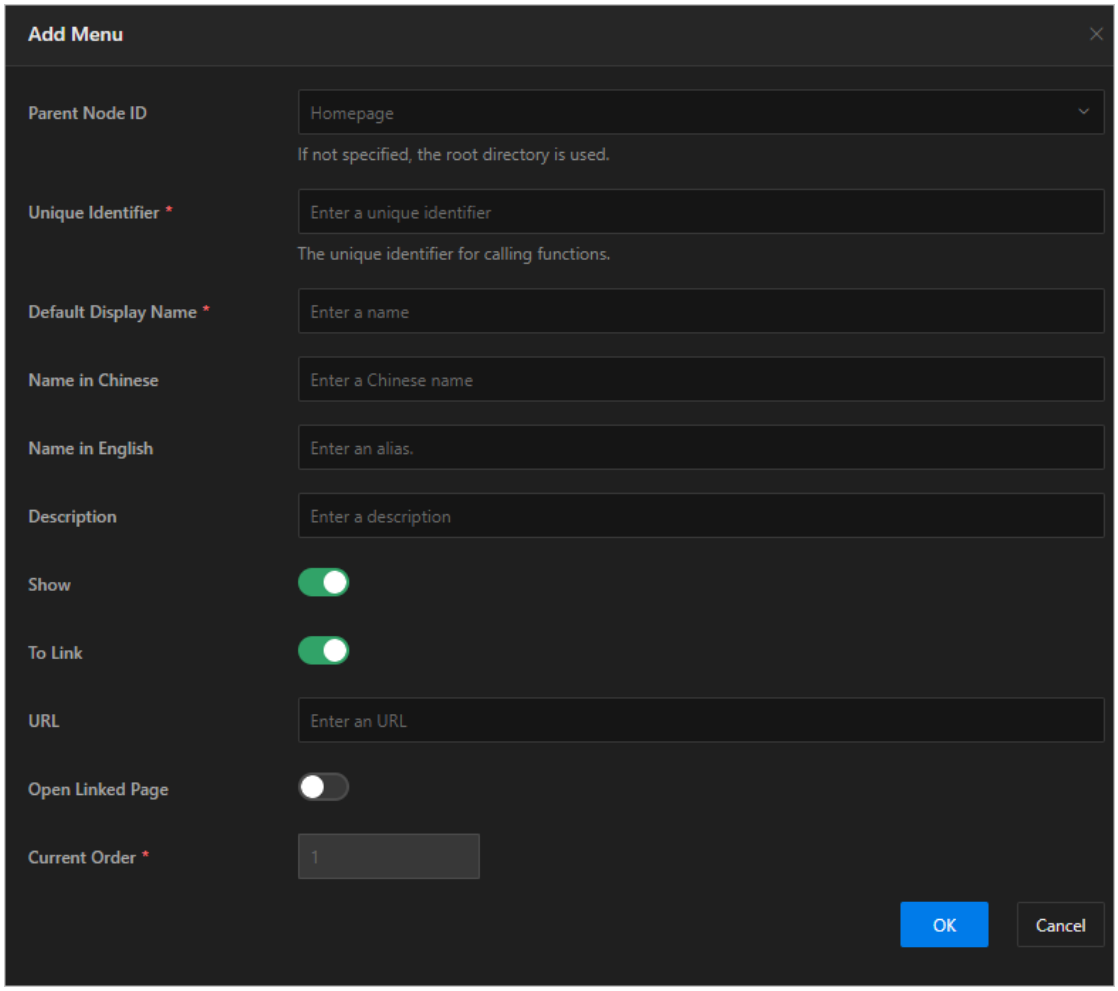

The following table describes the parameters.

Parameter	Description
Parent Node ID	The menu to which the submenu belongs.
Unique Identifier	The unique identifier used to call functions. It can consist only of letters and can be 5 to 20 characters in length.
Default Display Name	The default display name of the submenu.
Name in Chinese	The submenu name in Chinese. In the Chinese language environment, if the Chinese name of the submenu is specified, the default display name of the submenu is the specified Chinese name.
Name in English	The submenu name in English. In the English language environment, if the English name of the submenu is specified, the default display name of the submenu is the specified English name.
Description	The description of the submenu.

Parameter	Description
Show	Specifies whether to show the submenu after it is added. You can turn on or off **Show**. By default, Show is turned on.
To Link	Specifies whether to go to another page when you click the submenu. You can turn on or off **To Link**. By default, To Link is turned off.
URL	This parameter appears only when **To Link** is turned on. Set this parameter to the URL to go to when you click the submenu. ■ If the URL of a page within the current system is used, enter the absolute path or a relative path of the page. Example: /aso/aso-alarm/dashboard. ■ If the URL of a third-party system is used, enter the absolute path of the page. Example: http://example.com/TaskManageTool/#/taskView.
Open Linked Page	Specifies whether to open a new page for the URL to go to after you click the submenu. You can turn on or off **Open Linked Page**. By default, the switch is off.
Menu Type	The type of the menu. When you create a submenu, you do not need to configure this parameter.
Current Order	The order of the submenu under the selected menu. You cannot configure the order in the panel. You can modify the configuration on the **Menu Settings** page after you create the submenu.

iii. Click **OK**.
After you add a submenu, you can view it under the corresponding parent menu in the menu list and in the left-side navigation pane.

Note We recommend that you create a menu hierarchy of no more than five levels.

1.1.3.5.4.3. Hide a menu

This topic describes how to hide a menu.

Prerequisites

Notice Only custom menus and submenus can be hidden. After a menu or submenu is hidden, submenus beneath it are also hidden.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Menu Settings**.

4. In the menu list, find the menu or submenu that you want to hide and click **Modify** in the **Actions** column.
5. In the Modify Menu panel, turn off **Show** and click **OK**.

1.1.3.5.4.4. Modify a menu

After you add a menu or submenu, you can modify its configurations and sorting.

Prerequisites

Notice Only custom menus and submenus can be modified. Built-in menus and submenus can only be sorted.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Menu Settings**.
4. In the menu list, find the menu or submenu that you want to modify and click **Modify** in the **Actions** column.
5. In the Modify Menu panel, modify the configurations and click **OK**.
6. In the **Actions** column, click **Move Up** or **Move Down** to change the order of the menu.

1.1.3.5.4.5. Delete a menu

This topic describes how to delete menus or submenus that are no longer needed.

Prerequisites

Notice Only custom menus and submenus can be deleted.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Menu Settings**.
4. In the menu list, find the menu or submenu that you want to hide and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.3.5.5. Two-factor authentication

To make user logons more secure, you can configure two-factor authentication for users.

Context

The Apsara Uni-manager Operations Console supports only Google two-factor authentication.

This authentication method is 2-step verification and uses a password and mobile app to provide a two-layer protection for accounts. You can obtain the logon key after you configure users in the Apsara Uni-manager Operations Console, and then enter the key in the Google Authenticator app on your mobile phone. The app dynamically generates a verification code for your logon based on the time and key.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Two-factor Authentication**.
4. On the Two Factor Authentication page, perform the following operations:
   - Google two-factor authentication
     a. Set **Current Authentication Method** to **Google Two-Factor Authentication**.
     b. In the upper-right corner of the page, click **Add User**. In the Add User dialog box, enter a username and click OK. The added user is displayed in the user list.
     c. Find the username for which you want to enable Google two-factor authentication and click **Create Key** in the **Actions** column. When the **Added** message appears, the **Show Key** button appears in the **Actions** column. Click **Show Key**. The key is displayed in plaintext in the Key column.
     d. Enter the key in the Google Authenticator app on your mobile phone. The app dynamically generates a verification code for your logon based on the time and key. While two-factor authentication is enabled, you are required to enter the verification code on your app whenever you log on to the system.

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

     e. To disable two-factor authentication, click **Delete Key** 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.

1.1.3.5.6. Department management

The Department Management module allows administrators to create, modify, delete, and search for departments, as well as create users or user groups for departments.

Context

After the Apsara Uni-manager Operations Console is deployed, a root department is automatically generated. You can create other 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Department Management**.

   On the **Department Management** page, you can view the tree structure of all created departments and the information about all users in each department.
4. (Optional)In the upper-left corner of the page, enter a department name in the search box and click the search icon to find the department that you want to manage.
5. Perform the following operations:
   - Add a department

     In the left catalog tree, select the department to which you want to add sub-departments and click **Add Department**. In the **Add Department** dialog box, set **Department Name** and **Role**, and then click **OK**. Then, you can view the created department in the left catalog tree.
   - Modify a department

     In the left catalog tree, select the department that you want to modify and click **Modify Department**. In the **Modify Department** dialog box, set **Department Name**, **Department Administrator**, and **Role**, and 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 that you want to delete and click **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 Users section on the right, click **Create User**. In the **Add User** dialog box, set **Username**, **Password**, **Confirm Password**, and optional parameters such as **Display Name**. Click **OK**.

     Then, you can choose **System Settings > User Management** to view the added user on the **Users** tab.
   - 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 Group section on the right, click **Create User Group**. In the **Create User Group** dialog box, enter a user group name and click **OK**.

     Then, you can choose **System Settings > User Group Management** to view the added user group.

1.1.3.5.7. Region management

In multi-region scenarios, the system administrator can bind a department to a region. After you bind a department to a region, 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 administrators log on to the Apsara Uni-manager Operations Console, each administrator can manage only resources in the region that they are authorized to manage.

Relationship between departments and regions:

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Region Management**.
4. (Optional)In the upper-left corner of the page, enter a department name in the search box and click the search icon to find the department that you want to bind.
5. In the left catalog tree, click a department and select one or more regions in the **Regions** list on the right.
6. Click **Update Association**.

1.1.3.5.8. Operating system logs

The Operating System Logs module allows you to collect statistics on the number of logs and identify events related to the operating system.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Operating System Logs**.
4. On the **Operating System Logs** page, select **Time** and set **Range** by specifying IP, HostName, or MachineGroup. Select the **Search by Keyword** check box, enter **Keyword**, and then click **Search**. View the number of logs, the number of events, and the logs generated around the time point when the event occurred.

   Specify **EventType**, **DataSource**, and **LogLevel** to further filter information. The default value of each parameter is **Total**.

Parameter	Description
EventType	The exception category, which can be kernel, security, or system.
DataSource	The system application from which logs are generated. It mainly is used to indicate a system application.
LogLevel	The log error level, which can be debug, info, notice, warning, err, or crit.

5. In the section below the graphs, click the  icon and select an event. Click the  icon before the event to show the detailed log information around the time point when the event occurred.

   Click **Export** to export the full log information within the specified time range.

1.1.3.5.9. Operation logs

You can view logs to view the resource usage and running status of all modules on the platform.

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, view call details, and export selected logs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Operation Logs**.
4. On the **Log Management** page, perform the following operations:
   - Manage logs

     In the upper part of the page, enter **User Name** and select a period of time for **Time Period**. Click **Search** to view related logs in the list below.
   - Delete logs

     Select the logs that you want to delete and click **Delete**. In the message that appears, click **OK**.
   - Export logs

     Select the logs that you want to export and click the icon. If you do not select logs, when you click the icon, all displayed logs are exported.

     > **Note** If the number of logs to be exported is greater than 10,000, only the first 10,000 logs are exported.

1.1.3.5.10. View authorization information

The Authorization module allows customers, field engineers, and operations engineers to query services that are experiencing authorization problems and troubleshoot the problems.

Prerequisites

The logon user has the 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 only when you have the administrator permissions.

If you are not granted the administrator permissions, a message appears indicating that you have insufficient permissions when you access this tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.

3. In the left-side navigation pane, choose **System Settings > Authorization**.

   The **Authorization Details** tab appears.

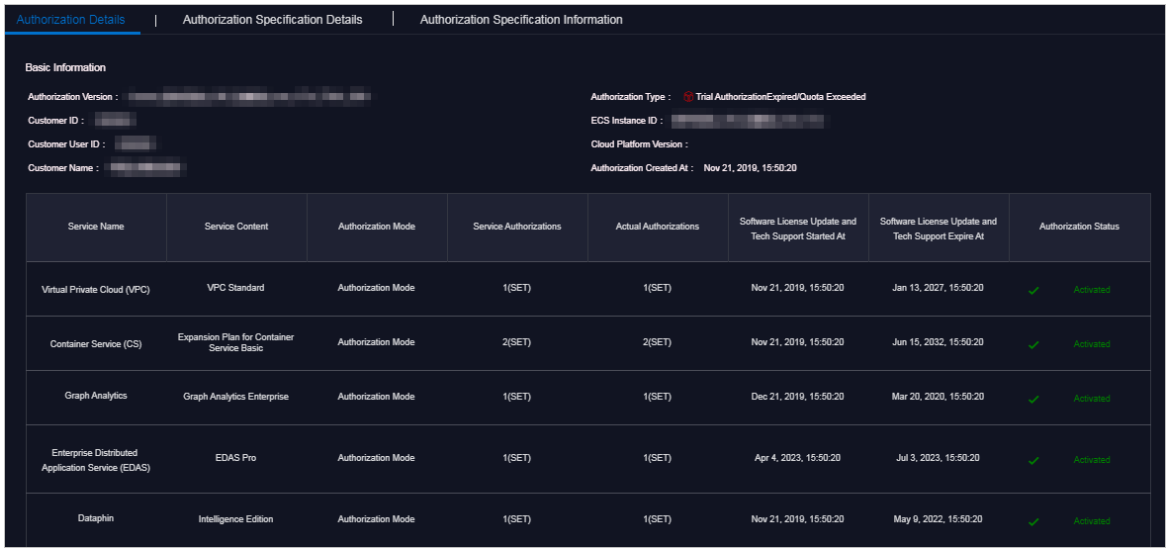

4. Perform the following operations to view the authorization information.

   > **Note** For formal authorization, you must enter the authorization code to view the authorization information. You can obtain the authorization code from the authorization letter appended to the project contract or by contacting the business manager (CBM) of your project.

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

     You can view authorization information including the 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 cloud services in different data centers.

     The following table describes the detailed authorization information.

Item	Description
Authorized Version	You can use the BP number in the version to associate a project or contract. The following parameters in the Input Parameters section must be specified: ■ **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 formal authorization. The authorization information of the service comes from the signed contract.

Item	Description
Authorization Type	Indicates the current authorization type and authorization status. ■ The following authorization types are available: ■ **Trial Authorization** ■ **Formal Authorization** ■ The following authorization states are available: ■ **Not Activated** ■ **Expire Soon** ■ **Activated** ■ **Expired** ■ **Expired/Quota Exceeded**
Customer information	The information about the customer, including the username, UID, and user ID.
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.
Apsara Stack Product Authorization Details (Data Center)	The authorization information of cloud services within different regions, including 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, take note of the following items: ■ **RENEW Service Expired** Indicates that the customer must renew the subscription as soon as possible. Otherwise, field operations services (including ticket processing) are terminated. ■ **Specification 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 Specifications Details** tab to view the authorization specification information of services across different data centers or regions.

  The following table describes the authorization specification information.

Item	Description
Service Name	The name of an authorized service.
Authorization Specification Name	The specification name of an authorized service.
Authorized Specifications	The total number of current authorizations of a specification for a service.
Service Authorizations	The authorization quota of a specification for a service.
Authorization 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.

  Set **Licensing Specification Level**, **Region ID** or **Data Center ID**, **Service Name**, and time range, and then click **Search**. You can view the authorization specification information of a service in the current environment. Such information includes the maximum and minimum number of specifications and their occurrence time as well as the average number of specifications within the specified time range.

  In the **Authorization Specification Information** or **Excess Authorization Specification 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.

1.1.3.5.11. Multi-cloud management

The multi-cloud management module allows you to use the same platform to perform O&M on different data centers.

1.1.3.5.11.1. Add multi-cloud configurations

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

Prerequisites

Before you add multi-cloud configurations, make sure that the following requirements are met:

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.

3. In the left-side navigation pane, choose **System Settings > Multi-cloud Management**.
4. In the upper part of the page, click **Add**.
5. In the dialog box that appears, add the console link of the new data center and click **OK**.

Parameter	Description
Cloud Name	The name of the new data center.
Central Region Console URL	The console URL of the new data center. Make sure that the console URL is valid. Otherwise, an error message is returned.
Longitude	Optional. The geographic longitude value of the new data center.
Latitude	Optional. The geographic latitude value of the new data center.

After you add multi-cloud configurations, you can log on to the Apsara Uni-manager Operations Console by using a shared account to switch to different data centers and perform related operations.

1.1.3.5.11.2. Modify multi-cloud configurations

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Multi-cloud Management**.
4. Find the data center and click **Modify** in the **Actions** column.
5. In the dialog box that appears, modify the multi-cloud configurations and click **OK**.

1.1.3.6. Personal Settings

The Personal Settings module allows you to modify the logon password and logon settings of the current account.

1.1.3.6.1. Change the logon password

The Logon Settings module allows you to change the password that you use to log on to the Apsara Uni-manager Operations Console.

Context

For security reasons, we recommend that you change your logon password on a regular basis.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.

3. In the left-side navigation pane, choose **Personal Settings > Personal Information**.
4. View the personal information of the current user, such as **Username** and **Department**.
5. Click **Change Password** to change the password that you use to log on to the Apsara Uni-manager Operations Console.
6. In the **Change Password** dialog box, specify **Current Password**, **New Password**, and **Confirm Password**, and then click **OK**.

1.1.3.6.2. Modify logon settings

The Logon Settings module allows you to configure whether to allow multi-terminal logon and modify the logon timeout period, maximum allowed password retries, logon policy, and validity period of the account you are using.

Context

To make your system more secure, you can modify the logon settings based on your scenario.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **Personal Settings > Logon Settings**.
4. On the **Logon Settings** tab, modify the following parameters.
   - **Timeout Period (Minutes)**: Set the logon timeout period of the current account. If the logon time exceeds the specified time period, the system prompts you that the logon timed out and you can try to log on again.
   - **Multi-Terminal Logon Settings**: Set whether to allow multi-terminal logon on the current account. You can select **Multi-Terminal Logon Allowed**, **Forbid Multi-Terminal Logon in ASO**, or **Forbid Multi-Terminal Logon in O&M**.
     - **Multi-Terminal Logon Allowed**: The current account is allowed to log on to the Apsara Uni-manager Operations Console from multiple terminals at the same time.
     - **Forbid Multi-Terminal Logon in ASO**: The current account is not allowed to log on to the Apsara Uni-manager Operations Console from multiple terminals at the same time. The current account is allowed to go to another console from the Apsara Uni-manager Operations Console.

       For example, User A uses the current account to go to another console from the Apsara Uni-manager Operations Console. At the same time, User B uses the current account to log on to the Apsara Uni-manager Operations Console. The system disables the logon of User A only after User A returns to the Apsara Uni-manager Operations Console.
     - **Forbid Multi-Terminal Logon in O&M**: The current account is not allowed to log on to the Apsara Uni-manager Operations Console or the console redirected from the Apsara Uni-manager Operations Console from multiple terminals.
   - **Maximum Allowed Password Retries**: Set the maximum number of password retries before the account is locked. When the number of retries reaches the specified number, the account is locked. After the account is locked, you must use the system administrator account to unlock it.
   - **Logon Policies**: Set the logon policy of the current account. You can select **Blacklist** or **Whitelist**. For more information about how to add logon policies, see Logon policies.

- **Blacklist**: If this option is selected, you cannot use the IP addresses configured in logon policies to log on to the Apsara Uni-manager Operations Console.
- **Whitelist**: If this option is selected, you can use the IP addresses configured in logon policies to log on to the Apsara Uni-manager Operations Console.

5. Click **Save**.
6. Click the **Account Validity Period** tab and set **Account Validity (Days)**.

> **Note** When your account expires, you must use the system administrator account to unlock it.

7. Click **Save**.

1.1.4. Resources

1.1.4.1. Products

1.1.4.1.1. Product overview

On the Product Overview page, you can view information about each product and its clusters, server roles, machines, alerts, and final-state status.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Products**.
3. On the **Product Overview** page, view the information about clusters, server roles, machines, alerts, and final-state statuses.

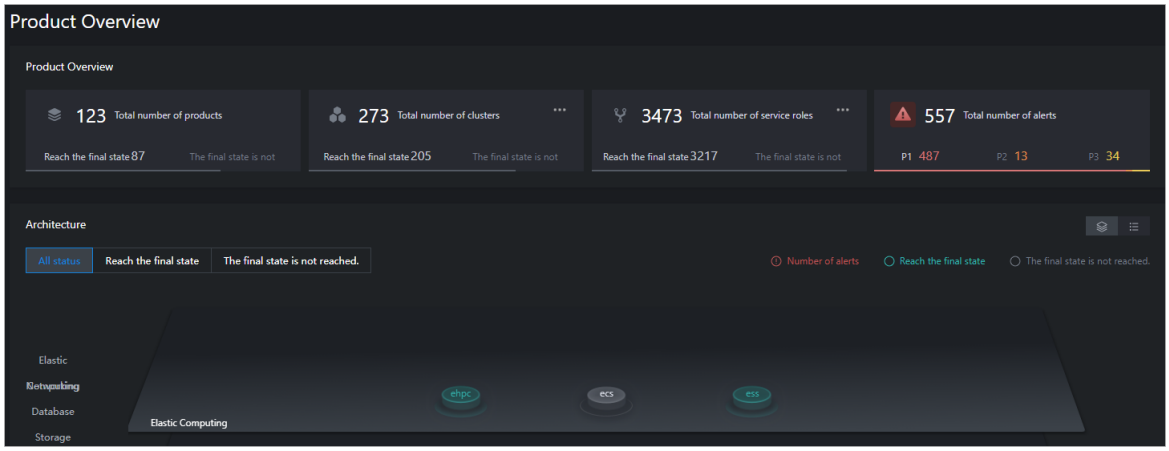

   - In the **Product Overview** section, view the total numbers of products, clusters, server roles, and alerts as well as the final-state status of these products, clusters, and server roles.
     - Click the ••• icon next to **Total Clusters** and view the information about the clusters on the **Clusters** page.

- Click the icon next to **Total Server Roles** and view the information about the server roles on the **Server Roles** page.

- In the **Architecture** section, view the product information.
  - Click **All status**, **Desired State**, or **Not Desired State** to view the information about the products in the corresponding state.
  - Click a product type in the left-side product type list to view the final-state status and alerts of the products of that type.
  - Click the icon in the upper-right corner to view the cluster status and the numbers of clusters, server roles, machines, and alerts in a list form.

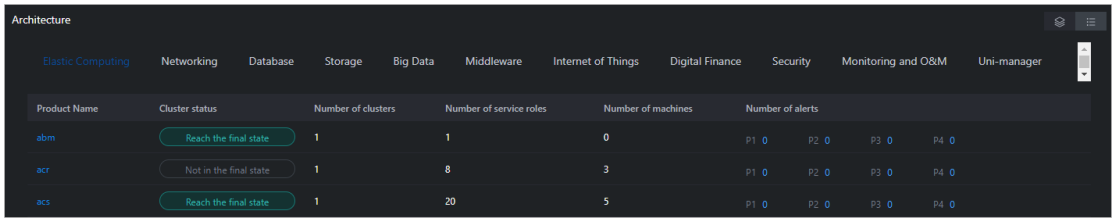

  - Click a product type in the upper product type list to view the information about the products of that type.
  - Click a product name in the Product Name column to view the details of the product on the **Product Details** page.

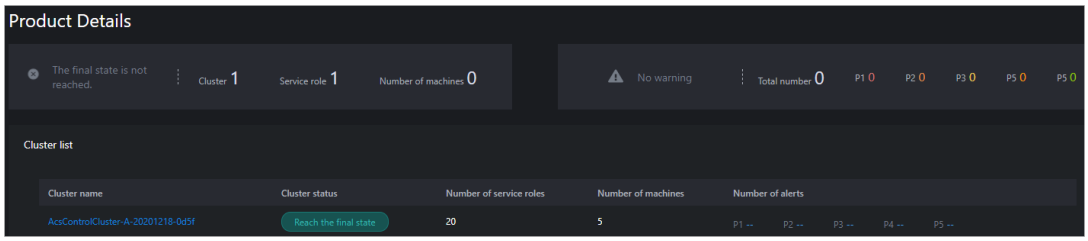

    a. Click a cluster name in the Cluster Name column to view the details of the cluster on the **Cluster Details** page.

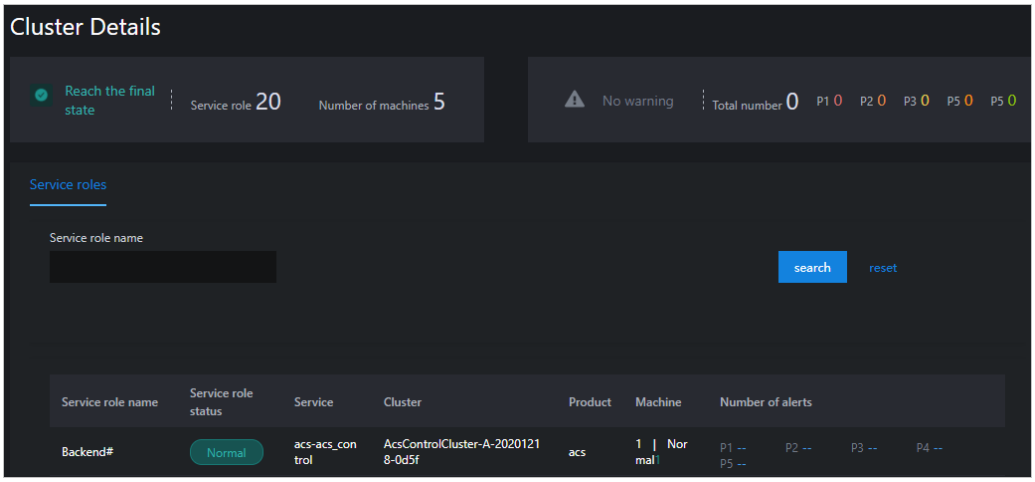

    b. Enter a server role name in the **Server Role Name** search box and click **Search** to view the details about the server role.

    c. (Optional) Click **Reset** to clear the search conditions.

- Click the icon in the upper-right corner to view the information about the products in a graphical form.

1.1.4.1.2. Clusters

You can view the status and alerts of all the deployed clusters and their server roles on the Cluster list page.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Products**.
3. In the left-side navigation pane, click **Clusters**.
4. On the **Cluster list** page, select a product from the **Product Name** drop-down list, select a cluster from the **Cluster Name** drop-down list, and then select a state from the **Cluster Status** drop-down list. Then, click **Search** to view the search results.

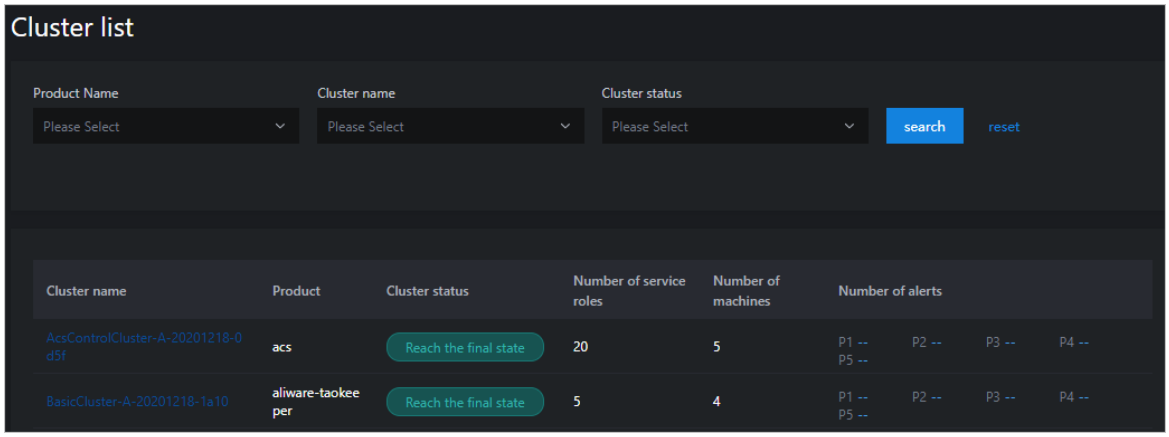

5. (Optional) Click **Reset** to clear the filter conditions.
6. Click a cluster name in the Cluster Name column to view the details of the cluster on the **Cluster Details** page.

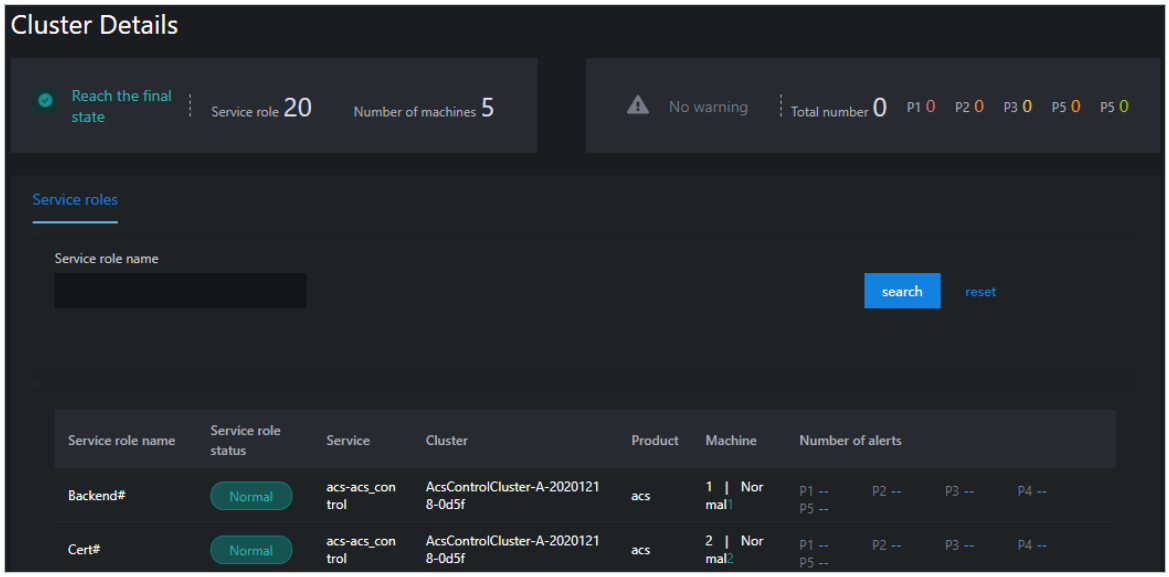

7. Enter a server role name in the **Server Role Name** search box and click **Search** to view the details about the server role.
8. (Optional) Click **Reset** to clear the filter conditions.

1.1.4.1.3. Server roles

On the Server Roles page, you can view the server role of a specific product or cluster and the status and alerts of the cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Products**.
3. In the left-side navigation pane, click **Server Roles**.
4. On the **Server Roles** page, select a product from the **Product Name** drop-down list, select a cluster from the **Cluster Name** drop-down list, and then select a state from the **Server Role Status** drop-down list. Then, click **Search** and view the search results displayed below.

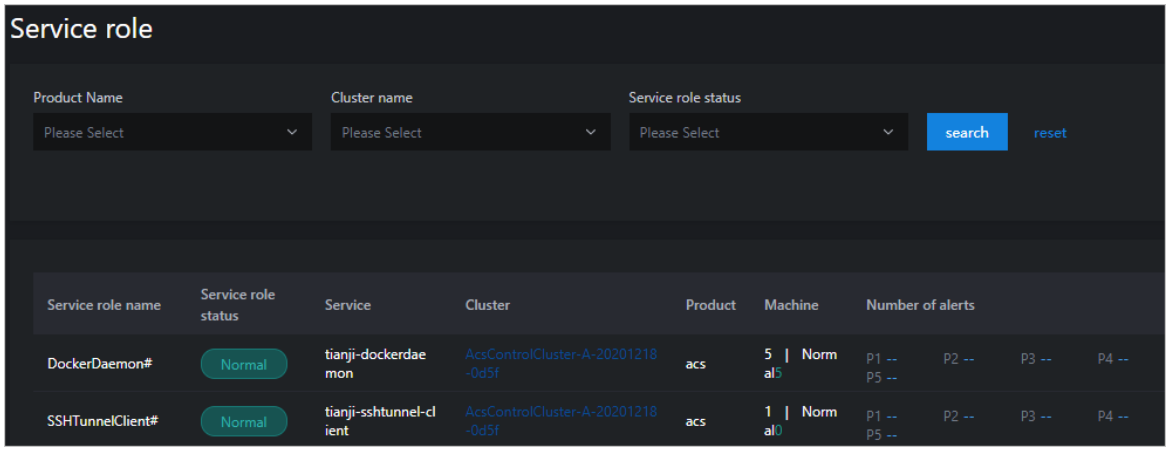

5. (Optional) Click **Reset** to clear the search conditions.
6. Find a server role and click the number next to **Normal** or **Exception** in the **Machine** column. In the dialog box that appears, find a virtual machine and click **Security Operations** on the right side to go to the CLI window of the virtual machine. For more information, see Log on to the host where a server role is deployed.
7. Click a cluster name in the Cluster column to view the details of the cluster on the **Cluster Details** page.

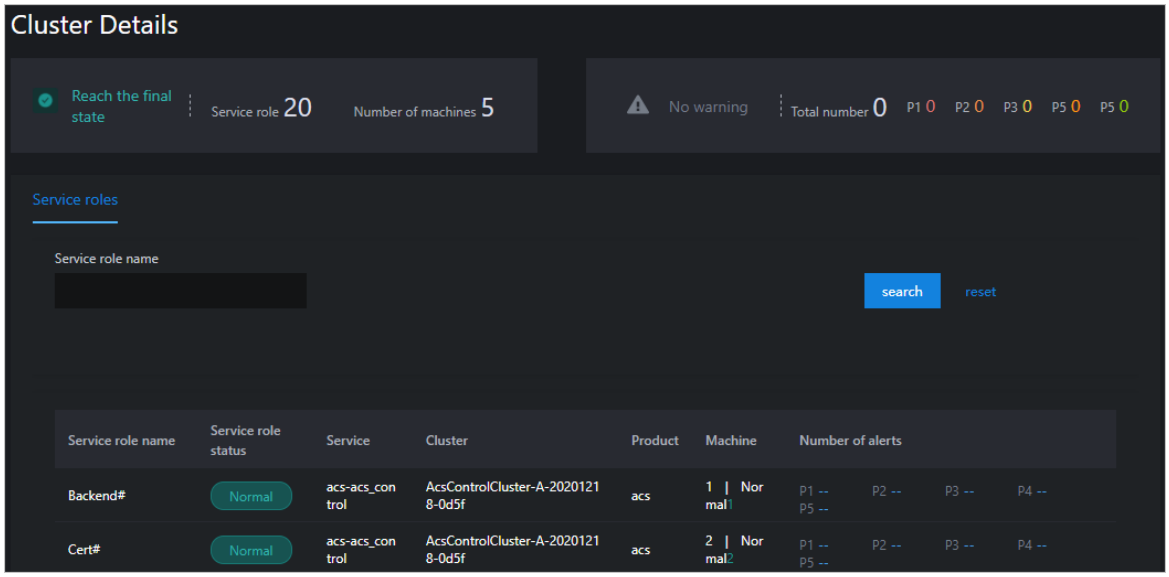

8. Enter a server role name in the **Server Role Name** search box and click **Search** to view the details about the server role.
9. (Optional) Click **Reset** to clear the search conditions.

1.1.4.2. Network

The network module provides the cloud service interconnection and network resource management features for the hybrid cloud network. The user is the network administrator of Apsara Stack.

1.1.4.2.1. Cloud service interconnection

The cloud service interconnection feature provides network access to cloud services, configuration of VIPs and DNS, and mutual access between cloud services in multiple clouds, IDCs, and Apsara Stack.

1.1.4.2.1.1. Dynamic VIP

The dynamic VIP feature expands the network capabilities of cloud service interconnection. It allows you to create dynamic VIPs by using VIP resources applied for by the cloud services in Apsara Infrastructure Management Framework as templates. You can modify the type, tunnel_type, and ipprotocol properties when you create a dynamic VIP. Other properties are the same as those in the template VIP resources defined by the cloud products. During scaling, upgrades, and downgrades of products, dynamic VIP resources and template VIP resources can be simultaneously updated by linking with Apsara Infrastructure Management Framework.

Context

- In scenarios that require large storage capacities, the dynamic VIP feature can be used to scale out multiple OSS Intranet VIP resources to bypass the single-VIP network speed of 5 Gbit/s.
- In hybrid cloud scenarios, you must access the Apsara Uni-manager Operations Console over the Internet and use the dynamic VIP feature to create Internet VIP resources for the Apsara Uni-manager Operations Console.
- The dynamic VIP tunnel type can be set only to classic_to_any_tunnel.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. On the **Dynamic VIP** page, view the dynamic VIPs.
   - Enter a VIP name in the **VIP Address** search box, enter a template VIP name in the **Template VIP Name** search box, and then select a type from the **SLB Instance Type** drop-down list. Then, click **Search** to view the dynamic VIPs. Click **Details** in the Actions column corresponding to a dynamic VIP to view the details of the dynamic VIP.

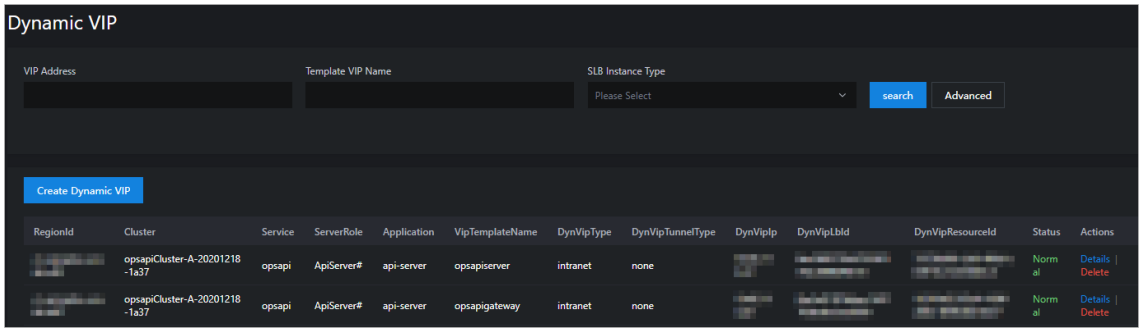

   - Use more filter conditions to view the information of dynamic VIPs.
     a. Click **Advanced**, select a type from the **Tunnel Type** drop-down list (you can select multiple types), and enter a region name in the **Region** search box. Then, click **Search** to view the dynamic VIPs.
     b. (Optional) Click **Reset** to clear the filter conditions.
     c. (Optional) Click **Collapse** to hide the **Tunnel Type** and **Region** options.
4. Create, update, and delete dynamic VIPs.
   - Click **Create Dynamic VIP** above the dynamic VIP list. In the Create Dynamic VIP dialog box, configure the parameters. Then, click **OK** to create a dynamic VIP.

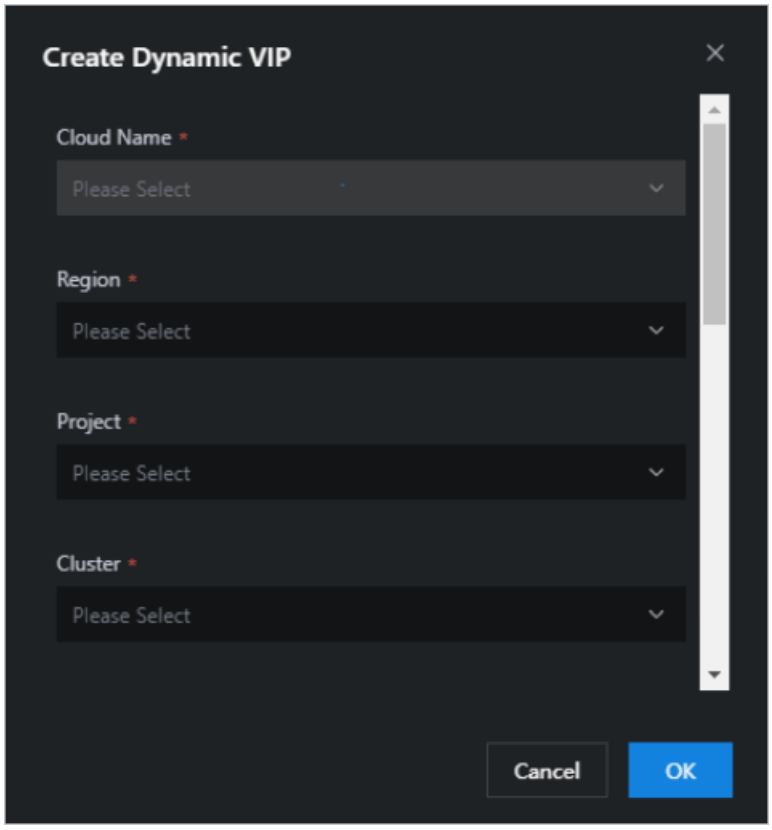

The parameters of a dynamic VIP:

- Cloud Name: the name of the cloud instance.
- Region: the ID of the region.
- Project: the name of the product.
- Cluster: the name of the cluster.
- Service: the name of the service.
- ServerRole: the name of the server role.
- Application: the name of the application.
- Template VIP Name: the name of the template VIP.
- DynVipType: the type of the instance.
- DynVipTunnelType: the type of the tunnel.
- Ipprotocol: the type of the dynamic VIP.

- Update a dynamic VIP.

  If the basic properties of a dynamic VIP are inconsistent with those of a template VIP, or if an operation on the dynamic VIP fails, **Abnormal** is displayed in the **Status** column corresponding to the dynamic VIP. When the system detects an exception in the dynamic VIP status, the system automatically checks and deletes DNS records associated with the abnormal dynamic VIP to reduce the impact on your business.

  The O&M personnel must check whether the abnormal dynamic VIP status is caused by resource changes during product upgrades or scaling and perform the following operations accordingly:

- If the abnormal dynamic VIP status is caused by resource changes during product upgrades or scaling, click **Update** in the Actions column corresponding to the dynamic VIP. The system queries the template VIP parameters through calls to Apsara Infrastructure Management Framework API operations and combine these parameters with the variable properties of the dynamic VIP to update the dynamic VIP.
- If the abnormal dynamic VIP status is not caused by resource changes during product upgrades or scaling, submit a ticket to contact Apsara Stack technical support.

- To delete a dynamic VIP, click **Delete** in the Actions column corresponding to the dynamic VIP. In the message that appears, click **OK**.

> **Notice** Risks may arise if you delete dynamic VIPs. Proceed with caution. Before you delete a dynamic VIP, make sure that the VIP does not have sessions and that DNS records can no longer be resolved to the dynamic VIP.

1.1.4.2.1.2. Dynamic DNS

Dynamic DNS works with the dynamic VIP feature. It allows you to implement horizontal network expansion by resolving DNS resources in cloud services to multiple VIPs (one static VIP and multiple dynamic VIPs).

Context

Dynamic DNS is applicable to scenarios such as multi-cloud geo-disaster recovery and hybrid clouds. The following features are supported:

- Adds domain names for zones corresponding to ops-dns intranet-domain and internet-domain. By default, ops-dns returns the resolution records in rotation mode.
- Adds the forwarding records of a tenant DNS to forward domain name resolution requests in a specified zone to external DNS servers such as DNS servers in Apsara Stack, heterogeneous clouds, and Alibaba Cloud.

> **Warning** Before you resolve DNS domain names in the Apsara Infrastructure Management Framework console, make sure that the parameters of the dynamic VIPs and the static VIP are exactly the same, and VIP listening on access is normal. Otherwise, network exceptions may occur.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Dynamic DNS**. On the **Dynamic DNS** page, view the information about dynamic DNS.
   - Select a type from the **DNS Type** drop-down list, select a cloud ID from the **Remote Cloud ID** drop-down list, and then enter a domain name in the **ZoneName** search box. Then, click **Search** to view the dynamic DNS records.

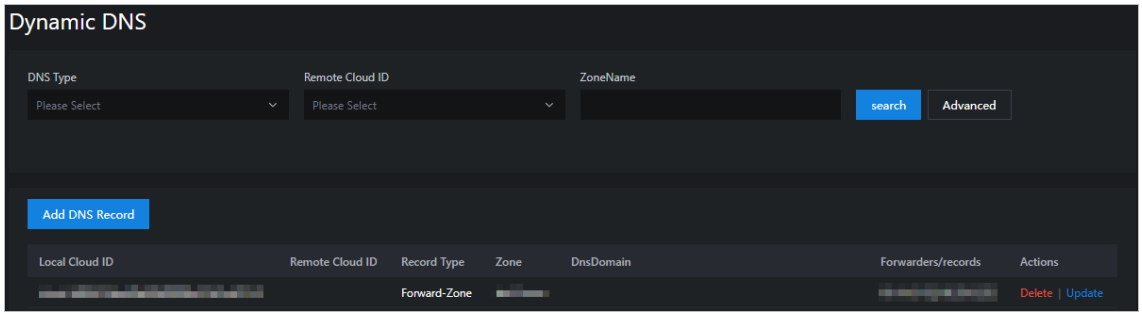

- Use more filter conditions to search for the dynamic DNS records.
  a. Click **Advanced**, enter a domain name in the **DnsDomain** search box, and then enter a DNS record in the **DNS Record** search box. Then, click **Search** to view the dynamic DNS records.

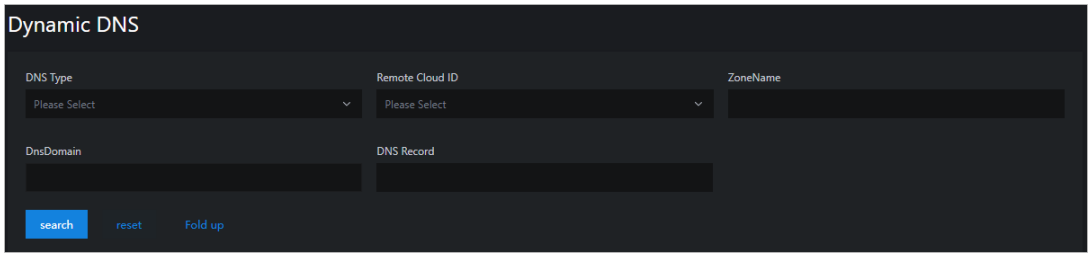

  b. (Optional) Click **Reset** to clear the filter conditions.
  c. (Optional) Click **Collapse** to hide the **DnsDomain** and **DNS Record** options.

4. Update and delete the dynamic DNS records or create new dynamic DNS records.
   - Click **Add DNS Record** above the dynamic DNS record list. In the Add DNS Record dialog box, configure the parameters. Then, click **OK** to create a dynamic DNS record.

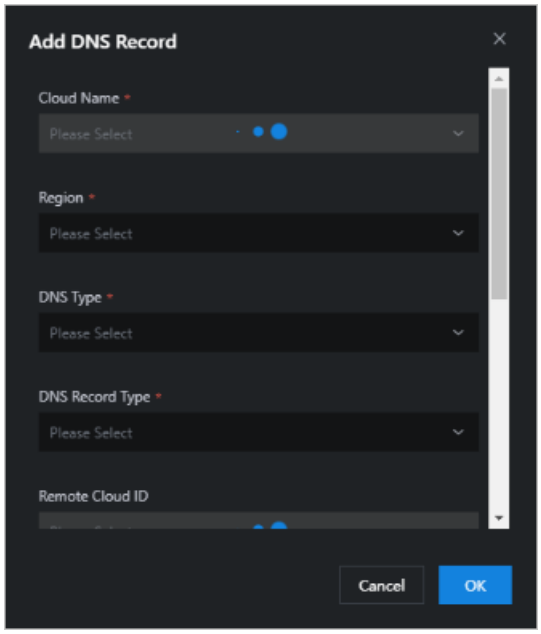

The parameters of a DNS record:

- Cloud Name: the name of the cloud instance.
- Region: the ID of the region.
- DNS Type: the DNS type. Valid values: dns product and ops dns.

- DNS Record Type: the type of the DNS record. Valid values: Forward-Zone and A.
- Remote Cloud ID: the ID of the cloud instance in the current cloud.
- DnsZone: the DNS zone. Separate multiple zones with commas (,).
- Forwarders: the forwarding IP addresses. Separate multiple IP addresses with commas (,).
- DnsDomain: the DNS endpoint.
- DnsRecord: the IP address of the DNS record.

- Click **Update** in the Actions column corresponding to the DNS record. In the dialog box that appears, modify the Remote Cloud ID and DnsRecord parameters. Then, click **OK**.
- To delete a DNS record, click **Delete** in the Actions column corresponding to the DNS record. In the message that appears, click **OK**.

> **Warning** Risks may arise if you delete DNS records. Proceed with caution. Before you delete a DNS record, make sure that an alternate DNS record is available for the domain name or that no other clients have access to this domain name.

1.1.4.2.1.3. Cross-cloud access

Cross-cloud access allows you to manage data for network access of cloud services in hybrid clouds.

Context

Before you enable cross-cloud access, you must configure IP network routing based on the access matrix data and grant the network access permissions.

Cross-cloud access is applicable to the following scenarios:

- Connection between Apsara Stack Enterprise and public cloud VPCs through dedicated lines: Public cloud VPCs and Apsara Stack VPCs can be connected by using dedicated lines.
- Connection between Apsara Stack Enterprise and public cloud VPCs over the Internet: Apsara Stack provides Internet egresses, so that public cloud VPCs and Apsara Stack VPCs can be connected over the Internet.
- Connection between Apsara Stack Agility ZStack and public cloud VPCs over the Internet: Apsara Stack Agility is connected to the Internet and ZStack VPCs are connected to public cloud VPCs over the Internet by using the IPSec VPN.
- Connection between Apsara Stack Enterprise and public cloud services through dedicated lines: Apsara Stack is connected to the public cloud by using dedicated lines to implement cloud management and network connection.
- Connection between Apsara Stack Enterprise, Apsara Stack Agility, and all-in-one cloud services (management) over the Internet: Apsara Stack Enterprise, Apsara Stack Agility, and all-in-one cloud services provide Internet egresses and can access cloud services on the Internet based on the NAT capabilities provided by user-managed firewalls.
- Hybrid clouds for multi-cloud remote disaster recovery: Hybrid clouds are connected to implement remote disaster recovery across regions.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.

3. In the left-side navigation pane, click **Cross-cloud Access**. On the **Cross-cloud Access** page, view the information about access matrices.
   i. Enter a product name in the **Product Name** search box and click **search**. The access matrices are displayed in the lower part of the page.

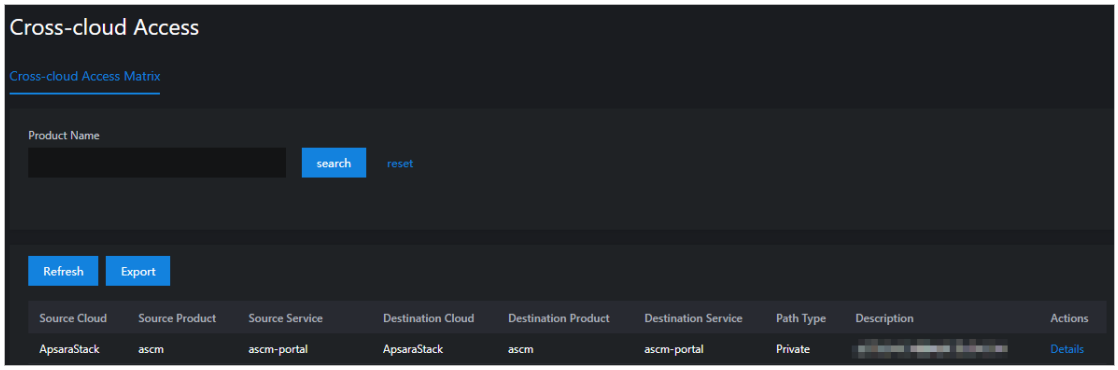

ii. Click **Details** in the Actions column corresponding to an access matrix to view the details of the access matrix.

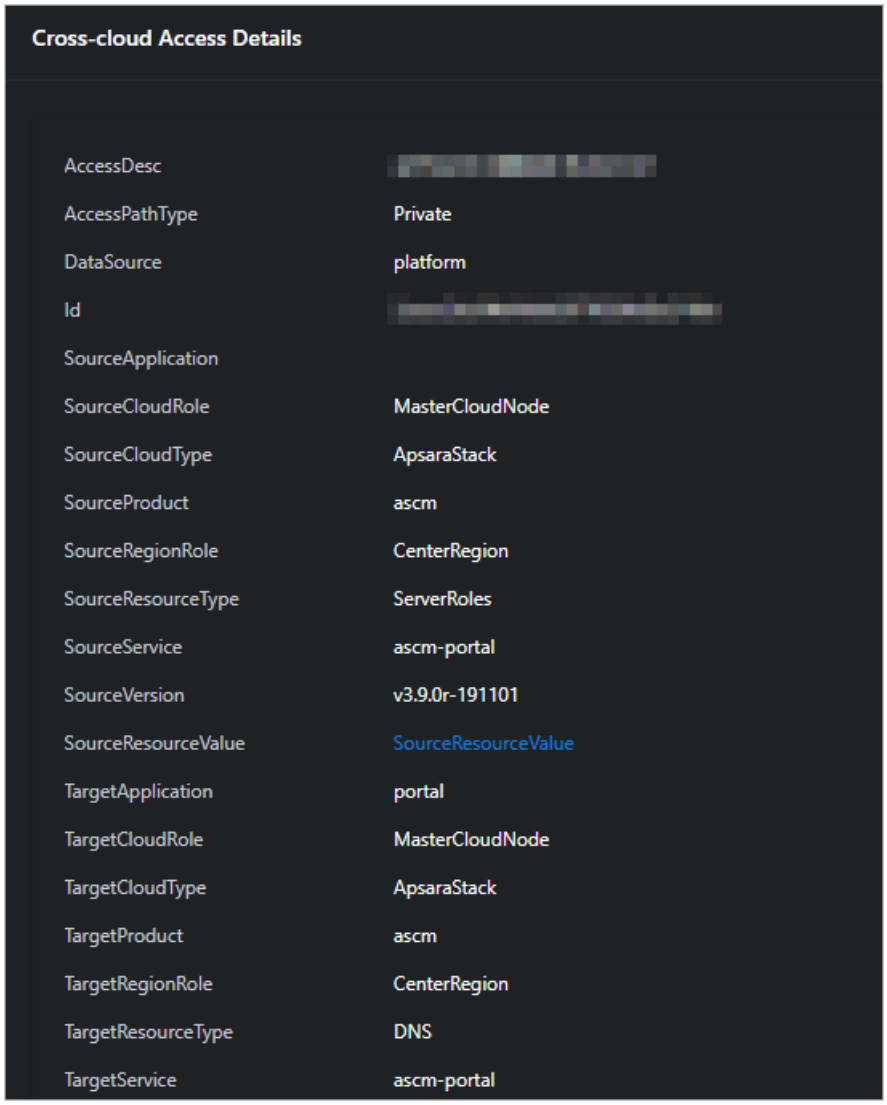

> **Note** If the cloud instance has multiple product clusters deployed or manages multiple lower-level cloud instances, each of the **SourceResourceValue** and **TargetResourceValue** parameters has multiple values. You must configure the network based on the source to the destination full mash.

iii. (Optional) Click **reset** to clear the filter conditions.

4. Click **Refresh**. Apsara Infrastructure Management Framework API is called to query the VIP, DNS, and SR resource values to refresh the **SourceResourceValue** and **TargetResourceValue** values.
5. Click **Export** to download the access matrix to your computer.

1.1.4.2.2. Hybrid cloud resources

You can manage hybrid cloud network resources such as physical network devices, network topology, and IP addresses in a centralized manner.

1.1.4.2.2.1. Physical topology

You can view the physical network topologies of hybrid clouds from multiple perspectives on the Physical Topology page.

Context

You can view the following types of physical network topologies:

- Standard topology: the physical network topology of the Apsara Stack data center. The initial data comes from Deployment Planner.
- IDC topology: the physical network topology of a self-managed IDC. The data comes from user-defined network devices or links.
- Global topology: the physical network topology of multiple data centers, including the standard topology and all IDC topologies. Data of interconnection links across data centers comes from user-defined links.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, choose **Hybrid Cloud Resources > Physical Topology**. On the **Physical Topologies** page, view the physical topologies.

i. Select the required options from the **Topology Type** and **Data Center** drop-down lists and click **Search** to view the physical topologies.

Note Action icons:

- : proceeds with the next step.
- : goes back to the previous step.
- : zooms in.
- : zooms out.
- : scales in proportion to the original aspect ratio.
- : scales to fit canvas.

ii. Click a node in the physical topology graph and the network device information of the node appears on the right side of the page.

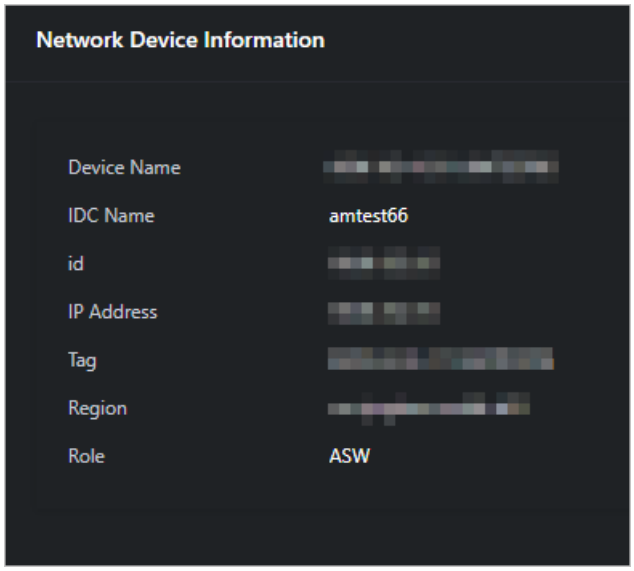

iii. (Optional) Click **Reset** to clear the filter conditions.

4. On the **Physical Topologies** page, click **Custom Link**. In the Custom Link dialog box, configure the parameters and click **OK** to set the link.
5. On the **Physical Topologies** page, click **Port Monitoring** to view the status of the ports.

i. On the Port Monitoring page, configure the parameters and click **Search** to view the port status trend chart.

> **Note** You can select multiple options for **Monitoring Metric** at a time.

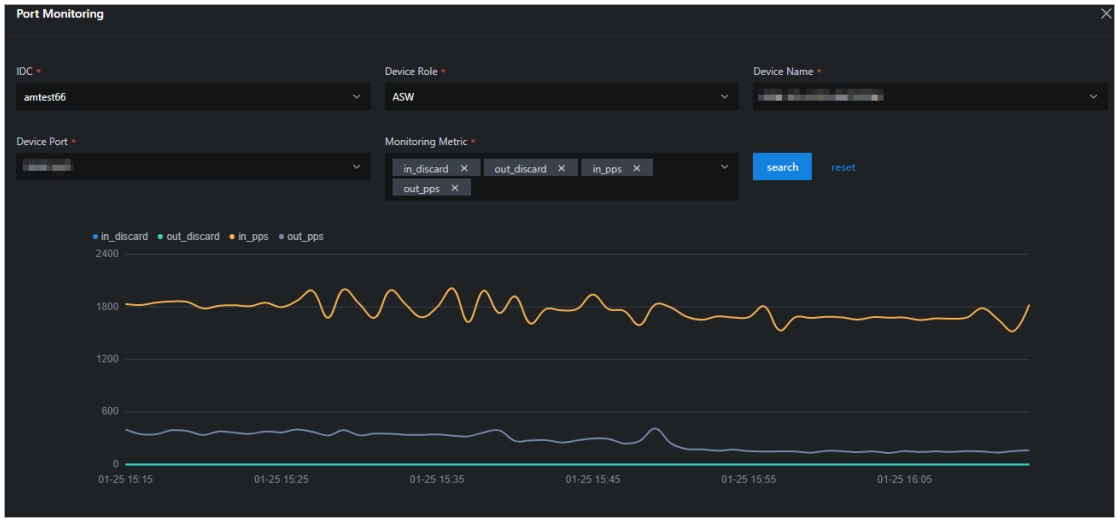

Monitoring metrics:

- in_discard: the inbound packet loss rate
- out_discard: the outbound packet loss rate
- in_pps: the inbound packet rate
- out_pps: the outbound packet rate
- in_bps: the inbound byte rate
- out_bps: the outbound byte rate
- in_error: the inbound packet error rate
- out_error: the outbound packet error rate

ii. (Optional) Click **Reset** to clear the filter conditions.

6. If you move a node, click **Save Topology** to save the new coordinates of the node.
7. Click **Refresh Topology** to generate coordinates for the nodes in the background.

1.1.4.2.2.2. Network element management

You can manage Apsara Stack data centers and user-managed data centers as well as their network element devices on the Network Element Management page.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, choose **Hybrid Cloud Resources > Network Element Management**. On the **Network Element Management** page, view the details of the network element devices.

i. Enter a keyword in the **Device Name/Management IP Address** search box and click **Search** to view the information about the network element device.

> **Note** The current Apsara Stack version supports status monitoring only for network elements in the Apsara Stack data center.

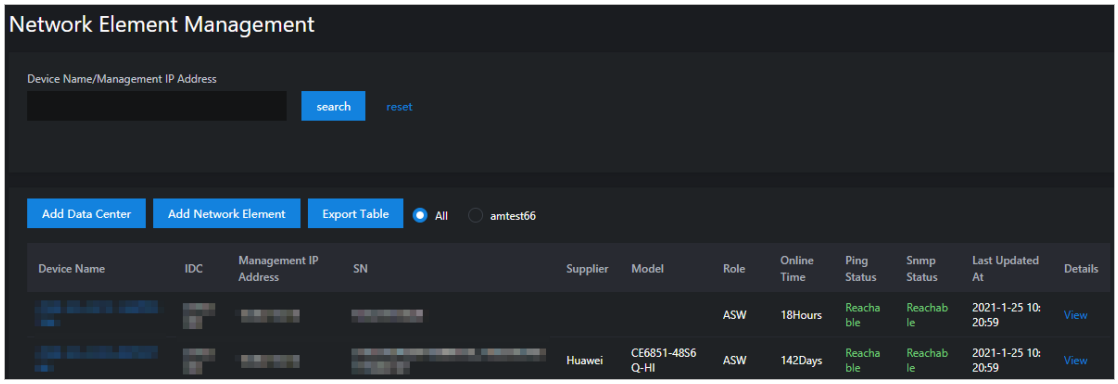

ii. Click **View** in the Details column corresponding to a network element device. On the **Network element details** page, enter a port name in the **Port Name** search box and click **Search** to view the port information. Click **View** in the Actions column to view the details of the port.

> **Note** The current Apsara Stack version supports management and monitoring only for network element ports in the Apsara Stack data center.

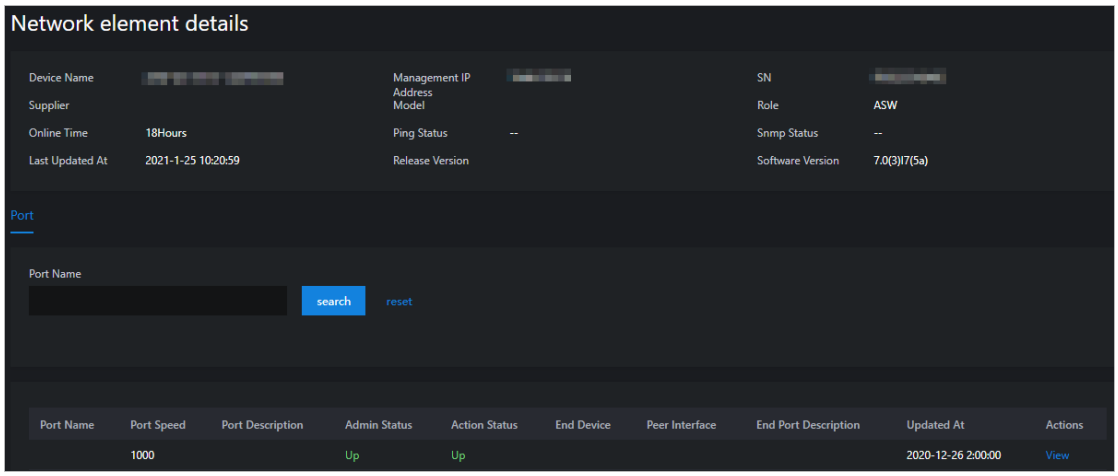

iii. (Optional) Go back to the **Network Element Management** page and click **Reset** to clear the filter conditions.

4. Click **Add Data Center**. In the Add Data Center dialog box, configure the parameters and click **OK** to add a data center.

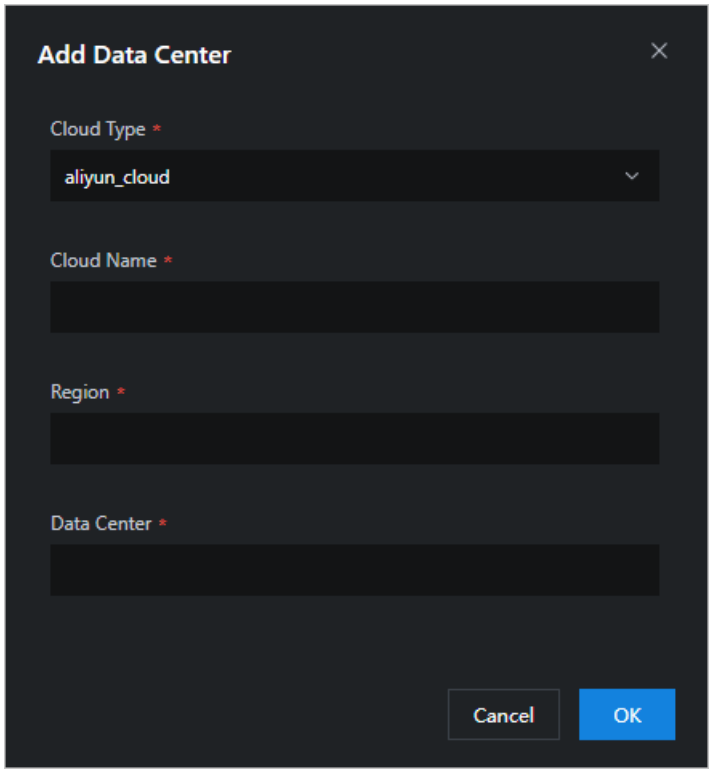

Parameter	Description	Example
Cloud Type	The type of the cloud instance.	apasara_stack
Cloud Name	The name of the cloud instance.	gddgzwy
Region	The region where the cloud instance is deployed.	cn-qingdao-envxxx
Data Center	The name of the data center.	amtestxx

5. Add a network element.
   i. Click **Add Network Element**. In the Add Network Element panel, configure the parameters.

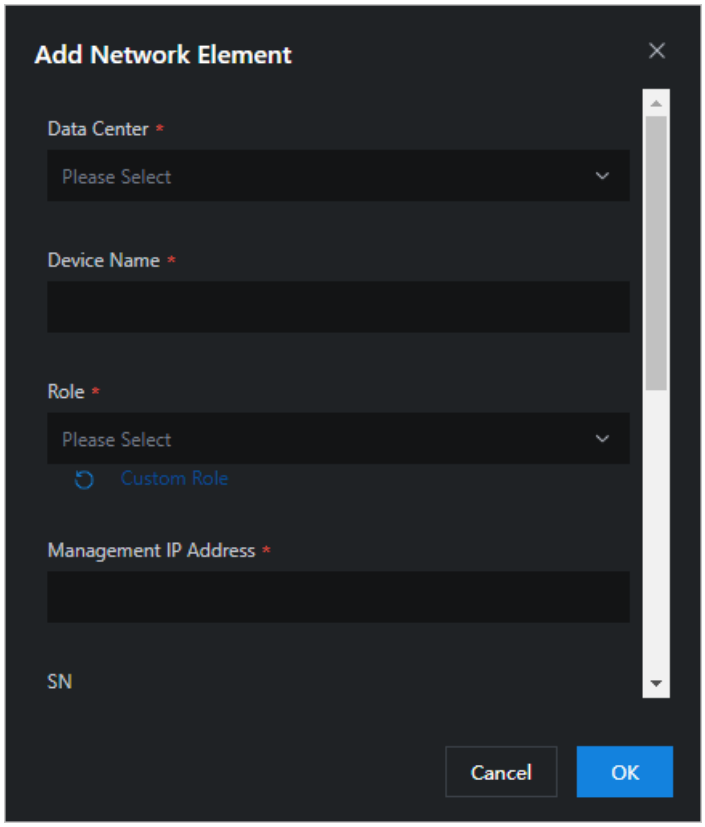

Parameter	Description	Example
Data Center	The name of the data center where the network element device is located.	amtest11
Device Name	The name of the network element device.	DSW-VM-G1-P-1.xxxx
Role	The role of the network element device.	ASW
Management IP Address	The management IP address of the network element device.	10.66.1.1
SN	The serial number of the network element device.	FDO23511111
Software Version	The version number of the software run by the network element device.	7.1.070
Release Version	The release version of the network element device.	V200R002C50SPC800
Supplier	The supplier of the network element device.	Ruijie

Parameter	Description	Example
Model	The model of the network element device.	S6510-48VS8CQ

ii. Click **Custom Role** next to **Role**. On the **Add a custom role** page, enter a device role name in the **Device Role** search box and click **Add**.

iii. In the **Add Network Element** panel, click **OK** to add the network element.

6. Click **Export Table** to download the information about the network element device to your computer.

1.1.4.2.2.3. IP address pools

Multiple cloud services may be deployed on a single cloud instance and you may need to view the IP addresses of the cloud services in scenarios such as network changes and security audits. You can view the IP address pool of each cloud service on a cloud instance on the IP Address Pools page.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, choose **Hybrid Cloud Resources > IP Address Pools** to view the IP address pools.
   - Select the required options from the **IP Address Pool ID**, **Cloud Name**, and **Product Name** drop-down lists, and click **search** to view the information about the IP address pool.

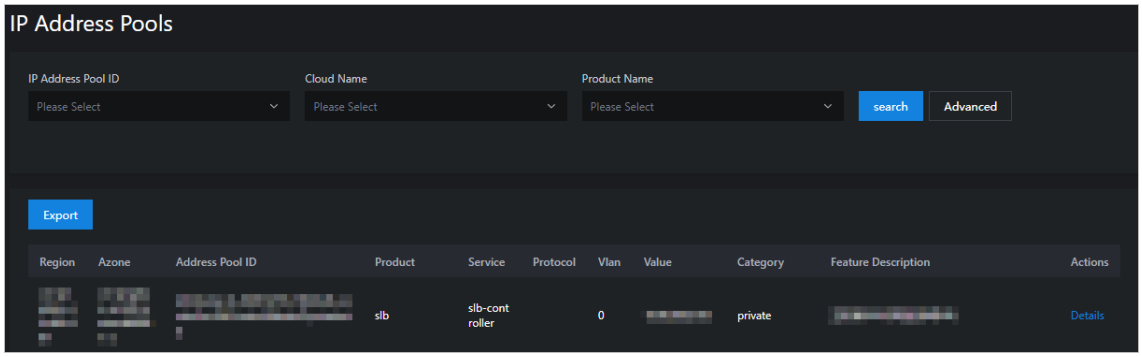

   - Use more filter conditions to view the information about the IP address pool.

a. Click **Advanced**. Select the required options from the additional **IP Resource Name**, **IP Address Type**, **CIDR Block of the IP Address Pool**, **Protocol**, and **Data Source** drop-down lists, and click **search** to view the information of the IP address pool.

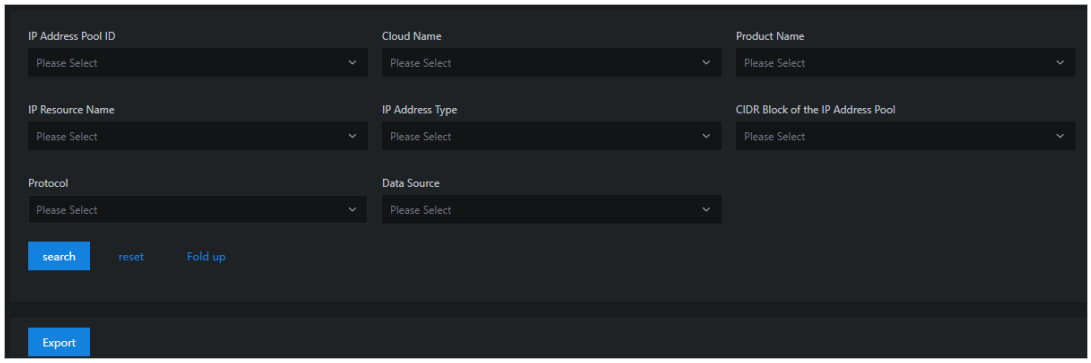

b. (Optional) Click **reset** to clear the filter conditions.

c. (Optional) Click **Fold up** to hide the advanced filter options.

4. Click **Details** in the Actions column to view the details of the IP address pool.

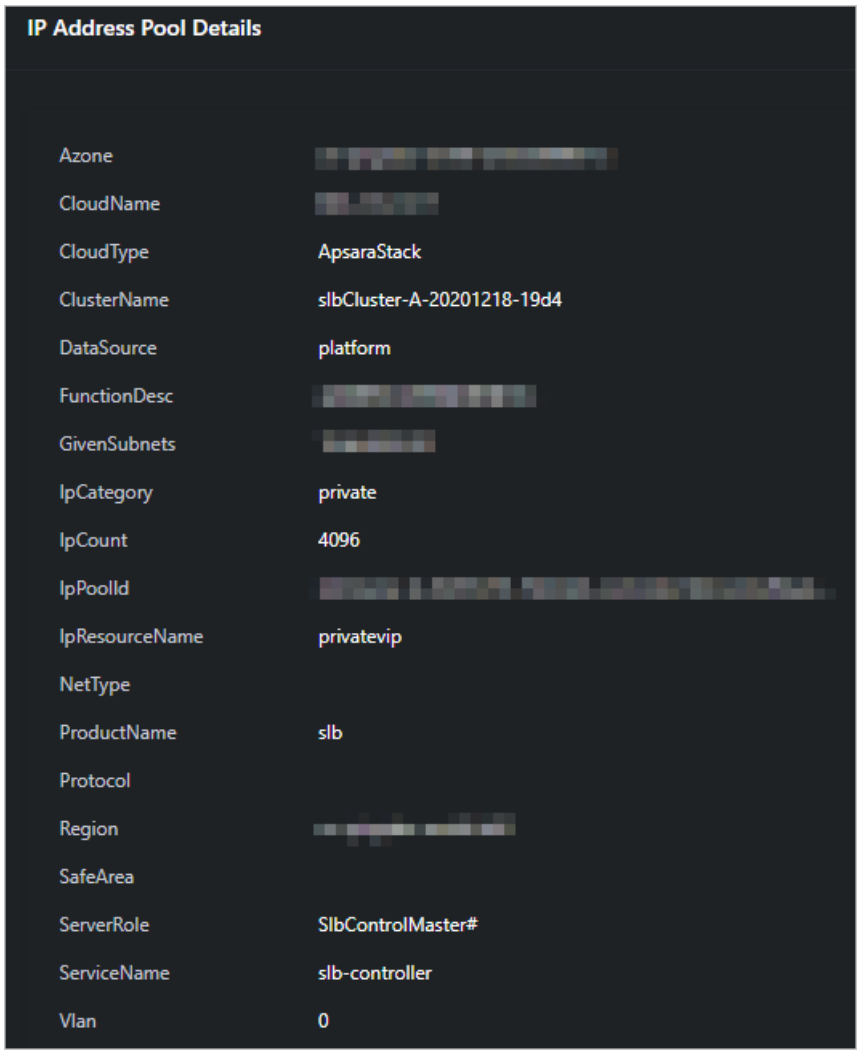

5. Click **Export** to download the information of the IP address pool to your computer.

1.1.4.2.3. Network service provider

1.1.4.2.3.1. View access gateway instances

You can view information of access gateway instances, such as the access gateway name, IBGP role, and creation time, on the Instance Management tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **Instance Management** tab.
5. 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 enter the ID of another region.

6. Click **Display Device List** to view the list of access gateway devices in the current environment.

   > **Note** If new devices are added, click **Scan for New Devices** and then click **Display Device List**.

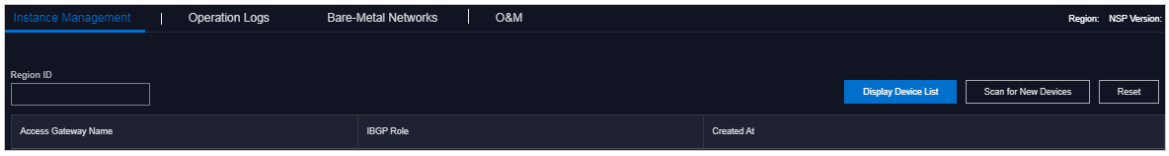

Column	Description
Access Gateway Name	The access gateway name in the current system.
IBGP Role	The role of the access gateway in the environment. Take note of the following items: ○ **RR-Active**: indicates that the role of the current gateway device is RR active device. ○ **Client**: indicates that the role of the current gateway is not RR active device.
Management IP Address	The management IP address of the current HSW vSwitch.
Created At	The time when the current vSwitch began to act as an access gateway instance.
Authorization Status	Indicates whether the access gateway instance is authorized.

1.1.4.2.3.2. View operation logs

You can view the API operation logs of bare metal instances based on your O&M needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **Operation Logs** tab.
5. Set filter conditions such as vSwitch ID, bare metal instance name, access gateway name, and time range, and click Search to search for the operation logs that meet the filter conditions.

   The following table describes some of the filter conditions.

Parameter	Description
VSwitch ID	The ID of the vSwitch when the bare metal instance is applied for or released in the VPC.
Bare Metal Name	The name of the bare metal that was applied for or released in the VPC. The serial number is used to identify the bare metal instance as a unique one in the region.
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 set the filter conditions again.

The following table describes the fields in the query result.

Column	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 instance in the VPC. ○ **add** indicates that a bare metal instance is applied for in the VPC. ○ **del** indicates that a bare metal instance 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	Description
VSwitch ID	The ID of the vSwitch when the bare metal instance 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 instance belongs.
Bare Metal Name	The name of the bare metal instance that is applied for or released in the VPC. To identify the bare metal instance as a unique one in the region, the serial number of the bare metal instance 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.

6. Find an operation log in the search results and click **View Details** in the **Detail** column to view the details of the API operation.

1.1.4.2.3.3. View network information of bare metal instances in a VPC

You can view the information of bare metal instances that are added to a VPC on the Bare-Metal Networks tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **Bare-Metal Networks** tab.

   By default, the network information of bare metal instances in the current system are displayed by page.
5. Configure the filters such as bare metal name, VPC ID, vSwitch ID, VBR ID, BD ID, access gateway name, and time range, and click Search to search for the bare metal instances that meet the filter conditions.

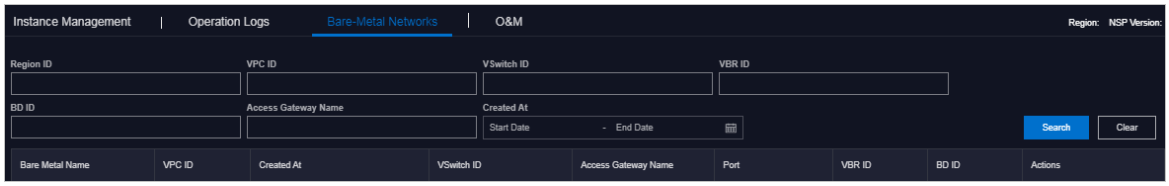

Filter	Description
Bare Metal Name	The name of the bare metal instance that was applied for or released in the VPC. The serial number of the is used to identify the bare metal instances as a unique one in the region.
VPC ID	The ID of the VPC to which the bare metal instances belongs.
Vswitch ID	The ID of the vSwitch to which the bare metal instances belongs.
VBR ID	The VBR ID of the physical connection created on HSW by the VPC to which the bare metal instances belongs.
BD ID	The value of the hardware bridge-domain (BD) to which the bare metal instances is added.
Access Gateway Name	The name of the access gateway to which the bare metal instances belongs.
Created At	The time range within which the bare metal instances is allocated to the VPC.

> **Note** To modify the filter conditions, click **Clear** in the upper-right corner of the tab and configure the filters again.

6. Find a bare metal instances in the search result and click **View Details** in the **Detail** column to view the details of the bare metal instances.

1.1.4.2.3.4. O&M configurations

Apply for a bare metal instance in the VPC

In O&M emergency scenarios, you can use this feature to add the physical port of the access gateway associated with a bare metal instance 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, take note of the following items:

- Typically, you cannot use this feature to apply for a bare metal instance 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 instance to the access gateway port but cannot be used to perform operations on the bare metal instance. To configure the network port IP address and routing information of the bare metal instance, contact the corresponding product team for guidance.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.

2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Apply for Bare Metal in VPC** from the drop-down list in the upper-left corner of the tab.

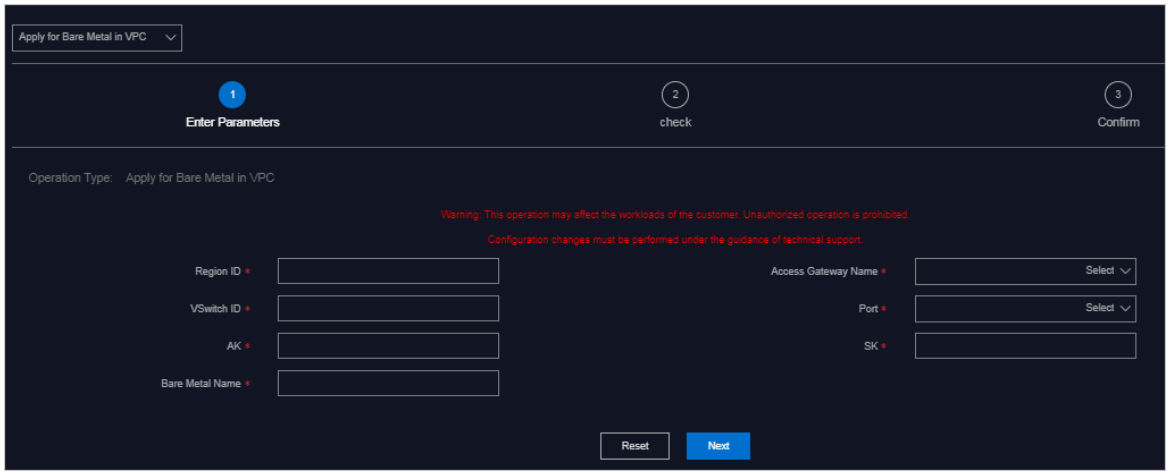

6. Configure the parameters. The following table describes the parameters.

Parameter	Description
Region ID	The name of the region in the current environment.
Access Gateway Name	The name of the access gateway to which the bare metal instance is connected.
Vswitch ID	The ID of the vSwitch to which the bare metal instance is to be added. You can obtain the vSwitch ID from the VPC console.
Port	The port of the access gateway to which the bare metal instance is connected.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained from the **Organizations** page of the Apsara Uni-manager Management Console based on the organization to which the VPC belongs.
Bare Metal Name	The name of the bare metal instance. In this case, enter the serial number of the bare metal instance.

> **Note** If the specified values are incorrect, click **Reset** in the lower part of the tab and set the parameters again.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system starts to push the configurations. After the configurations are pushed, the **Result: Successful** message appears.

After the configurations are pushed, you can search for the bare metal instance based on the bare metal instance name on the **Bare-Metal Networks** tab. If the bare metal instance is displayed, it is added to the VPC.

Release a bare metal instance in the VPC

In O&M emergency scenarios, you can use this feature to disconnect the physical port of the bare metal instance from the VPC.

Prerequisites

Before you use this feature, take note of the following items:

- 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 be used only to disconnect the bare metal instance from the access gateway port but cannot be used to perform operations on the bare metal instance. To configure the network port IP address and routing information of the bare metal instance, contact the corresponding product team for guidance.

Context

Warning This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Release Bare Metal in VPC** from the drop-down list in the upper-left corner of the tab.

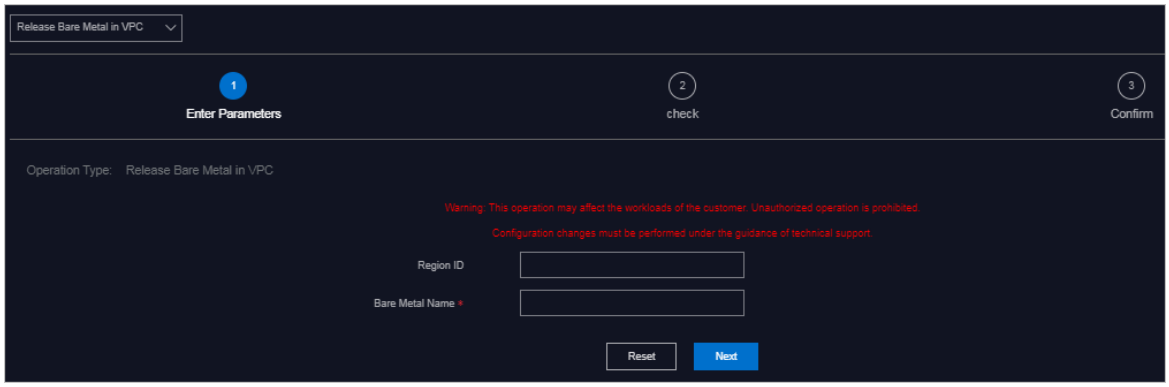

6. Configure the parameters described in the following table.

Parameter	Description
Region ID	The name of the region in the current environment.
Bare Metal Name	The name of the bare metal instance that you want to release. Enter the serial number of the bare metal instance.

> **Note** If the specified values are incorrect, click **Reset** in the lower part of the tab and set the parameters again.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system starts to push the configurations. After the configurations are pushed, the **Result: Successful** message appears.

   After the configurations are pushed, you can search for the bare metal instance based on the bare metal instance name on the **Bare-Metal Networks** tab. If the bare metal instance is not displayed, it is released.

Delete a VPC route table entry

In O&M emergency scenarios, you can use the VPC route table entry deletion feature to delete route table entries that point to the bare metal subnet in the VPC.

Prerequisites

Before you use this feature, take note of the following items:

- Typically, you cannot use this feature to delete a VPC route table entry. This operation is only used for emergency situations.
- You can perform this operation to delete only a single route table entry at a time. To delete multiple route table entries, you must perform this operation multiple times.

Context

> **Warning** This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Delete Route Table Entry** from the drop-down list in the upper-left corner of the tab.

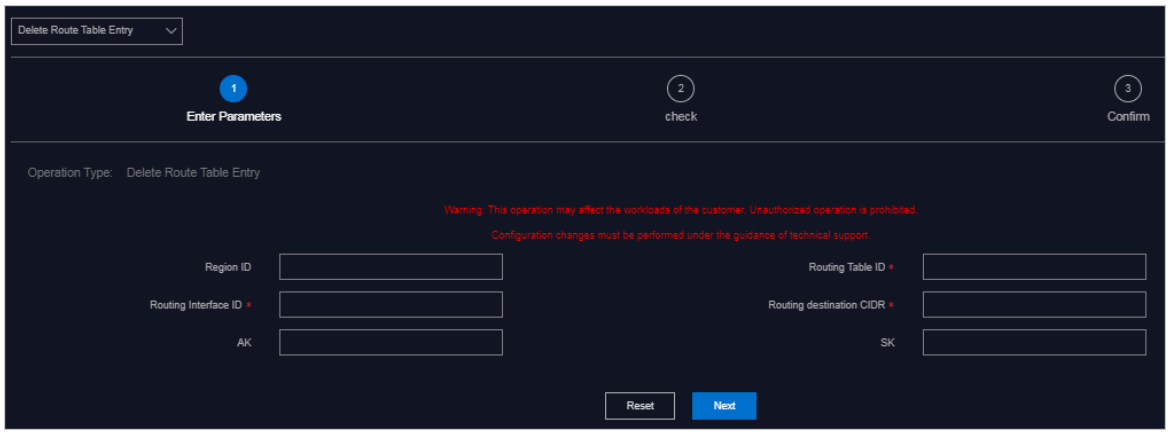

6. Configure the parameters described in the following table.

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*.
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 from the **Organizations** page of the Apsara Uni-manager Management 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 set the parameters again.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system begins to push the configurations. After the configurations have been pushed, the **Result: Successful** message appears.

   After the configurations are pushed, you can log on to the VPC console and view the route table entry of the specified destination CIDR block. If the route table entry is not displayed, it is deleted.
9. (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.

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.

Context

> **Warning** This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.

5. Select **Delete Route Table Entry** from the drop-down list in the upper-left corner of the tab.
6. Configure the parameters described in the following table.

Parameter	Description
Region ID	The name of the region in the current environment.
Routing Table ID	The VBR route table ID. If the bare metal instance 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 instance name, and then click View Details. The VBR Route Table ID in the details is the value of this parameter. If the bare metal instance involved with the VBR is not added to the VPC, specify the bare metal instance name and creation time to search for the 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 instance 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 instance name, and then click View Details. The VBR RI in the details is the value of this parameter. If the bare metal instance involved with the VBR is not added to VPC, specify the bare metal instance name and creation time to search for the 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 Uni-manager Management 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.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system begins to push the configurations. After the configurations are pushed, the **Result: Successful** message appears.

After the configurations are pushed, you can choose **Products > Product List** in the left-side navigation pane and click **Apsara Network Intelligence**. 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.

Delete a VPC router interface

In O&M emergency scenarios, you can use this feature to delete a VPC router interface.

Context

Warning This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Delete Router Interface** from the drop-down list in the upper-left corner of the tab.

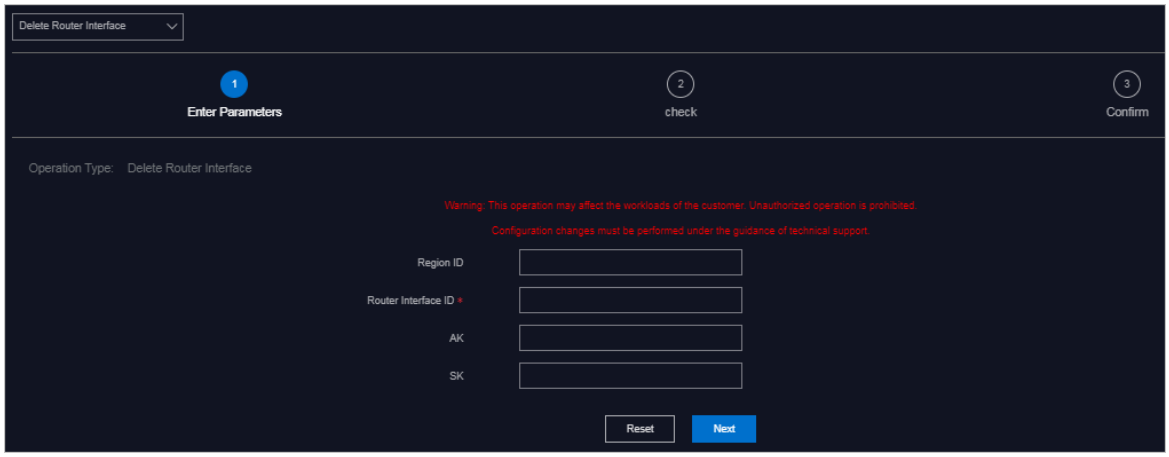

6. Configure the parameters described in the following table.

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 instance name and creation time to search for the 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 from the **Organizations** page of the Apsara Uni-manager Management 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 set the parameters again.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system begins to push the configurations. After the configurations are pushed, the **Result: Successful** message appears.

   After the configurations are pushed, you can choose **Products > Product List** in the left-side navigation pane and 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.

Delete a VBR router interface

In O&M emergency scenarios, you can use this feature to delete a VBR router interface.

Context

> **Warning** This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Delete Router Interface** from the drop-down list in the upper-left corner of the tab.
6. Configure the parameters described in the following table.

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 instance involved with the VBR is added to VPC, you can search for the bare metal instance on the **Bare-Metal Networks** tab based on the bare metal instance name, and then click View Details. The VBR RI in the details is the VBR router interface ID. If the bare metal instance involved with the VBR is not added to VPC, specify the bare metal instance name and creation time to search for the 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.

Parameter	Description
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained from the **Organizations** page of the Apsara Uni-manager Management 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 set the parameters again.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system starts to push the configurations. After the configurations have been pushed, the **Result: Successful** message appears.

   After the configurations are pushed, you can choose **Products > Product List** in the left-side navigation pane and 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.

Delete a VBR

In O&M emergency scenarios, you can use this feature to delete a VBR.

Context

> **Warning** This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Delete VBR** from the drop-down list in the upper-left corner of the tab.

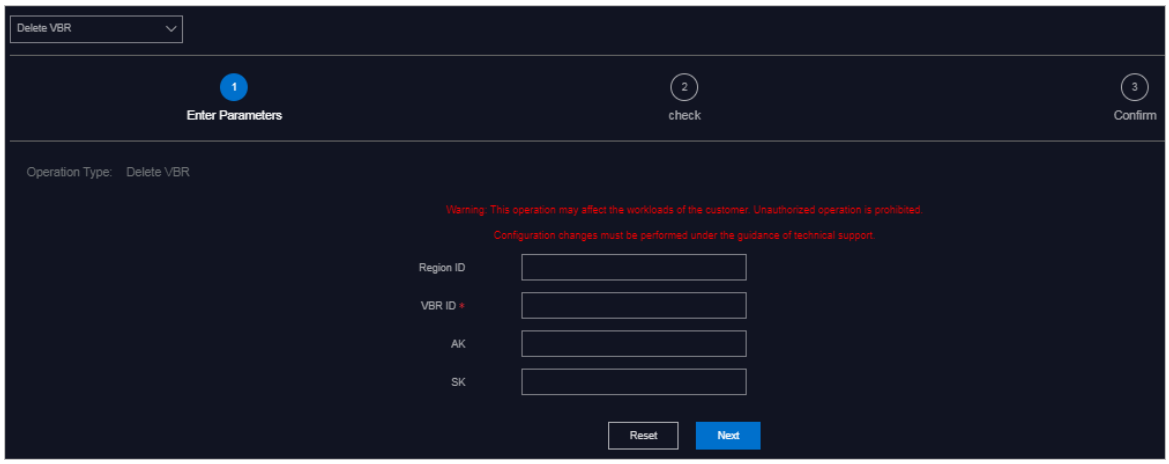

6. Configure the parameters described in the following table.

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 instance name and creation time to search for the 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 Uni-manager Management 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.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system begins to push the configurations. After the configurations are pushed, the **Result: Successful** message appears.

   After the configurations are pushed, you can choose **Products > Product List** in the left-side navigation pane and 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.

Delete a physical connection

In O&M emergency scenarios, you can use this feature to delete a physical connection.

Context

> **Warning** This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Delete Express Connect Circuit** from the drop-down list in the upper-left corner of the tab.

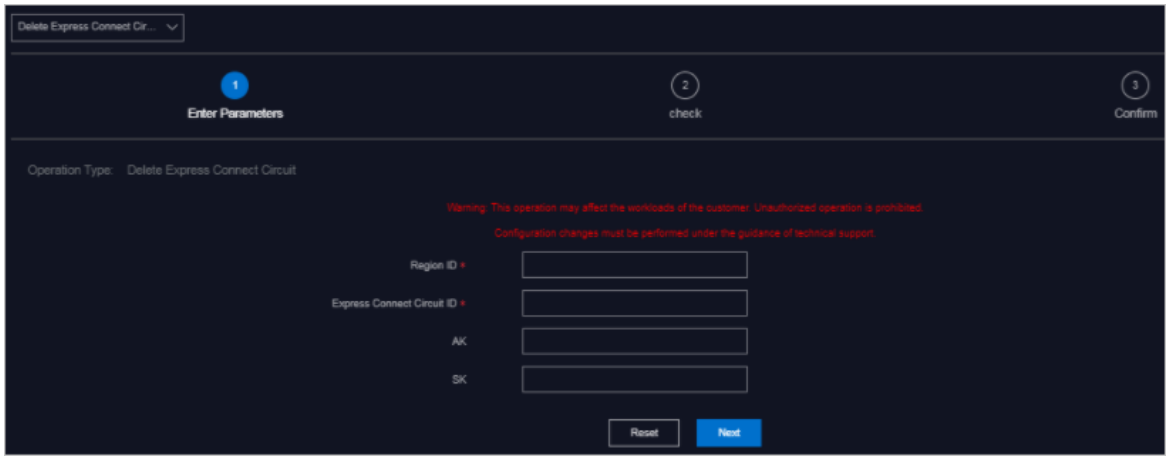

6. Configure the parameters described in the following table.

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 instance name and creation time to search for the 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.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained on the **Organizations** page of the Apsara Uni-manager Management 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.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system begins to push the configurations. After the configurations are pushed, the **Result: Successful** message appears.

   After the configurations are pushed, you can choose **Products > Product List** in the left-side navigation pane and click **Apsara Network Intelligence**. On the homepage of Apsara Network Intelligence, enter the physical connection ID to search for the physical connection. If no search result appears, the physical connection is deleted.

Delete all resources

In O&M emergency scenarios, you can use this feature to delete all resources, including the VPC route table entries, VBR route table entries, VPC router interfaces, VBR router interfaces, VBRs, and physical connections.

Context

Warning This operation is only for emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Delete ALL Resources** from the drop-down list in the upper-left corner of the tab.

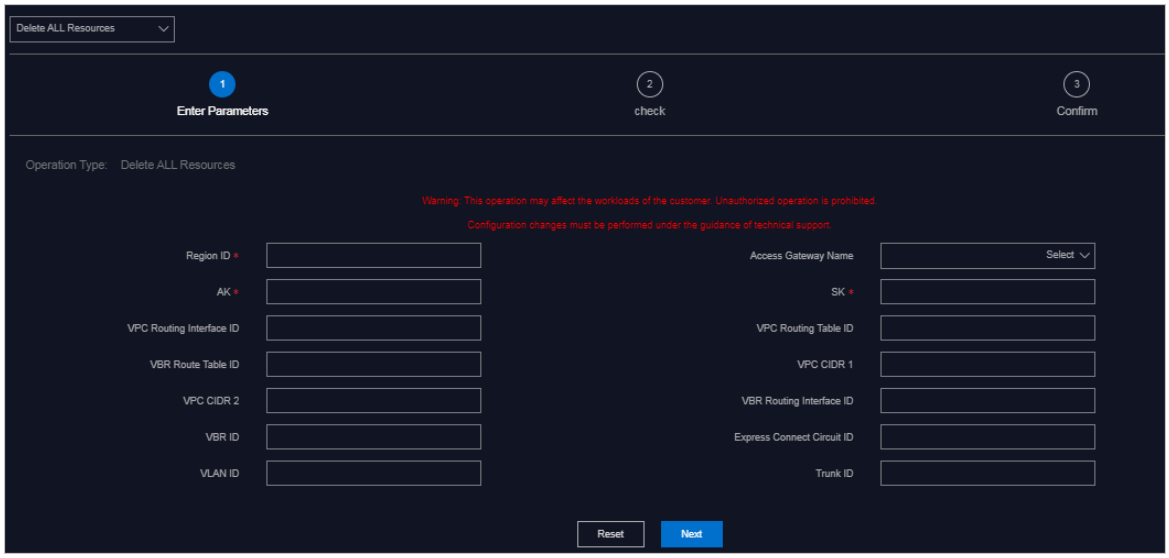

6. Configure the parameters described in the following table.

Parameter	Description
Region ID	The name of the region in the current environment.
Access Gateway Name	The name of the access gateway to which the bare metal instance is connected.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained from the **Organizations** page of the Apsara Uni-manager Management Console based on the organization to which the VBR belongs.
VPC Routing Interface ID	The VPC router interface ID, which can be obtained from the VPC console. For more information about how to obtain the VPC router interface ID, see the *VPC User Guide*.
VPC Routing Table ID	The VPC route table ID, which can be obtained from the VPC console. For more information about how to obtain the VPC route table ID, see the *VPC User Guide*.

Parameter	Description
VBR Route Table ID	The VBR route table ID. If the bare metal instance involved with the VBR has already been added to the VPC, you can search for the bare metal instance on the **Bare-Metal Networks** tab based on the bare metal instance name, and click **View Details**. The VBR Route Table ID in the details is the value of this parameter. If the bare metal instance involved with the VBR is not added to VPC, specify the bare metal instance name and creation time to search for the 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.
CPC CIDR1	The destination CIDR block 1 to which the VPC points, which can be obtained from the VPC console. For more information about how to obtain the VPC CIDR block 1, see the *VPC User Guide*.
VPC CIDR2	The destination CIDR block 2 to which the VPC points, which can be obtained from the VPC console. For more information about how to obtain the VPC CIDR block 2, see the *VPC User Guide*.
VBR Routing Interface ID	The VBR router interface ID. If the bare metal instance involved with the VBR has already been added to the VPC, you can search for the bare metal instance on the **Bare-Metal Networks** tab based on the bare metal instance name, and then click **View Details**. The VBR RI in the details is the value of this parameter. If the bare metal instance involved with the VBR is not added to VPC, specify the bare metal instance name and creation time to search for the 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.
VBR ID	The ID of the VBR to be deleted. On the **Operation Logs** tab, specify the bare metal instance name and creation time to search for the 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 instance name and creation time to search for the 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.
Trunk ID	This parameter is not required.

> **Note** If the specified values are incorrect, click **Reset** in the lower part of the tab and configure the parameters again.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system starts to push the configurations. After the configurations are pushed, the **Result: Successful** message 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.

View the physical connection bandwidth

You can view the physical connection bandwidth when the access gateway is connected to a VPC in the system based on your O&M needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **View Express Connect Bandwidth** from the drop-down list in the upper-left corner.

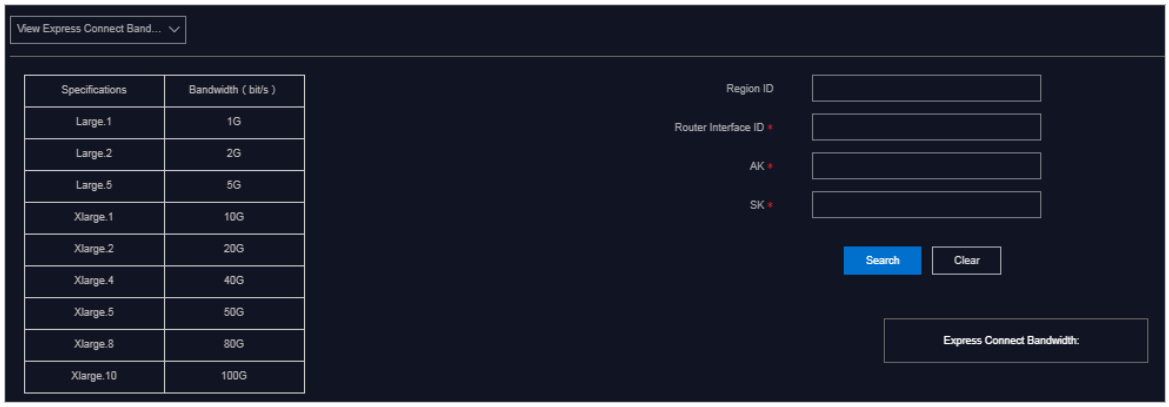

6. Configure the filter conditions and click **Search**.

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 click View Details. VBR RI 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 Uni-manager Management Console based on the organization to which the VBR belongs.

The information about the physical connection bandwidth that meets the filter conditions is displayed.

The bandwidth information describes the specifications of the physical connection bandwidth on the HSW of the current VPC. View the table on the left and obtain the bandwidth (bit/s) based on the specification.

Modify the physical connection bandwidth

In O&M emergency scenarios, you can use this feature to modify the physical connection bandwidth.

Context

Warning This operation is used only in emergency situations and must be performed under the guidance of technical personnel. Otherwise, normal business operations may be affected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **Change Express Connect** from the drop-down list in the upper-left corner.

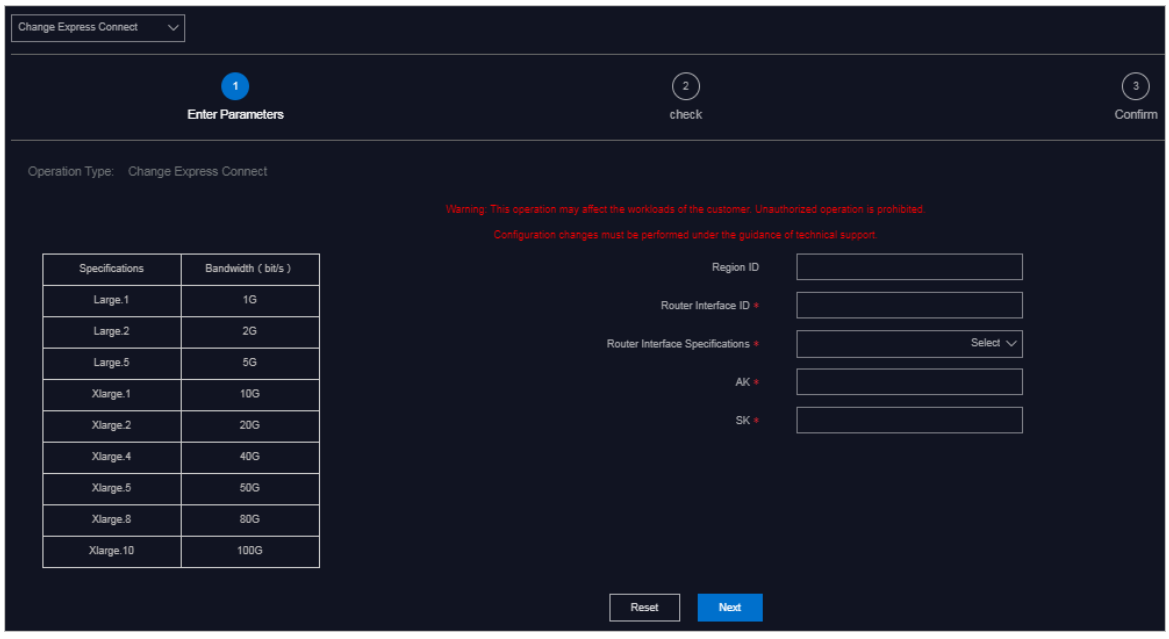

6. Configure the parameters described in the following table.

Parameter	Description
Region ID	The name of the region in the current environment.
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 instance. Click **View Details** and the VBR RI in the details is the value of this parameter.
Router Interface Specifications	The specification of the physical connection bandwidth.
AK and **SK**	The organization AccessKey ID and AccessKey secret, which can be obtained from the **Organizations** page of the Apsara Uni-manager Management 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.

7. Click **Next**.
8. Check the information. If the information is correct, click **Confirm**.

   The system starts to push the configurations. After the configurations have been pushed, the **Result: Successful** message appears.

   After the configurations are pushed, check whether the physical connection bandwidth has been modified. For more information, see View the physical connection bandwidth.

View BD usage

You can view BD usage to learn about the BD configuration distribution in a timely manner.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **View BD Usage** from the drop-down list in the upper-left corner.
6. Configure the filter conditions and click **Search**.

View BM VPN usage

You can view the information of the BM VPN that is assigned to the access gateway instance.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.

4. Click the **O&M** tab.
5. Specify **Access Gateway Name**, **vxlan id**, and **BM VPN Name**, select a state from the **Status** drop-down list, and then click **Search** to view the BM VPNs assigned to all the HSW vSwitches.

View trunk usage

You can view trunk usage to learn about the usage of hardware ports in a timely manner.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Network**.
3. In the left-side navigation pane, click **Network Service Provider**.
4. Click the **O&M** tab.
5. Select **View Trunk Usage** from the drop-down list in the upper-left corner.

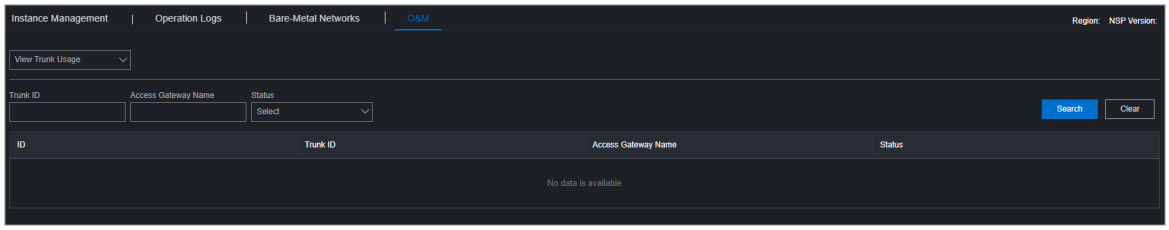

6. Specify the trunk ID and access gateway name, select a trunk state, and then click **Search**. The Trunk Status options include Idle, Used, Creating, and Deleting.

   Set **Trunk ID** 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 the filter conditions, click **Clear** in the upper-right corner and set the filter conditions again.

1.1.4.3. Data centers

Operations personnel can monitor and view the physical servers where products are located.

1.1.4.3.1. View physical server information

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

Go to the Data Centers page

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Data Centers**.

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

Product tab

1. On the Product tab, perform the following operations to view physical server information:
   - Expand the left-side hierarchy tree by selecting a region, a product, and a 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 a 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.
- Find a product and click **Details** in the **Operation** 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 between the tabs to view the monitoring details and alert information.

  Monitoring details include the CPU utilization, system load, disk usage, memory usage, network throughput, and disk I/O. When you view the monitoring details, you can select a monitoring item in the upper-right corner of each monitoring graph and then select a time range to view the monitoring value within the time range.

  In the upper-right corner of the CPU utilization, system load, disk usage, memory usage, network throughput, and disk I/O 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 computer.
  - Click the icon to manually refresh the monitoring data.
  - Click the icon. The icon turns green. The system refreshes the monitoring data at 10-second intervals. To disable the auto-refresh feature, click the icon again.

Server tab

1. Click the **Server** tab.
2. On the Server tab, perform the following operations to view the physical server list:
   - Expand the left-side hierarchy tree by selecting an IDC and a rack in sequence to view the physical server list in a rack.
   - Enter a rack name in the left-side search box and press the Enter key or click the icon to search for and view the list of all physical servers in the rack.
3. To view the details of a physical server, enter the hostname, IP address, device role, or serial number (SN) in the right-side search box and press the Enter key.
4. Find the physical server that you want to view and click **Details** in the **Operation** column. 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 details and alert information.

   Monitoring details include the CPU utilization, system load, disk usage, memory usage, network throughput, and disk I/O. When you view the monitoring details, you can select a monitoring item in the upper-right corner of each monitoring graph and then select a time range to view the monitoring value within the time range.

   In the upper-right corner of the CPU utilization, system load, disk usage, memory usage, network throughput, and disk I/O 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 computer.
- Click the  icon to manually refresh the monitoring data.
- Click the  icon. The icon turns green. The system refreshes the monitoring data at 10-second intervals. To disable the auto-refresh feature, click the icon again.

Physical View of Device tab

1. Click the **Physical View of Device** tab.
2. On the **Physical View of Device** tab, expand the left-side hierarchy tree by selecting an IDC and a rack in sequence to view the corresponding rack information on the right. In addition, the rack details panel 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 running normally.

   In the upper-right corner, you can view the alert legend. By default, the check box on 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 panel. Clear the check box on 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 panel.
3. To view the details of a physical server, perform the following operations:
   i. Find the physical server that you want to view in the left-side hierarchy tree or right-side rack graph of the tab.
   ii. In the rack details panel that appears, click the color block corresponding to 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 details and alert information.

Monitoring details include the CPU utilization, system load, disk usage, memory usage, network throughput, and disk I/O. When you view the monitoring details, you can select a monitoring item in the upper-right corner of each monitoring graph and then select a time range to view the monitoring value within the time range.

In the upper-right corner of the CPU utilization, system load, disk usage, memory usage, network throughput, and disk I/O 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 computer.
- Click the icon to manually refresh the monitoring data.
- Click the icon. The icon turns green. The system refreshes the monitoring data at 10-second intervals. To disable the auto-refresh feature, click the icon again.

1.1.4.3.2. Export physical server information

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

Go to the Data Centers page

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Data Centers**.

Product tab

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

1. 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 computer.

Server or Physical View of Device tab

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

1. Click the **Server** or the **Physical View of Device** tab.
2. In the upper-right corner of the **Server** tab, click the icon to export all the information of

physical servers to your computer.

3. In the upper part of the **Physical View of Device** tab, click the icon to export all the information of physical servers to your computer.

1.1.4.3.3. Security operations

You can log on to a virtual machine and perform remote operations.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Data Centers**.

   The **Product** tab appears.
3. Find a virtual machine and click **Safe operation and maintenance** in the **Operation** column to log on to the virtual machine.
4. Perform remote operations on the virtual machine.
   i. Enter Linux commands in the CLI window to perform related operations.

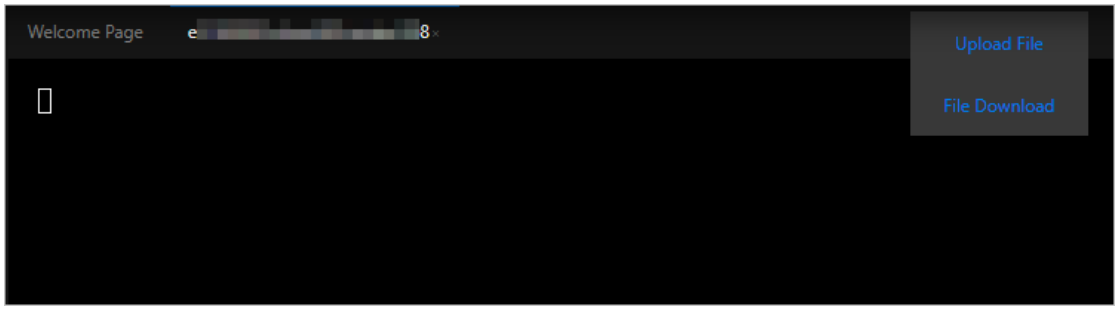

   ii. Click **Upload File** in the CLI window. The **Upload File** dialog box appears. You can upload an object in one of the following ways:
   - Click the dotted box. In the dialog box that appears, select the file to be uploaded, click **Open**. Click **Upload** in the **Upload File** dialog box.
   - Drag the file to the dotted box and then click **Upload**.

   iii. Click **File Download**. The **File Download** dialog box appears. Set **File Name** and **File Directory** and then click **Download** to download the file to the default download directory of the local browser.

> **Note** The uploaded and downloaded files cannot exceed 200 MB in size.

1.1.4.4. Full stack

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.

1.1.4.4.1. 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 the ilogtail reporting feature is enabled on each SLB service node, 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **Full Stack Log Monitoring**.
4. Click the **ECS**, **SLB**, or **All in ECS** tab.
5. Enter a keyword in the **Query** search box, select a time range in **Time**, and then click **Search**.

   > **Note** You can enter a string in the Query search box as the filter condition, such as the instance ID, request ID, or the keyword **error**.

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

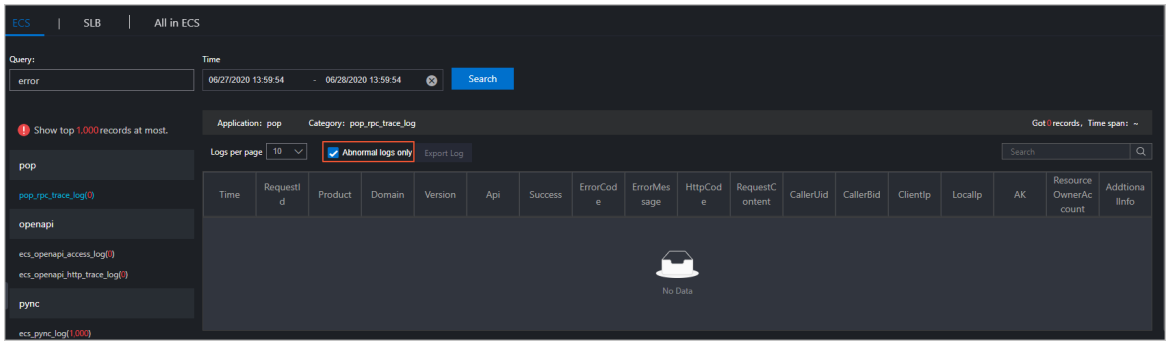

7. Select **Abnormal logs only** to view only the exceptional logs.

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

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

1.1.4.4.2. SLA console

The SLA console provides availability monitoring, forward diagnostics, and reverse influence analysis features and allows you to view various monitoring data.

1.1.4.4.2.1. Product availability dashboard

The Product Availability Dashboard page displays the product availability sequence diagram, product fault diagnosis, and product dependency information.

Access Product Availability Dashboard

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.

   The **Product Availability Dashboard** page appears.

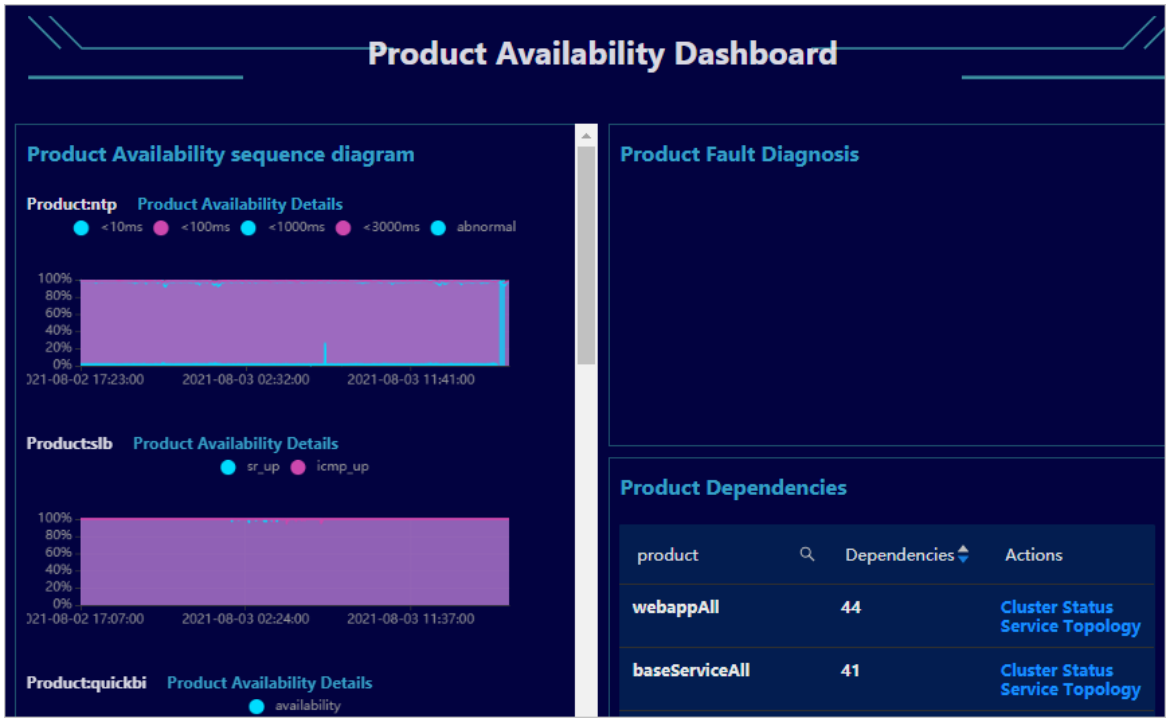

View availability sequence diagrams

1. View availability sequence diagrams.

   Availability sequence diagrams show the availability details of products and the base in the SLA console. For different products, the dots above sequence diagram represent different meanings. For more information, see View product availability.

2. Move the pointer over the sequence diagram of a product. The system displays the availability details of the product.
3. (Optional)After you click a dot above the sequence diagram, the information that the dot indicates disappears in the sequence diagram.

> **Note** If you want to view the information that the dot indicates in the sequence diagram, click the dot again.

4. Click **Product Availability Details**. The **Product Availability** or Base Availability page appears. For more information, see View product availability.

View product fault diagnostics

1. In the **Product Fault Diagnosis** section, find a product and click **Details** next to the product.
2. Check the states of components of the product on the page. They are in one of the two states: **Normal** and **Exceptions**.
3. Move the pointer over the **Exceptions** area of a component. A dialog box appears, displaying the alerts for the exceptions.
4. (Optional)After you click **Refresh** on the page, the system updates the diagnostic details of the product.

View dependencies of a product

1. In the **Product Dependencies** section, find a product and click **Cluster Status** in the **Actions** column to view the cluster status.

> **Note** A cluster with its status in red is exceptional. A cluster with its status in green is normal.

2. Find a product and click **Service Topology** in the **Actions** column. The **Product** tab on the Forward Diagnostics & Demarcation page for the product appears. For more information, see View diagnostic information of products.

1.1.4.4.2.2. SLA platform

The SLA platform allows you to view product availability, base availability, and diagnostic results.

View product availability

The Product Availability page contains the Instance and Service tabs and allows you to view product availability details and download availability reports.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **SLA Platform > Product Availability**.
5. Click the **Instance** or **Service** tab. On the tab, click the product that you want to view.
6. (Optional)Select a time range and click **Download All Reports** to export the availability details of the base products within the time range.
7. Find the product and click **View Details** in the **Actions** column.
8. To view the availability details, you must perform the following operations:

i. (Optional)Select the gradient and time range from the **Gradient** and **Time Range** drop-down lists, and then click **Search**.

ii. In the **SLI Details** section, move the pointer over the sequence diagram to view the availability data of the current point in time.

The dots above the sequence diagram are the legend. After you click a dot, the information that the dot indicates disappears in the sequence diagram. The following table describes the meanings of the dots for different products.

Product	SLI description
ECS	■ vm_up: the ratio of the period when the VM is up to the period when the VM is connected. ■ io_nohang: the ratio of the period when the VM does not have IO hang events to the period when the VM is connected.
SLB	■ icmp_up: the SLI ICMP is pinged. ■ sr_up: the ratio of the period when the three SRs on which SLB depends reach the desired state to the period when they do not reach the desired state.
OSS	2xx: the success rate of requests.
Tablestore	availability: the service availability rate.
RDS	availability: the service availability rate.
Redis	availability: the service availability rate.
MongoDB	availability: the service availability rate.
MySQL 3.0	availability: the service availability rate.
MySQL 2.0	availability: the service availability rate.
PostgreSQL	availability: the service availability rate.
MaxCompute	availability: the service availability rate.
Realtime Compute	availability: the service availability rate.
Dataworks	availability: the service availability rate.
Quick BI	availability: the service availability rate.

Product	SLI description
ODPS	■ availability: the service availability rate. ■ sr_up: the ratio of the period when the three SRs on which ODPS depends reach the desired state to the period when they do not reach the desired state.
Elasticsearch	availability: the service availability rate.
PAI	availability: the service availability rate.
DataQ	availability: the service availability rate.
DataHub	availability: the service availability rate.
EDAS	■ CoDeploy: the success rate of application deployments. ■ CoStart: the success rate of application startups. ■ CoStop: the success rate of application shutdowns. ■ CoScaling: the success rate of auto scaling operations. ■ CoCreateApp: the success rate of application creations. ■ CoRollback: the success rate of application rollbacks. ■ CoReset: the success rate of application resets. ■ CoScaleOut: the success rate of application scale-outs. ■ CoScaleIn: the success rate of application scale-ins. ■ CoUpdateContainer: the success rate of container upgrades. ■ CoDeleteApp: the success rate of delete application operations. ■ CoChangeGroup: the success rate of change group operations. ■ CoBindSlb: the success rate of bind SLB operations. ■ CoModifyConfig: the success rate of modify configuration operations.
CSB	availability: the service availability rate.

Product	SLI description
MQ	■ send_success: the success rate of send operations. ■ subscribe_success: the success rate of subscriptions.
SLS	2xx: the success rate of requests.
API	2xx: the success rate of requests.
KMS	2xx: the success rate of requests.

iii. In the **Unavailable Events** section, the list contains the Instance ID, Start time, End time, State, Note, and Label columns.

View base availability

The Base Availability page contains the Instance and Service tabs and allows you to view base availability details and download availability reports.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **SLA Platform > Base Availability**.
5. Click the **Instance** or **Service** tab. On the **Service** tab, click the product that you want to view. On the **Instance** tab, you do not need to select a product.
6. (Optional)Select a time range and click **Download All Reports** to export the availability details of the base products within the time range.
7. Find the base product and click **View Details** in the **Actions** column.
8. To view the availability details, you must perform the following operations:
   i. (Optional)Enter the instance ID in the **Instance ID** field, select a time range and gradient, and then click **Search**.

ii. In the **SLI Details** section, view the SLI sequence diagram.

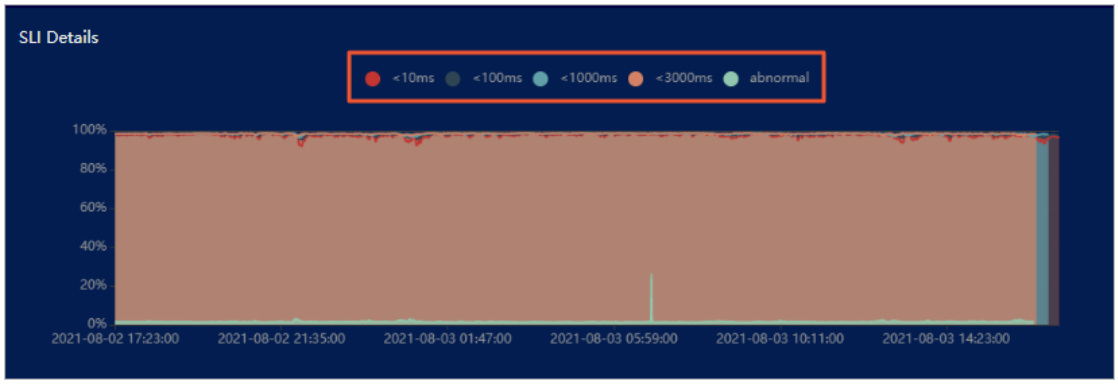

> **Note** The dots above the sequence diagram are the legend. For different base products, the dots above sequence diagram represent different meanings. After you click a dot, the information that the dot indicates disappears in the sequence diagram. The meanings of the dots for different base products:
>
> - **ntp: abnormal** represents the SLI with abnormal clock synchronization. **<10ms**, **<100ms**, **<1000ms**, and **<3000ms** represent the SLIs in different clock difference intervals.
> - **dns: dns** indicates the success rate of the DNS probe.

iii. In the **Unavailable Events** section, the list contains the Instance ID, Start time, End time, State, Note, and Label columns.

> **Note** Values of the State column: **resolved** and **firing**. **resolved** indicates that an event has been resolved and **firing** indicates that an event is in progress.

View availability reports

The Availability Report page allows you to view the service availability of all instances in a region by month or year.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **SLA Platform > Availability Report**.
5. Select **By-month** or **By-year** from the drop-down list in the upper-left corner and select the month or year.

> Note
>
> - Region: The region to which the product is deployed.
> - Total Instances (Services): the total number of instances or services.
> - Total SLA-incompliant Instances (Services): the total number of instances or services whose service availability is less than 100%.
>
> You can click the icon in a column to change the sorting order.

6. Find the product and click **View Details** in the **View Details** column. The system displays the report on the service availability details of the product. The details list contains the following columns: Product Code, Month, Region, Instance ID, and Service Availability.

1.1.4.4.2.3. End-to-end diagnostics & demarcation

You can view and diagnose the status and exception causes of products deployed on the cloud, and analyze root causes of exceptions.

Forward diagnostic and demarcation information of products

You can view the number of resources on the cloud by product, view the dependencies of specified products, services, and server roles, and view diagnostic information of products such as status and root cause analysis.

View diagnostic information of products

You can view the diagnostic information of products, including exceptional products, exceptional services, abnormal service roles, dependencies, and root cause analysis.

Prerequisites

The Apsara Uni-manager Operations Console and the NetworkBaseServiceCluster for BaseServiceAll are deployed and reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **End-to-end Diagnostics > End-to-end Diagnostics & Demarcation**.

   The **Product** tab appears.
5. You can view the total number of products, total number of services, and total number of server roles in the upper part of the tab.
6. View diagnostic information of a product.

i. Move the pointer over a product. The system displays the dependencies between the product and other products.

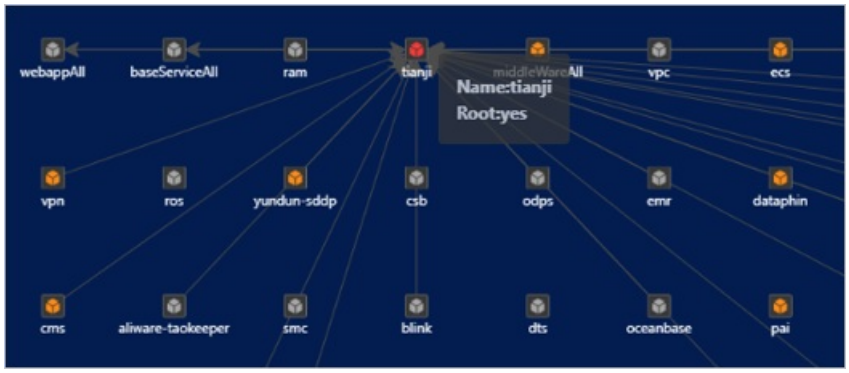

Note

- By default, the system displays the dependencies of all products.
- The product that the arrow points is the dependent product.
- **Exceptional**: indicates that the product has exceptions.
- **Exceptional - Root**: indicates that the product is the root cause for other exceptional products.
- You can click the style buttons or icons in the upper-left corner to change the display style of the topology.
- Fuzzy search is supported.

ii. Right-click a product and select **Root Cause Analysis**. The topology of components of the product is displayed. You can find the causes for internal exceptions.

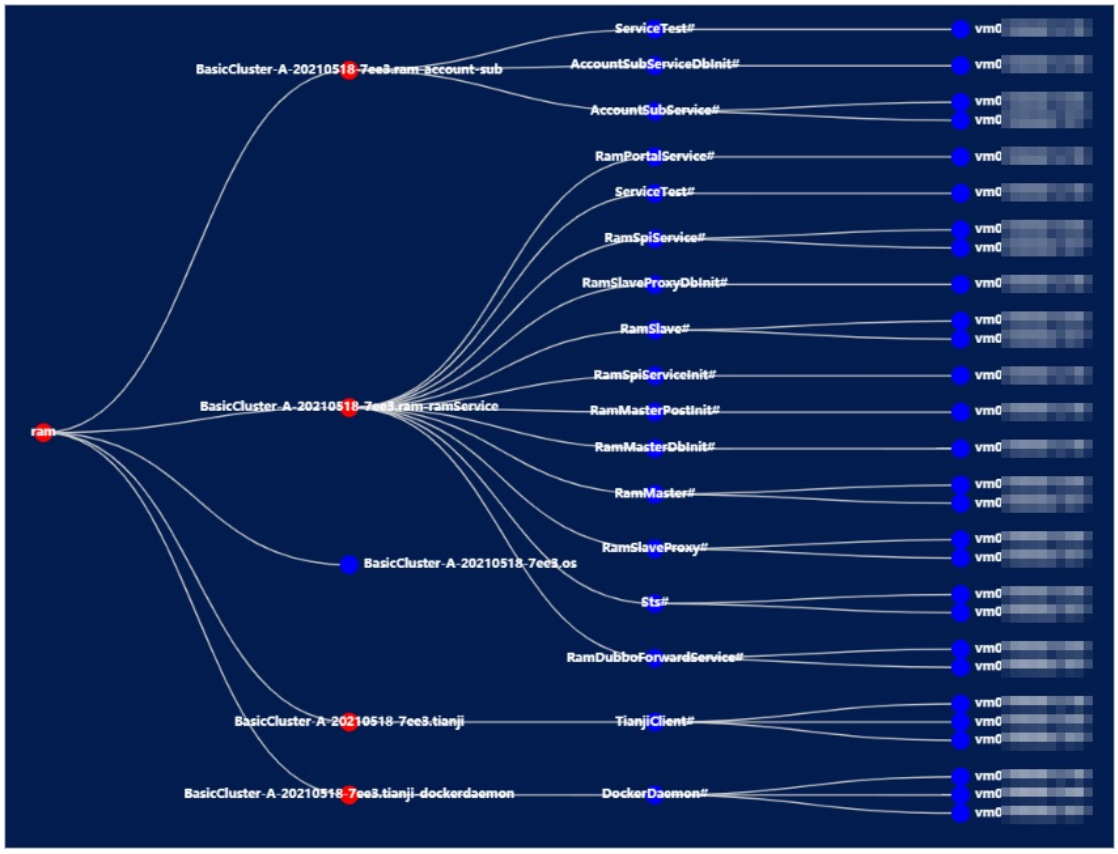

> **Note**
>
> - The four columns in the figure indicate the product, services, server roles, and hosts from left to right.
> - A red circle indicates an exceptional item, and a blue circle indicates a normal item.

iii. Click a red circle. A dialog box appears to display relevant component alerts and dependency alerts.

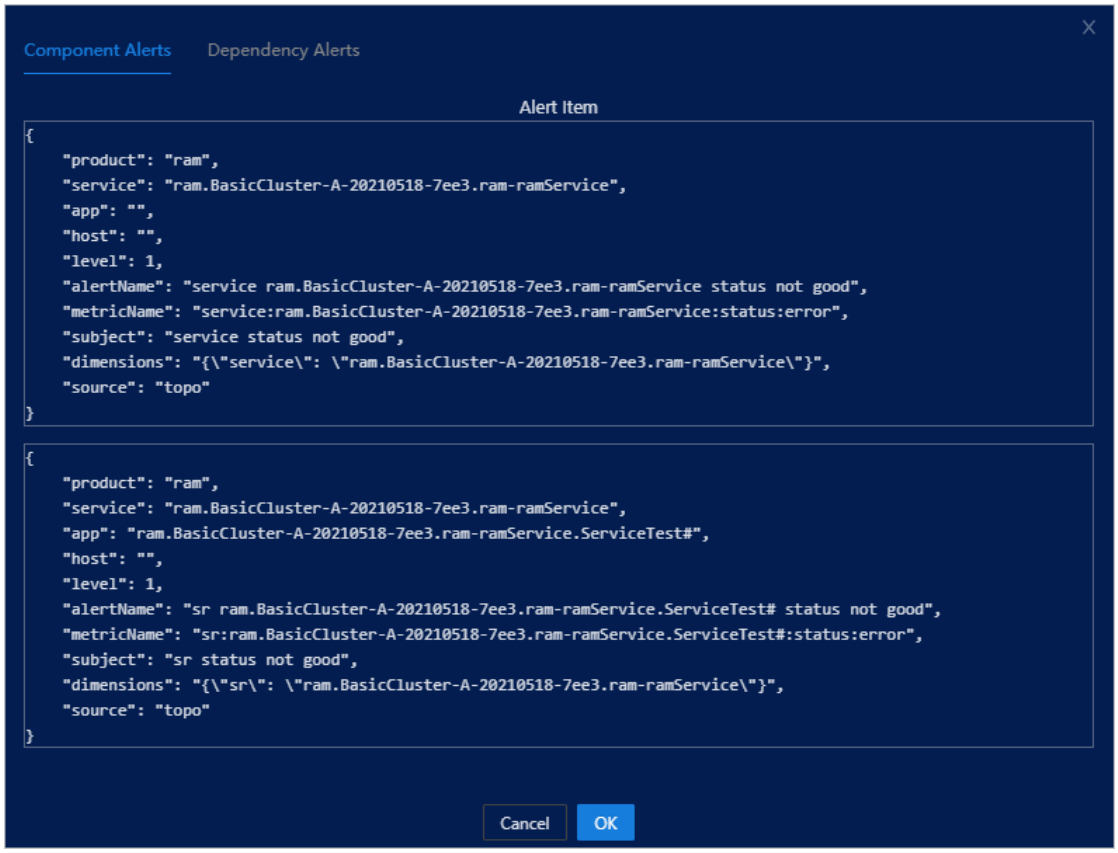

7. View diagnosis information of a service.
   i. Select the product from the **Product** drop-down list or enter the product name in the field to find the product. The system displays the status and dependencies of all services in the product.
   ii. Move the pointer over a service. The system displays the dependencies between the service and other services.

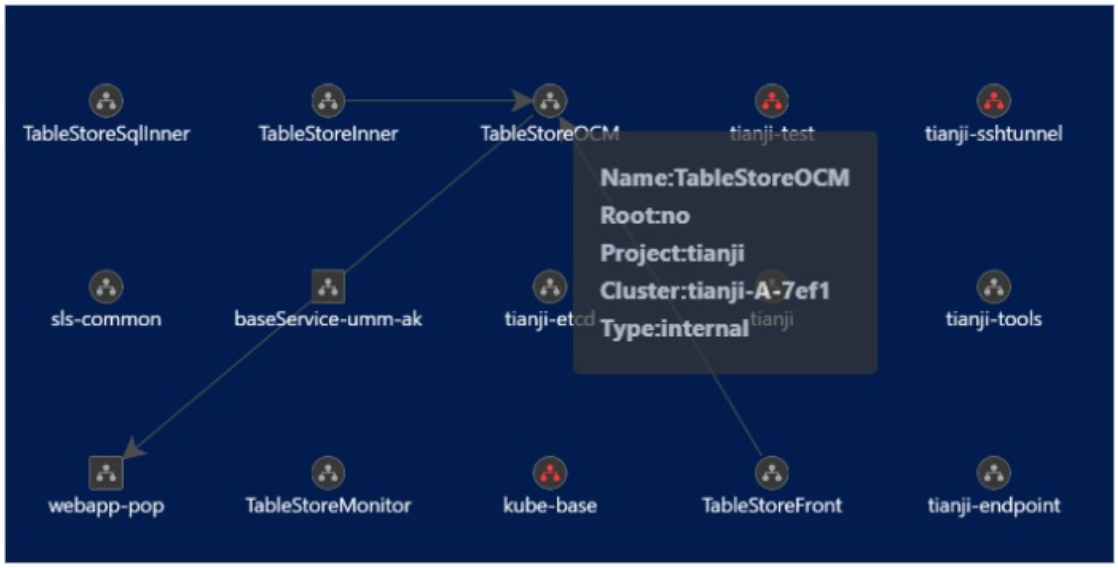

iii. Right-click a service and select **Root Cause Analysis**. The topology of components of the service is displayed. You can find the causes for internal exceptions.

iv. Right-click a service and select **Resources**. The system displays a list of the resources related to the service. The list contains the following columns: type, cluster, service, and details.

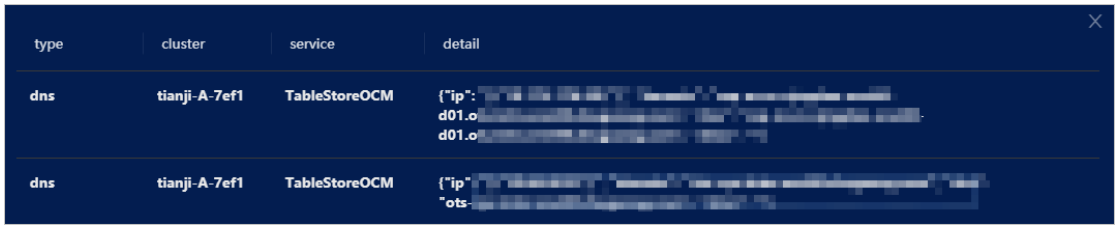

8. View diagnosis information of a server role

i. Click the service or select the service from the **Service** drop-down list or enter the service name in the field to find the service. The system displays the status and dependencies of all server roles in the service.

ii. Move the pointer over a server role. The system displays the dependencies between the server role and other server roles.

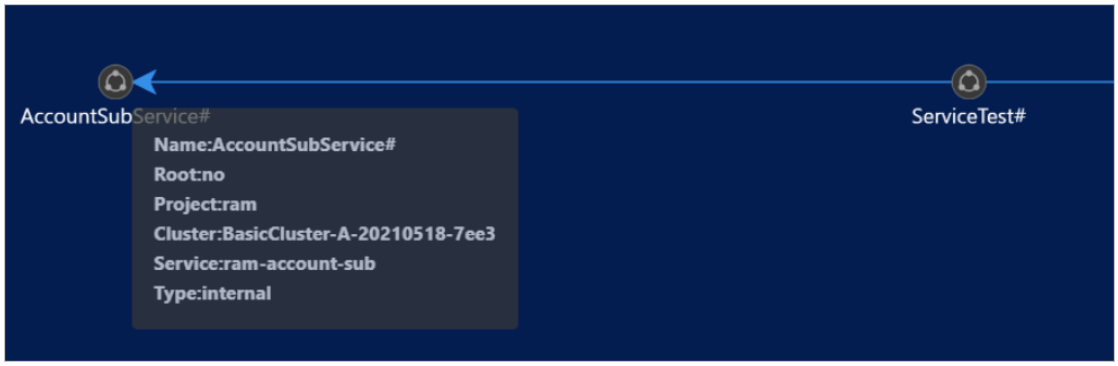

iii. Right-click a server role and select **Root Cause Analysis**. The topology of components of the server role is displayed. You can find the causes for internal exceptions.

View physical server information

You can view the information about the physical servers where the server roles are deployed.

Prerequisites

The Apsara Uni-manager Operations Console and the NetworkBaseServiceCluster for BaseServiceAll are deployed and reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **End-to-end Diagnostics > End-to-end Diagnostics & Demarcation**.

   The **Product** tab appears.
5. Select the product, service, and server role from the **Product**, **Service**, and **ServerRole** drop-down lists or enter values in these fields to find the server role. The system displays the physical

server information of the server role.

> **Note**
>
> - The physical server information includes the cluster, service, server role, machine (virtual machine or physical server) ID (**id**), machine IP address (**ip**), physical server ID (**nc_id**), and physical server IP address (**nc_ip**).
> - Fuzzy search is supported.

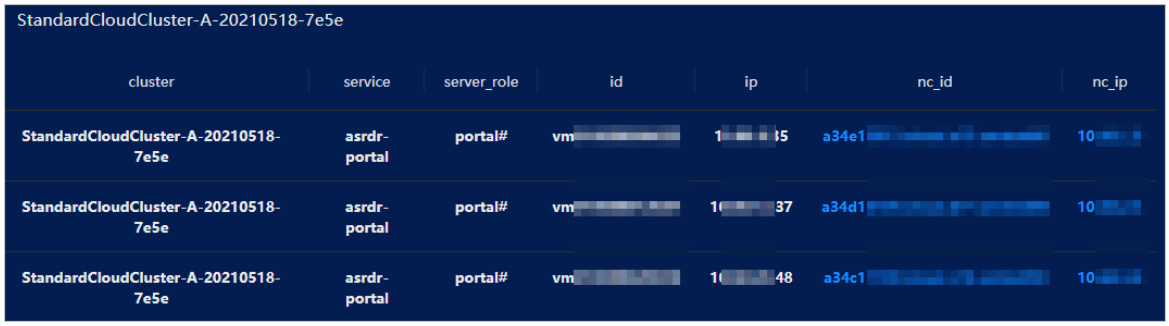

6. Click the ID or IP address of the physical server or select the ID from the **NC** drop-down list or enter the ID in the field to find the physical server. The system displays the network topology information of the physical server.

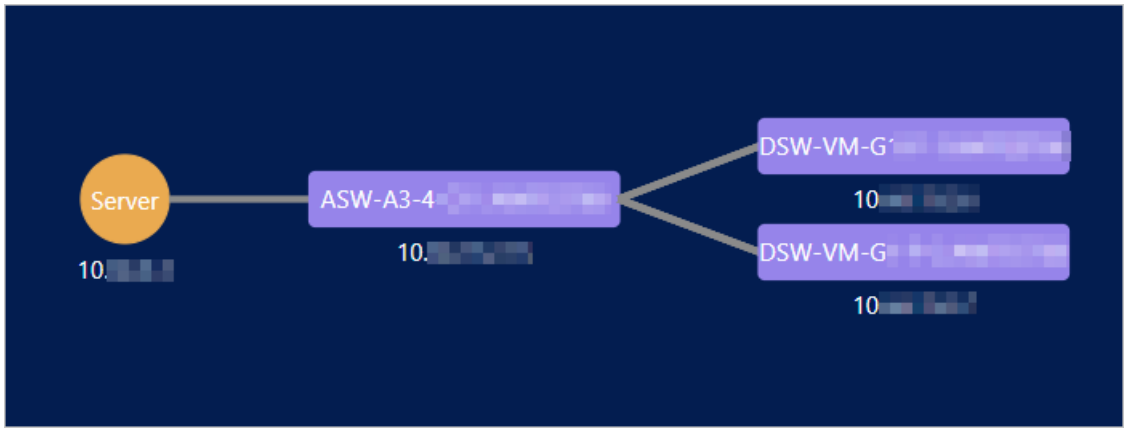

7. Click a link in the preceding figure. The system displays the port mapping information of the link, including the local device, local port, remote device, and remote port.

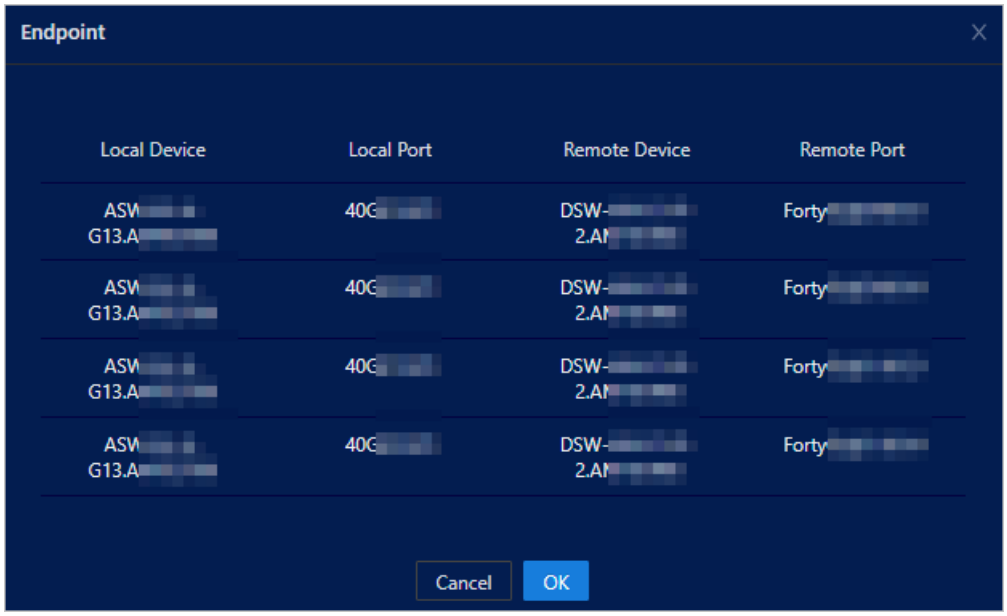

Quick diagnostics

This feature is used to diagnose products deployed on the cloud and view their diagnostic information.

Prerequisites

The Apsara Uni-manager Operations Console and the NetworkBaseServiceCluster for BaseServiceAll are deployed and reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **End-to-end Diagnostics > End-to-end Diagnostics & Demarcation**.

   The **Product** tab appears.
5. In the upper-right corner of the page, click **Quick Diagnostics**. A dialog box appears to show the diagnostic results, including exceptional products and services.

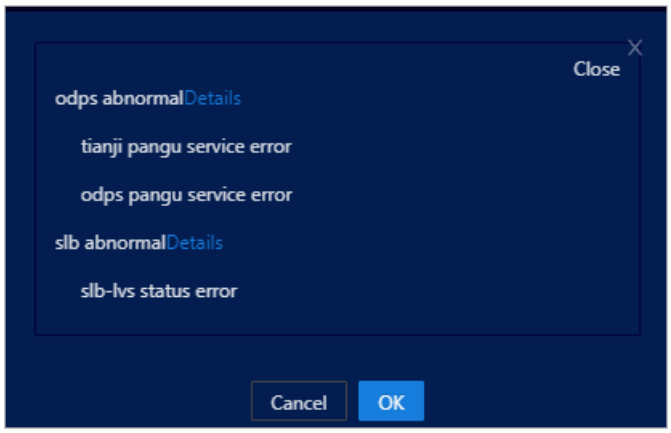

6. Click **Details** next to an exceptional product. A dialog box appears and shows the diagnostic

details of the product.

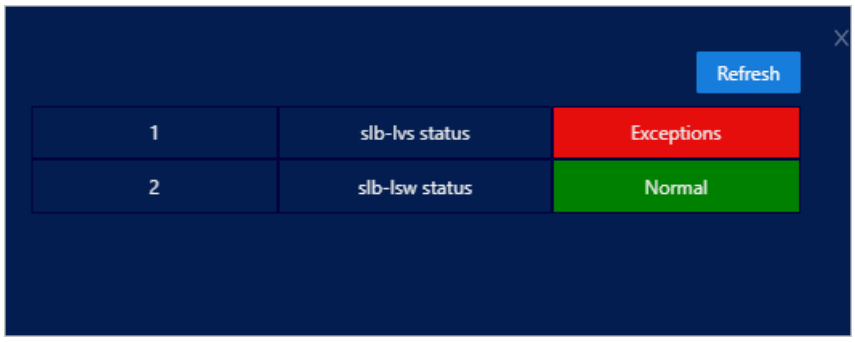

Note The second column is the name of the exceptional product, service, server role, or physical server. If you move the pointer over a red area, the alert information is displayed.

7. (Optional)After you click **Refresh** in the upper-right corner, the system updates the diagnostic details of the product.

Forward diagnostic and demarcation information of resources

You can view the number of resources on the cloud by resource, and view diagnostic information of specified resources such as services, service status, dependencies, and root cause analysis.

Prerequisites

The Apsara Uni-manager Operations Console and the NetworkBaseServiceCluster for BaseServiceAll are deployed and reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **End-to-end Diagnostics > End-to-end Diagnostics & Demarcation**.
5. Click the **Resource** tab.
6. You can view the total number of various resources in the upper part of the tab.
7. Select the resource type and resource ID from the two drop-down lists next to **Resource** or enter values in these fields and then click **Search**. The system displays the services that are related to the specified resource in green, the dependencies between the specified resource and other services, and the diagnostic information of the specified resource.

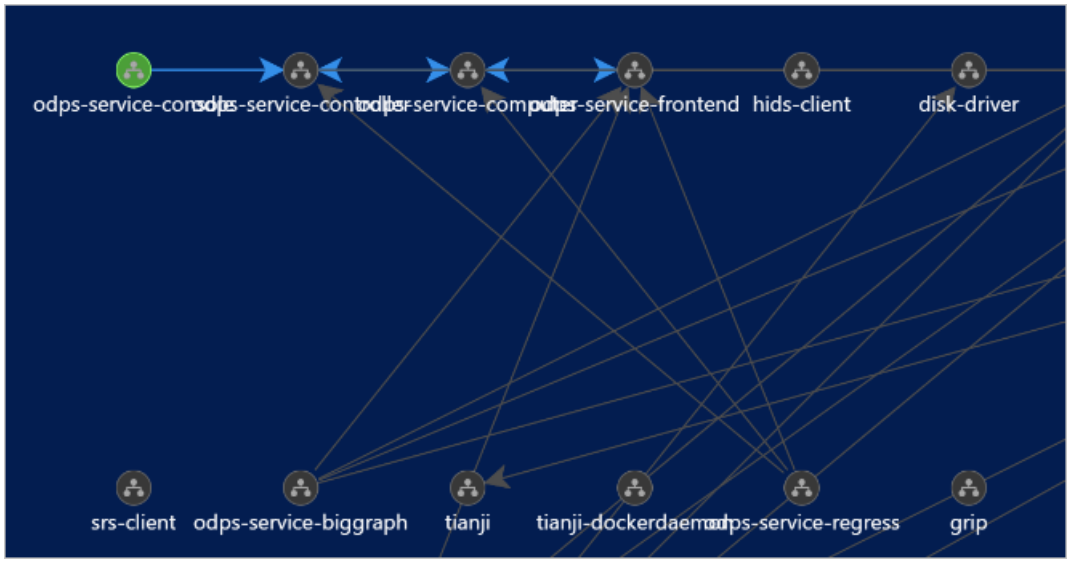

> **Note**
> - Move the pointer over a service. The system displays the dependencies between the service and other services.
> - If you select **vip** or **dns** from the first drop-down list next to the Resource field, the attached backend services are displayed.
> - If you select **db** from the first drop-down list next to the Resource field, the services that use the **db** instance are displayed.
> - Click **Reset** to clear the value that you set.
> - Fuzzy search is supported.

8. You can view detailed diagnostic information and physical server information about services and server roles. For more information, see View diagnostic information of products and View physical server information.

Reverse influence analysis of products

This feature is used to view dependencies between a product and other products that depend on the product.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **End-to-end Diagnostics > Reverse Influence Analysis**.

> **Note** The **Product** tab appears. By default, this tab displays the dependencies between all products. If you move the pointer over a product, the dependencies of the product are displayed.

5. Select the product from the **Product** drop-down list or enter the product name in the field to find

the product. The system displays the dependencies between the product and other products that depend on the product.

> Note
>
> - You can click the style buttons or icons in the upper-left corner to change the display style of the topology.
> - Click **Reset** to clear the value that you set.
> - Fuzzy search is supported.

Reverse influence analysis of services

This feature is used to view dependencies between a service and other services that depend on the service.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Full Stack**.
3. In the left-side navigation pane, click **SLA Console**.
4. In the left-side navigation pane, choose **End-to-end Diagnostics > Reverse Influence Analysis**.
5. Click the **Service** tab.
6. Select the service, product, and cluster from the drop-down lists from left to right or enter values in these fields and then click **Search** to find the service. The system displays the dependencies between the service and other services that depend on the service.

Note

- You can click the style buttons or icons in the upper-left corner to change the display style of the topology.
- Click **Reset** to clear the value that you set.
- Fuzzy search is supported.

1.1.5. Alerts

1.1.5.1. Dashboard

This topic describes how to view alerts within each region of a multi-region scenario.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the **Alerts for Regions** section, view the distribution of alerts in different regions.

   Move the pointer over a region with alerts, and the specific number of alerts is displayed.

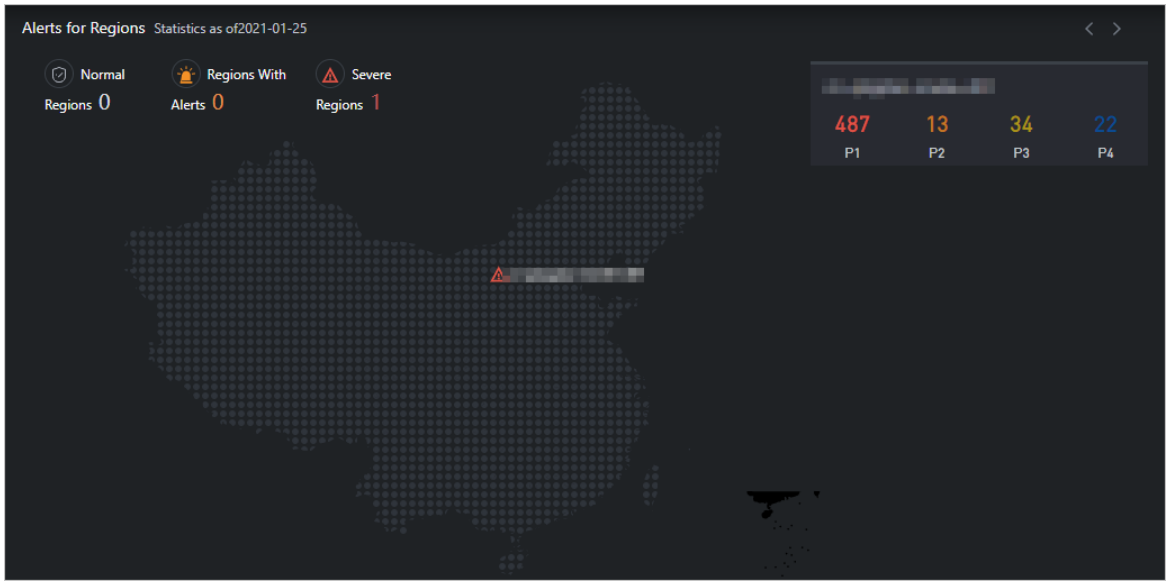

> **Note** P1, P2, P3, and P4 have the following meanings:
>
> - P1: urgent alerts
> - P2: major alerts
> - P3: minor alerts
> - P4: reminder alerts

4. In the **Alert Statistics** section, view the statistical data of alerts in the region.

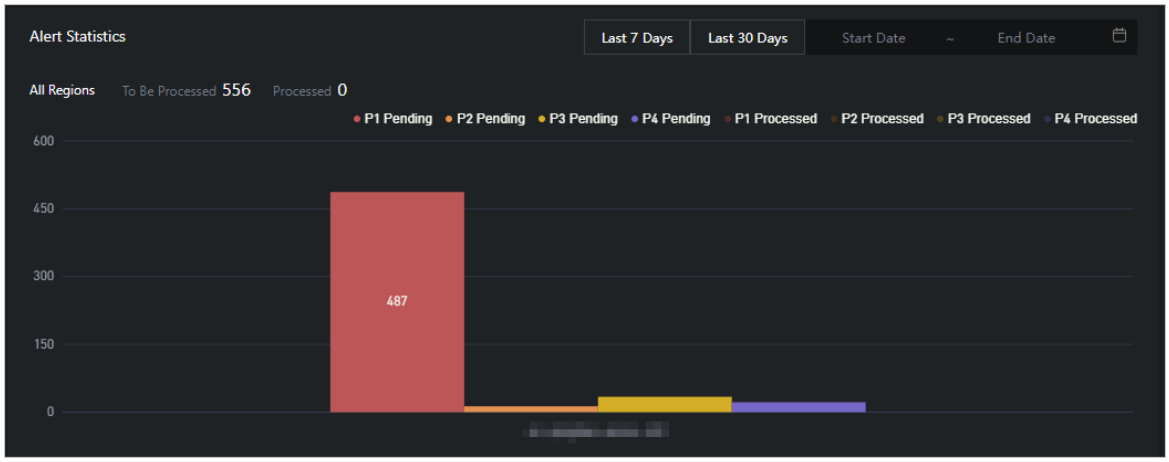

   - Click **Last 7 Days**, **Last 30 Days**, or specify **Start Date** and **End Date** to view the statistical data of alert handling within the period.
   - Move the pointer over a column chart, and the corresponding statistical data is displayed.
   - Select the region and priority level, and then click **Search** to view the alert details.
   - Specify more filter conditions to query alerts.
     - Click **Advanced**. Select the required options from the **Status** and **Resource** drop-down lists, and click **Search** to view the alert details.
     - (Optional) Click **Reset** to clear the filter conditions.
     - (Optional) Click **Collapse** to hide the **Status** and **Resource** options.
5. Click **Export Report** to download the alert details to your computer.

1.1.5.2. View alerts

This topic describes how to view the alerts for each cloud service, basic service, and hardware in the system.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, click **Alerts** to view the distribution of alerts in the services.
   - Click the **Critical Alerts**, **Existing Alerts**, and **Alert History** tabs to view the information about different types of alerts.
   - The different colors of alerts indicate different ranges of quantities.

- Drag the scroll bar to view more alerts.
- Move the pointer over a color block, and alerts of the corresponding service are displayed.

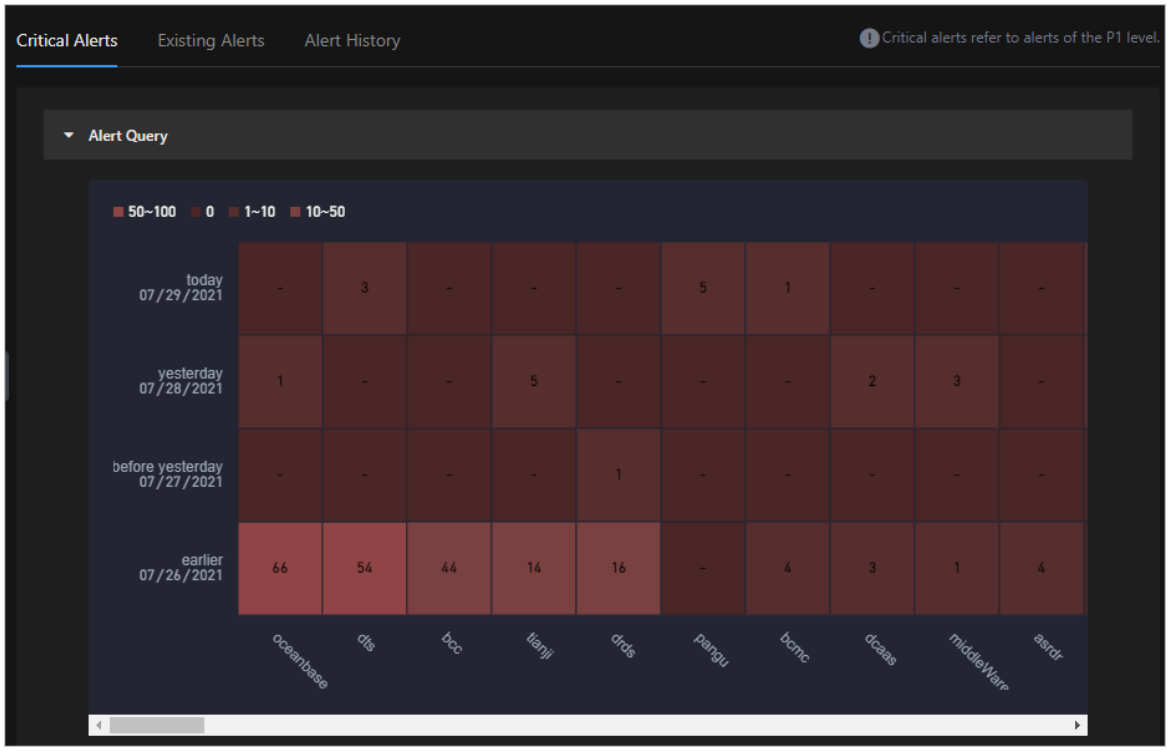

4. View alerts.
   - Enter an alert resource ID in the **Resource With Alerts** search box, select the required options from the **Resource Owner** and **Alert Status** drop-down lists, and then click **Search** to view the alerts.

     Move the pointer over an alert resource or a piece of alert information, and the full description of the alert information is displayed.

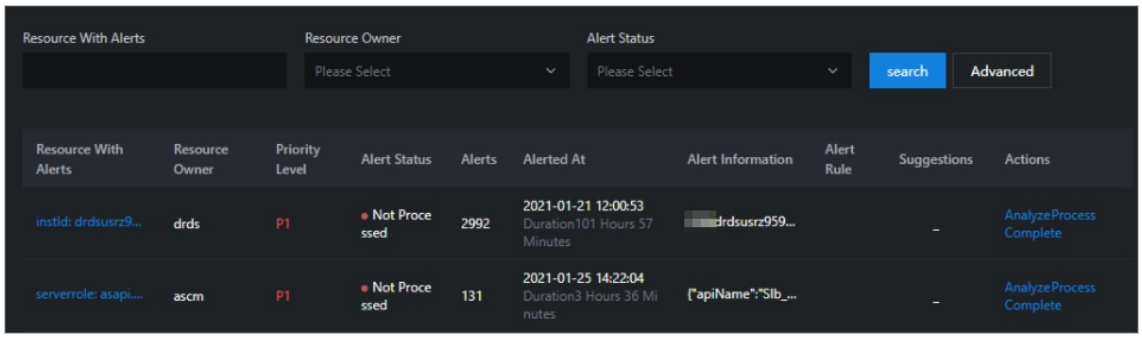

   - Specify more filter conditions to view the alerts.

a. Click **Advanced**. Enter an alert resource ID in the **Resource With Alerts** search box, select the required options from the **Resource Owner**, **Alert Status**, and **Alert Level** drop-down lists, and then select a start date and an end date to specify **Time Range**. Then, click **Search** to view the alert information.

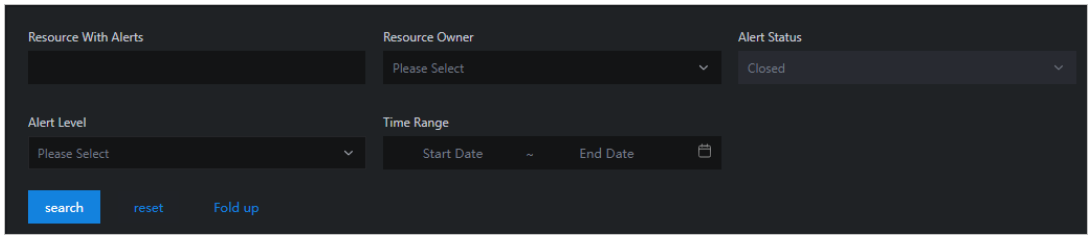

b. (Optional) Click **Reset** to clear the filter conditions.

c. (Optional) Click **Collapse** to hide the **Alert Level** and **Time Range** options.

5. Analyze the alert information.

i. Click **Analyze** in the Actions column corresponding to the alert information. On the **Alert Analysis** page, view the service role to which the alert belongs, the dependency link diagram of the service, or the logical topology of the server.

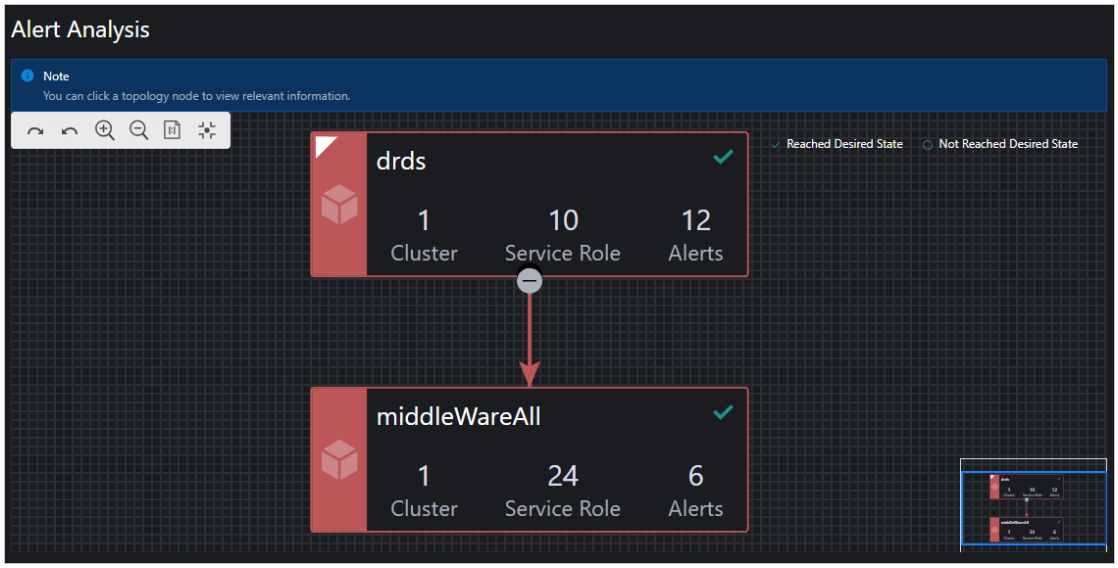

Note Icons:

- : indicates the server.
- : indicates the cluster.
- : Indicates the service.
- : indicates the service role.

ii. Click a node in the topology. The details panel of the node appears on the right.

The panel shows the elements related to this node. You can search for the elements by entering a keyword in the search box.

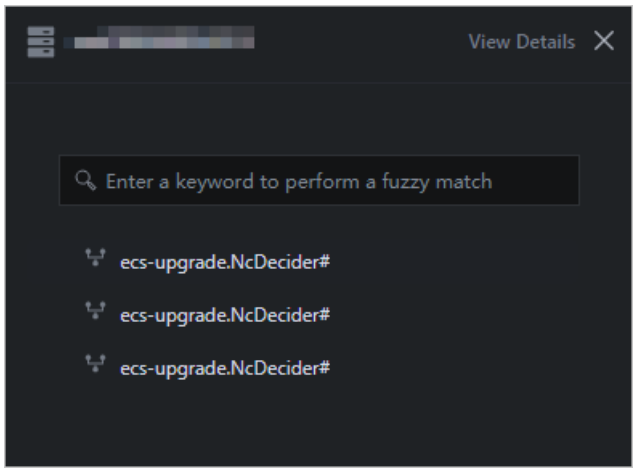

iii. Click **View Details** in the upper-right corner to view the details of the node.

Select a start time and an end time to view the monitoring details within the time range.

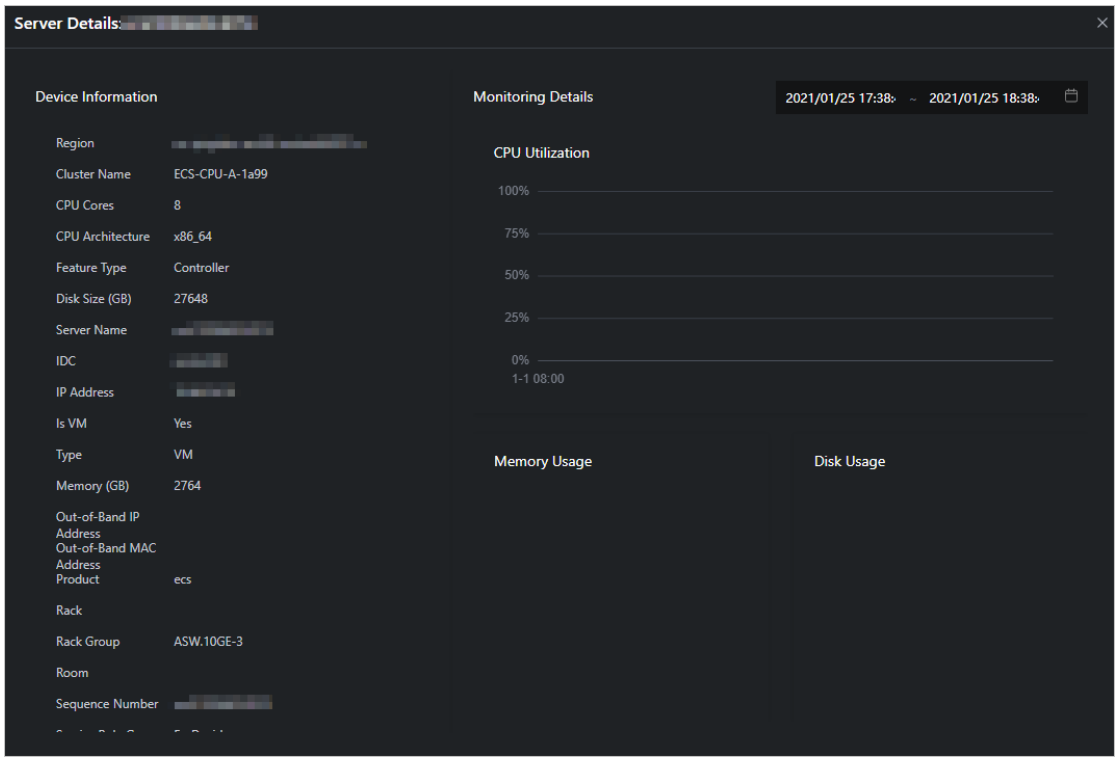

iv. In the node details panel, click an element related to this node to view the details of this element.

> **Note** In this example, the node is a server, and the element related to the node is service role.

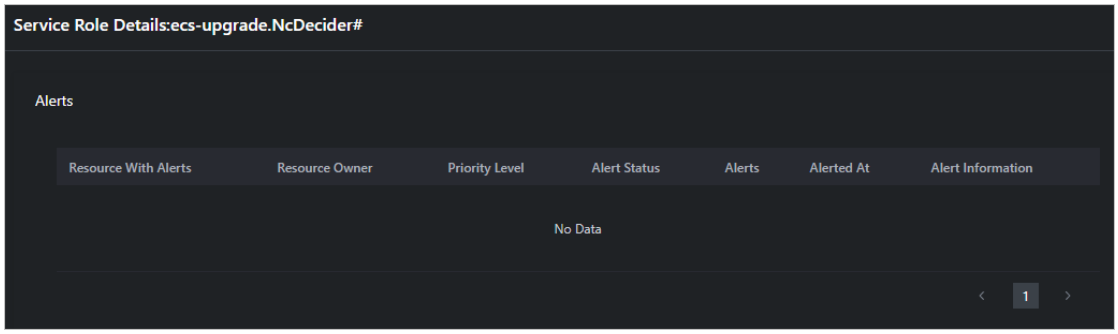

6. Click **Process** in the Actions column corresponding to the alert information. In the message that appears, click **OK** to mark the alert as being processed.
7. Click **Processed** in the Actions column corresponding to the alert information. In the message that appears, click **OK** to mark the alert as processed.

1.1.5.3. Alert settings

1.1.5.3.1. Policy management

The Policy Management module allows you to manage contacts and contact groups, and configure static parameters.

1.1.5.3.1.1. Alert contacts

You can query, add, modify, or delete alert contacts based on your business needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Policy Management**.
4. On the Contacts tab, perform the following operations:
   - Query alert contacts

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

     In the upper-left corner of the tab, click **Add**. In the **Add Contact** panel, configure the parameters. Then, click **OK**.
   - Modify an alert contact

     Find the alert contact whose information you want to modify and click **Modify** in the **Actions** column. In the **Modify Contact** panel, modify the relevant information and click **OK**.

- Delete an alert contact

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

1.1.5.3.1.2. Alert contact groups

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Policy Management**.
4. Click the **Contact Groups** tab.
5. Perform the following operations:
   - Query an alert contact group

     In the upper-left corner of the tab, enter a group name in the search box and click **Search**. The information about the alert contact group that meets the filter condition is displayed.
   - Add an alert contact group

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

     Find the contact group that you want to modify and click **Modify** in the **Actions** column. In the **Modify Contact Group** panel, modify the group name, description, contacts, and notification method. Then, click **OK**.
   - Delete one or more alert contact groups

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

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

1.1.5.3.1.3. Configure static parameters

You can configure alert-related static parameters to suit 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Policy Management**.

4. Click the **Static Parameter Settings** tab.
5. (Optional)In the upper-left corner of the tab, enter a parameter name in the search box and click **Search** to query static parameter configurations.
6. Find the static parameter that you want to modify and click **Modify** in the **Actions** column.
7. In the **Modify Static Parameter** panel, modify the parameters described in the following table.

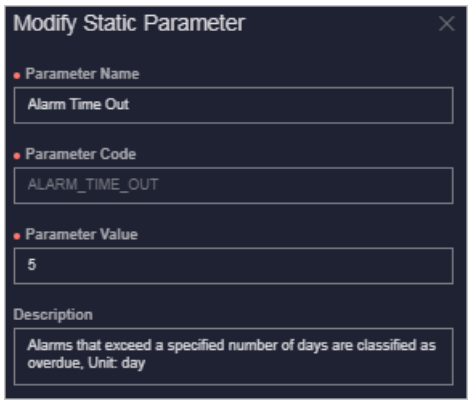

Parameter	Description
Parameter Name	Enter a parameter name related to the configuration.
Parameter Value	Enter a parameter value. The default value is 5, indicating five days. After you complete the configurations, you can choose **Alert Monitoring > Alert Events** and then click the **Timeout Alert**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

8. Click **OK**.

1.1.5.3.2. Alert templates

For Ant Financial Service products deployed on the PaaS platform, you can upload alert templates to configure or modify the rules that trigger alerts.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Alert Template**.
4. On the **Alert Template** page, select the required options from the **Product**, **Cluster**, and **Service** drop-down lists, and click **Search** to view the details of the service.

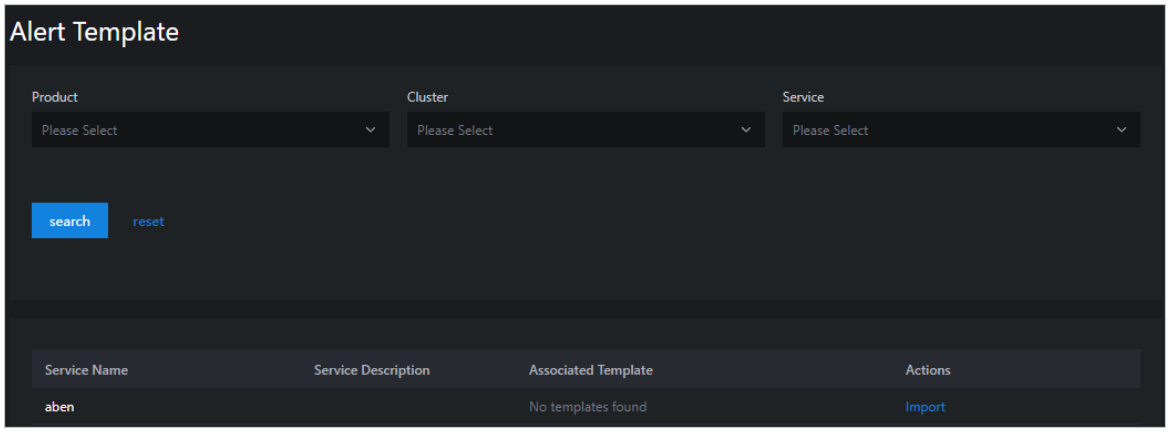

5. (Optional) Click **Reset** to clear the filter conditions.
6. Download Alert Templates.

> Note : For Ant Financial Service products deployed on the PaaS platform, use the simple_template.json template.

7. Click **Import** in the Actions column corresponding to an entry. In the **Import Template** dialog box, click **Upload and Parse File**. Select the template and click **Open**. After the template is uploaded, click **OK**.

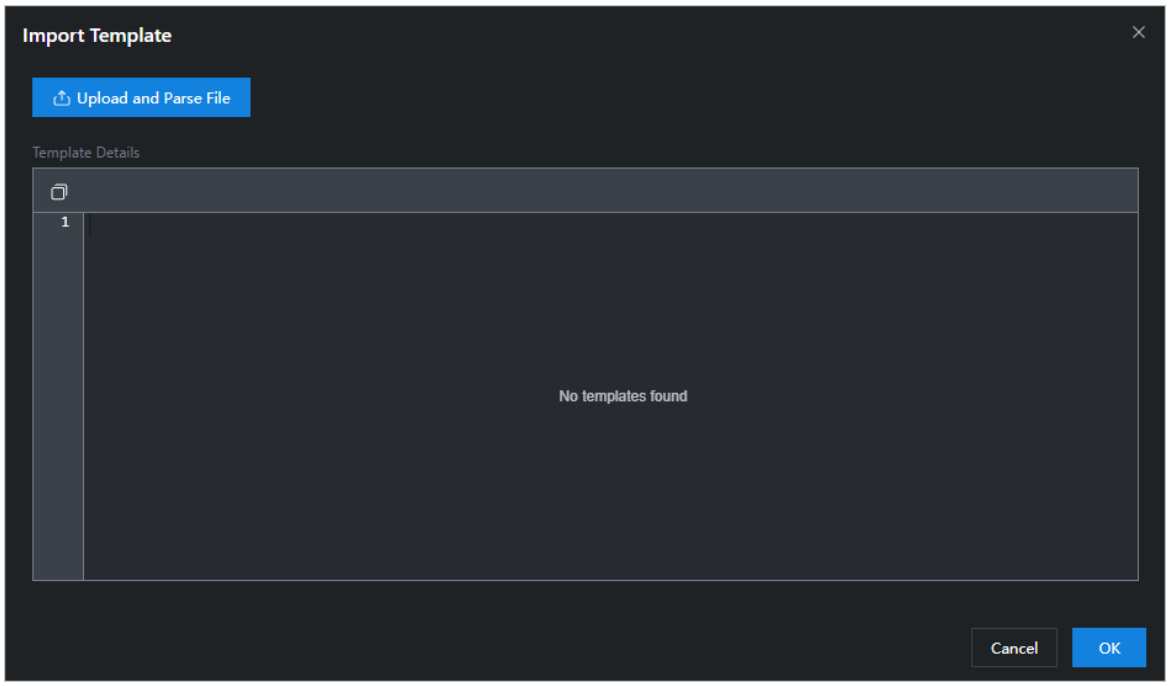

1.1.5.3.3. Notification management

The notification management feature allows you to configure alert notification channels and then push alerts to O&M engineers.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.

2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Notification Management**.

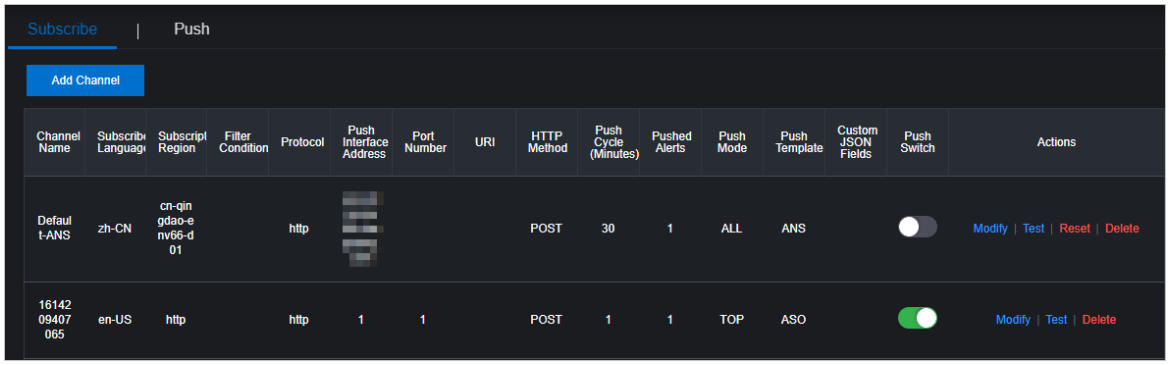

4. On the **Subscribe** tab, click **Add Channel**.
5. In the **Add Subscription** panel, configure the parameters described in the following table.

Parameter	Description
Channel Name	The name of the subscription channel.
Subscribed Language	The subscription language. Valid values: Chinese and English.
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 at which to push alerts. Unit: minutes.
Pushed Alerts	The number of alerts pushed each time.

Parameter	Description
Push Mode	The mode used to push alerts. Valid values: ○ **ALL**: All alerts are pushed in each push cycle. ○ **TOP**: Only high priority alerts are pushed in 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 configure only one channel of this type. **Note** A preset ANS template exists if the system is already connected to ANS. To restore the initial configurations of the template, click **Reset** in the **Actions** column corresponding to the channel.
Custom JSON Fields	The push receiver 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 panel is not turned on, you can enable the push feature in the **Push Switch** column after you configure the subscription channel.

6. Click **OK**.

   To modify or delete a channel, click **Modify** or **Delete** in the **Actions** column corresponding to the channel.

7. (Optional)The newly added channel is displayed in the list. Click **Test** in the **Actions** column to test the connectivity of the push channel.

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

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

1.1.5.3.4. Alert masking

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

1.1.5.3.4.1. Add a masking rule

Masking rules allow you to mask alerts that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Alert Masking**.
4. In the upper part of the page, click **Add**.
5. In the **Add** panel, configure the parameters to filter the alerts to be masked.

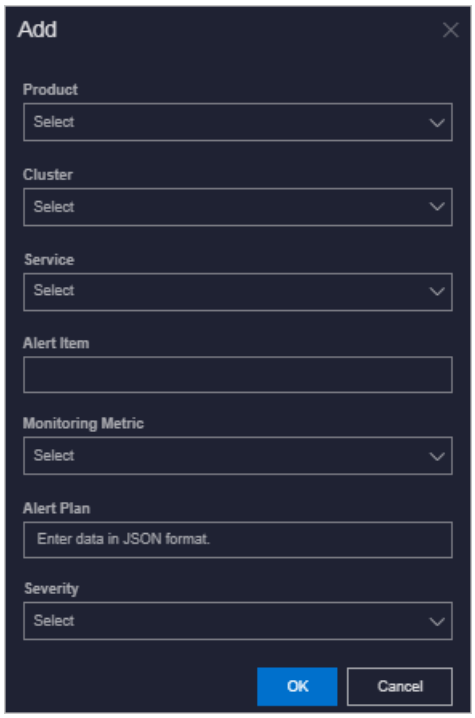

Parameter	Description
Product	Optional. The product to which the alerts to be masked belong.
Cluster	Optional. The cluster to which the alerts to be masked belong.
Machine	Optional. The host to which the alerts to be masked belong.
Service	Optional. The service to which the alerts to be masked belong.
Alert Item	Optional. The name of the alerts to be masked. **Note** It may take an extended period of time if the number of alerts is large when you configure **Alert Item**.
Metric	Optional. The monitoring metric to which the alerts to be masked belong.

Parameter	Description
Alert Plan	Optional. Details about the alerts to be masked. Example: {"serverrole":"ecs-yaochi.ServiceTest#","machine":"vm01001******","level":"error"}
Class	Optional. The severity level of the alerts. Valid values: ○ **P0**: indicates the alerts that have been cleared. The **Alert Level** of these alerts is **Restored** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P1**: indicates critical alerts. The **Alert Level** of these alerts is **P1** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P2**: indicates major alerts. The **Alert Level** of these alerts is **P2** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P3**: indicates minor alerts. The **Alert Level** of these alerts is **P3** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P4**: indicates reminder alerts. The **Alert Level** of these alerts is **P4** in **Monitoring > Alert History** of Apsara Infrastructure Management Framework. ○ **P5**: indicates system alerts.

6. 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 **Alerts > Alerts**.

1.1.5.3.4.2. Disable masking

You can disable masking for masked alerts.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Alerts**.
3. In the left-side navigation pane, choose **Alert Settings > Alert Masking**.
4. (Optional)Specify a product, service, or an alert item, and click **Search**.
5. Find the alert masking rule that you want to disable and click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

Result

After masking is disabled, the unmasked alerts are displayed in **Alerts > Alerts**.

1.1.6. O&M

1.1.6.1. Automated O&M

The automated O&M feature automates O&M for data centers. A web-based method is provided to implement O&M operations for resources at scale, simplify O&M management of IT resources, and support full-stack automated O&M of the infrastructure, the Apsara Stack environment, operating systems, and the application layer.

1.1.6.1.1. View host resources

You can view the information about hosts such as physical machines or Docker virtual machines.

Context

Before you execute a script or an O&M job on a host, you can view the specific information about the host to ensure that the script or job can be effectively executed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Host Resources**.
4. On the **Host Resources** page, enter a hostname, project name, or cluster name in the upper-left search box and click **search**. Fuzzy match is supported.

   You can view the information about the hosts that meet the filter condition, including the hostname, IP address, project name, cluster name, operating system, and IDC.
5. (Optional) Click **reset** to clear the filter conditions.

1.1.6.1.2. View Docker resources

You can view the information about Docker containers.

Context

Before you execute a script or an O&M job on a Docker container, you can view the specific information about the Docker container to ensure that the script or job can be effectively executed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Docker Resources**.
4. On the **Docker Resources** page, enter a server role name, project name, cluster name, or service

name in the upper-left search box and click **search**. Fuzzy match is supported.

You can view the information about the Docker containers that meet the filter conditions, including the server role name, type, hostname, host IP address, project name, cluster name, and service name.

5. (Optional) Click **reset** to clear the filter conditions.

1.1.6.1.3. Manage scripts

The script library is used to store scripts for implementing various features and is the basis for automated O&M. All O&M commands are run by using scripts. The system provides some common built-in default scripts and supports custom scripting. You can create, import, view, modify, export, and delete scripts.

1.1.6.1.3.1. Create a script

You can create a script and test whether it can be properly executed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Script Library**.
4. Click **Create Script**.
5. Configure the parameters for the script.

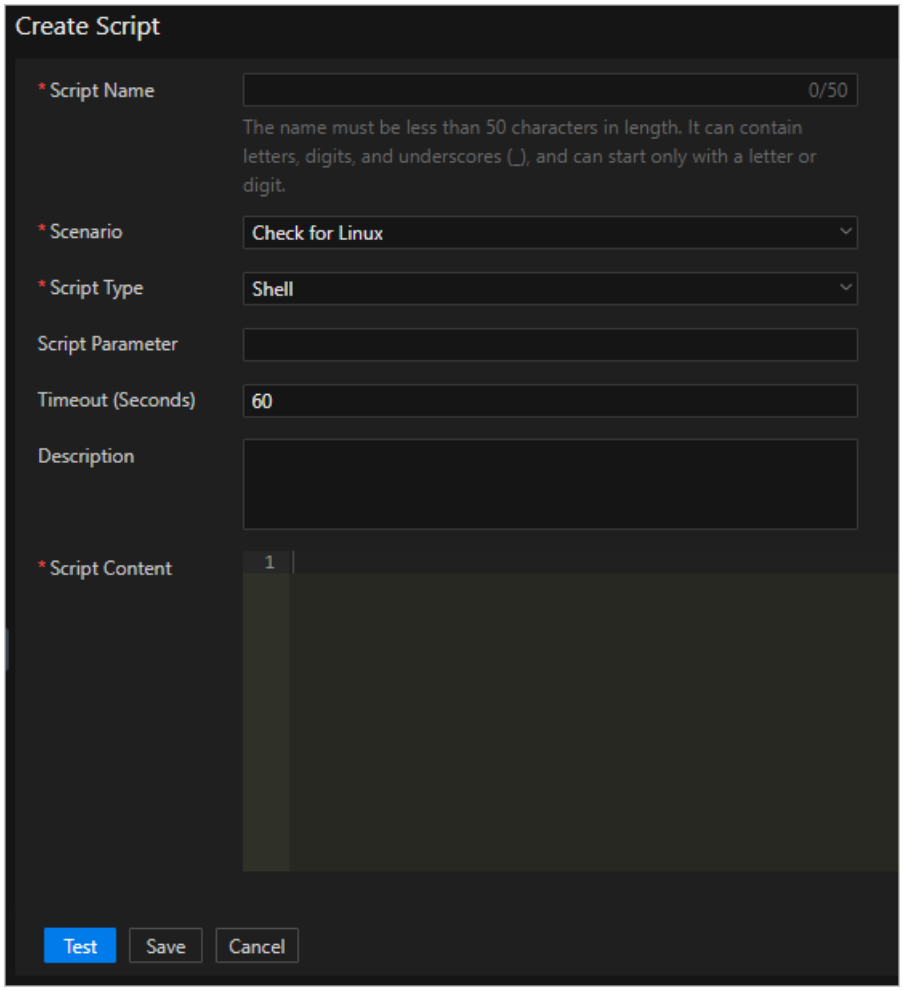

The following table describes the parameters.

Parameter	Description
Script Name	The name of the script.
Scenario	The application scenario of the script. Valid values: **Check for Linux**, **Run Command**, and **Install Software**.
Script Type	The type of the script. Valid values: Shell and Python.
Script Parameter	The parameters passed in when the script is executed. Separate multiple parameters with spaces.
Timeout (Seconds)	The timeout period for script execution. After the specified number of seconds, the script stops executing and the execution timeout result is returned.
Description	The description of the script.

Parameter	Description
Script Content	The content of the script. When you write a script, you must add a script interpreter. For example, for a Shell script, you must enter `#!/bin/bash`. For a Python script, you must enter `#!/usr/bin/python`. The path of the interpreter may vary with the execution resources and environments.

6. Click **Test** to test whether the script can be properly executed. If you confirm that the script can be executed, you can directly click **Save**.
   i. After you click **Save** or **Test**, the system checks the script content. If the **The script has high-risk commands. Do you want to continue?** message appears, check whether the script content is correct.
      - Correct: Click **OK**.
      - Incorrect: Modify the script and test again.
   ii. After you click **Test**, click **Host Resources** or **Docker Resources** in the **Script Test** dialog box, select one or more host resources or Docker resources, and then click **Execute**.

      > **Note** The SSH protocol is used to copy files to the host or Docker container. Therefore, the test execution process may be slow.

   iii. In the **Test Results** dialog box, view the test result of the script.
      - Click **OK** to exit the dialog box and click **Save** to save the script.
      - Click **Re-select Resources** to select another host or Docker container to test the script.

1.1.6.1.3.2. Import a script

You can import a script on your computer to the script library.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Script Library**.
4. Click **Import Script**.
5. In the **Upload Script File** dialog box, click **Click Here to Upload** to upload a script on your computer to the script library.

   > **Note** Only JSON files can be uploaded. The file to be uploaded cannot exceed 500 KB in size.

1.1.6.1.3.3. View scripts

You can view scripts in the script library.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Script Library**.
4. In the **Script Name** search box, enter the name of the script that you want to view and click **search**. Fuzzy match is supported.

   You can view the information about the scripts that meet the filter conditions, including the script name, script type, scenario, parameter, modification time, description, user who updates the script, and whether the script is a default script.
5. (Optional) Click **reset** to clear the filter conditions.

1.1.6.1.3.4. Modify a script

After a script is created or imported, you can modify the script to suit your requirements.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Script Library**.
4. Find the script that you want to modify and click **Edit** in the **Actions** column.
5. Modify the parameters and click **Test** to test whether the script can be executed. If you confirm that the script can be executed, you can click **Save**.
   i. After you click **Save** or **Test**, the system checks the script content. If the **The script has high-risk commands. Do you want to continue?** message appears, check whether the content of the script is correct.
      - Correct: Click **OK**.
      - Incorrect: Modify the script and test again.
   ii. After you click **Test**, click **Host Resources** or **Docker Resources** in the **Script Test** dialog box, select one or more host resources or Docker resources, and then click **Execute**.

      > **Note** The SSH protocol is used to copy files to the host or Docker container, which may result in slow execution speeds.

   iii. In the **Test Results** dialog box, view the test result of the script.
      - Click **OK** to exit the dialog box and click **Save** to save the script.
      - Click **Re-select Resources** to select another host or Docker container to test the script.

1.1.6.1.3.5. Export a script

You can export a script to your computer.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.

3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Script Library**.
4. Select one or more scripts that you want to export and click **Export Script** to export the scripts to your computer.

> **Note** If you export multiple scripts at a time, the content of the scripts is stored as a single JSON file.

1.1.6.1.3.6. Delete a script

You can delete scripts that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Script Library**.
4. Select one or more scripts that you want to delete and click **Delete Script** in the lower part of the page, or click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.6.1.4. Manage software

The software repository is used to manage software, including uploading, viewing, downloading, and deleting software. The term software used in this topic is in its broad sense, including compressed packages, JAR packages, images, and files. Only software uploaded to the software repository can be used in subsequent jobs.

1.1.6.1.4.1. Upload software

You can upload software to the software repository.

Context

When an O&M job is executed in an on-site environment, software is downloaded from the software repository and deployed on the host or Docker container. You must upload the software to the software repository before you can use it in subsequent jobs.

> **Note** Delete software that is no longer needed to free up storage space.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Software Repository**.
4. Click **Upload Software**.
5. In the **Upload Software** dialog box, enter a software name in the **Software Name** field and click

Click Here to Upload to upload an on-premises file. If you do not enter a software name, the software name is the same as the file name.

> **Note** The file to be uploaded cannot exceed 500 MB in size.

1.1.6.1.4.2. View software

You can view software in the software repository.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Software Repository**.
4. In the **Software Name** search box, enter the name of the software that you want to view and click **search**. Fuzzy match is supported.

   You can view information about software that meets the filter conditions, including the software name, file name, file size, upload time, and upload user.
5. (Optional) Click **reset** to clear the filter conditions.

1.1.6.1.4.3. Download software

You can download software from the software repository to your computer.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Software Repository**.
4. Select the software that you want to download and click **Download** in the **Actions** column.

1.1.6.1.4.4. Delete software

To save storage space, you can delete software that is no longer needed after O&M jobs are executed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Software Repository**.
4. Select one or more software that you want to delete and click **Delete Software** in the lower part of the page, or click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.6.1.5. Manage processes

Process orchestration is one of the core features of automated O&M. It is used to manage processes, including creating, importing, viewing, exporting, modifying, running, and deleting processes. You can define a process to combine a series of logical actions into a task and automate O&M.

1.1.6.1.5.1. Create a process

You can create a process to visually orchestrate the O&M process and automate O&M.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.
4. Click **Create Process**.
5. On the **Create Process** page, click **Process Settings**.
6. In the **Process Settings** dialog box, configure the following parameters:
   - **Process Name**: Enter a name for the process.
   - **Process Description**: Enter a description for the process.
   - **Trigger Method**: Select **Manual** or **Scheduled**.
     - **Manual**: The process must be manually triggered.
     - **Scheduled**: The process is triggered at the specified time.
   - **Timing Rule**: This parameter is available only when **Trigger Method** is set to **Scheduled**. Set the time to trigger the process.
     - **Once**: The process is triggered only once at the specified time. Select a date and time to trigger the process.
     - **Daily**: The process is triggered once at the specified time every day. Select a time to trigger the process every day.
     - **Monthly**: The process is triggered at the specified date and time every month. Select a date and time to trigger the process. For example, if you set **Days** to 10 and **Time** to 09:00:00, the process is triggered at 09:00:00 on the tenth day of every month.
7. Click **OK**.
8. Drag nodes on the left side to the right side and add lines between the nodes. The nodes that can be added to a process include the Start, Task, Judgement, Manual, Wait, and Notification nodes.
   -  : the Start node. Each process has only one Start node, which represents the start of the process. The Start node has no parameters and can have only one line to connect to another node.
   - : the Task node. The Task node is the major node for process execution and can execute one script or job. Click the Task node and configure the following parameters in the **Node Properties** dialog box.

Tab	Parameter	Description
Specify Node	**Node Name**	The name of the node.
	Description	The description of the node.
Specify Parameters	**Input Parameters**	Input parameters are the output parameters of the previous node. If a previous node has no output parameters, the node has no input parameters. Input parameters follow a script in sequence when the script is executed. Example: `./test.sh params1 params2`. **Note** If you set Operation to **Execute Job** on the Select Operation tab, you do not need to specify the input parameters.
	Output Parameters	Output parameters take effect when only one script and execution resource is selected for the node. Output parameters come from the execution result of the script. Therefore, the script must return a fixed result. For example, if the execution result is `echo "CPU=22,MEM=30"`, click **Add** to configure the output parameters. Specify **Output Parameter Name** and enter CPU and MEM in the **Parsed Key Value** field. **Note** If you set Operation to **Execute Job** on the Select Operation tab, you do not need to specify the output parameters.

Tab	Parameter	Description
Select Operation	**Operation**	Select **Script** or **Execute Job** and select a script from the Script drop-down list, or select an execute job from the Execute Job drop-down list.
	Resource Type	Select **Host** or **Docker**. ■ **Host**: Click **Select Host**. In the **Select Host** dialog box, select one or more hosts and click **OK**. You can also enter a hostname, project name, or cluster name in the **Host** search box and press the Enter key to search for the hosts that you want to select. Fuzzy matching is supported. ■ **Docker**: Click **Select Docker**. In the **Select Docker** dialog box, select one or more Docker containers and click **OK**. You can also enter a server role name, project name, or cluster name in the **Docker** search box and press the Enter key to search for the Docker containers that you want to select. Fuzzy matching is supported. You can click the icon to delete the specified host or Docker container.
		Select **None**, **Automatically Executed**, or **Stop Waiting**. ■ **None**: The scripts or jobs are executed on all the selected hosts or Docker containers in one batch.

Tab	Parameter	Description
Select Execution Resources	Phased Execution Settings	■ **Automatically Executed**: The scripts or jobs are executed in two batches on the selected hosts or Docker containers. One batch is executed at a time. After the first batch is executed, the second batch starts to be executed. If the first batch fails to be executed, the second batch is not executed. **Note** If you set Resource Type to Docker, the selected Docker containers are automatically divided into batches based on the server role name such that Docker containers with the same server role are in different batches. ■ **Stop Waiting**: The scripts or jobs are executed in two batches on the selected hosts or Docker containers. The first batch is executed first. After the first batch is executed, a process approval is sent. The approver can click **Passed** or **Stop** in **Process Review**. When the approval is passed, the second batch starts to be executed. If the first batch fails to be executed, the execution stops and no process approval is sent. **Note** If you set Resource Type to Docker, the selected Docker containers are automatically divided into batches based on the server role name such that Docker containers with the same server role are in different batches.

Tab	Parameter	Description

-  : the Judgement node, which is used to judge the process routes of different directions. The previous node of the Judgement node must be a Task node that has output. Otherwise, the Judgement node is of no use. Click the Judgement node. In the **Node Properties** dialog box, configure the following parameters.
  - **Node Name**: Enter a name for the node.
  - **Description**: Enter a description for the node.
  - **Judgement Condition**: Click **Add**. In the **Add** dialog box, set **Output Parameter Name** of the previous node and set the judgement condition by specifying **Judge** and **Value**. Click **OK** to save the judgement condition.
  - **Judgement Type**: If you set multiple judgement conditions, you must select **Or** or **And**.
    - **Or**: The judgement result is yes if one judgement condition is met. The judgement result is no if no judgement conditions are met.
    - **And**: The judgement result is no if one judgement condition is not met. The judgement result is yes if all the judgement conditions are met.

  You must click the lines coming out from the Judgement node and set **Yes** or **No** to define the execution of the subsequent process.
-  : the Manual node. When the process is executed to this node, the process is suspended for manual approval. If the approval is passed, the process continues execution. If the approval is not passed, the subsequent process is not executed. If a timeout period is specified and manual approval is not performed after this period expires, the subsequent process is automatically stopped.
-  : the Wait node. When the process is executed to this node, the process waits a specified period of time before the subsequent process is executed. For example, the previous node executed the script for service startup, but usually the service startup takes some time. In this case, you can wait 1 minute and then check whether the service is normal in the next node.
-  : the Notification node. Notifications are sent by email. Select **Email** from the **Notification Type** drop-down list. Then, set Recipient, Notification Title, and Notification Content. Separate multiple email addresses with commas (,).

9. Click **Save** in the upper-right corner.

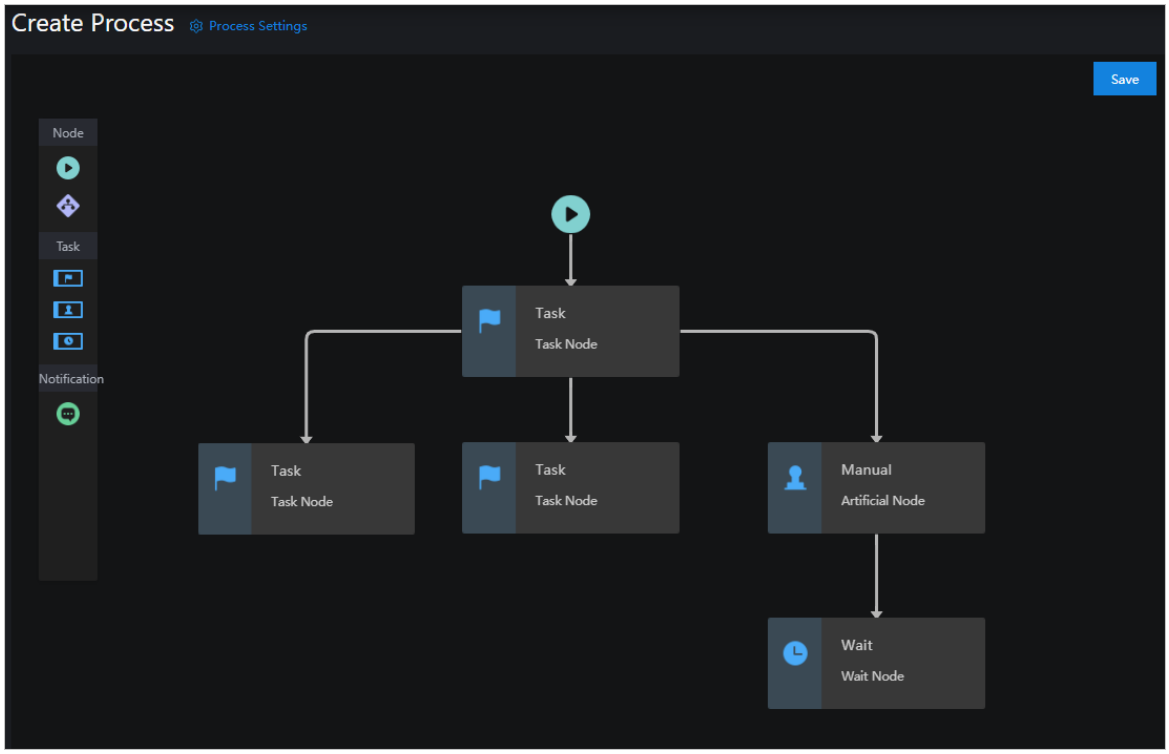

1.1.6.1.5.2. Import a process

You can import existing processes.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.
4. Click **Import Process**.
5. In the **Upload Process** dialog box, click **Click Here to Upload** to upload an existing process.

> **Note** Only JSON files can be uploaded. The file to be uploaded cannot exceed 500 KB in size.

1.1.6.1.5.3. View processes

You can view existing processes.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.

4. In the **Name** search box, enter the name of the process that you want to view and click **search**. Fuzzy match is supported.

   You can view the information of processes that meet the filter conditions, including the process name, description, execution method, modification time, execution history, latest execution time, and latest execution status.

5. (Optional) Click **reset** to clear the filter conditions.
6. Click the process name to go to the **Process Details** page.

   On the **Process Details** page, you can perform the following operations:

   - Click the **Process Details** tab to view the structure of the process.
   - Click the **Execution History** tab to view the execution history of the process. You can also click **View Details** to view the details of the execution history, including the node name, node type, start time, end time, task status, and execution information.
   - Click **Run** in the upper-right corner to manually run the process.
   - Click **Modify** in the upper-right corner to modify the process.
   - Click **Delete** in the upper-right corner to delete the process.

1.1.6.1.5.4. Export a process

You can export processes to your computer.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.
4. Select one or more processes that you want to export and click **Export Process** to export the processes to your computer.

   > **Note** If you export multiple processes at a time, the content of the processes is stored in one JSON file.

1.1.6.1.5.5. Modify a process

After you have created or imported a process, you can modify the process to suit your requirements.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.
4. Select the process that you want to modify and click **Modify** in the **Actions** column.
5. Modify the process and click **Save** in the upper-right corner.

1.1.6.1.5.6. Run a process

You can manually trigger processes.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.
4. Select the process that you want to run and click **Run** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.6.1.5.7. Delete a process

You can delete processes that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > Process Orchestration**.
4. Select the process that you want to delete and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.6.1.6. Manage O&M jobs

O&M jobs is one of the core features of automated O&M and can be used to independently complete O&M tasks such as software distribution, patch upgrade, and program update. You can create, import, view, export, modify, run, and delete O&M jobs.

Each O&M job is a collection of features for O&M resources, software, and scripts. Scripts are used to implement features and are executed on different hosts or Docker instances in a specified order to reduce the workloads of O&M personnel.

1.1.6.1.6.1. Create an O&M job

You can create an O&M job to independently complete an O&M task to implement automated O&M.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.
4. Click **Create O&M Job**.
5. On the **Create O&M Job** page, configure the parameters described in the following table.

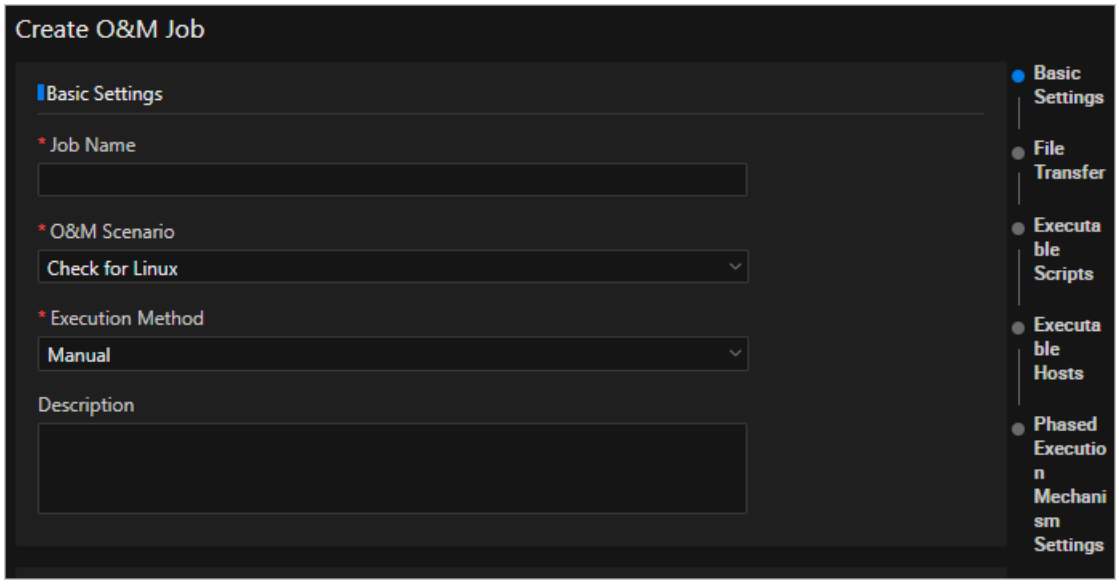

Parameter	Description
Job Name	The name of the O&M job.
O&M Scenario	The scenario of the O&M job. You can select **Check for Linux**, **Run Command**, or **Install Software**.
Execution Method	Select **Manual** or **Scheduled**. ○ **Manual**: The O&M job must be manually executed. ○ **Scheduled**: The O&M job is executed at the specified time.
Timing Rule	This option is available only when **Execution Method** is set to **Scheduled**. Set the time to execute the O&M job. ○ **Once**: The O&M job is executed only once at the specified time. Select a date and time to execute the O&M job. ○ **Daily**: The O&M job is executed once at the specified time every day. Select a time to execute the O&M job every day. ○ **Monthly**: The O&M job is executed at the specified date and time every month. Select a date and time to execute the O&M job. For example, if you set **Days** to 10 and **Time** to 09:00:00, the O&M job is executed at 09:00:00 on the tenth day of every month.
Description	The description of the O&M job.

Parameter	Description
File Transfer	Click **Add File**. In the **Add File** dialog box, select the file that you want to transmit to the host or Docker container, enter an absolute path in the Transmission Path field, and then click **OK**. You can add multiple files or click the icon to delete the files that are no longer needed. **Note** The files that you can select all come from the software repository.
Executable Scripts	Add one or more scripts that you want to execute on the host or Docker container, and the system executes them in sequence. Click **Add Executable Script** and select a script from the script library, or click **Add Script** to create a script. You can click the icon to modify the script. The modification does not change the original script content in the script library. You can click the or icon to change the order of execution of the script. You can also click the icon to delete the scripts that are no longer needed.
Executable Hosts	Set **Resource Type** to **Host** or **docker**. Click **Add Host resources** or **Add Docker Resources** to add one or more hosts or Docker containers. These hosts or Docker containers are where all files are transferred to and where the scripts are executed. **Note** You can specify only one resource type. You cannot add both host and Docker resources at the same time.
	You can select the following options to set the phased execution rules: ○ **None (You do not need to specify this parameter, and all target hosts are executed directly)**: The script is executed on all the selected hosts or Docker containers in a single batch. You can select this option if you are selecting only a few hosts or Docker containers, or if you have confirmed that you

Parameter	Description
Phased Execution Rules	do not have problem with script execution. ◦ **Automatically Execute (You do not need to specify batches. The system executes jobs in two batches. After the execution of the first batch is complete, the second batch is automatically executed)**: The script is executed in two batches on the selected hosts or Docker containers. The first batch is executed first. If the first batch is executed, the second batch start to be executed. If the first batch fails to be executed, the second batch is not executed. **Note** If you set Resource Type to docker, the selected Docker containers are automatically divided into batches based on the service role name so that Docker containers with the same service role are in different batches. ◦ **Stop Waiting (You do not need to specify batches. The system executes jobs in two batches. After the execution of the first batch is complete, the process suspends. The second batch is executed after confirmation)**: The script is executed in two batches on the selected hosts or Docker containers. The first batch is executed first. If the first batch is executed, a job approval is sent. The approver clicks **Passed** or **Stop** in **Job Review**. When the approval is passed, the second batch is executed. If the first batch fails to be executed, the execution is stopped and no job approval is sent. **Note** If you set Resource Type to docker, the selected Docker containers are automatically divided into batches based on the service role name such that Docker containers with the same service role are in different batches. ◦ **Phased execution rules apply to physical servers by default but are not applicable to VMs or Docker containers (The default phased execution rule is a cluster-based algorithm and is executed automatically on multiple hosts in parallel)**: This option is applicable only to physical machines, but not to VMs or Docker containers. The script is executed in batches based on the phased execution rules provided in Apsara Infrastructure Management Framework. After a batch is executed, a job approval is sent. The approver clicks **Passed** or **Stop** in **Job**

Parameter	Description
	Review. When the approval is passed, the next batch is executed. If the batch fails to be executed, the execution is stopped and no job approval is sent. Jobs are executed by cluster based on the following rules on the machines in each cluster: ■ For clusters in SLB, VPC, Apsara Infrastructure Management Framework, ApsaraDB RDS, MiniRDS, OSS, and Blink, the job is executed machine by machine. ■ For clusters other than the preceding ones and that contain 10 or fewer machines, the job is executed on the machines in the following order: 1 machine, 1 machine, 2 machines, 3 machines, and then the remaining machines. ■ For clusters other than the preceding ones and that contain more than 10 machines, the job is executed on the machines in the following order: 1 machine, 3 machines, 5 machines, N/3-1 (rounded down) machines, and N/3-1 machines until the job is executed on all the machines. N is the number of machines in the cluster. **Note** If a cluster contains both physical machines and VMs, the job is executed on all the VMs in the last batch. ○ **Custom**: You can set the batches on your own to execute the job. In the **Batch Settings** drop-down list, select the hosts or Docker containers to add. You can click the  icon to add batches or click the  icon to delete batches. You can add up to three batches. Set **Phased Execution Condition** to **Automatic Execution per Batch** or **Waiting for Review and Confirmation**. ■ **Automatic Execution per Batch**: A batch is executed first. If the batch is executed, the next batch start to be executed. If the batch fails to be executed, the next batch stops execution. ■ **Waiting for Review and Confirmation**: A batch is executed first. If the batch is executed, a job approval is sent. The approver clicks **Passed** or **Stop** in **Job Review**. When the approval is passed, the next batch starts execution. If the batch fails to be executed, the execution is stopped and no job approval is sent.

Parameter	Description

6. Click **Create**.

1.1.6.1.6.2. Import an O&M job

You can import existing O&M jobs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.
4. Click **Import O&M Job**.
5. In the **Upload O&M Job** dialog box, click **Click Here to Upload** to upload an existing O&M Job.

> **Note**
> - Only JSON files can be uploaded. The file to be uploaded cannot exceed 500 KB in size.
> - An imported O&M job cannot be directly executed. You must select a host before you can execute the job.

1.1.6.1.6.3. View O&M jobs

You can view existing O&M jobs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.
4. In the **O&M Job Name** search box, enter the name of the O&M job that you want to view and click **search**. Fuzzy match is supported.

   You can view information of the O&M job that meets the filter conditions, including the job name, O&M scenario, description, execution method, modification time, execution history, latest execution time, latest execution status, update user, and whether the job is set to default.
5. (Optional) Click **reset** to clear the filter conditions.
6. Click the number in the **Execution History** column to view the execution history of the job, including the start time, end time, execution method, execution result, and job information.
7. On the **Execution History** page, you can perform the following operations:
   - Click **View** in the **Details** column to view the execution details of each step on each host or Docker container.

- Click **Snapshot Records** in the **Details** column to view the job history snapshot.
- Click **Proceed** in the **Details** column to continue to execute the job.

  > **Note** If an O&M job is executed based on the phased execution rules and if you want to execute a subsequent batch when the previous batch fails to be executed, you can click **Proceed**.

- Select one or more execution history entries and click **Delete Execution History** above the list, or click **Delete** in the **Details** column corresponding to the entries. In the message that appears, click **OK**.

1.1.6.1.6.4. Export an O&M job

You can export existing O&M jobs to your computer.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.
4. Select one or more O&M jobs that you want to export and click **Export O&M Job** to export the O&M jobs to your computer.

   > **Note** If you export multiple O&M jobs at a time, the content of the jobs is stored in a single JSON file.

1.1.6.1.6.5. Modify an O&M job

After an O&M job is created or imported, you can modify the O&M job to suit your needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.
4. Find the O&M job that you want to modify and click **Modify** in the **Actions** column.
5. On the Modify O&M Job page, modify the settings and click **Save**.

1.1.6.1.6.6. Execute an O&M job

You can manually execute O&M jobs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.

4. Select the O&M job that you want to execute and click **Execute** in the **Actions** column.
5. In the Execute O&M Job message, click **OK**.

1.1.6.1.6.7. Delete an O&M job

You can delete O&M jobs that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > O&M Resources > O&M Jobs**.
4. Select one or more O&M jobs that you want to delete and click **Delete Job** in the lower part of the page, or click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.6.1.7. Review jobs

If an O&M job is executed based on the phased execution rules, the system enables job review. After a batch is executed, the next batch does not start to be executed until the previous batch passes the review. You can pass or stop an O&M job.

Context

During the execution of an O&M job, the system can only judge whether the O&M job is executed, but cannot know the execution result. If the O&M personnel want to confirm the execution results of one batch before they execute the next batch, the O&M personnel can set **Phased Execution Rules** to **Stop Waiting (You do not need to specify batches. The system executes jobs in two batches. After the execution of the first batch is complete, the process suspends. The second batch is executed after confirmation.)** to review the O&M jobs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > Audit Management > Job Review**.
4. Select the O&M job that you want to review and click **Passed** in the **Actions** column to execute the next batch, or click **Stop** to stop the execution of the next batch.
5. In the message that appears, click **OK**.

1.1.6.1.8. Review processes

If a Manual node or a Task node for which **Phased Execution Settings** is set to **Stop Waiting** is available when a process is running, the system initiates process review. You can pass or stop the process.

Context

When a process is running, the system can judge whether the task is complete but cannot determine whether the task is correctly executed. If a Manual node or a Task node for which **Phased Execution Settings** is set to **Stop Waiting** is available in a process, the O&M personnel can view the task execution result to determine whether to pass or stop the process.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > Audit Management > Process Review**.
4. Select the process that you want to review and perform the following operations in the **Actions** column:
   - Click **Passed** to pass the process or run the next batch.
   - Click **Stop** to stop the process or stop the execution of the next batch.
5. In the message that appears, click **OK**.

1.1.6.1.9. Review OOS executions

If an approval process is configured in the template for an OOS execution, the system administrator of the Apsara Uni-manager Operations Console must review and approve the execution before the execution can be processed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > Audit Management > OOS Review**.
4. Find the execution and click **View Details** in the **Task Details** column.
5. In the upper-right corner of the details page, click **Confirm Task**.
6. Approve the execution.
   - If you confirm that the execution is approved, select **Deny Task** in the panel that appears and click **OK**.
   - If you confirm that the execution is not approved, select **Deny Task** in the panel that appears and click **OK**.

1.1.6.1.10. View O&M logs

You can view the logs of various automated O&M operations.

Context

You can view the type, time, user, and details of automated O&M operations to help you perform subsequent audits.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Automated O&M > Audit Management > Operations Logs**.
4. Select an O&M operation type from the **Type** drop-down list, select the start time and end time of the O&M operation, and then click **Search**.

   You can view the operation logs that meet the filter conditions, including the type, time, user, and details of the operation.
5. (Optional) Click **Reset** to clear the filter conditions.

1.1.6.2. STM

The STM module can monitor the availability of the cloud platform by simulating user operations and testing and inspecting product features. The module also can provide early warning for event alerts.

1.1.6.2.1. Check health overview

You can check the inspection results in the STM module, including product health overview, current events, and product inspection details.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Health Overview**.
4. In the **Overview** section, you can view health score, the numbers of abnormal products, risky products, normal products, and no-data products, and the top 5 abnormal products.
5. In the **Current Events** section, you can view details of current events in a list. The list contains the Event Title, Level, Status, Created At, and Event Summary.

   > **Note** Click **View All** in the upper-right corner. The **Event Center** page appears. You can view information of all events on the page.
6. In the **Product Details** section, you can view the inspection details of the modules in each product.
7. Find a module and click **Details** next to the module. In the dialog box that appears, you can view the inspection results of the module.

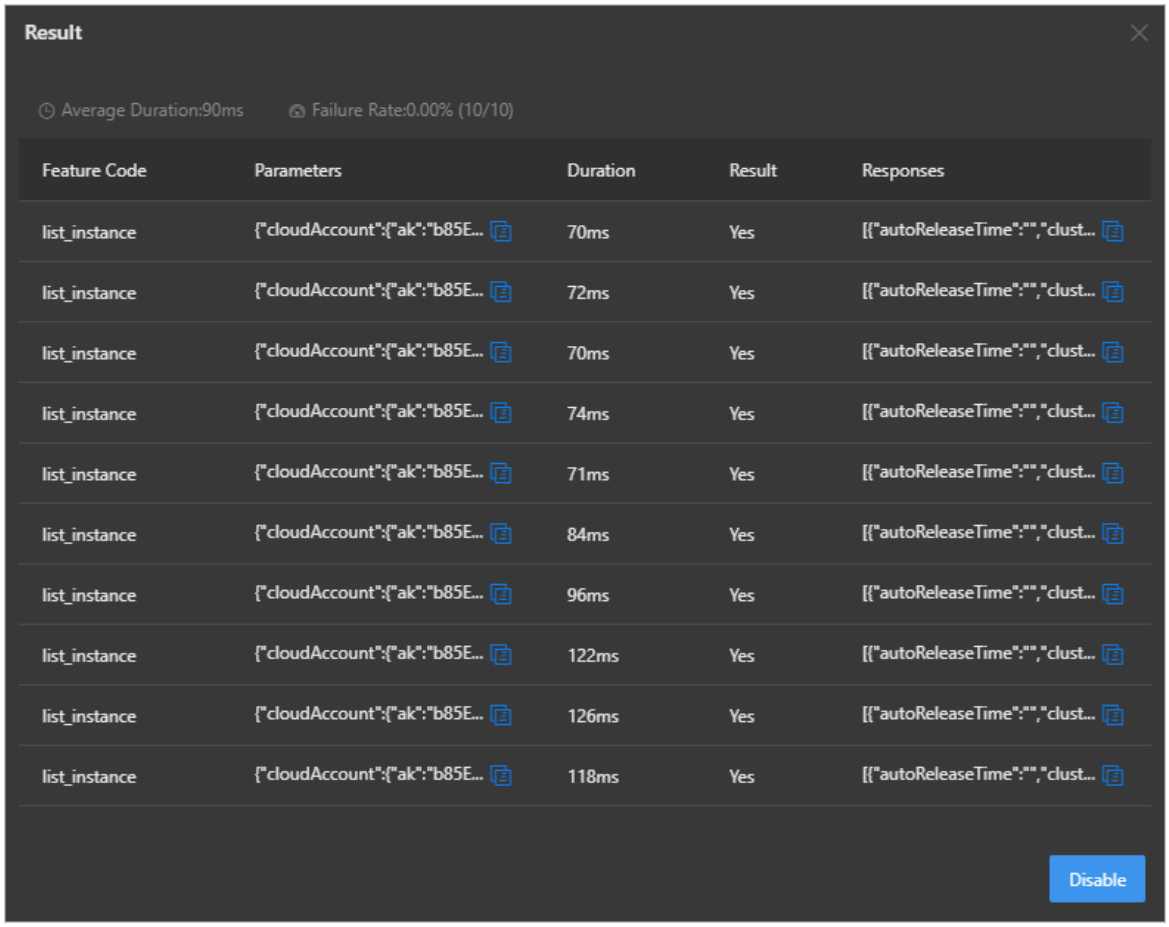

1.1.6.2.2. Event center

The Event Center feature manages STM events and provides root cause analysis. The feature allows you to quickly discover, locate, and resolve alerts or issues of the cloud platform and can improve the health of the cloud platform. You can use the Event Center feature to view the details, status, and level of events.

1.1.6.2.2.1. View events

On the Event Center page, you can view the level, summary, status, and handler of all alert events to learn about the health status of the cloud platform.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Event Center**.
4. (Optional)In the upper-left corner, enter a keyword in the search bar and click **Search**.

   > **Note** Click **Advanced** in the upper-right corner. Select the product, level, status, creation time and handler from the **Product**, **Event Level**, **Current Status**, **Created At**, and **Handled By** drop-down lists, and then click **Search** to run exact search. If you do not enter keywords or perform exact search, all events are displayed in multiple pages.

5. View the event list. The following table describes the columns of the list.

Column	Description	Example
Event Title	The name of the alert event.	test
Event Level	The severity level of the alert event. Valid values: ○ Abnormal: Business may be affected. The alert is relatively serious and must be handled in a timely manner. ○ Risky: You must continuously observe and determine whether the event affects business.	Abnormal
Event Summary	A brief description of the alert event.	testContent
Current Status	The current status of the event. Valid values: ○ New: the initial state. ○ Being Handled: The event is being handled and will be resolved. ○ Resolved: The event is resolved. ○ Canceled: The event is automatically restored in the initial state and with the intervention of a handler. ○ Unresolved: The event is in the New or Being Handled state.	New
Handled By	The current user in the Apsara Uni-manager Operations Console to handle the event.	xt****

1.1.6.2.2.2. View event details

On the Event Details page, you can view the basic information, dynamic information, and related alerts of a warning event to learn the details of the event in a timely manner.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Event Center**.
4. Find the event and click **Details** in the **Actions** column.

> **Note** You can also click the title of an event to view its details.

5. On the **Event Details** page, view the information. The following table describes the parameters.

Parameter	Description
Basic Information	○ ID: the ID of the event. ○ Product: the name of the product. ○ Status: the current status of the alert event. The event status is automatically generated. Usually, you do not need to modify it. To modify the status, select one of two values from the **Status** drop-down list. ■ Being Handled: The current event is being handled. ■ Resolved: The incident has been handled and resolved. ○ Assign To: If you want to transfer the event, you can select another user from the **Assign To** drop-down list. ○ Event Summary: a brief description of an alert event. ○ Event Content: the details of the event. ○ Root Cause Analysis: the diagnosis results for the related product when the event occurs, such as related alerts, product status, resource connectivity, host load, Pangu usage, and product inventory.
Comments	All system logs and comments for the alert event. You can also enter a comment in the box and click **Publish** to add a comment.
Related Alerts	The alerts that are highly relevant to the event. The alerts are obtained from the Alerts module of the Apsara Uni-manage Operations Console. The system also associates the event with OOS plans and recommends solutions. You can use the solutions to solve the event.

1.1.6.2.3. STM settings

1.1.6.2.3.1. Authorize account resource monitoring

In addition to monitoring the availability of the cloud platform, STM can also monitor the availability of cloud resources. You can add Apsara Stack accounts so that STM can monitor cloud resources that belong to the accounts.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > STM Settings**.
4. Click **Apsara Stack Account Authorization** in the upper-left corner. On the page that appears, you can view, add, or delete Apsara Stack accounts.
   - View: You can view the AccessKey pair information of the STM account in the list.
   - Add: Click **Create Apsara Stack Account** in the upper-left corner. In the dialog box that appears, enter AccessKey pair information of the STM account and then click **OK**.

     > **Note**
     > - AK: the AccessKey ID of the Apsara Stack account, used to identify the user.
     > - SK: the AccessKey Secret of the Apsara Stack account, used to verify the user.

   - Delete: Find the account and click **Delete** in the **Actions** column. In the message that appears, click **OK**.

1.1.6.2.3.2. Enable or disable STM inspection

You can enable or disable STM inspection for a product.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > STM Settings**.
4. Find a product and turn on or off the switch in the **STM Inspection** column to enable or disable STM inspection.

1.1.6.2.3.3. Authorize product resource monitoring

You can configure whether to monitor the cloud resources of a product.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > STM Settings**.
4. Find a product and turn on or off the switch in the **Authorization** column to enable or disable monitoring on its cloud resources.

1.1.6.2.4. Alert thresholds

1.1.6.2.4.1. View alert threshold settings

You can view alert threshold settings for features of products.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Alert Thresholds**.
4. You can view product names, feature names, RT thresholds, and success rate thresholds in the list.

1.1.6.2.4.2. Add alert threshold settings

You can add alert threshold settings for features of products.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Alert Thresholds**.
4. Click **Add Configuration** in the upper-left corner of the page. In the dialog box that appears, select a product and a feature, enter the RT threshold and success rate threshold, and then click **OK**.

   The following table describes the parameters.

Parameter	Description
Product	The product to which the alert threshold settings belong.
Feature	The feature to which the alert threshold settings belong.
RT Threshold (ms)	If this value is exceeded, the feature is regarded to be at risk.
Success Rate Threshold	If the call success rate of a feature is lower than this value, the feature is regarded to be abnormal.

1.1.6.2.4.3. Modify alert threshold settings

You can modify existing alert threshold settings.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Alert Thresholds**.
4. Find an alert threshold rule and click **Modify** in the **Actions** column. In the dialog box that appears, modify the parameters and click **OK**.

> **Note** For more information about the parameters, see Add alert threshold settings.

1.1.6.2.4.4. Delete alert threshold settings

You can delete existing alert threshold settings. After the alert threshold settings of a feature are deleted, the feature uses the common thresholds.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **STM > Alert Thresholds**.
4. Find an alert threshold rule and click **Delete** in the **Actions** column. In the message that appears, click **OK**.

1.1.6.3. Network Operation Center

The Network Operation Center (NOC) is a comprehensive operations tool platform that covers the entire network (virtual and physical).

NOC provides operations capabilities such as the visualization of network-wide monitoring, automated implementation, automated fault location, and network traffic analysis to enhance the efficiency of network operations engineers, reduce operations risks, and improve the quality of Apsara Stack services.

1.1.6.3.1. Dashboard

1.1.6.3.1.1. View the dashboard

You can view the status of the current devices, network, and traffic on the Dashboard tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Dashboard**.
4. On the **Dashboard** tab, view the dashboard information.

Item		Description
	Device Overview	The model distribution of the network devices in use.

Item		Description
Device Management	Ports Usage	○ **Ports Utilization**: the proportion of the number of ports in use to the total number of ports in the network devices. ○ **Error Packets by Port Top 5**: the total number of error packets generated by the device ports within a specified time range, of which the top 5 are displayed.
	Configuration Management	○ **Automatic Backup**: shows the proportions of **Backup Completed**, **Connection Failed**, and **Out-of-Scope** data sources. Move the pointer over the corresponding section and the details are displayed. ○ **Configuration Sync**: shows the proportions of **Configuration Synchronized**, **Connection Failed**, and **Out-of-Scope** data sources. Move the pointer over the corresponding section and the details are displayed.
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.
Traffic Dashboard	SLB Overview	The bandwidth usage of SLB clusters.
	XGW Overview	The bandwidth usage of XGW clusters.

1.1.6.3.1.2. View the network topology

You can view the physical network topology on the **Network Topology** tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Dashboard**.
4. Click the **Network Topology** tab.
5. On the **Network Topology** tab, view the physical network topology of a physical data center.

   You can set **Topology Type** to **Standard Topology** or **Dynamic Topology**.

If an offset exists between the dynamic topology and the standard topology, a message appears when you go to the **Network Topology** tab 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 normally.
> - Red: The connection has an error.
> - Grey: The connection is inactive.

By default, if **Topology Type** is set to **Standard Topology**, the **Refresh Alert** switch is turned on. You can turn off **Refresh Alert** to stop receiving new alerts that are triggered for the devices or connection statuses within the topology.

If **Topology Type** is set to **Dynamic Topology**, **Refresh Alert** is turned off.

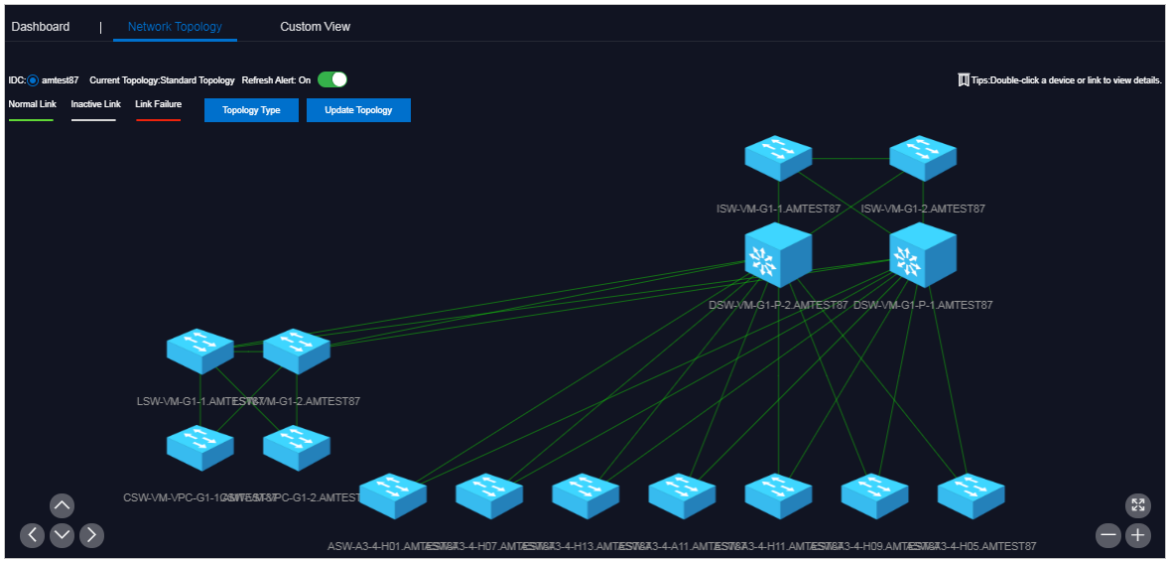

6. In the topology, double-click a connection between two devices to view the connection and alerts between the two devices.
7. In the topology, double-click a physical network device to view the basic information and node alerts of the device on the right.

1.1.6.3.1.3. Manage custom views

You can create a custom view to configure how to show the independent monitoring data set. You can configure the content and rules to display in the view to summarize and demonstrate the monitoring data and graph information you are interested in.

Go to the Dashboard page

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Dashboard**.

Create a view

1. Click the **Custom View** tab.
2. Create a view.
   i. In the upper part of the tab, click **Create View**.

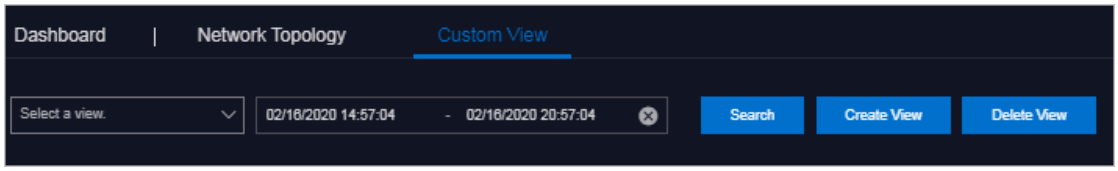

   ii. In the dialog box that appears, set View Name and Description and click **OK**.

   The view name cannot be the same as the name of an existing view. If the **A view with the same name already exists** message appears, you must change the view name to a unique one and then click **OK**.

3. Add a subview.

   By default, no subviews exist in a view after you create the view.

   i. Select the created view from the drop-down list in the upper part of the page, select a start time and an end time, and then click **Search**.

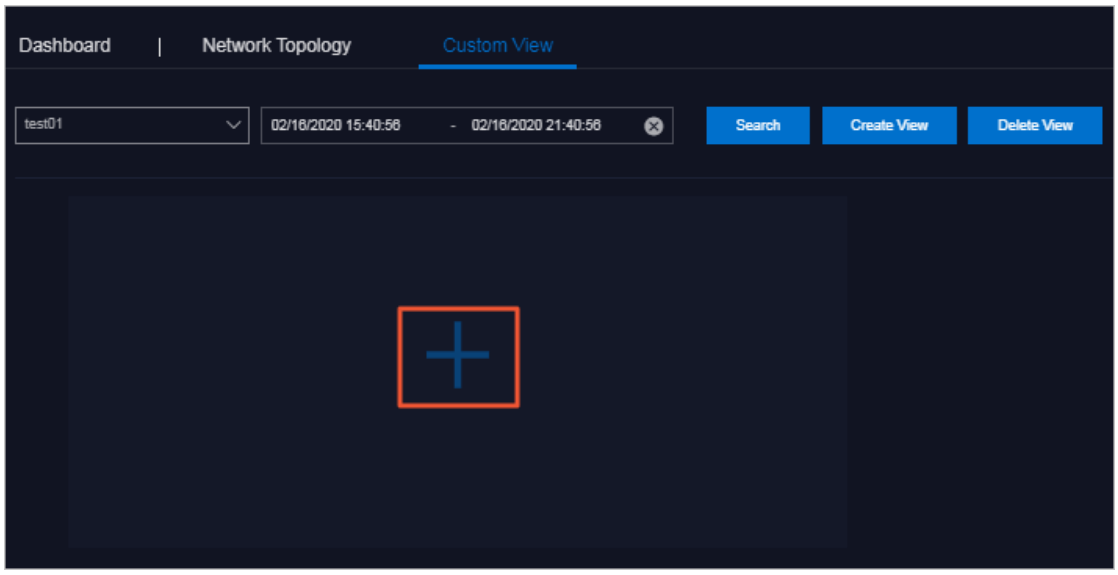

   ii. Click the icon.

iii. In the panel that appears, configure the parameters described in the following table.

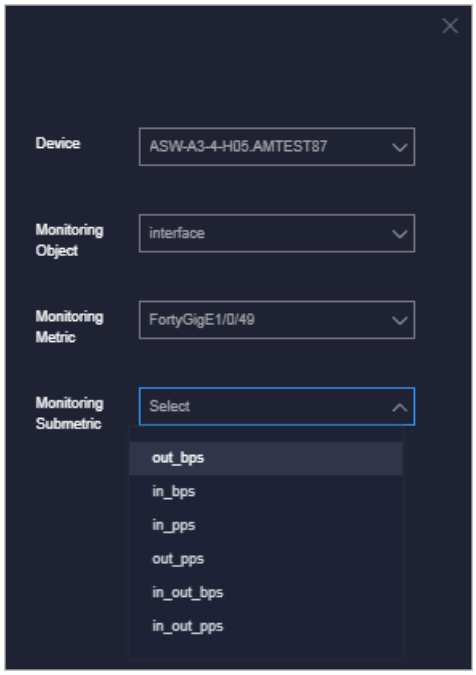

Parameter	Description
Device	Required. Select the device to be monitored from the drop-down list.
Monitoring Object	Required. Select the monitoring object from the drop-down list. ■ **interface**: the switch interface, including the water level, packet error, and packet loss of the interface. ■ **hardware**: the switch hardware, including the memory usage and CPU usage. ■ **capacity**: others, which is not supported.
Monitoring Metric	Required. Select the corresponding monitoring metric from the drop-down list.
Monitoring Submetric	Optional. Select the corresponding monitoring submetric from the drop-down list.

iv. Click **OK**.

After the subview is added, the system automatically shows the subview on the view to which the subview belongs.

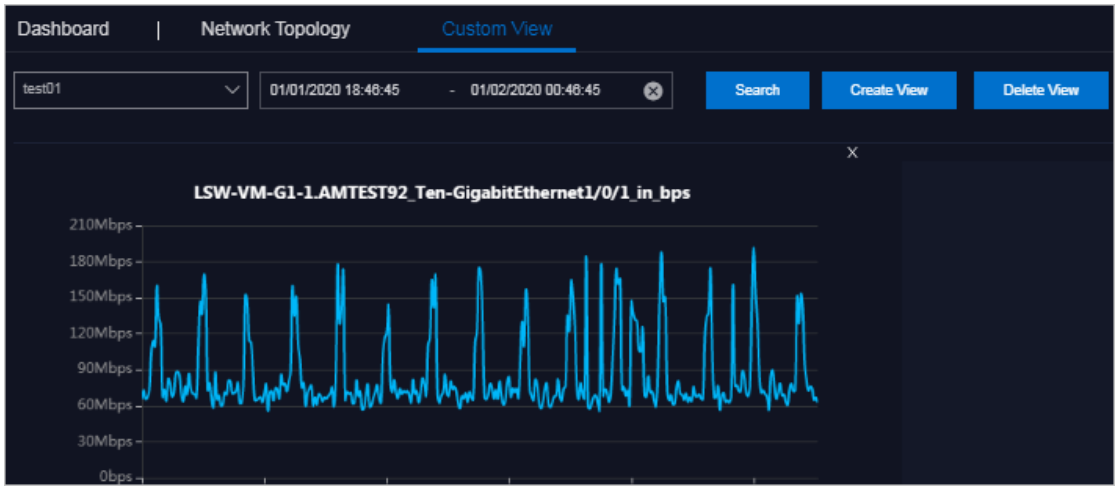

v. You can add other subviews.

Delete a subview

1. Click the **Custom View** tab.
2. Select the view to which the subview that you want to delete belongs from the drop-down list in the upper part of the page, select a start time and an end time, and then click **Search**.
3. Click the x icon in the upper-right corner of the subview.

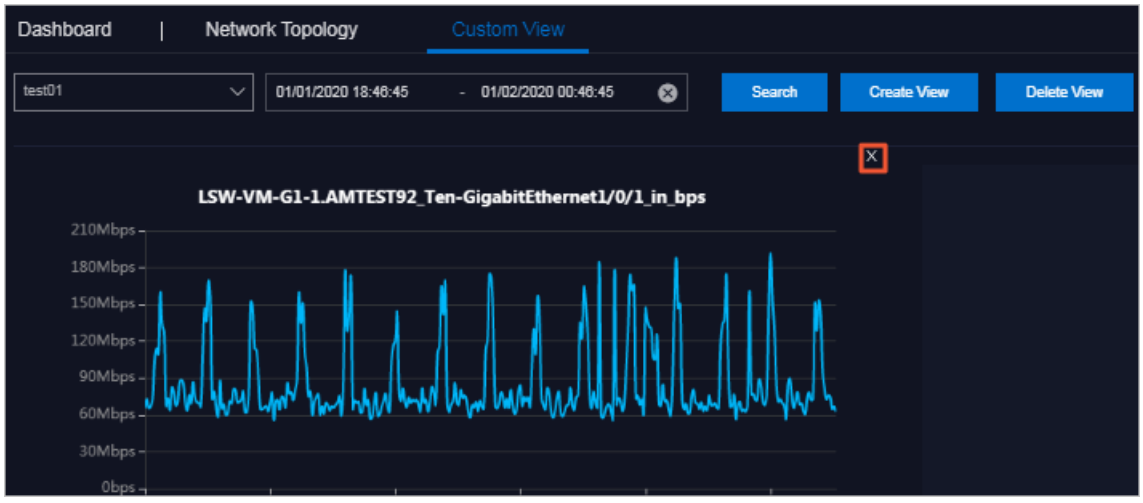

4. In the message that appears, click **OK**.

Delete a view

> **Notice** If you delete a view, all subviews of the view are also deleted. Proceed with caution.

1. Click the **Custom View** tab.
2. Select the view that you want to delete from the drop-down list in the upper part of the page, select a start time and an end time, and then click **Search**.

3. Click **Delete View** in the upper part of the tab.
4. In the message that appears, click **OK**.

1.1.6.3.2. Network element management

Network elements are network devices such as vSwitches and routers. The Network Element Management module shows the basic information and running status of physical network devices. The module also provides configuration management operations for physical network devices, including device management, password management, and configuration comparsion.

1.1.6.3.2.1. Device management

The Device Management module shows the basic information, running status, traffic monitoring information, and logs of physical network element devices. The module also allows you to configure the collection settings of network devices.

View network monitoring information

The Network Monitoring tab allows you to view the basic information, running status, and traffic monitoring information of Apsara Stack physical network devices and check the health status of network devices in a timely manner.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. On the **Device Management** tab, click the **Network Monitoring** tab.
5. 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 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 O&M personnel must troubleshoot the problem.
   - In the upper-right corner of the tab, enter the device name or IP address in the search box 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 link information of the device.

   a. In the upper-right corner of the **Port** tab, search for the port that you are about to view by using the search box. Click **View** in the **Details** column corresponding to the port.

   b. Select a time range on the right and click **Search** to view the traffic in the selected time range.

      You can select 5MIN, 30MIN, 1H, 6H for **Quick Query** to view the traffic within the last 5 minutes, 30 minutes, 1 hour, or 6 hours.

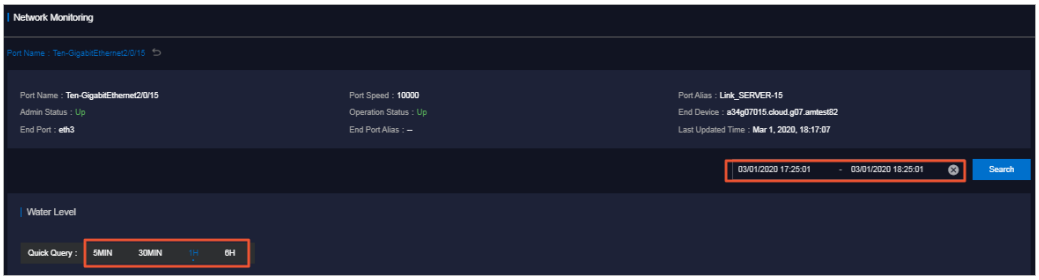

c. On the **CPU Utilization** tab, view all the CPU utilization information of the device.

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, you must closely monitor 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.

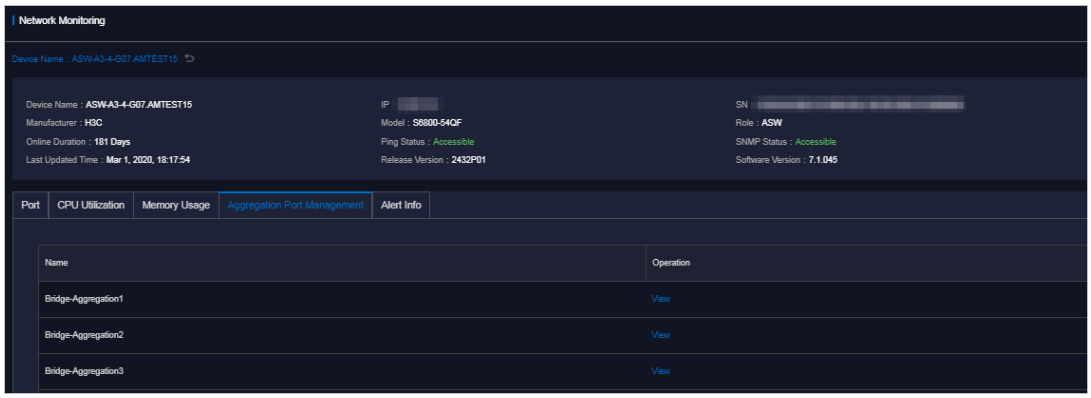

View logs

The Syslogs tab allows you to view logs of physical network element devices and provides necessary data for fault location and diagnosis information collection.

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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. On the **Device Management** tab, click the **Syslogs** tab.
5. 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 for system logs generated by the device within 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 are returned.

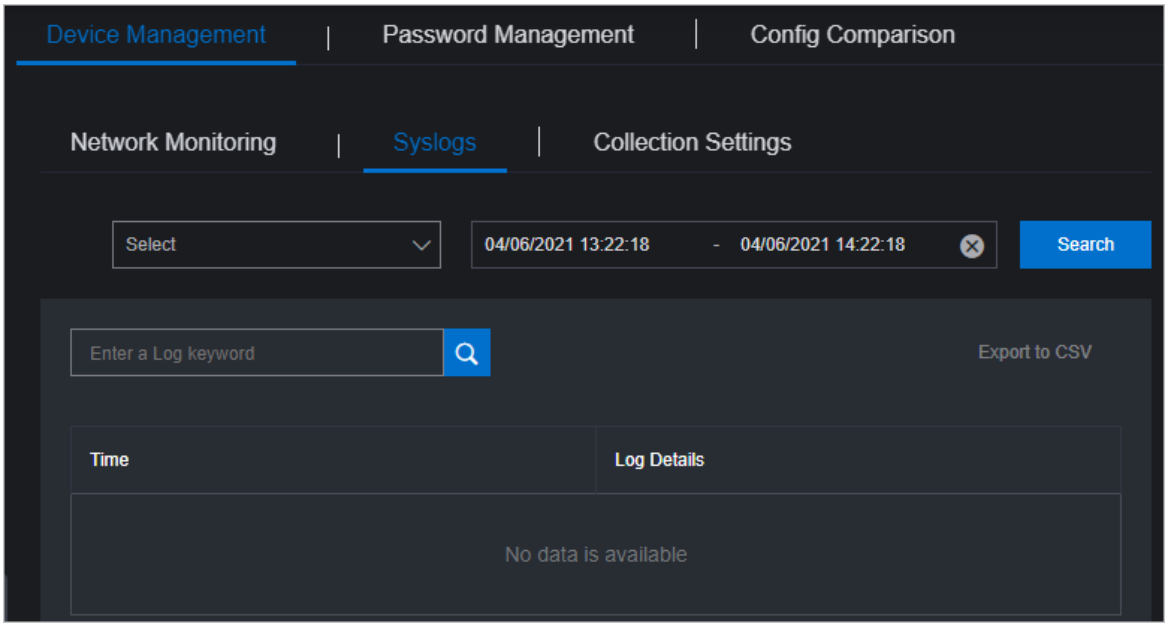

6. (Optional)You can filter the search results based on log keywords.
7. (Optional)Click **Export to CSV** in the upper-right corner to export the search results to your computer.

Collection settings

The Collection Settings tab allows you to set the collection interval of physical network element devices and manage OOB network segments.

Configure the collection interval

Before you collect network device information, you must configure a collection interval.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. On the **Device Management** tab, click the **Collection Settings** tab.

5. In the **Collection Interval Settings** section, configure the auto scan interval, device scan interval, port scan interval, and link scan interval.

   If you have no special requirements, we recommend that you use the initial default value.

6. Click **Submit**.

   Then, the system collects the device information based on your configurations.

Modify the collection interval

This topic describes how to modify the interval at which network device information is collected.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. On the **Device Management** tab, click the **Collection Settings** tab.
5. In the **Collection Interval Settings** section, modify the parameter values.

   > **Note** To cancel your modification before you submit it, click **Reset** in the upper-right corner to reset the collection interval to the previous version.

6. Click **Submit**.

   The modified collection interval of the network device information takes effect after 1 minute.

Add an OOB network segment

If this is the first time you are using the Network Elements feature of Network Operations Center (NOC), you must add the device loopback network segment planned by the current Apsara Stack network device, which is typically the network segment of the netdev.loopback field in the IP address planning list.

Context

The OOB Network Segments section is used to configure the management scope of a physical network element device. Typically, O&M engineers must add the loopback network segment in which the network device to be managed resides.

In the Apsara Stack scenario, a loopback network segment is used to configure the management scope of a physical network element device. To expand the network and the loopback network segment, you must add the network segment involved in the expansion to the management scope. The procedure to add an expanded network segment is the same as that used to add the loopback network segment for the first time. Then, you can search for the network segment of the managed device on this page.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the left-side navigation pane, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. On the **Device Management** tab, click the **Collection Settings** tab.

5. In the lower part of the **OOB Network Segments** section, click **Add Network Segment**.
6. In the Add Network Segment dialog box, enter the network segment that contains the mask information and a subnet mask and select an IDC.

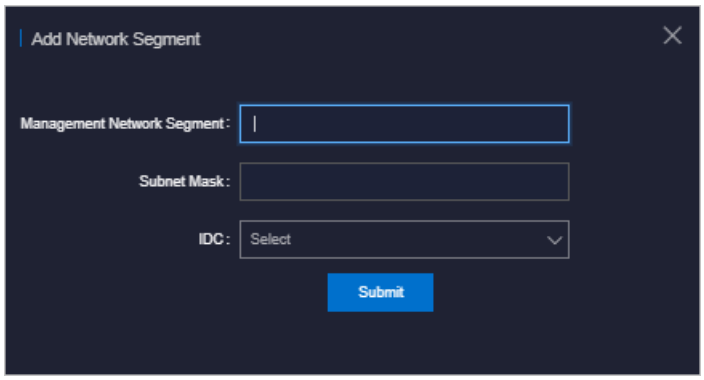

7. Click **Submit**.

   The initial data entry is complete.

   To modify or delete an OOB network segment, find it in the list and click **Edit** or **Delete** in the **Actions** column.

View the OOB network segment information

You can search for and view the network segment information of your managed devices.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. On the **Device Management** tab, click the **Collection Settings** tab.
5. In the **OOB Network Management** section, click **Refresh** in the upper-right corner of the section.

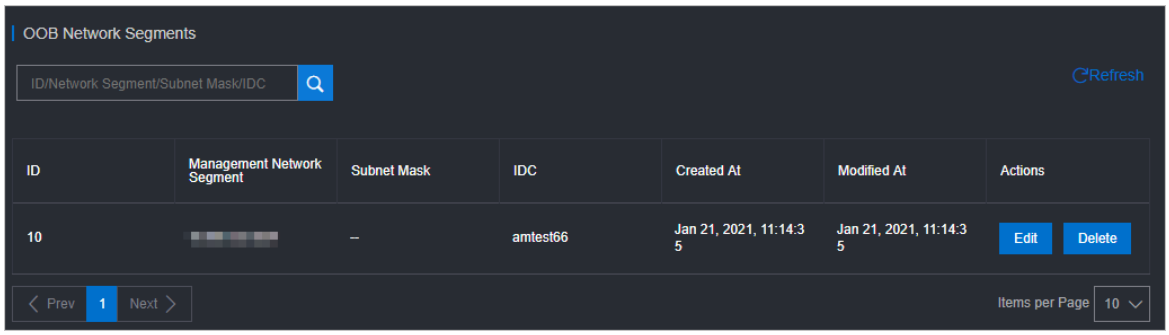

6. In the list, view the network segment information of your managed devices.

> **Note** You can search for the information of a specific network segment by entering keywords in the search box.

1.1.6.3.2.2. Modify the device password

You can modify the passwords of physical network devices.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. Click the **Password Management** tab.
5. (Optional)Enter the name of the device for which you want to modify the password in the search box of the **Devices on Live Network** section and click **Search**.

   To search for another device, you can click **Reset** to reset the previous search conditions.
6. Select one or more devices and 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.
7. The system must verify the old password before you modify it. Enter **Username** and **Old Password** in the lower-right corner and click **Verify**.

   You must verify the old passwords for all the devices in the **Target Devices** section.
8. After the verification is passed, you can modify the password for one or more devices.
   - Modify the password of a device
     a. Add a device to the Target Devices section at a time. Or choose **Manage > Delete** in the **Actions** column corresponding to a device for which you do not modify the password to remove the device.
     b. In the lower part of the **Target Devices** section, click **Modify**.
     c. In the dialog box that appears, enter and confirm the new password, and click **OK**.
   - Modify the passwords of all devices
     a. In the lower part of the **Target Devices** section, click **Modify**.
     b. In the dialog box that appears, enter and confirm the new password, and then click **OK**. The passwords of all the devices that are added to the **Target Devices** section and that are in the **Accessible** and **Verified** are modified.

1.1.6.3.2.3. Compare device configurations

You can compare the current configuration of a device with its configuration on startup and check whether they are consistent.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.
4. Click the **Config Comparison** tab.

5. (Optional)Enter the name of the device whose configurations you want to compare in the **Device Name** search box and click **Search**.

   To search for more devices, you can click **Reset** to reset the configured search condition.

6. Select devices and click **Compare Configuration**.

   After you compare the configurations, click **Refresh** and click **Export Results**.

1.1.6.3.3. SLB cluster management

After you add SLB cluster tags, you can easily select clusters in the tenant console. This improves the ease of use of SLB clusters.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > SLB Cluster Management**.
4. Select a cluster and click **Search** to view the cluster information.
5. Find the cluster and click **Modify Tag** in the **Actions** column. In the dialog box that appears, modify the tag name and click **OK**.

   > **Note** The default cluster tag **default** cannot be modified.

1.1.6.3.4. SLB management

The SLB Management module contains the Cluster Monitoring and Instance Monitoring tabs and shows the basic information, running status, and usage of SLB network products.

1.1.6.3.4.1. View 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > SLB Management**.
4. The **Cluster Monitoring** tab appears.
5. Select the cluster that you want to view from the drop-down list and click **Search**.

   Information of all the device nodes in the cluster is displayed.

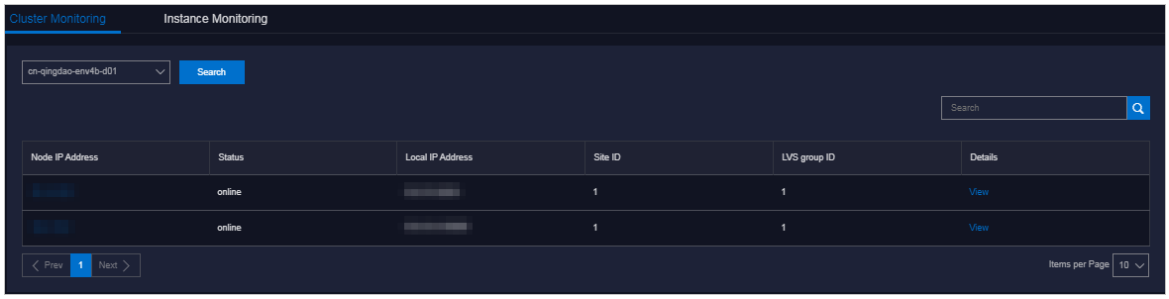

6. Find a device node and click **View** in the **Details** column.
7. On the **Node Message** page, view the basic information, inbound limit (bit/s), outbound limit (bit/s), inbound limit (PPS), outbound limit (PPS), active connection limit, and inactive connection limit of the device node.

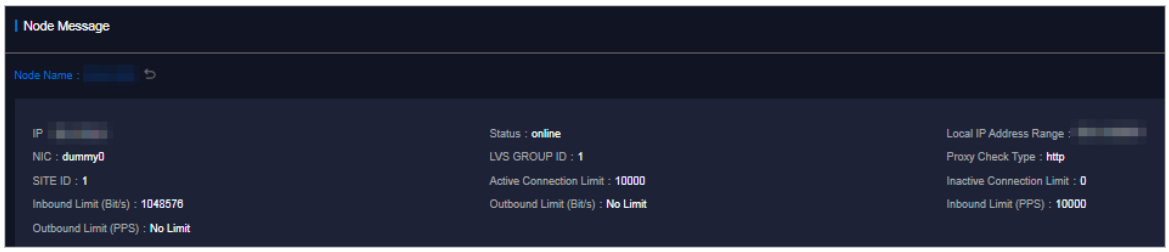

8. On the **Node Message** page, view the usage information of the device node.

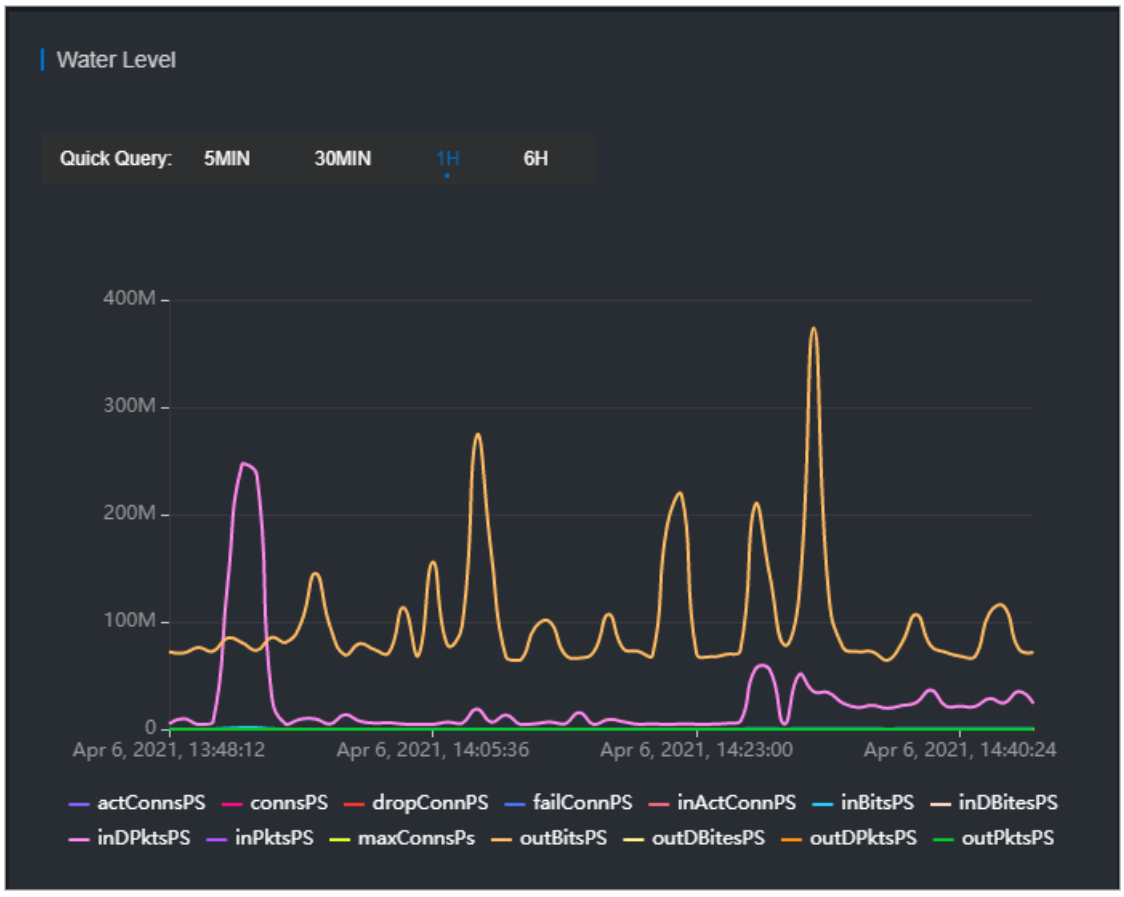

- In the upper-right corner, you can set the start and end time and click **Search** to view the usage data within the specified time range.
- You can also click **5MIN**, **30MIN**, **1H**, or **6H** in the upper-left corner to query the usage data

within the corresponding time range.

- Click a metric in the lower part of the chart, and the curve corresponding to the metric disappears from the chart. Click the metric again, the curve appears.
- Move the pointer over a point in time to display the values of all metrics for that point in time.

Metric	Description	Example
actConnsPS	The number of active connections.	0
connsPS	The number of new connections.	1
dropConnsPS	The number of connections dropped per second.	0
failConnPS	The number of connections failed per second.	0
inActConnPS	The number of inactive connections.	1
inBitsPS	The amount of inbound data per second.	248
inDBitesPS	The amount of inbound data dropped per second.	0
inDPktsPS	The number of inbound packets dropped per second.	0
inPktsPS	The number of inbound packets per second.	1
maxConnsPs	The total number of connections.	1
outBitsPS	The amount of outbound data per second.	208
outDBitesPS	The amount of outbound data dropped per second.	0
outDPktsPS	The number of outbound packets dropped per second.	0
outPktsPS	The number of outbound packets per second.	1

1.1.6.3.4.2. 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > SLB Management**.
4. Click the **Instance Monitoring** tab.
5. Select the cluster where the instance that you want to view is located from the Cluster drop-down list. Enter the ID or VIP address of the instance that you want to query in the search box and click **Search**.
6. In the query result, view the monitoring information of the instance.
   - The first section shows the basic information of the SLB instance, which allows O&M 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 click **Search**, or select 5MIM, 30MIN, 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.

1.1.6.3.5. Collect IP addresses

The system routinely collects the IP addresses of all physical networks within the current Apsara Stack environment at a specified collection interval. You can use a CIDR block or IP address, and subnet mask to search for the information of corresponding devices and ports.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IP Address Collection**.
4. 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.

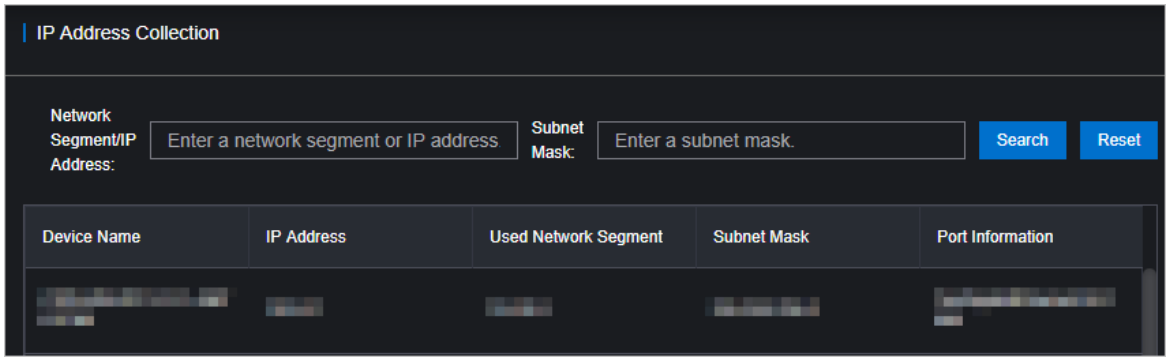

1.1.6.3.6. IP address range management

The IP Address Ranges module allows you to manage planning information in the Apsara Stack environment, including the network architecture and address planning. You can modify, import, and export the planning information.

1.1.6.3.6.1. Import the planning file

No data is imported at the time 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 table is obtained from the Apsara Stack deployment planner. If you have not obtained the allocation table, contact your account manager or submit a Apsara Stack ticket.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IP Address Ranges**.
4. Click **Import** in the upper-right corner.
5. In the dialog box that appears, click **Select a file to upload** or **Browse** and select the IP address allocation list.
6. Click **Import**.

1.1.6.3.6.2. Manually add the IP address pool information

You can also manually add new IP address pool information to the Apsara Uni-manager Operations Console for centralized management.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Console > IP Address Ranges**.
4. Click **Add**.
5. In the dialog box that appears, configure the IP address pool information.

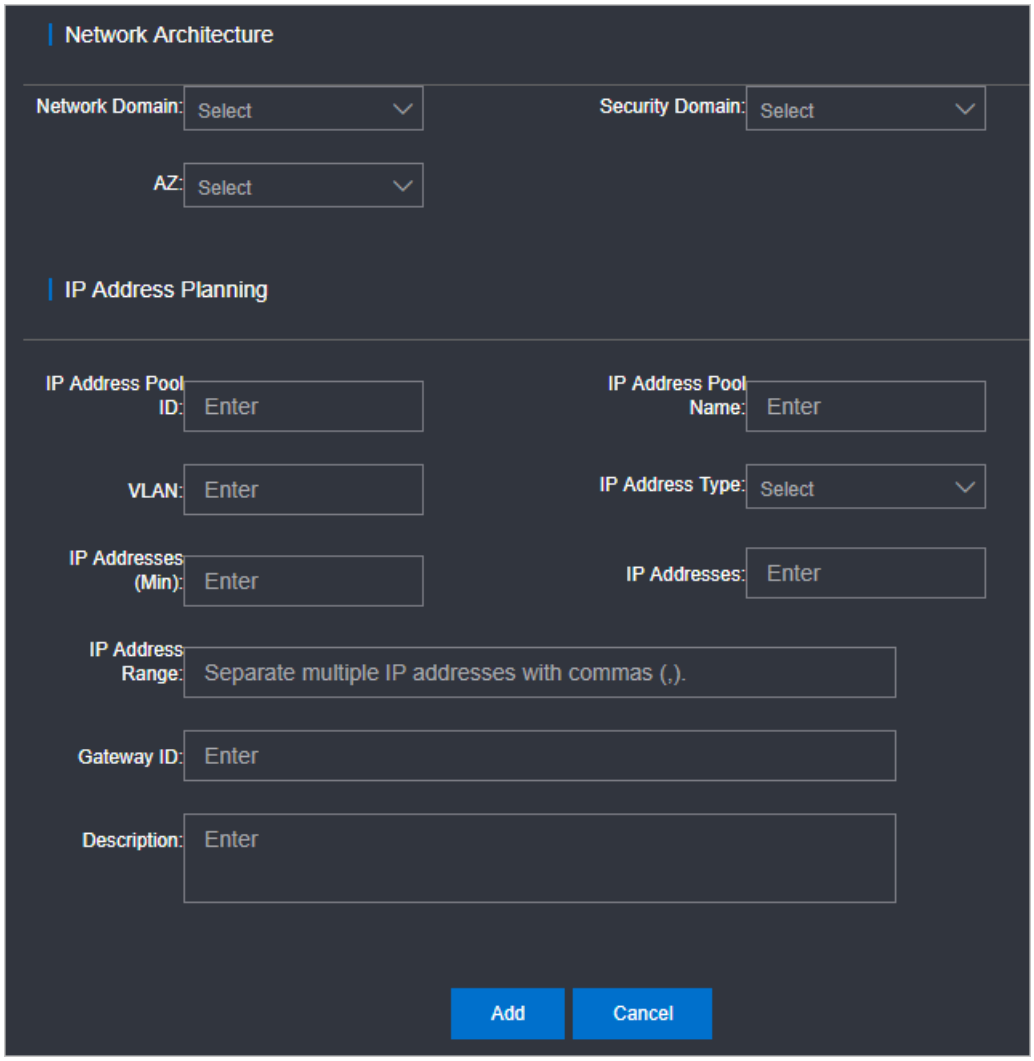

6. Click **Add**.

1.1.6.3.6.3. Modify the IP address pool information

If an IP address range is changed, you can modify the IP address pool information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IP Address Ranges**.
4. (Optional)On the **IP Address Ranges** page, configure the search conditions and click **Search**.

   > **Note** To reset the search conditions, you can click **Reset** to clear your configurations with one click.

5. Find the IP address pool for which you want to modify the configurations and choose **Manage > Modify** in the **Actions** column.
6. In the dialog box that appears, modify the network architecture and IP address planning.

7. Click **Edit**.

1.1.6.3.6.4. Export the IP address pool information

You can export the IP address pool information to your computer and then view the information offline.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IP Address Ranges**.
4. Select IP address pools that you want to export and click **Export**.

1.1.6.3.6.5. Delete an IP address pool

You can delete IP address pools that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IP Address Ranges**.
4. Find the IP address pool that you want to delete and choose **Manage > Delete** in the **Actions** column.

1.1.6.3.7. View Anytunnel information

You can view the Anytunnel information to see the Anytunnel resources registered by projects within the current environment or whether a project has Anytunnel registered. The system allows you to query the registration information of Anytunnel resources based on the project, cluster, service instance, and server role. You can use the global query feature to query the usage of all the Anytunnel resources in the current environment.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Anytunnel Management**.
4. Query all information or specific information. Perform the following steps:
   - In the upper part of the page, click **Query Details** to view all the Anytunnel information in the environment.

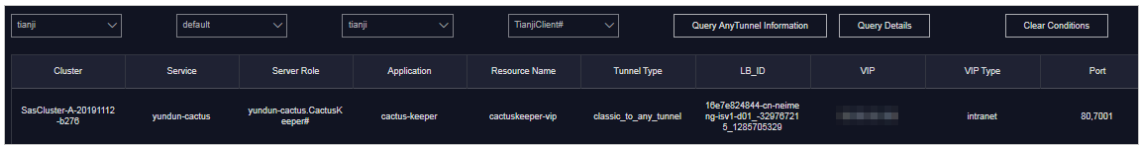

   - In the upper part of the page, select the project, cluster, service instance, or service role, and

then click **Query AnyTunnel Information** to view the AnyTunnel information that meets the search conditions.

> **Note** You can click **Clear Conditions** and modify the search conditions.

1.1.6.3.8. XGW management

The XGW Management module allows you to manage VPC gateways and view the usage information of device nodes and service instances.

1.1.6.3.8.1. View node information

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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > XGW Management**.
4. Select the cluster that you want to view from the drop-down list and click **Search**.

   The system shows the basic information and usage information of aggregated traffic of all device nodes within the selected cluster. By default, the usage information of the last one hour is displayed. You can select 1 hour, 3 hours, 6 hours, or one day as the time range, or customize the time range to search for the usage information.

5. Find a device node and click **View** in the **Details** column.
6. On the page that appears, view the traffic usage information of the device node.

1.1.6.3.8.2. View the instance monitoring information

The Instance Monitoring tab allows you to view information such as bps, pps, drop_speed, fin_speed, ratelimit_drop_speed, rst_speed, and syn_ack_speed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > XGW Management**.
4. Click the **Instance Monitoring** tab.
5. Select the cluster where the instance you want to view is located from the drop-down list, enter the instance ID (EIP address) in the search box, and then click **Search**.

   By default, the data information within the last one hour is displayed. You can select a time range such as 5 minutes, 30 minutes, 1 hour, or 6 hours.

Metric	Description
in_bps	The data receiving rate. Unit: bit/s.
in_drop_speed	The inbound packet loss rate. Unit: PPS.
in_fin_speed	The inbound Fin packet forwarding rate. Unit: PPS.
in_pps	The inbound packet forwarding rate. Unit: PPS.
in_ratelimit_drop_speed	The inbound throttling packet loss rate. Unit: PPS.
in_rst_speed	The inbound RST packet forwarding rate. Unit: PPS.
in_syn_ack_speed	The inbound SYN/ACK packet forwarding rate. Unit: PPS.
out_bps	The data transmission rate. Unit: bit/s.
out_drop_speed	The outbound packet loss rate. Unit: PPS.

Metric	Description
out_fin_speed	The outbound FIN packet forwarding rate. Unit: PPS.
out_pps	The outbound packet forwarding rate. Unit: PPS.
out_ratelimit_drop_speed	The outbound throttling packet loss rate. Unit: PPS.
out_rst_speed	The outbound RST packet forwarding rate. Unit: PPS.
out_syn_ack_speed	The outbound SYN/ACK packet forwarding rate. Unit: PPS.

1.1.6.3.9. CGW management

1.1.6.3.9.1. View node information

The CGW Management module allows you to view the basic information, running status, aggregated traffic, and usage of each device node of CGW network products.

Prerequisites

The NetworkBaseServiceCluster cluster for BaseServiceAll has been deployed and has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > CGW Management**.
4. Select the cluster that you want to view from the drop-down list and click **Search**.

   > **Note** The system shows the basic information and usage information of aggregated traffic of all device nodes within the selected cluster. By default, the usage information of the last one hour is displayed. You can also select **3 Hours**, **6 Hours**, or **1 Day** as the time range, or customize the time range to search for the usage information.

5. Find a device node and click **View** in the **Details** column. On the page that appears, view the traffic usage information of the device node.

   The following table describes the parameters.

Parameter	Description
InByteRate	The data receiving rate. Unit: bit/s.
OutByteRate	The data transmission rate. Unit: bit/s.
PacketLossRateIn	The inbound packet loss rate. Unit: PPS.

Parameter	Description
PacketLossRateOut	The outbound packet loss rate. Unit: PPS.
PacketRateIn	The inbound packet rate. Unit: PPS.
PacketRateOut	The outbound packet rate. Unit: PPS.

1.1.6.3.9.2. View instance information

You can view the usage information of instance metrics.

Prerequisites

The NetworkBaseServiceCluster cluster for BaseServiceAll has been deployed and has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > CGW Management**.
4. Click the **Instance Monitoring** tab. Select a cluster from the drop-down list, enter the instance ID(Tunnel ID), and click **Search**. You can view the usage charts of the instance.

   > **Note** By default, the usage information of the last one hour is displayed. You can also select **5MIN** (5 minutes), **30MIN** (30 minutes), **1H** (1 hour), and **6H** (6 hours) as the time range to search for the usage information.

   The following table describes the parameters.

Parameter	Description
InBps	The data receiving rate. Unit: bit/s.
InDropSpeed	The inbound packet loss rate. Unit: PPS.
InFinSpeed	The inbound Fin packet forwarding rate. Unit: PPS.
InPps	The inbound packet forwarding rate. Unit: PPS.
InRateLimitDropSpeed	The inbound throttling packet loss rate. Unit: PPS.
InRstSpeed	The inbound RST packet forwarding rate. Unit: PPS.
InSynAckSpeed	The inbound SYN/ACK packet forwarding rate. Unit: PPS.
OutBps	The data transmission rate. Unit: bit/s.

Parameter	Description
OutDropSpeed	The outbound packet loss rate. Unit: PPS.
OutFinSpeed	The outbound FIN packet forwarding rate. Unit: PPS.
OutPps	The outbound packet forwarding rate. Unit: PPS.
OutRateLimitDropSpeed	The outbound throttling packet loss rate. Unit: PPS.
OutRstSpeed	The outbound RST packet forwarding rate. Unit: PPS.
OutSynAckSpeed	The outbound SYN/ACK packet forwarding rate. Unit: PPS.

1.1.6.3.10. Firewall management

If the cloud firewall is deployed in your environment, you can use the firewall 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 firewall feature to isolate or restore the firewall.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Cloud Firewall Management**.
4. Select the operation type, firewall type, and data center from the corresponding drop-down list and click **Confirm**.

   Supported operations:

   - **Isolate Firewall**: physically isolates the firewall from the network structure. If an exception occurs on the cloud firewall service, the system removes the firewall device from the network forwarding path to ensure that the normal business traffic forwarding is not affected.
   - **Restore Firewall**: restores the firewall from the network isolated state to the normal state. After the exception on the cloud firewall is resolved, the system restores the firewall device back to the network forwarding path to ensure that the firewall is restored to the initial online status.

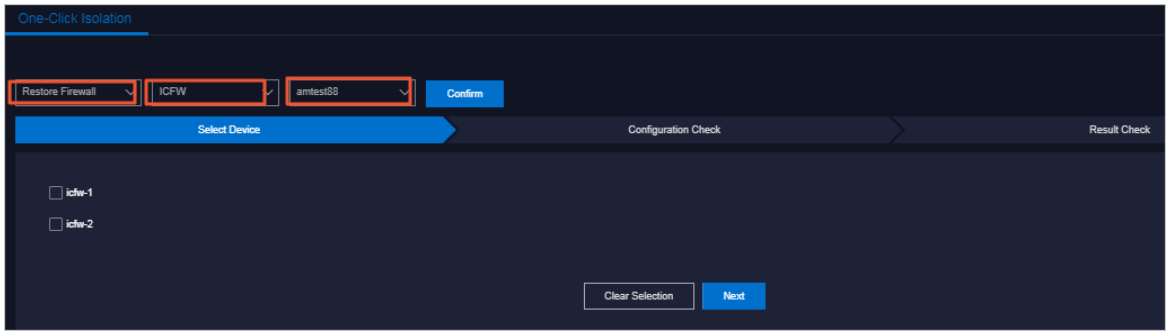

5. In the **Select Device** step, select devices and click **Next**.
6. In the **Configuration Check** step, check the selected devices and template information. If the information is correct, click **Confirm**.

7. 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.
8. In the **Result Check** step, click **Details** in the **Details** column corresponding to each device to view the corresponding result.

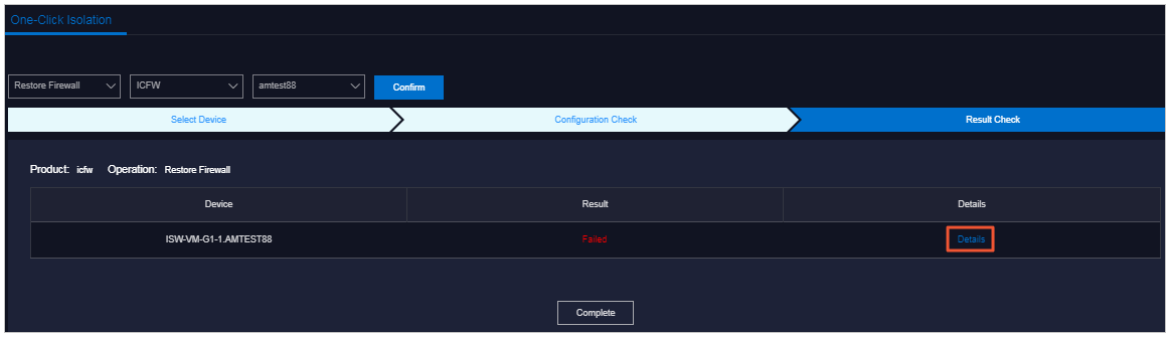

9. Click **Complete**.

1.1.6.3.11. Alerts

1.1.6.3.11.1. View and process current alerts

You can view and process current alerts on the Current Alerts tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Alert Dashboard**.
4. Click the **Current Alerts** tab.
5. In the upper-right corner, enter a keyword in the search box and click **Search**.

   Alerts that meet the search conditions are displayed.

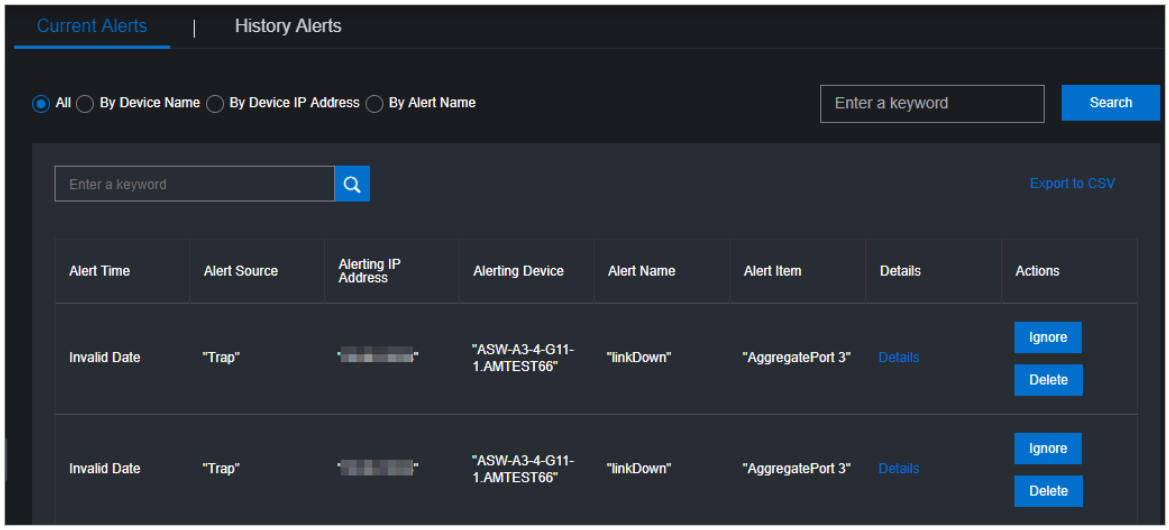

6. (Optional)Filter the search results by the device name, device IP address, or alert name.
7. Find an alert and move the pointer over **Details** in the **Details** column to view the detailed alert information.
8. Find the reason why the alert is triggered and then process the alert.
   - If the alert does not affect the operation of the system, you can click **Ignore** in the **Actions** column corresponding to the alert. In the **Confirm Operation** message, click **OK** to ignore the alert.
   - If the alert is no longer significant, you can click **Delete** in the **Actions** column corresponding to the alert. In the **Confirm Operation** message, click **OK** to delete the alert.

     After the alert is deleted, you can query it on the **History Alerts** tab.
9. (Optional)Click **Export to CSV** to export the alert information to your computer.

1.1.6.3.11.2. View historical alerts

You can view historical alerts on the History Alerts tab.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Alert Dashboard**.
4. Click the **History Alerts** tab.
5. 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 click **Search**.

Alerts that meet the search conditions are displayed.

6. Click **Details** in the **Details** column corresponding to an alert to view detailed information about the alert.
7. (Optional)Click **Export to CSV** to export the alert information to your computer.

1.1.6.3.12. Alert settings

1.1.6.3.12.1. Add a trap

If the initially configured trap subscription does not meet the monitoring requirements, you can add a trap for monitoring match.

Context

Simple Network Management Protocol (SNMP) traps are used in this topic. SNMP trap is a part of SNMP and a mechanism that enables devices being managed (here refers to network devices such as switches and routers) to send SNMP messages to NOC monitoring servers. If an exception occurs on the device being monitored or the switch monitoring metrics have an exception, the SNMP agent running in the switch sends an alert event to the NOC monitoring server.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Alert Settings**.
4. On the **Alert Settings** page, click **Configure Trap**.
5. In the **Configure Trap** dialog box, configure the parameters.

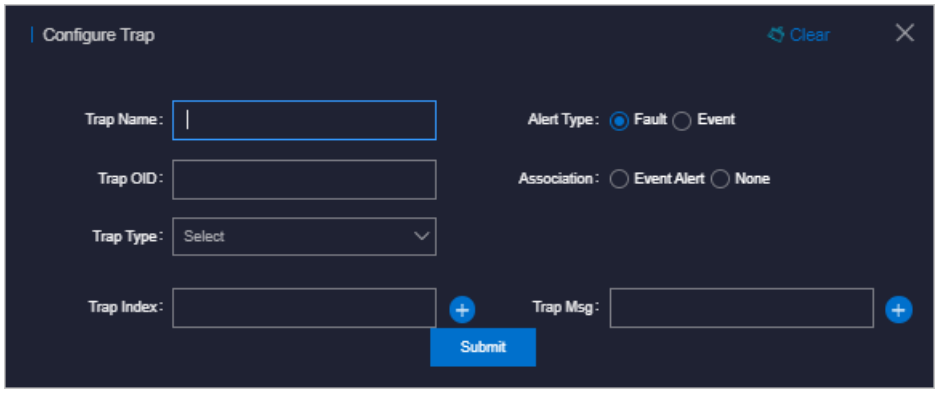

The following table describes the parameters.

Parameter	Description	Example
Trap Name	The name of the alert event.	linkdown or BGPneighbor down. You can customize the value.

Parameter	Description	Example
Trap OID	The OID of the alert event.	.1.3.6.1.4.1.25506.8.35.12.1.12 The value must be configured based on the device document and cannot be customized.
Trap Type	The type of the alert event.	N/A
Trap Index	The index ID of the alert item. The KV information in trap messages, which is used to identify alert objects. Typically, the value of this parameter can be an API name, protocol ID, or index ID. You can configure multiple values for this parameter. The value must be configured based on the device document and cannot be customized.	N/A
Trap Msg	The message of the alert item. The KV information in trap messages, which is used to identify the alert data. Typically, the value of this parameter 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. You can configure multiple values for this parameter. The value must be configured based on the device document and cannot be customized.	N/A
Alert Type	Specifies whether the alert is of the fault type or the event type.	N/A

Parameter	Description	Example
Association	Specifies whether the alert has an event alert. **Note** If **Event** is selected, you must enter the associated alert trap configurations.	N/A

6. 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 trap is configured.

   After the trap is added, the alert events of the configured Trap OID are monitored and are displayed on the **Current Alerts** and **Alert History** tabs of the **Alert Management** module.

1.1.6.3.12.2. View traps

You can view traps configured in the current system.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Alert Settings**.
4. Enter a keyword in the search box in the upper-right corner and click **Search**.

   Note You can click **Export to CSV** in the upper-right corner of the list to export the trap information to your computer.

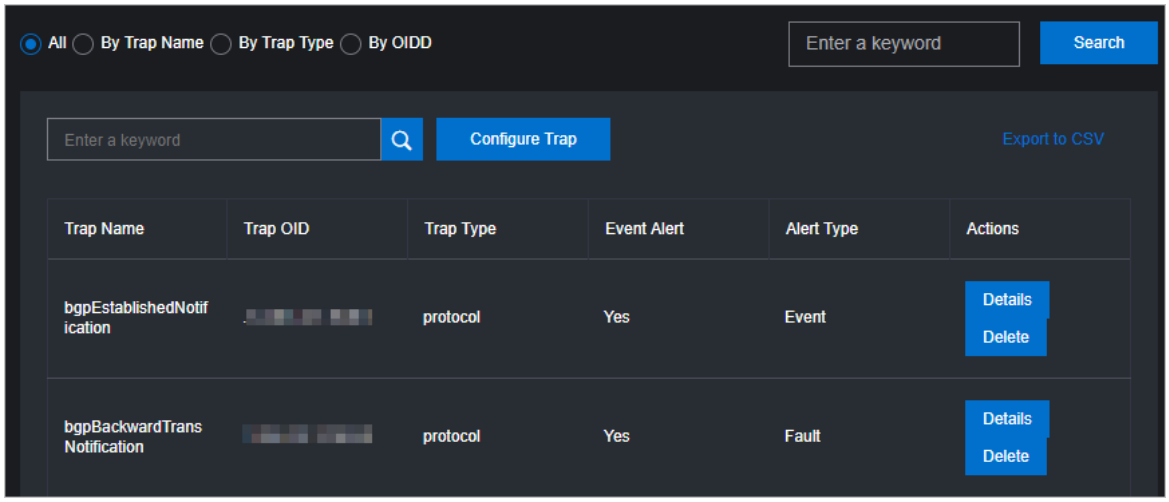

5. (Optional)Filter the search results by trap name, trap type, or OID.
6. Move the pointer over **Details** in the **Actions** column corresponding to a trap to view detailed information about the trap.

> **Note** If a trap is no longer needed, you can click **Delete** in the **Actions** column.

1.1.6.3.13. Physical network integration

The Physical Network Integration module allows network operations engineers to automate the integration of physical networks in the Apsara Uni-manager Operations Console by specifying the integration parameters. Network Operations Center (NOC) automatically generates and issues the configurations to specified devices and then performs the network integration test.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Physical Network Integration**.
4. Enter a project name and click **Create**.

   You must create a project file for this change to store the parameters related to the change. In the **History** section, you can choose **Manage > Import** in the Actions column.
5. In the upper-right corner, click **Save Project** to save the project details. Click **Next**.
6. 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 click **Search**.

      After you add a device, you can click **Reset** to clear the search condition and search for other devices to add to the list.
   ii. Find a device and click **Add** in the Actions column to add it to the Target Device section on the right.

      To remove a device from the Target Device section, choose **Manage > Delete** in the Actions column corresponding to the device. You can also choose **Manage > Set the username and password** in the Actions column 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.
   iv. Click **Next**.
7. Configure the interface parameters.
   i. In the **Configure Interfaces** step, click **Edit** in the Actions column. The **Configure Interfaces** section appears.
   ii. In the **Configure Interfaces** section, configure the parameters and click **Add**.

      You can choose **Manage > Edit** or **Manage > Delete** to modify or delete the added interface.
   iii. In the upper-right corner, click **Save Project** to save the configurations.
   iv. Click **Next**.
8. Configure the route parameters.
   i. In the **Configure Routes** step, click **Edit** in the Actions column. The **Configure Routes** section appears.

ii. In the **Configure Routes** dialog box that appears, configure the parameters and click **Add**.

You can choose **Manage > Edit** or **Manage > Delete** to modify or delete the added route.

iii. In the upper-right corner, click **Save Project** to save the information.

iv. Click **Next**.

9. Configure the route policies.

i. In the **Configure Route Policies** step, click **Edit** in the Actions column. The **Configure Route Policies** section appears.

ii. In the Configure Route Policies section, configure the parameters and click **Add**.

You can choose **Manage > Delete** or **Manage > Delete** in the Actions column to modify or delete the added route policy.

iii. In the upper-right corner, click **Save Project** to save the information.

iv. Click **Next**.

10. In the **Generate Integration Configuration** step, click **Generate**.

The system generates the integration configuration commands and rollback commands for all of the devices that have parameters configured.

O&M engineers can generate configurations for each device based on the configured parameters. After the configurations are generated, 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 computer.

1.1.6.3.14. ASW scale-up

This topic describes how to use NOC to automatically scale up ASW devices. 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > ASW Scale-up**.
4. Select devices.

i. In the **Select Device** step, enter a device name in the search box of the **Devices on Live Network** section and click **Search**.

After you add a device, you can click **Reset** to clear the search condition and search for other devices to add to the list.

ii. Find a device and click **Add** in the Actions column to add it to the Target Device section on the right.

To remove a device from the Target Device section, choose **Manage > Delete** in the corresponding column. You can also modify the logon username and password of a device by choosing **Manage > Set the username and password** in the Actions column.

5. Click **Next**.
6. Disable the DSW port.
   i. In the **Disable DSW Port** step, find the device for which you want to disable the DSW port and click **Port Settings**.
   ii. Disable the port and click **Implement**.
   iii. In the message that appears, click **OK** to run the script.
7. Click **Next**.
8. Configure the DSW port.
   i. In the **Configure DSW Port** step, find the device for which you want to configure the DSW port and click **Edit** in the Actions column. The **Interface Parameter Configuration** section appears.
   ii. In the **Interface Parameter Configuration** section, set **Display Ports**, **Port Description**, **IP Address**, and **Subnet Mask**, and then click **Add**.

   You can choose **Manage > Edit** or **Manage > Delete** to modify or delete the added interface.
   iii. After you add the interface, click **Implement** in the Actions column corresponding to the device.
   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 previous version.
9. Click **Next**.
10. Configure BGP.
    i. In the **Configure BGP** step, find the device for which you want to configure BGP and click **Edit** in the Actions column. The **Interface Parameter Configuration** section appears.
    ii. In the **Interface Parameter Configuration** section, set **Group Name**, **Peer ASN**, **Peer IP Address**, and **Local Port Name**, and then click **Add**.

    You can choose **Manage > Edit** or **Manage > Delete** in the Actions column to modify or delete the added interface.
    iii. After you add the interface, click **Implement** in the Actions column corresponding to the device.
    iv. In the message that appears, click **OK** to run the script.

    If an exception occurs after the implementation, you can click **Back** to roll back to the previous version.
11. Click **Next**.
12. In the **Upload ASW Configurations** step, upload the new ASW configurations.
13. Click **Next**.
14. Enable the DSW ports.

i. In the **Enable DSW Port** step, find the device for which you want to enable the DSW port and click **Port Settings** in the Actions column.

ii. Enable the port and click **Implement** in the Actions column.

iii. In the message that appears, click **OK** to run the script command.

15. Click **Next**.
16. Perform the scale-up test.
    i. In the **Test Scale-up** step, find the device for which you want to perform the scale-up test and click **Select** in the Actions column. The route table is displayed on the right.
    ii. In the **ASW IP Address** search box, enter the IP address to be tested and then click **Add**.
    iii. Click **Test**. The system returns the test result.

1.1.6.3.15. 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 generates the IPv6 configurations and pushes the configurations to specified devices.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IPV6 Configuration**.
4. In the **Initialize Device** step, select one or more devices and configure the parameters to complete the initialization.
    i. In the device list, find a device and click **Check** in the **Check** column to check whether the device is accessible.

    You can check multiple devices at a time.

    ii. Select one or more devices whose **Status** is **Accessible**.

iii. Configure the parameters on the right.

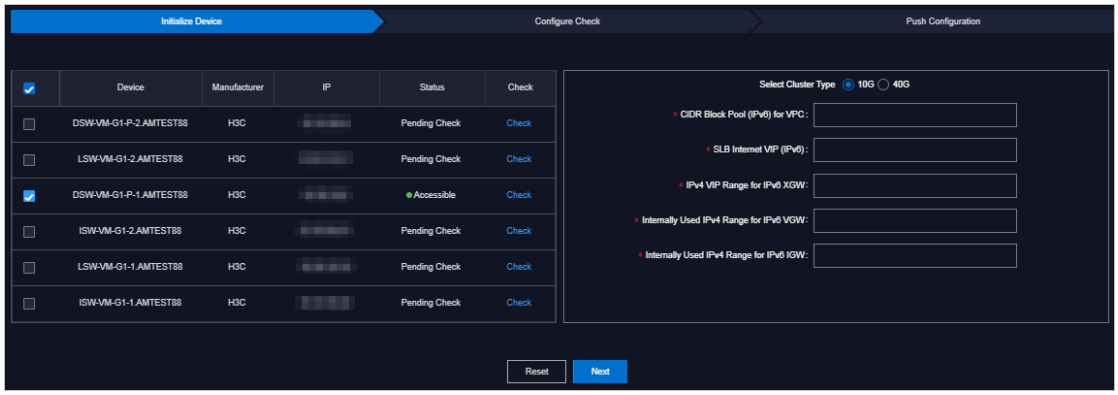

Parameter	Description
Select Cluster Type	The type of the cluster. Valid values: **10G** and **40G**. Select a value based on the planned cluster type.
CIDR Block Pool (IPv6) for VPC	The VPC CIDR block pool in the IPv6 format.
SLB Internet VIP (IPv6)	The SLB public VIP address in the IPv6 format.
IPv4 VIP Range for IPv6 XGW	The XGW VIP CIDR block in the IPv4 format.
Internally Used IPv4 Range for IPv6 VGW	The internally used CIDR block for VGW in the IPv4 format.
Internally Used IPv4 Range for IPv6 IGW	The internally used CIDR block for IGW in the IPv4 format.

5. Click **Next**.
6. In the **Configure Check** step, check the configurations.

   During the configuration check, the system checks the current configurations of the selected devices and generates the IPv6 configuration script based on the check results. Click **View** on the right of the script file to view the generated configuration script, or click **Download** to download the configuration script to your computer.

   > **Note** If you select multiple devices in the **Initialize Device** step, you can click **Batch Download** to download multiple configuration scripts to your computer at a time.

   One of the following results may occur 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.**
7. Click **Next**.

The system checks whether the configuration pushing feature is enabled. If not, the **Contact the onsite manager to enable the feature before you continue** message appears. If yes, check whether 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.**, or **Failed to check the configuration. The IPv6 configuration already exists.** in the previous step, a dialog box appears and the system does not push the configurations.

8. 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 whether the IPv6 configurations are pushed.

1.1.6.3.16. Check IP address conflicts

The IP Address Conflicts module allows you to check whether the current Apsara Stack environment contains conflicting IP addresses.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > IP Address Conflicts**.

   After the **IP Address Conflict Check** page appears, the system checks whether the current Apsara Stack environment contains conflicting IP addresses. If it does, 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.

1.1.6.3.17. Leased line discovery

You can automate the leased line discovery for devices in the ASO console. After network operations engineers configure the discovery parameters, Network Operations Center (NOC) automatically generates the discovery configuration, pushes the configuration to a specific device, and then automatically performs the discovery test.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Leased Line Discovery**.
4. Select a discovery source.

i. In the **Select Sources** step, enter a device name in the search box of the **Devices on Live Network** section and click **Search**.

After you add a device, you can click **Reset** to clear the search condition, and then search for another device and add it to the Devices for Discovery list.

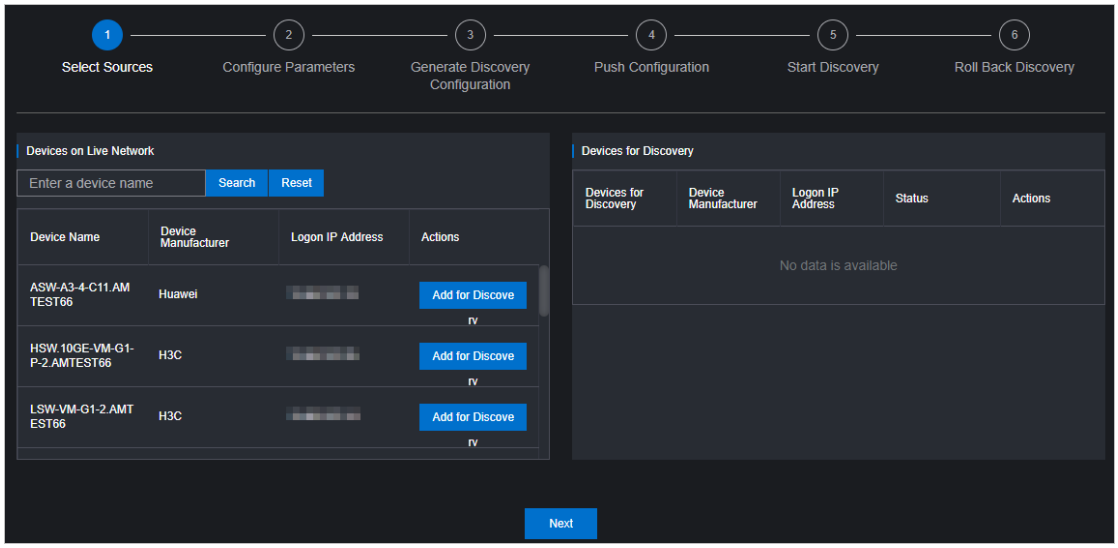

ii. Click **Add for Discovery** in the Actions column corresponding to a device to add the device on the live network to the Devices for Discovery list on the right.

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.

iii. Click **Next**.

5. Configure the discovery parameters.

i. In the **Configure Parameters** step, click **Edit**. The **Configure Parameters** section is displayed.

ii. Set **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** in the Actions column to modify or delete the discovery parameters.

iii. Click **Next**.

6. In the **Generate Discovery Configuration** step, click **Generate** to generate the discovery configuration and roll back commands of all devices that have discovery parameters configured.

i. Click **View** in the **Actions** column. The corresponding commands are displayed on the left.

ii. 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 computer.

iii. Click **Next**.

7. In the **Push Configuration** step, click **Push Configurations**.

i. In the message that appears, click **Continue** to push the discovery configuration commands to the corresponding device.

ii. After the configuration is pushed, you can click **View Logs** to view detailed push logs.

iii. Click **Next**.

8. In the **Start Discovery** step, click **Started** in the Actions column corresponding to a device to perform the leased line discovery test.

   After the test is complete, click **Next**.

9. 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.

   i. After the rollback is complete, you can click **View Logs** to view detailed rollback logs.

   ii. Click **Completed**.

1.1.6.3.18. Baseline configuration audit

The Baseline Configuration Audit module allows you to compare the baseline and the current running configurations of devices.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Baseline Configuration Audit**.
4. Select one or more devices in the device list and click **Audit**. The system begins to audit the baseline configurations of the selected devices.

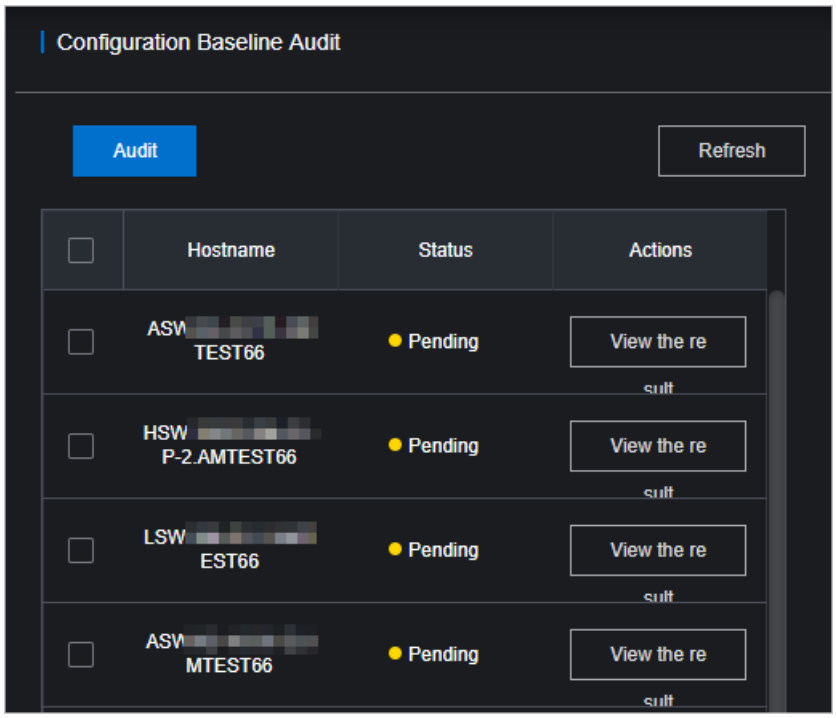

The following table describes the audit status.

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.

5. After the audit is complete, click **Refresh** to update the audit results.
6. Click **View the result** in the **Actions** column of the device. The audit result is displayed on the right.

1.1.6.3.19. Inspection dashboard

The Inspection Dashboard module allows you to view the inspection data and the last 10 inspection records.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Dashboard**.
4. Perform the following operations:
   - View the inspection statistics of the current day and the last 10 inspection records. The inspection statistics include the number of successful tasks, failed tasks, and scheduled tasks for the current day, as well as the progress.
   - 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.

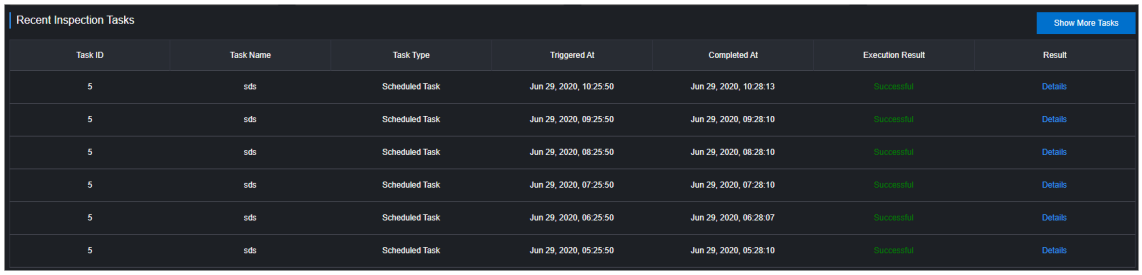

In the **Inspection Details** section, move the pointer over **Details** in the **Rollback** column corresponding to a subtask. The inspection result of the inspection subtask is displayed.

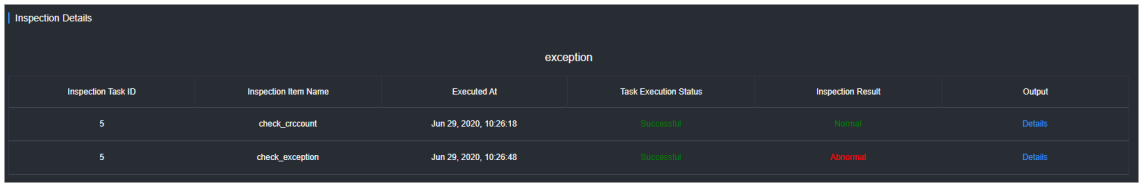

- Click **Show More Tasks** to go to the **Inspection History** page to view the inspection history.

1.1.6.3.20. Inspection history

You can query the inspection history and view detailed inspection records 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection History**.

   By default, all inspection records in the last 24 hours are displayed.
4. Select the inspection type (**All**, **One-time Task**, or **Scheduled Task**), specify the time range, and then click **Search**.
5. View inspection records that meet the query conditions.
6. 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.
7. In the **Inspection Details** section, move the pointer over **Details** in the **Rollback** column corresponding to a subtask. The inspection result of the inspection subtask is displayed.

1.1.6.3.21. Inspection management

The Inspection Management module allows you to create, view, modify, start, suspend, and delete inspection tasks.

1.1.6.3.21.1. 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Management**.
4. Click the **One-time Task** icon.

   The configuration wizard for the inspection task appears.
5. In the Specify Inspection Task Name step, enter the inspection task name and click **Next**.
6. In the **Add Device for Inspection** step, select one or more devices from the drop-down list and click **Next**.

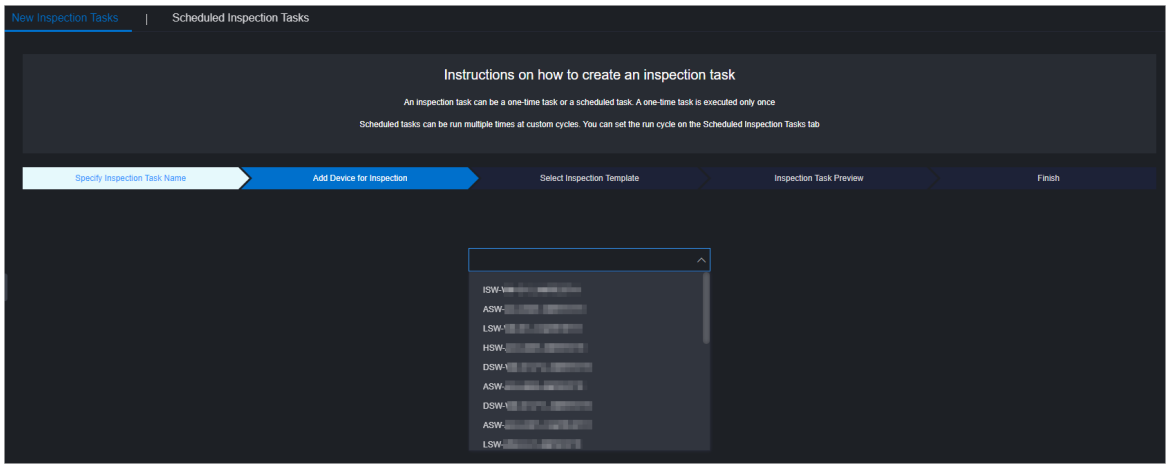

7. In the **Select Inspection Template** step, select an existing template from the drop-down list or click **Create Temporary Inspection Template**.

   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 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.

8. Click **Next**.
9. In the **Inspection Task Preview** step, confirm the inspection task information and click **Next**.

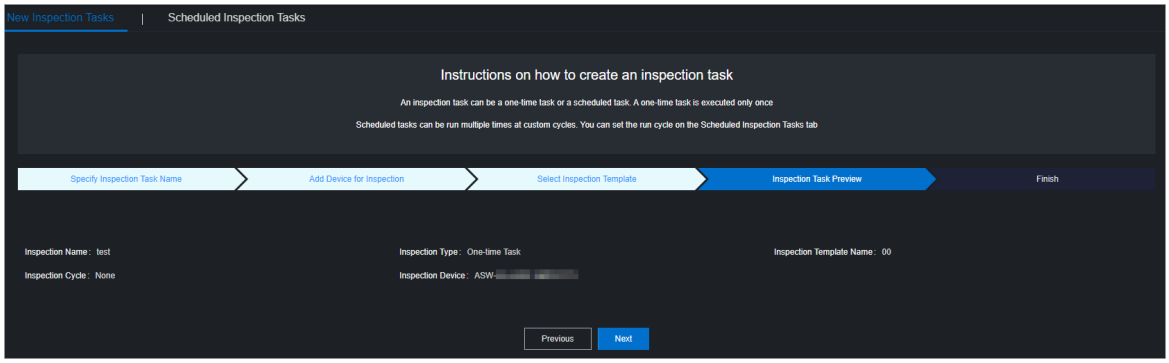

10. Click **Finish**.

The message **Created** is displayed. You can choose **Network Operations Center > Inspection Management** to view the created one-time inspection task on the **Scheduled Inspection Tasks** tab.

1.1.6.3.21.2. 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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Management**.
4. Click the **Scheduled Task** icon.

   The configuration wizard for the inspection task appears.
5. In the Specify Inspection Task Name step, enter the inspection task name and click **Next**.
6. In the **Add Device for Inspection** step, select one or more devices from the drop-down list and click **Next**.
7. 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 click **OK**.
8. Click **Next**.
9. Specify the inspection cycle and the time point when the task is triggered and click **Next**.

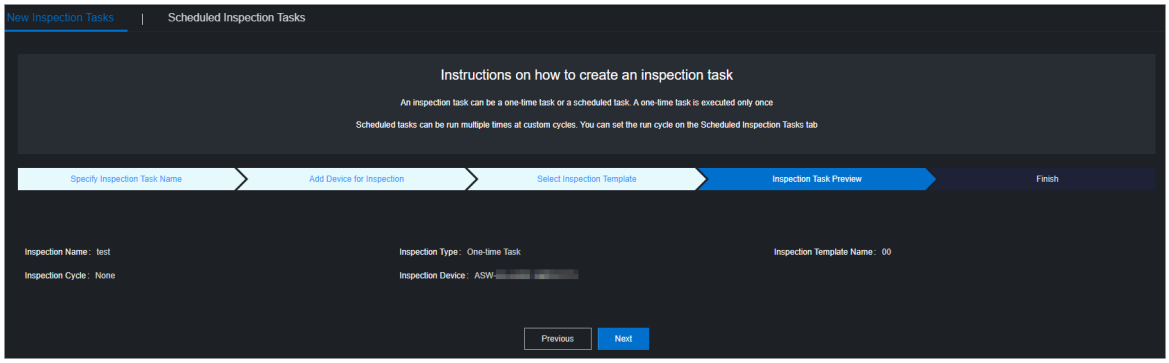

10. In the **Inspection Task Preview** step, confirm the inspection task information and click **Next**.
11. Click **Finish**.

    The message **Created** is displayed. You can choose **Network Operations Center > Inspection Tasks** to view the newly created scheduled task on the **Scheduled Inspection Tasks** tab.

1.1.6.3.21.3. Manage scheduled inspection tasks

After an inspection task is created, you can view, modify, start, suspend, or delete the task.

Go to the scheduled inspection task management page

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Management**.
4. Click the **Scheduled Inspection Tasks** tab.

View tasks

1. 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 task list, find the task that you want to modify and click **Modify** in the **Actions** column.
2. In the dialog box that appears, modify the task parameters.
   - 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.
3. 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 task list, find a task and click **Start** or **Suspend** in the **Actions** column.

   > **Note** After a one-time task is executed, it automatically enters the **Suspended** state. You can click **Start** to execute it again.

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

Delete a task

1. In the task list, find the task that you want to delete and click **Delete** in the **Actions** column.
2. In the message that appears, click **OK**.

1.1.6.3.22. Inspection templates

The Inspection Templates module allows you to manage, create, view, modify, and delete inspection templates.

1.1.6.3.22.1. Create a template

You can create a common inspection template to facilitate routine inspection task creation.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.

2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Templates**.
4. Click **Create Template**.
5. 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 Manufacturer	The manufacturer of the device.
Template Tag	The tag added to the template to make it easier to differentiate.
Template Inspection Item Collection	The collection of inspection items associated with the template.

6. Click **OK**.

   After the template is created, you can view the new template in the template list.

1.1.6.3.22.2. View template details

Before you use an inspection template, you can view its details to determine whether it meets your requirements.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Templates**.
4. In the template list, find the template that you want to view and click **Details** in the **Details** column.

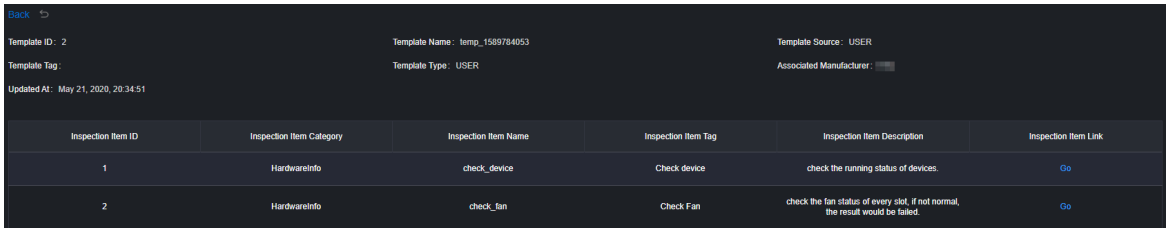

5. View the basic information about the template and the inspection items related to the template.
6. (Optional)To manage an inspection item in the template, click **Go** in the **Inspection Item Link** column to go to the inspection item management page.

   For other management operations that can be performed on inspection items, see **Network operations > Network management and operations > Network automation > Configure templates** in *Apsara Stack Enterprise Operations and Maintenance Guide*.

> **Note** Typically, no other management operations are required for inspection items.

1.1.6.3.22.3. Modify a template

After you create a template, you can modify its information based on your needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Templates**.
4. In the template list, find the template that you want to modify and click **Modify** in the **Actions** column.
5. In the dialog box that appears, modify the template name, associated manufacturer, template tag, and template inspection item collection.
6. Click **OK**.

1.1.6.3.22.4. Delete a template

You can delete inspection templates that are no longer needed for routine O&M.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Inspection Templates**.
4. In the template list, find the template that you want to delete and click **Delete** in the **Actions** column.

   > **Notice** When you delete a template, its associated inspection tasks and records are also deleted. Exercise caution when you delete templates.

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

1.1.6.3.22.5. 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. Log on to the Apsara Uni-manager Operations Console. For more information, see Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, click **Network Operations Center > Inspection Templates**.
4. Click the **Inspection Items** tab.

5. View the information of all inspection items in the system.
6. 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 other management operations that can be performed on inspection items, see **Network operations > Network management and operations > Network automation > Configure templates** in *Apsara Stack Enterprise Operations and Maintenance Guide*.

   > **Note** Typically, no other management operations are required for inspection items.

1.1.6.3.23. Use cases

1.1.6.3.23.1. Troubleshoot network failures

This topic uses a typical case to describe how to use the Network Operations Network module to troubleshoot network failures.

Scenario

If the visit latency and retransmission time of a cloud service increase, you must determine whether this is caused by network failures.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Dashboard**.
4. Click the **Network Topology** tab.
5. On the tab that appears, click **Topology Type** and select **Standard Topology**.

   Wait five seconds. After the page loads, 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, device icons are blue, links between devices are green, and device names are 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.

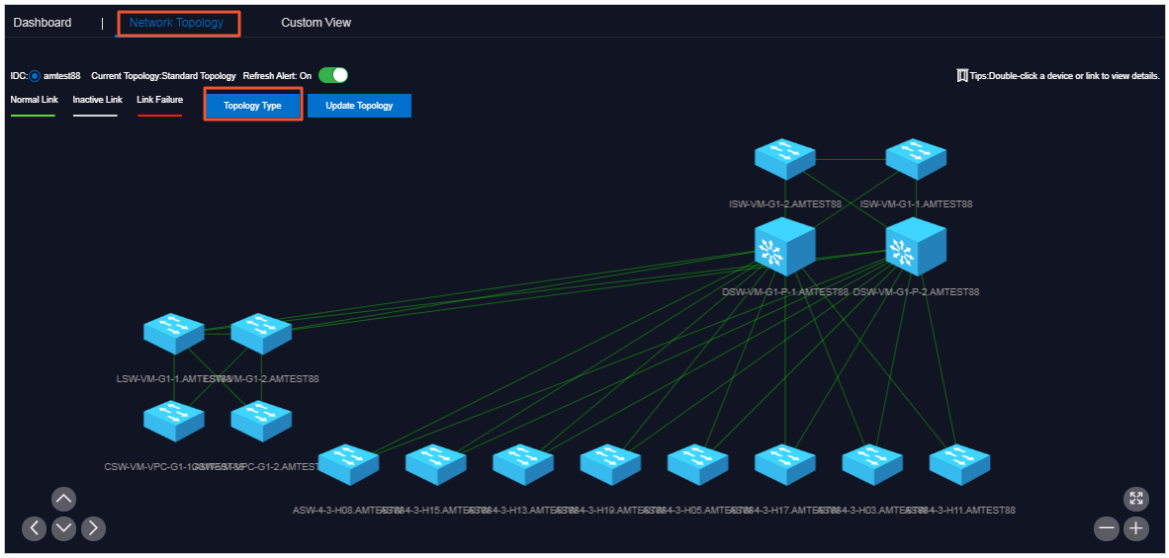

6. If a device name or link in the topology becomes red, alerts are detected in the network device or link port. Double-click the icon of a red device name. In the panel that appears, you can view the basic information of the device and the network alert information related to the device.

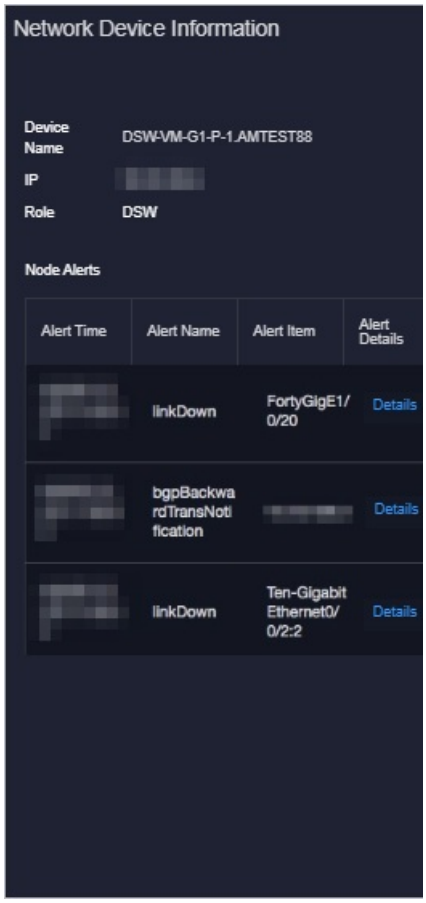

In the preceding figure, the port that is connected to the DSW has a **linkDown** alert and a bgp peer alert. An ASW is identified based on the IP address of the BGP peer. This allows you to determine that a problem in a link between DSW and ASW exists, which caused the port to go down and triggered the alerts.

7. Click the red link in the topology. In the panel that appears, you can view one or more physical links

contained in the logical link and the alert information of the link between devices.

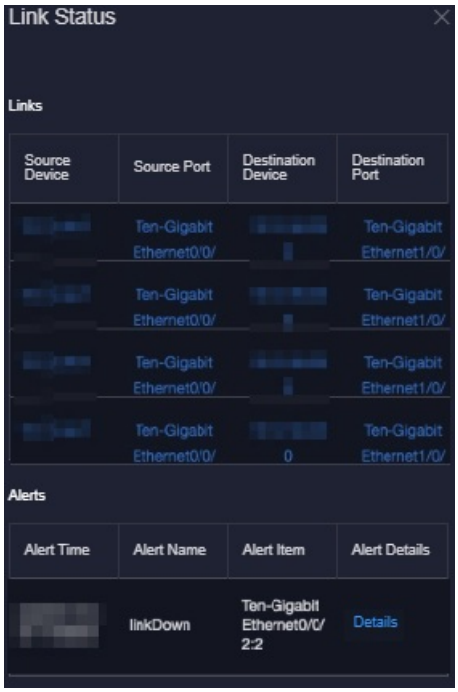

In the preceding figure, the logical link that connects the two devices contains four physical end-to-end links. The port 0/0/2:2 has a port **linkDown** alert. Then, you can proceed to log on to the device and check whether this is caused by the low optical power or damaged modules.

8. When the problem corresponding to the previous alerts is solved, the system updates the alert information. When the fault is repaired, the alerts automatically disappear, the topology is restored to the normal state, and no device names or links remain red.

Use the Alert Management module as a supplement to troubleshoot problems

If a device name or link in the topology becomes red, alerts are detected in the network device or link. You can choose **Network Operations Center > Alerts** 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 problems

If a device name or link in the 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 feature, you must view the detailed exception logs. You can use the syslog log query tool of vSwitches to search for logs.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Network Operations Center > Network Element Management**.

4. Click the **Syslogs** tab.
5. In the upper-right corner, select the device that you want to query, specify the time range, and then click **Search**. Logs generated within the specified time range are displayed.

   By default, you can query a maximum of 1,000 logs.
6. In the upper-left corner, enter a keyword in the search box and click **Search**.
7. After the query is complete, if you want to export logs to submit a ticket or submit logs to device vendors for troubleshooting, click **Export to CSV** in the upper-right corner. Logs are stored in your computer as a .csv file.

1.1.6.4. 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.

1.1.6.4.1. Product list

On the Product List page, you can go to the O&M page or ISV page corresponding to a product 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

When you use accounts that have different permissions to log on to the Apsara Uni-manager Operations Console, the product O&M icons and ISV icons on the **Product List** page are displayed in different ways. 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 separate for each product to allow the system to dynamically assign different permissions based on different roles.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products**.

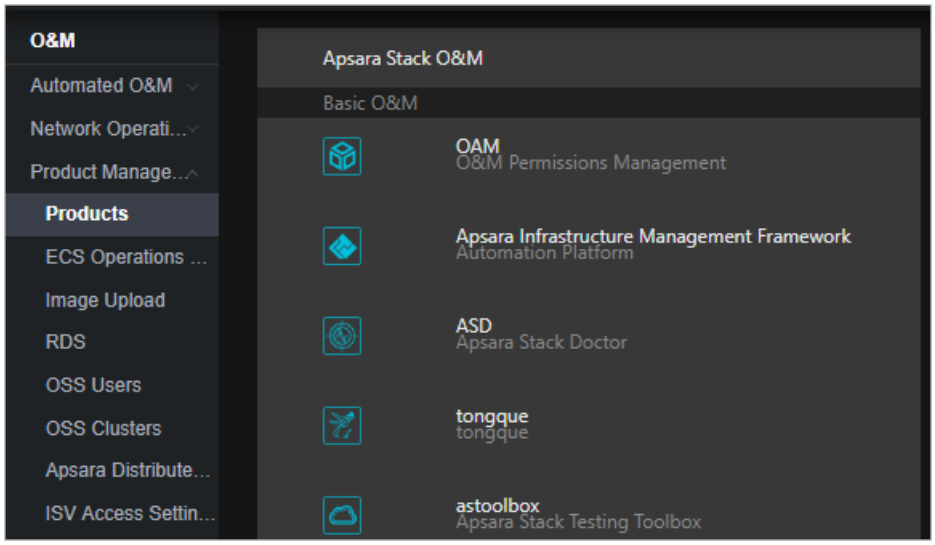

4. On the **Product List** page, you can view the O&M icons of different products and ISV icons based on your permissions.

1.1.6.4.2. SRS

Security Reinforce Service (SRS) is a security hardening component for internal access of cloud services. The Apsara Uni-manager Operations Console provides an entry point for SRS.

SRS provides security isolation and allows you to view the IP address whitelist information, SRS VIPS information, information about the services deployed on a host, SRS security isolation status configurations, SRS IP address whitelist configurations, and SRS isolated and isolation-free product configurations.

1.1.6.4.2.1. SRS management

You can configure SRS security isolation status, SRS IP whitelists, and SRS isolation-free products. You can also view SRS IP address whitelists and SRS VIPS information.

Configure SRS isolation status

To perform operations involving IP address changes, you must first manage the SRS switch. You can configure SRS security isolation status.

Prerequisites

- SRS has reached the desired state.
- After all base services have reached the desired state, the data in Skyline is manually updated.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > SRS Management**.

   The **Configuration Management** tab appears.
4. In the **SRS Security Isolation Status Configuration** section, you can view and configure the SRS security isolation status.

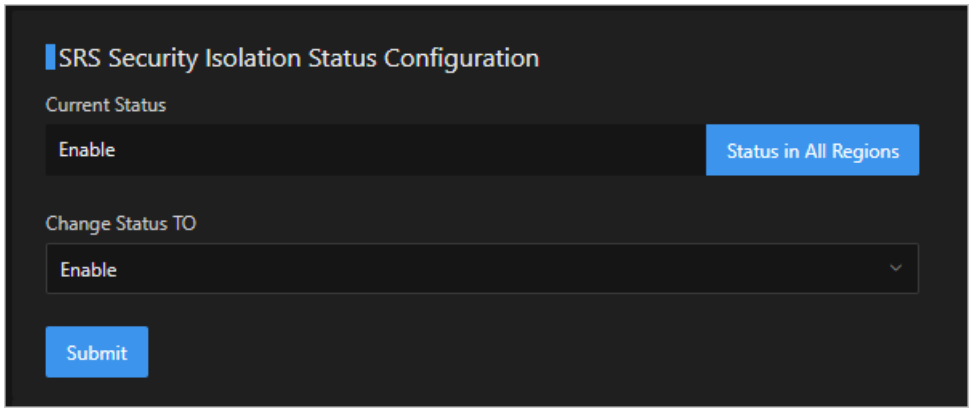

- View the SRS security isolation status: The **Current Status** field shows the current SRS security isolation status. Click **All Region Status** to view the SRS security isolation status of all nodes in the window that appears.
- Configure the SRS security isolation status: Select a parameter from the **Change Status TO** drop-down list and then click **Submit** to change the SRS security isolation status.
  - **Close**: disables SRS security isolation.

    Warning

    - After SRS security isolation is disabled, the cloud service ports that are originally protected by SRS may be detected and scanned by devices outside the cloud.
    - SRS security isolation is disabled by default. You can change it to Enable or DEBUG Mode. We recommend that you do not frequently change SRS isolation status. You can change SRS isolation status 10 minutes after you change it.
    - SRS uses the unit deployment method. To disable SRS for the entire cloud instance, you must disable SRS for each region.
    - You can disable SRS in the following ways:
      - General scenarios: Disable SRS for the current region.
      - Multi-region scenario: First, disable SRS for the current region, central region, and associated regions. Then, wait for a maximum of 50 minutes (usually within 30 minutes) for the disable operation to complete. Then, start O&M tasks.

- **Enable**: enables SRS security isolation.

  > **Warning**
  >
  > - Strict conditions are required for enabling SRS security isolation. Contact Alibaba Cloud O&M engineers for evaluation first. Otherwise, the business traffic may be interrupted.
  >
  >   Preconditions for enabling SRS security isolation:
  >
  >   - Apsara Stack Enterprise Edition must be V3.15.0 or later.
  >   - Apsara Infrastructure Management Framework API, SLB control API, Skyline, Networkbase, and MetadataSync services are available.
  >   - The products and services connected to SRS cannot be involved in the network connection scenario. You can view the products and services connected to SRS on the **Isolation Configuration Management** page.
  >   - The products and services connected to SRS cannot be involved in the scenarios where VPCs and the classic network for base services can be accessed each other by using non-standard methods. You can view the products and services connected to SRS on the **Isolation Configuration Management** page.
  >   - The data in Skyline is correct.
  >   - In the zone-disaster recovery scenario, Apsara Stack Enterprise Edition and the SRS version must be consistent in the primary and secondary data centers.
  >
  > - SRS depends on whether topology and IP address information of services is properly synchronized. Therefore, you must first contact Alibaba Cloud O&M engineers to confirm that the services and data collection are normal. Otherwise, access to the services may fail.
  > - When you perform operations such as scale-out and upgrade, you must first disable SRS. Otherwise, the topology and IP address information cannot be synchronized to SRS during such operations. This may cause unexpected failure to access services.

- **DEBUG mode**: This mode is used in special scenarios. It is equivalent to the half-enabled state. The involved traffic is recorded, but not actually isolated.

  > **Warning** You can change the DEBUG Mode to the Close state, but not to the Enable state. You cannot change the Enable state to the DEBUG Mode. We recommend that you do not frequently change SRS isolation status. You can change SRS isolation status 10 minutes after you change it.

Configure an SRS IP whitelist

SRS only trusts internal IP addresses. After SRS is enabled, public IP addresses cannot access the services on the cloud that are connected to SRS over the classic network of the base for security isolation. You can add the public IP addresses allowed to access the services to an SRS IP whitelist.

Prerequisites

- SRS has reached the desired state.
- After all base services have reached the desired state, the data in Skyline is manually updated.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > SRS Management**.

   The **Configuration Management** tab appears.
4. In the **SRS IP Address Whitelist Configuration** section, you can add or delete an SRS IP whitelist.
   - Add an SRS IP whitelist: Click **Add**. In the dialog box that appears, enter one or more IP addresses. Separate multiple IP addresses with commas (,). Click **OK**.
   - Delete an SRS IP whitelist: In the **Actions** column, click **Delete**. In the message that appears, click **OK**.

Add isolation-free products

If you do not want a product that is connected to SRS for security isolation to be protected by SRS, you can add an isolation-free product.

Prerequisites

- SRS has reached the desired state.
- After all base services have reached the desired state, the data in Skyline is manually updated.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > SRS Management**.

   The **Configuration Management** tab appears.
4. In the **SRS Isolation-free Product Configuration** section, you can add or delete an isolation-free product.
   - Add an isolation-free product: Click **Add**. In the dialog box that appears, enter one or more product names. Separate multiple product names with commas (,). Click **OK**.
   - Delete an isolation-free product: In the **Actions** column, click **Delete**. In the message that appears, click **OK**.

View SRS isolation data

You can view the configured IP whitelist and SRS VIPS information.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > SRS Management**.

4. You can click the **Data Management** tab to view the configured IP whitelist and SRS VIPS information.

   > **Note**
   > - SRS IP whitelist: After SRS is enabled, public IP addresses cannot access the servers on the cloud that are connected to SRS over the classic network of the base for security isolation. You can add the public IP addresses allowed to access the services to an SRS IP whitelist.
   > - SRS VIPS information: The VIPs in the list use the DNAT forwarding mode.

5. (Optional)After you click **Refresh Data**, the configured isolation data will be refreshed.

1.1.6.4.2.2. Isolation configuration management

You can also add security isolation configurations for cloud services if they have special security isolation requirements.

View security isolation configurations

You can view the security isolation configurations of cloud services registered with SRS.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Isolation Configuration Management**.
4. Click **View Details** in the **Actions** column corresponding to the service to view the security isolation configurations.

Add a security isolation configuration rule

You can add security isolation configuration rules of cloud services registered with SRS.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Isolation Configuration Management**.

   > **Warning** Adding security isolation configurations is a high-risk operation. You must careful evaluate the impacts. Do not perform the operations unless necessary.

4. Click **Add Product Configurations**. In the dialog box that appears, select the product name and SR, enter the port and VIP name, and then click **Save**. After you confirm that the product isolation configurations are correct, click **OK**.

   > **Note**
   > - You can also click **Add** or **Delete** to add or delete an isolation configuration rule for the product.
   > - If the new product name is the same as an existing one, the original security isolation configurations are overwritten.

Modify a security isolation configuration rule

You can modify security isolation configuration rules of cloud services registered with SRS.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Isolation Configuration Management**.

   > **Warning** Modifying security isolation configurations is a high-risk operation. You must careful evaluate the impacts. Do not perform the operations unless necessary.

4. Find the product and click **Details** in the **Actions** column.
5. In the panel that appears, click **Modify**. In the dialog box that appears, scroll down the page and click **Modify**.
6. In the dialog box that appears, modify the isolation configurations and click **Save**. After you confirm that the product isolation configurations are correct, click **OK**.

   > **Note** You can also click **Add** or **Delete** to add or delete an isolation configuration rule for the product.

Delete a security isolation configuration rule

You can delete security isolation configuration rules of cloud services registered with SRS.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Isolation**

Configuration Management.

> **Warning** Deleting security isolation configurations is a high-risk operation. You must careful evaluate the impacts. Do not perform the operations unless necessary.

4. Find the product and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

Restore factory security isolation configurations

You can restore the factory security isolation configurations of a product.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Isolation Configuration Management**.

> **Warning** Restoring factory security isolation configurations is a high-risk operation. You must careful evaluate the impacts. Do not perform the operations unless necessary.

4. Click **Restore Products to Factory Settings** to restore the factory security isolation configurations of a product. In the message that appears, click **OK**.

1.1.6.4.2.3. Client status

View client status

You can view the host information, last heartbeat time, drop logs (including the packet information intercepted by SRS), and status of the SRS client node. You can also view information about the services deployed on a host.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Client Status**.
4. View the host information, last heartbeat time, and status of the SRS client node.

> **Note** You can enter **Hostname** and **SR**, and select **Only Hosts with Packet Loss** to filter client node information.

5. Find a host and click the ··· icon in the **Drop Logs** column to view SRS-intercepted packet

information in a new window.

6. Find a host and click **Details** in the **Actions** column to view information about the services deployed on the host in a new window.

Add an SRS IP blacklist

You can add an IP address to the SRS IP blacklist to stop the address from accessing the services on the cloud that are connected to SRS over the classic network of the base for security isolation.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Client Status**.

   Warning IP addresses in the SRS IP blacklist cannot access the services on the cloud that are connected to SRS over the classic network of the base for security isolation. Proceed with caution.

4. In the **SRS IP Blacklist Configuration** section, click **Add**. In the dialog box that appears, enter IP addresses. Separate multiple IP addresses with commas (,).
5. Click **OK**.

Delete an SRS IP blacklist

You can delete an IP address in the SRS IP blacklist to allow the address to access the services on the cloud that are connected to SRS over the classic network of the base for security isolation.

Prerequisites

SRS has reached the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > SRS > Client Status**.
4. In the **SRS IP Blacklist** section, find the IP address and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.1.6.4.3. ISV access settings

The ISV Access Settings module allows you to configure, modify, and delete the ISV access information.

1.1.6.4.3.1. Configure the ISV access information

This topic describes how to configure the ISV access information in the system to suit your business needs. Then, you can click an icon on the **Product List** page to access the corresponding ISV page.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > ISV Access Settings**.
4. Click **Add**.
5. In the **Add** panel, configure the ISV access information.

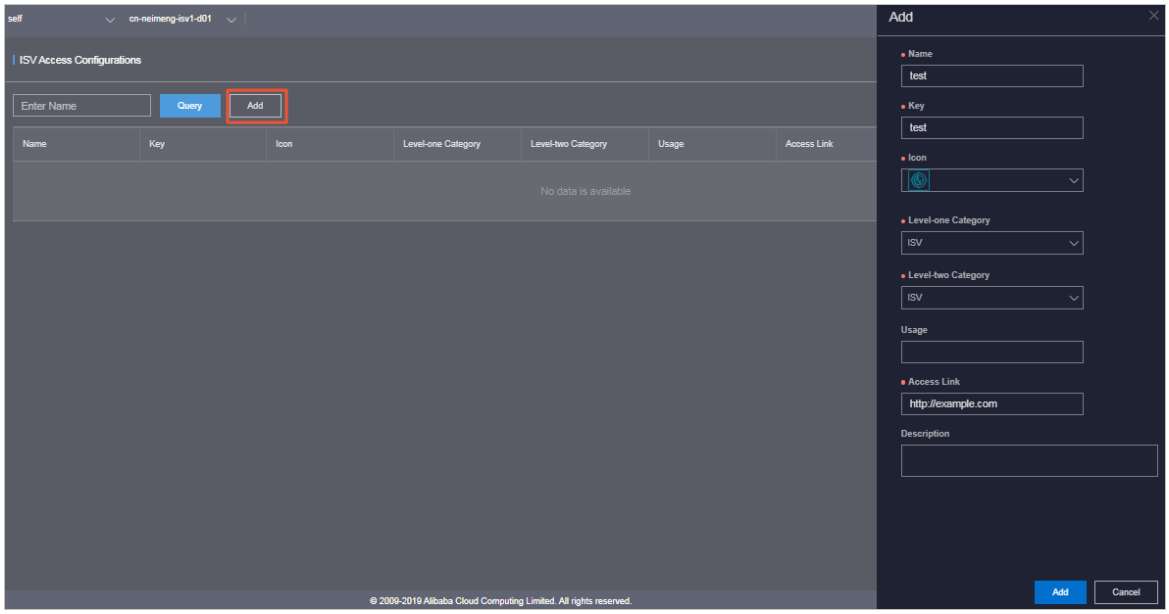

The following table describes the parameters.

Parameter	Description
Name	The name of the ISV to access.
Key	Set the value to an identifier related to the ISV business.
Icon	The icon displayed on the **Product List** page for the ISV to access.
Level-one Category and **Level-two Category**	The category to which the ISV to be accessed belongs on the **Product List** page.
Usage	The feature of the ISV to access.
Access Link	The address of the ISV to access.
Description	The description related to the ISV to access.

6. Click **Add**.

Result

You can view the added ISV icon on the Product List page by choosing **Product Management > Products**. Click the icon and then you can go to the corresponding page.

1.1.6.4.3.2. Modify the access information of an ISV

If the information of a third-party independent software vendor (ISV) is changed, you can modify the access information of the ISV.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > ISV Access Settings**.
4. (Optional)Enter the ISV name in the search box and click **Query**. Fuzzy search is supported.
5. Find the ISV for which you want to modify the access information and click **Modify** in the **Actions** column.
6. In the **Modify** panel, modify the name, key, icon, level-one category, level-two category, usage, access link, or description of the ISV.
7. Click **Edit**.

1.1.6.4.3.3. Delete the access information of an ISV

You can delete the access information of a third-party independent software vendor (ISV) from the system based on your business needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > ISV Access Settings**.
4. (Optional)Enter the ISV name in the search box and click **Query**. Fuzzy search is supported.
5. Find the ISV for which you want to delete the access information and click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

Result

In the left-side navigation pane, choose **Product Management > Products**, and the deleted ISV is no longer displayed on the Product List page.

1.1.6.5. Apsara Distributed File System Management

1.1.6.5.1. Dashboard

The Dashboard module allows you to view the overview information, health heatmap, and data of the top five clusters of a service.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.

3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Dashboard**.
4. 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 top five clusters of a service for the current date.

   - **Overview**

     The Overview section shows the storage space, server information, and health information of the service. In the **Health** column, values that are greater than 0 are displayed in red.

     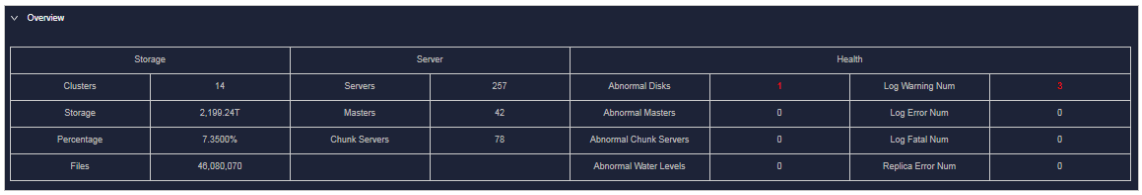

   - **Heatmap of Health**

     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 is working normally.
     - Yellow indicates that the cluster has an alert.
     - 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 Cluster Details page.

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

     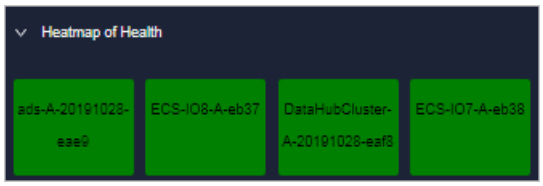

   - **Data of Top 5 Services**

     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 disk usage, abnormal masters, abnormal disks, and abnormal chunk servers. Click the name of a cluster to go to the Cluster Details page.

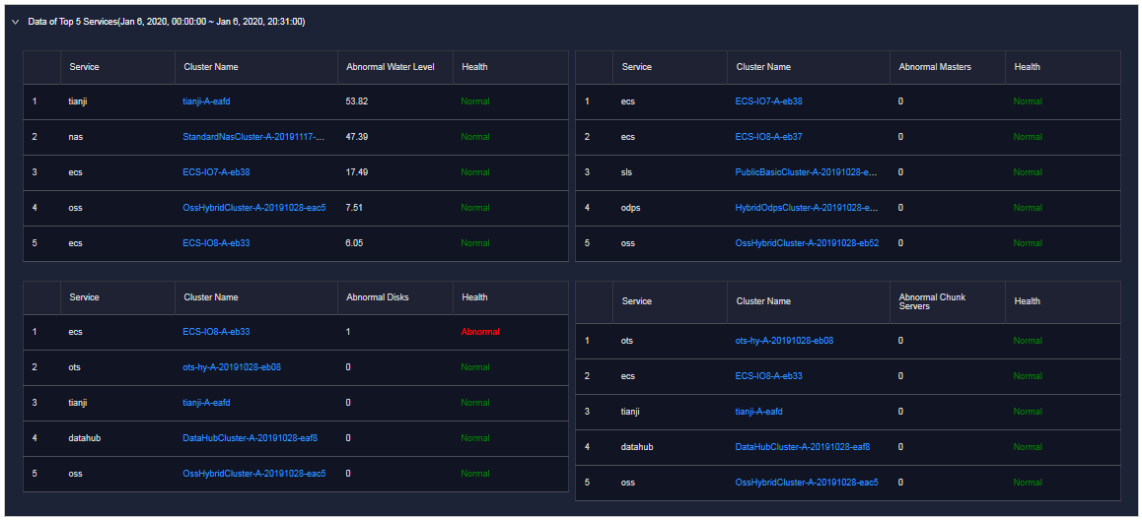

1.1.6.5.2. Clusters

The Clusters module allows you to view the overview information, alert monitoring information, replica information, trend charts, and rack information of a cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Clusters**.

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

   - **Overview**

     This section shows the storage space, device information, and health information of the specified cluster. In the **Health** column, values that are greater than 0 are displayed in red.

     

   - **Alarm Monitor**

     This section shows the alert information of the specified cluster, such as the level, machine, and server role. You can query data by keywords.

   - **Replica**

This section shows the replica information of the specified cluster.

- **Run Chart of Clusters**

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

  Predicted disk usage shows the run chart of the next seven days.

  > **Note** The disk usage can be predicted only when historical disk usage data exist. Some clusters may not have predicted disk usage data.

- **Rack Information**

  This section contains Storage and Servers in Rack.

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

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

1.1.6.5.3. Nodes

The Nodes module allows you to view the information about master and chunk servers within a cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Nodes**.

   On the page that appears, the data of the first cluster in the **Cluster Name** drop-down list is displayed, including the information about master and chunk servers.
4. 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 the enabled clusters in the current environment are displayed in the **Cluster Name** drop-down list.

   - **Master Info**

     This section shows the master node information of the specified cluster. You can click **Refresh** in the upper-right corner of the section to refresh the master node information of the cluster.

     

   - **Chunk Server Info**

     This section shows the chunk server information of the specified cluster. You can click **Refresh** in the upper-right corner of the section to refresh the chunk server information of the specified cluster. You can click the + icon in front of a server to view the disk and SSD cache information of the server. Fuzzy search is supported in this section.

     

1.1.6.5.4. Operations and maintenance

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Operations and Maintenance**.
4. Select a service from the **Service** drop-down list to view the cluster status of the service.

   Clusters are displayed in different colors based on their health status.
   - Green indicates that the cluster is running normally.
   - 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.

   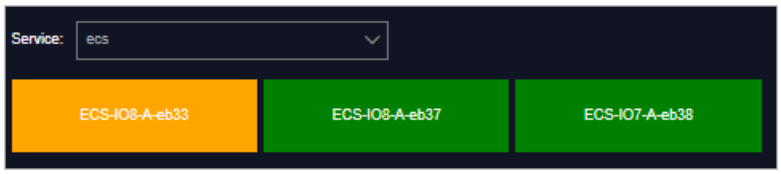

5. Move the pointer over a cluster name to view the service name, server name, and IP address of the cluster.

1.1.6.5.5. Modify cluster thresholds

By default, the thresholds for all clusters are configured by the system. You can modify these thresholds for storage usage, chunk server, and disk of each cluster based on your needs.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Settings**.
4. In the **Cluster Name** drop-down list, select a cluster for which you want to modify the thresholds.
5. In the lower part of the page, click **Modify** and configure the parameters.

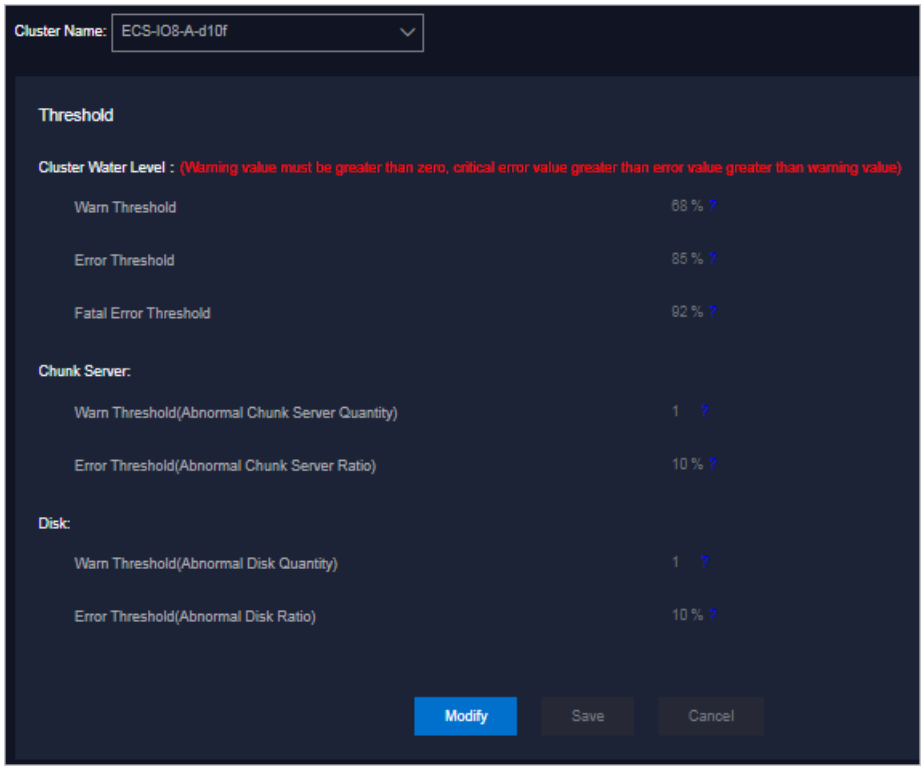

The following table describes the parameters.

Parameter		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]. The default threshold for the cluster storage usage is 65%. Note: The fatal error threshold value must be greater than the error threshold value, and the error threshold value must be greater than the warning threshold value.
	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]. The default threshold for the cluster storage usage is 85%. Note: The fatal error threshold value must be greater than the error threshold value, and the error threshold value must be greater than the warning threshold value.

Parameter		Description
	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]. The default threshold for the cluster storage usage is 92%. Note: The fatal error threshold value must be greater than the error threshold value, and the error threshold value must be greater than the warning threshold value.
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. The default threshold for the number of abnormal chunk servers is 1.
	Error Threshold (Abnormal Chunk Server Ratio)	If the ratio of abnormal chunk servers to all the chunk servers is greater than this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. The default threshold for the ratio of abnormal chunk servers to all the chunk servers is 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. The default threshold for the number of abnormal disks is 1.
	Error Threshold (Abnormal Disk Ratio)	When the ratio of abnormal disks to all the disks is greater than this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. The default threshold for the ratio of abnormal disks to all the disks is 10%.

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

6. Click **Save**.

1.1.6.5.6. Load information

The Load Information module allows you to view the NC information, VM information, and block device information.

1.1.6.5.6.1. View NC information

The Load Information module allows you to view the data overview information for each NC and real-time data such as the load, CPU utilization, memory information, sda usage, traffic, TCP information, network exception metrics, read and write rate, BPS, kernel status, IOPS, and latency.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Load Information**.

   The **NC Info** tab appears.
4. Select a cluster from the **Cluster Name** drop-down list and specify **Time Frame**. You can select **One Hour**, **Three Hours**, **Six Hours**, or **One Day**, or you can customize a time range. Click **Search** and view the following information:
   - Survey

     The **Survey** section shows all NCs in the current cluster. You can click the IP address of an NC to view the trend charts of real-time data of the NC in the current cluster.

     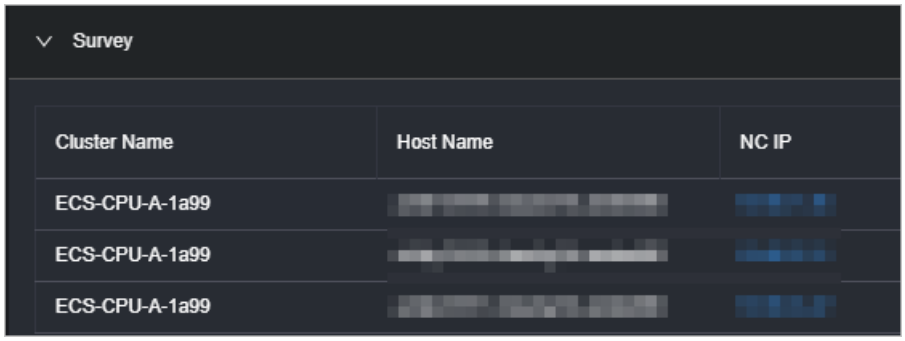

   - Real-time load

     In the **Survey** section, click the IP address of an NC. In the **Real Time Load** section, the load trend chart of the NC within the specified time range is displayed. By default, the real-time load trend chart corresponding to the NC in the first row of the **Data Overview** section is displayed.

     The following information is displayed in the real-time load trend chart:

     - **load_1**: the average load of the NC within the last minute.
     - **load_5**: the average load of the NC within the last 5 minutes.
     - **load_15**: the average load of the NC within the last 15 minutes.

     You can drag the slider below the chart to zoom in or out the chart.
   - Real-time CPU utilization

     In the **Survey** section, click the IP address of an NC. In the **Real-time CPU utilization** section, the CPU utilization trend chart of the NC within the specified time range is displayed.

     The following information is displayed in the real-time CPU utilization trend chart:

- **cpu_iowait**: the amount of time spent waiting for an I/O response.
- **cpu_guest**: the run time duration of vCPUs in a guest OS.
- **cpu_idle**: the duration for which vCPUs are available.
- **cpu_hardirp**: the amount of time spent handling hardware interrupts.
- **cpu_user**: the CPU time in user mode.
- **cpu_softirp**: the amount of time spent handling software interrupts.
- **cpu_steal**: the duration for which vCPUs are occupied by other VMs.
- **cpu_sys**: the CPU time in system mode.
- **cpu_nice**: the amount of CPU time consumed to process data for low-priority programs in user mode.

You can drag the slider below the chart to zoom in or out the chart.

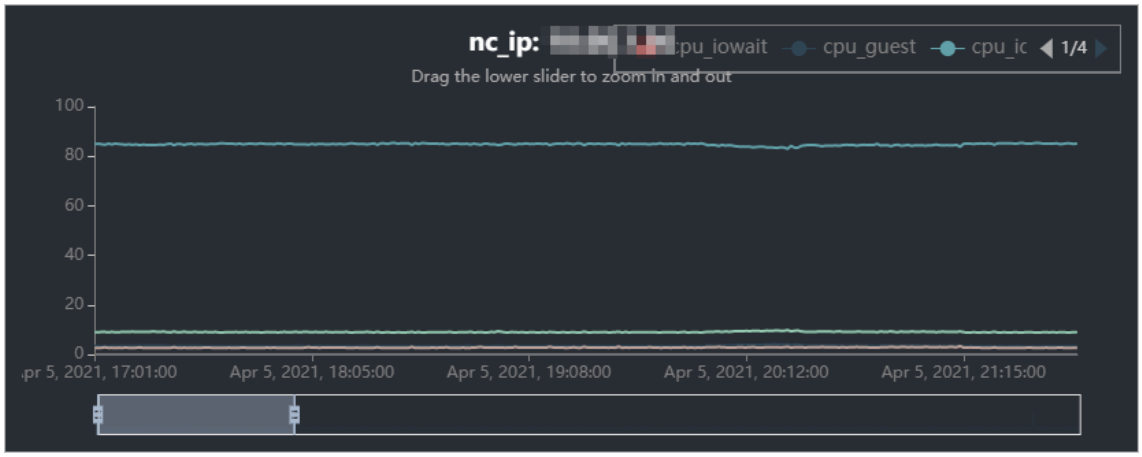

- Real-time memory information

  In the **Survey** section, click the IP address of an NC. In the **Real-time mem info** section, the memory usage trend chart of the NC within the specified time range is displayed.

  The following information is displayed in the real-time memory usage trend chart:

  - **mem_sunreclaim**: the amount of slab capacity in which the object is active and cannot be reclaimed.
  - **mem_cached**: the amount of physical memory that has been cached.
  - **mem_slab**: the total amount of the memory allocated by the slab allocator.
  - **mem_free**: the amount of available physical memory and swap space.
  - **mem_shmem**: the memory usage.
  - **mem_used**: the amount of physical memory and swap space that is occupied.
  - **mem_total**: the total amount of the physical memory and swap space of the system.
  - **mem_buffer**: the physical memory that has been buffered.
  - **mem_dirty**: the amount of dirty data, which is the data that is stored in the buffer zone but has not been written to physical disks.

  You can drag the slider below the chart to zoom in or out the chart.

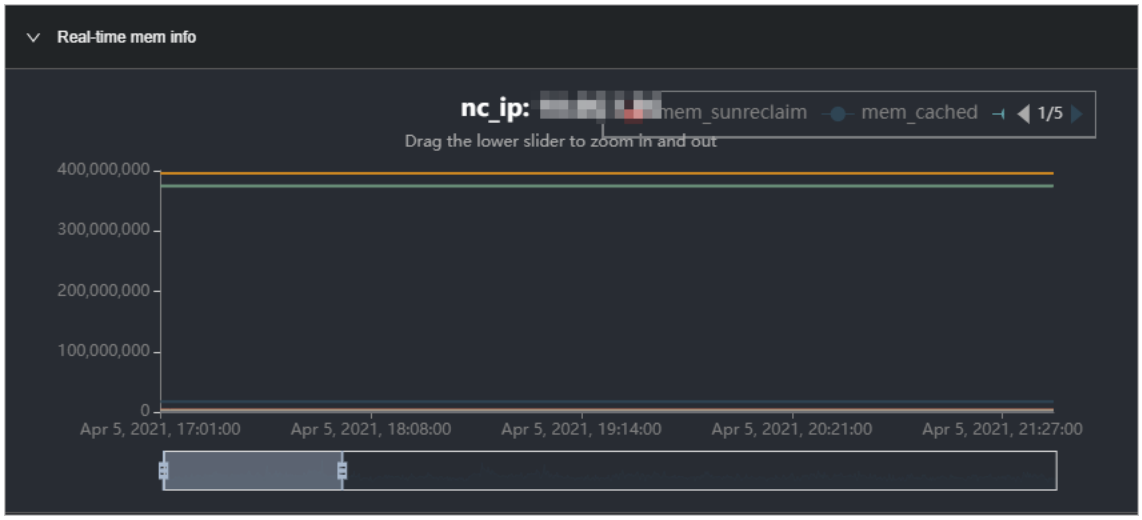

- Real-time sda usage

  In the **Survey** section, click the IP address of an NC. In the **Real-time SDA utilization** section, the disk usage trend chart of the NC within the specified time range is displayed.

  In the real-time sda usage trend chart, **sda_until** indicates the sda usage.

  You can drag the slider below the chart to zoom in or out the chart.

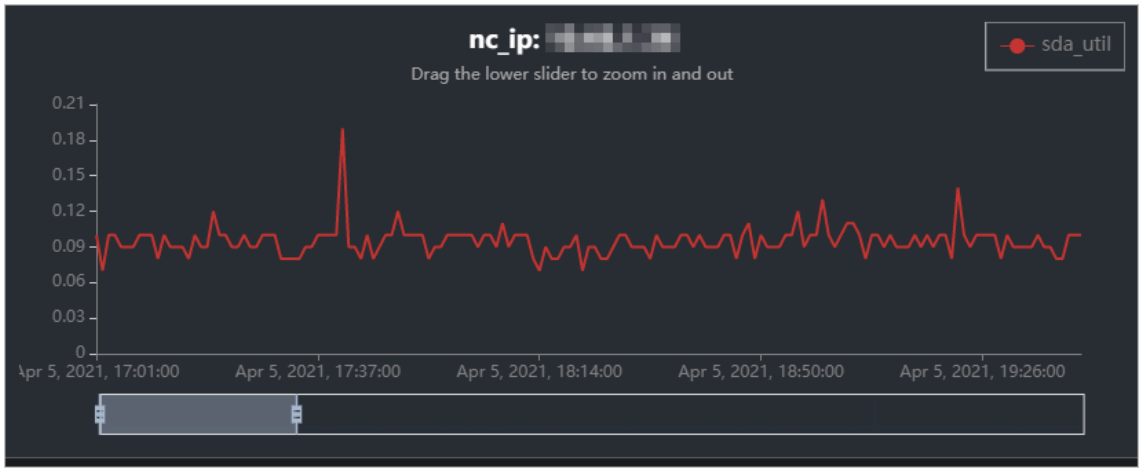

- Real-time traffic

  In the **Survey** section, click the IP address of an NC. In the **Real-time traffic** section, the traffic trend chart of the NC within the specified time range is displayed.

  The following information is displayed in the real-time traffic trend chart:

  - **traffic_pkterr**: the number of errors for transmission packets.
  - **traffic_pktdrp**: the number of transmission packets that are lost.
  - **traffic_pktin**: the number of input bytes during the traffic peak.
  - **traffic_bytesout**: the number of output bytes.
  - **traffic_bytesin**: the number of input bytes.
  - **traffic_pktout**: the number of output bytes during the traffic peak.

  You can drag the slider below the chart to zoom in or out the chart.

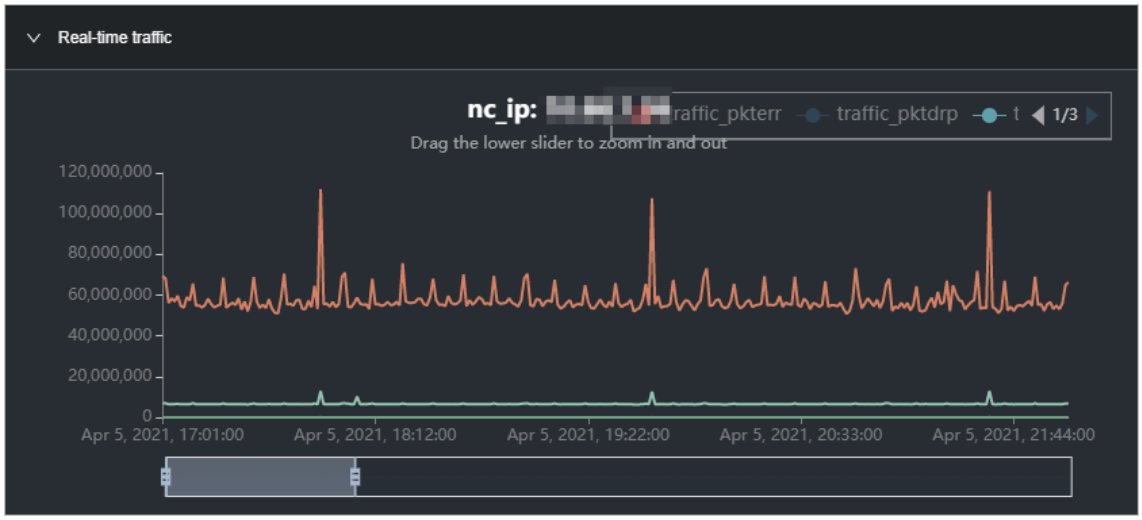

- Real-time TCP information

  In the **Survey** section, click the IP address of an NC. In the **Real-time tcp Info** section, the TCP connection trend chart of the NC within the specified time range is displayed.

  The following information is displayed in the real-time TCP connection trend chart:

  - **tcp_outseqs**: the number of TCP packets that have been sent.
  - **tcp_estab_resets**: the number of retry attempts of TCP connections that are in the ESTABLISHED state.
  - **tcp_opens**: the number of open TCP connections.
  - **tcp_inseqs**: the number of received TCP packets.
  - **tcp_attempt_fails**: the number of failed connection attempts.
  - **tcp_curr_estab**: the number of TCP connections that are in the ESTABLISHED state.

  You can drag the slider below the chart to zoom in or out the chart.

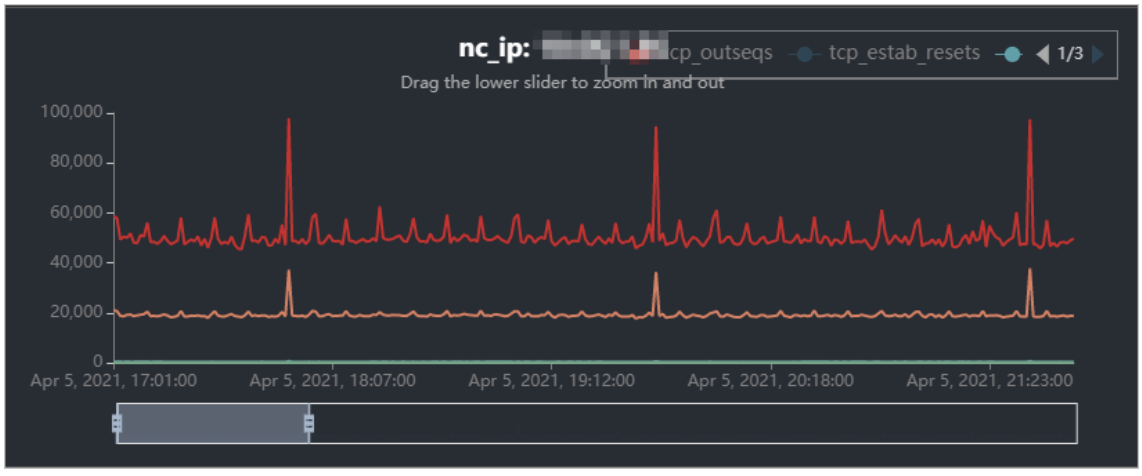

- Real-time network exception index

  In the **Survey** section, click the IP address of an NC. In the **Real-time network anomaly index** section, the trend chart of the network exception index of the NC within the specified time range is displayed.

The following information is displayed in the trend chart of the real-time network exception index:

- **traffic_pktdrp**: the number of transmission packets that are lost.
- **traffic_pkterr**: the number of errors for transmission packets.
- **traffic_retrans_ratio**: the number of retry attempts of transmission packets.

You can drag the slider below the chart to zoom in or out the chart.

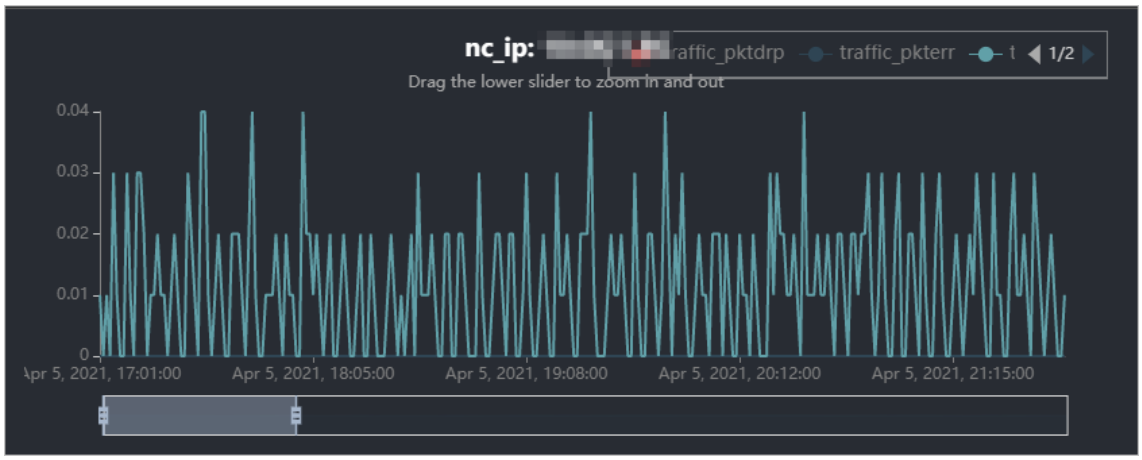

- Reading and writing B/S

  In the **Survey** section, click the IP address of an NC. In the **Reading and Writing B/S** section, the trend chart of the reading and writing rate of the NC within the specified time range is displayed.

  The following information is displayed in the trend chart of the read and write rate :

  - **bs_w**: the write rate in byte/s.
  - **bs_r**: the read rate in byte/s.

  You can drag the slider below the chart to zoom in or out the chart.

- Real-time BPS

  In the **Survey** section, click the IP address of an NC. In the **Real-time BPS** section, the BPS trend chart of the NC within the specified time range is displayed.

  The following information is displayed in the real-time BPS trend chart:

  - **bps_w**: the writing rate in bps.
  - **bps_r**: the reading rate in bps.

  You can drag the slider below the chart to zoom in or out the chart.

- Real-time kernel status

  In the **Survey** section, click the IP address of an NC. In the **Real-time Kernel State** section, the kernel status trend chart of the NC within the specified time range is displayed.

  In the real-time kernel status trend chart, **kernel_status** indicates the real-time kernel status.

  You can drag the slider below the chart to zoom in or out the chart.

- Real-time IOPS

In the **Survey** section, click the IP address of an NC. In the **Real-time iops** section, the IOPS trend chart of the NC within the specified time range is displayed.

The following information is displayed in the real-time IOPS trend chart:

- **iops_w**: the number of input operations per second.
- **iops_r**: the number of output operations per second.

You can drag the slider below the chart to zoom in or out the chart.

- Real-time latency

  In the **Survey** section, click the IP address of an NC. In the **Real-time latency** section, the latency trend chart of the NC within the specified time range is displayed.

  The following information is displayed in the real-time latency trend chart:

  - **latency_w**: the real-time latency of data writes
  - **latency_r**: the real-time latency of data reads
  - **latency_w_qos**: the QoS intelligent adjustment for real-time latency of data writes
  - **latency_r_qos**: the QoS intelligent adjustment for real-time latency of data reads

  You can drag the slider below the chart to zoom in or out the chart.

1.1.6.5.6.2. View virtual machine information

The Load Information module allows you to view the data overview information and trend charts of read and write rate, real-time BPS, real-time IOPS, and real-time latency of all virtual machines (VMs).

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Load Information**.

   The **NC Info** tab appears.
4. Click the **Virtual Machine Information** tab.
5. Select a cluster from the **Cluster Name** drop-down list, select an NC IP address from the **nc_Ip** drop-down list, and then set **Time Frame** to **One Hour**, **Three Hours**, **Six Hours**, **One Day**, or a customized time range. Click **Search** and view the following information:

   - Survey

     The **Survey** section shows a list of all VMs under an NC of the selected cluster. Click a VM name to view the trend charts of the real-time data of the VM.

   - Reading and Writing B/S

     Click a VM name in the **Survey** section. In the **Reading and Writing B/S** section, the trend chart of the reading and writing rate of the current VM within the specified time range is displayed.

     You can drag the slider below the chart to zoom in or out the chart.

     - bs_w: the instance writing rate
     - bs_r: the instance reading rate

- Real-time BPS

  Click a VM name in the **Survey** section. In the **Real-time BPS** section, the BPS trend of the current VM within the specified time range is displayed.

  You can drag the slider below the chart to zoom in or out the chart.

  - bps_w: the amount of data written per unit time
  - bps_r: the amount of data read per unit time

- Real-time IOPS

  Click a VM name in the **Survey** section. In the **Real-time iops** section, the IOPS trend chart of the current VM within the specified time range is displayed.

  You can drag the slider below the chart to zoom in or out the chart.

  - iops_w: the number of times per second data is written to the disk
  - iops_r: the number of times per second data is read from the disk

- Real-time latency

  Click a VM name in the **Survey** section. In the **Real-time latency** section, the latency trend of the current VM within the specified time range is displayed.

  You can drag the slider below the chart to zoom in or out the chart.

  - latency_w: the real-time latency of data writes
  - latency_r: the real-time latency of data reads
  - latency_w_qos: the intelligent QoS adjustment of real-time data writing latency
  - latency_r_qos: the intelligent QoS adjustment of real-time delay in reading data

1.1.6.5.6.3. View block device information

The Load Information module allows you to view the overview information and the trend charts of the data read/write rate, real-time BPS, real-time IOPS, and real-time latency for each device.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Load Information**.

   The **NC Info** tab appears.
4. Click the **Block Device Information** tab.
5. Select a cluster from the **Cluster Name** drop-down list, select an NC IP from the **nc_Ip** drop-down list, select an instance name from the **vmName** drop-down list, and then specify **Time Range** (**1 Hour**, **3 Hours**, **6 Hours**, **One Day** or a customized time range). Click **Search**. View the following information:

   - Survey

     The **Survey** section shows the information about all block devices in each VM of an NC in the current cluster. You can click a disk ID to view the real-time data trend charts for the device.

- Reading and Writing B/S

  In the **Survey** section, click a disk ID. In the **Reading and Writing B/S** section, the read/write rate trend chart of the current block device within the specified time range appears.

  You can drag the slider below the chart to zoom in or out the chart.

  - bs_w: the instance writing rate
  - bs_r: the instance reading rate

- Real-time BPS

  In the **Survey** section, click a disk ID. In the **Real-time BPS** section, the BPS trend chart of the current block device within the specified time range appears.

  You can drag the slider below the chart to zoom in or out the chart.

  - bps_w: the volume of data written per unit time
  - bps_r: the volume of data read per unit time

- Real-time IOPS

  In the **Survey** section, click a disk ID. In the **Real-time iops** section, the IOPS trend chart of the current block device within the specified time range appears.

  You can drag the slider below the chart to zoom in or out the chart.

  - iops_w: the number of times that data is written to the disk per second
  - iops_r: the number of times that data is read from the disk per second

- Real-time latency

  In the **Survey** section, click a disk ID. In the **Real-time latency** section, the latency trend chart of the current block device with the specified time range appears.

  You can drag the slider below the chart to zoom in or out the chart.

  - latency_w: the real-time latency of data writes
  - latency_r: the real-time latency of data reads
  - latency_w_qos: the QoS intelligent adjustment for real-time latency of data writes
  - latency_r_qos: the QoS intelligent adjustment for real-time latency of data reads

1.1.6.5.7. EBS dashboard

The EBS Dashboard module allows you to view the overview information and cluster usage trend charts of EBS clusters.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > EBS Dashboard**.

   On the page that appears, cluster overview information and cluster usage trend charts of all EBS clusters are displayed.

4. Select a cluster from the **Cluster Name** drop-down list.
5. 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 Disks**, **Abnormal Masters**, **Deleting Status**, **Abnormal Block GcWorkers**, or **Abnormal Block Servers** is greater than 0, the corresponding value is displayed in red.
   - The **Trend Chart of Cluster Usage** section shows the storage usage curve of the cluster for the last 30 days.

1.1.6.5.8. Block master operations

The Block Master Operations module shows the block master node information of Elastic Block Storage (EBS) clusters, including the IP addresses and roles. The module also allows you to switch the role of a node to leader as well as query and configure flags.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > 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.
4. Select a cluster from the **Cluster Name** drop-down list.
5. In the **Master List** section, 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.

     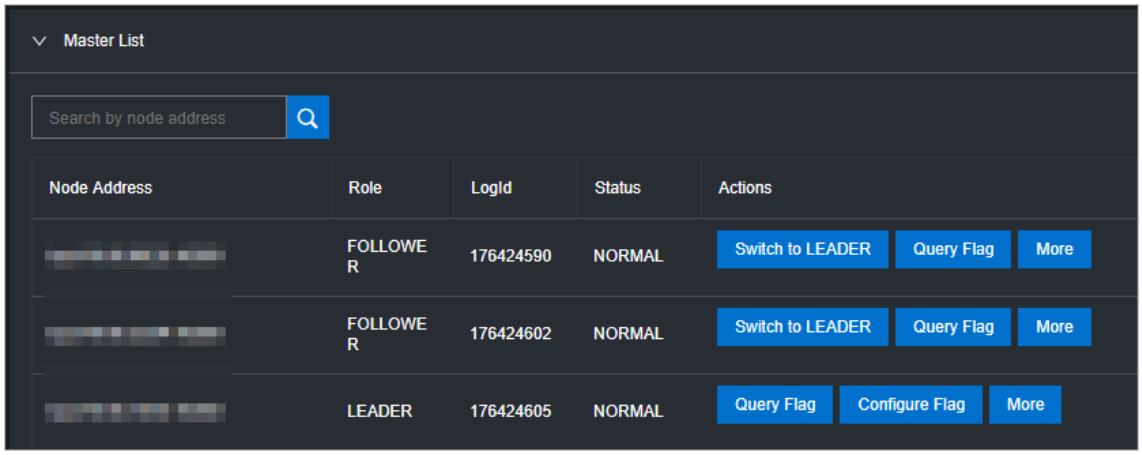

   - Switch to leader

     A leader role for a master node has the same features 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 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 list, find a node that assumes the leader role and click **Configure Flag** in the **Actions** column. In the dialog box that appears, configure the parameters and click **OK**.

  The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_blockmaster_flag.json* file.
flag_value	The custom flag value.
flag_type	The 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.

6. In the **Cluster Overview** section, you can query the disk size, number of segments, total storage size, and storage usage of the cluster.

1.1.6.5.9. Block server operations

The Block Server Operations module shows the block server node information of Elastic Block Storage (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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Block Master Operations**.

   On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
4. Select a cluster from the **Cluster Name** drop-down list.
5. In the **Server List** section, perform the following operations:
   - View the server node list

     You can view the 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 list, find a node and click **Query Flag** in the **Actions** column. In the dialog box that appears, set flag_key and 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. Click the **Configure** tab.

     e. 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 list, find a node and click **Configure Flag** in the **Actions** column. In the dialog box that appears, set flag_key and flag_value, select flag_type, and then click **OK**.

The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_blockserver_flag.json* file.
flag_value	The custom flag value.
flag_type	The flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the server node status

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

  The following table describes the server node status.

Status	Description
NORMAL	The node is running normally.
DISCONNECTED	The node is disconnected.
OFFLOADING	The node is being disabled.
OFFLOADED	The node is disabled.
UPGRADE	The node is upgraded.
RECOVERY	The node is restored.

- Query the version information

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

6. In the **Block Server Blacklist** section, perform the following operations:
   - 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 click **OK**.

     The block server node that is added to the blacklist is disabled and no longer provides services.
   - View the block server blacklist

In the **Block Server Blacklist** section, you can view all block server nodes that are added to the blacklist.

- 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 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 services.

1.1.6.5.10. Snapshot server operations

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > SnapShotServer**.

   On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
4. Select a cluster from the **Cluster Name** drop-down list.
5. Perform the following operations:
   - View the snapshot server node list

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

     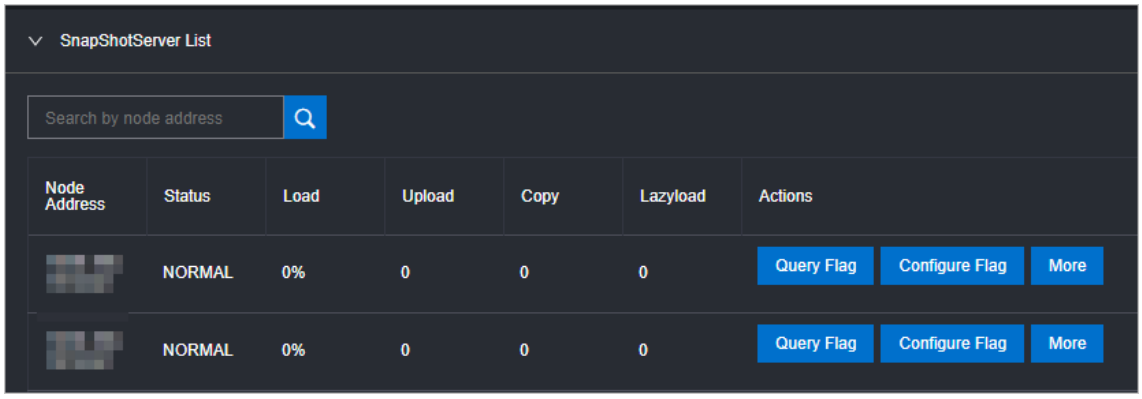

   - Query a flag

     In the snapshot server node list, find a node and click **Query Flag** in the **Actions** column. In the dialog box that appears, set flag_key and 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. Click the **Configure** tab.

e. 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, find a node and click **Configure Flag** in the **Actions** column. In the dialog box that appears, set flag_key, flag_value, and flag_type, and click **OK**.

The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_snapshotserver_flag.json* file.
flag_value	The custom flag value.
flag_type	The flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the snapshot server node status

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

The following table describes the snapshot server node status.

Status	Description
NORMAL	The node is running normally.
DISCONNECTED	The node is disconnected.
OFFLOADING	The node is being disabled.
OFFLOADED	The node is disabled.

- Query the version information

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

1.1.6.5.11. Block gcworker operations

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Block Gcworker Operations**.

   On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
4. Select a cluster from the **Cluster Name** drop-down list.
5. Perform the following operations:
   - View the gcworker node list

     You can view the IP addresses and status of the block gcworker nodes in the selected cluster.

     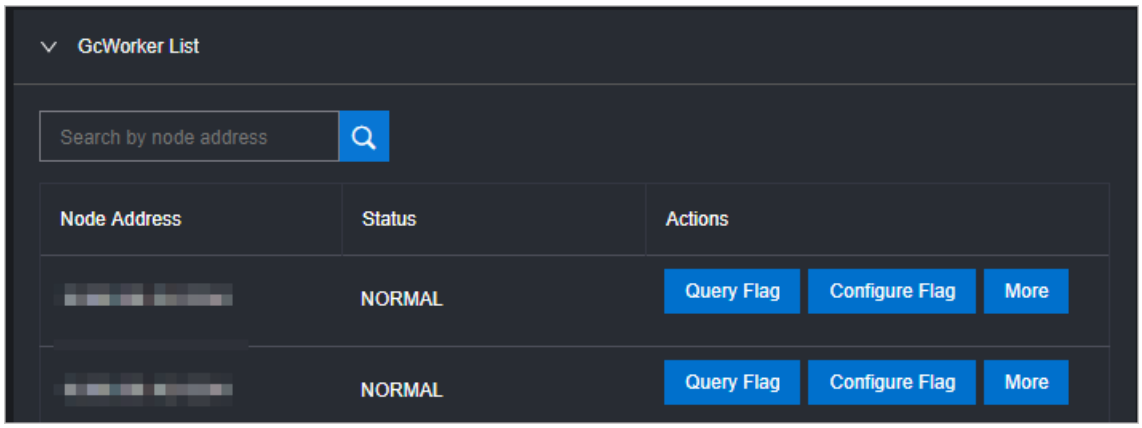

   - Query a flag

     In the gcworker node list, find a node and click **Query Flag** in the **Actions** column. In the dialog box that appears, set flag_key and 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, find a node and click **Configure Flag** in the **Actions** column. In the dialog box that appears, set flag_key, flag_value, and flag_type, and click **OK**.

  The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_blockgcworker_flag.json* file.
flag_value	The custom flag value.
flag_type	The flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the gcworker node status

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

  The following table describes the gcworker status.

Status	Description
NORMAL	The node is running normally.
DISCONNECTED	The node is disconnected.
OFFLOADING	The node is being disabled.
OFFLOADED	The node is disabled.

- Query the version information

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

1.1.6.5.12. Device operations

The Device Operations module allows you to view disk information in Elastic Block Storage (EBS) clusters such as the disk ID, state, capacity, and category. You can also perform flush operations, modify disk configurations, query segment information, and enable, disable, delete, or restore devices.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > EBS > Device Operations**.

   On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
4. Select a cluster from the **Cluster Name** drop-down list.
5. You can perform the following operations:
   - View the device list

     You can view the total number of devices, the total amount of logical space of devices, and information of each device in the cluster, including the device ID, state, logical capacity, number of segments, mode, and flags.
   - Check all segments in a cluster

     In the upper-right corner of the **Device List** section, click **Global Check Segment** to view all the segments in the selected cluster and their indexes and states.
   - Check disk status

     In the upper-right corner of the **Device List** section, click **Check Cloud Disk Status** to 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 disk information such as the disk ID, state, and capacity.
   - Delete a device

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

     After the device is deleted, its state changes to **DELETING** and the device becomes unavailable. While the device is in this state, operations such as enabling the device and modifying configurations cannot be performed on the device.
   - Restore a device

     In the device list, find a deleted device that is in the **DELETING** state and click **Restore** in the **Actions** column. In the message that appears, click **OK** to restore the deleted device to a normal state.

     After the device is restored, it becomes available and operations such as enabling the device and modifying configurations can be performed on the device.
   - Enable a device

     In the device list, find a device and choose **More > Turn On** in the **Actions** column. In the dialog box that appears, configure 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 used to enable a device.

Parameter	Description
client_ip	Optional. The IP address of the client on which to enable the disk. The client IP address is the IP address of the block server. If the client IP address is not specified, the IP address of your computer is used.
token	The string used as a token to disable the device.
mode	The disk mode. Valid values: ■ **ro**: read-only ■ **rw**: read and 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 when you disable a disk.

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

The following table describes the parameters used to disable a device.

Parameter	Description
client_ip	The IP address of the client on which to disable the disk. If the client IP address is not specified, the IP address of your computer is used.
token	The token to use to disable the device. This token was configured when the device is enabled. You can run the **dev -query** command on a server in the EBS cluster to query the token
open_ver	The current openversion of the device when the client IP address is not specified. If a client IP address is specified, you do not need to specify the openversion. You can run the **dev -query** command on a server in the EBS cluster to query the openversion.

- Flush a device

In the device list, find a device and choose **More > Flush** in the **Actions**. In the dialog box that appears, configure parameters and click **Submit** to flush the transaction logs of the disk or its segments.

The following table describes the parameters.

Parameter	Description
segment	The segment that you want to flush. If you do not specify this parameter, all segments are flushed.
ifnsw	Specifies whether to flush the index file. Valid values: ■ **0**: The index file is flushed. ■ **1**: The index file is not flushed.
dfnsw	Specifies whether to flush the data files. Valid values: ■ **0**: The data files are flushed. ■ **1**: The data files are not flushed.

- Perform a global flush operation for a cluster

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

  On the right of the **Device List** section, click **Global Flush**. In the dialog box that appears, select ifnsw and dfnsw and click **OK** to flush the transaction logs of all disks or segments in the current cluster.

- Query configuration status

  In the device list, find a device and choose **More > Query Configuration Status** in the **Actions** column. In the dialog box that appears, enter config_ver and click **OK**. You can determine whether the disk is configurable based on the check result.

  config_ver is the config_version parameter among the queried device information.

- Modify device configurations

  You can modify the configurations of a disk, such as the compression algorithm, storage mode, and whether data compression is enabled.

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

  The following table describes the parameters.

Parameter	Description

Parameter	Description
compress	Specifies whether to enable data compression. Valid values: ■ enable ■ disable
algorithm	The data compression algorithm. Valid values: ■ 0: No data compression algorithms are used. ■ 1: The snappy data compression algorithm is used. ■ 2: The lz4 data compression algorithm is used.
ec	Specifies whether to enable the EC storage mode. Default value: disable. Valid values: ■ enable ■ disable
data_chunks	The number of data chunks. Default value: 8.
parity_chunks	The number of parity chunks. Default value: 3.
packet_bits	The size of a single data block in EC mode. Default value: 15.
copy	The number of data replicas. Default value: 3.
storage_mode	The storage mode of the disk.
cache	Specifies whether to enable the cache mode. Default value: 0. Valid values: ■ 0: disables the cache mode. ■ 1: enables the cache mode.
storage_app_name	The data storage name.
simsuppress	Specifies whether to enable the delay simulation feature. Default value: disable. Valid values: ■ enable ■ disable
baselatency	The basic latency. Default value: 300.
consumespeed	The processing speed. Default value: 256 B/μs.
lat80th	The 80th percentile latency. The default value is 110% of the specified basic latency.

Parameter	Description
lat90th	The 90th percentile latency. The default value is 150% of the specified basic latency.
lat99th	The 99th percentile latency. The default value is 500% of the specified basic latency.

- Query segment information

  In the device list, find a device and choose **More > Segment Information** in the **Actions** column. In the dialog box that appears, view the information of the segments, such as the indexes and states.

- Check a segment

  In the device list, find a device and choose **More > Check Segment** in the **Actions** column. In the dialog box that appears, select a segment and click **Submit** to view the information of the segment such as the index and state.

1.1.6.5.13. Enable or disable Rebalance

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Rebalance**.
4. 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**.

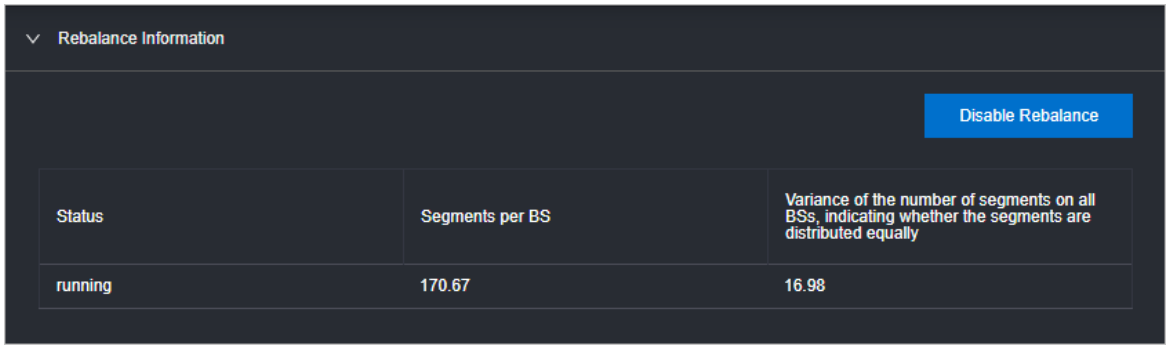

1.1.6.5.14. I/O hang fault analysis

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > IO HANG**.

   By default, the system shows the affected VM list, VM cluster statistics, and device cluster statistics for the last 24 hours.
4. Select a time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or a customized time range) and click **Search**. View the following information:
   - **Affected VM List**

     The **Affected VM List** section shows the I/O hang start time and recovery time of all the VMs, as well as the cluster name and user ID of the cluster to which these VMs belong.

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

     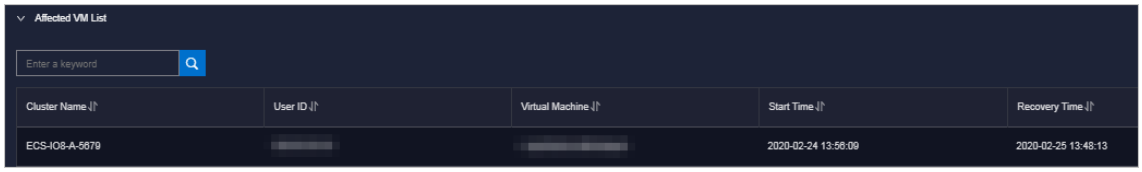

   - **VM Cluster Statistics**

     The **VM Cluster Statistics** section shows 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.

     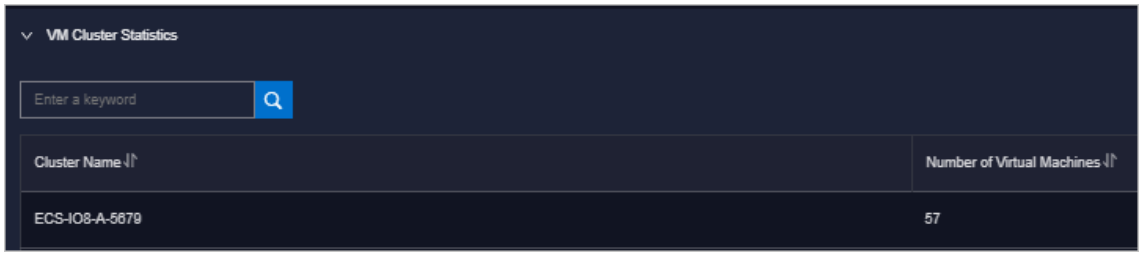

   - **Device Cluster Statistics**

     The **Device Cluster Statistics** section shows 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.

     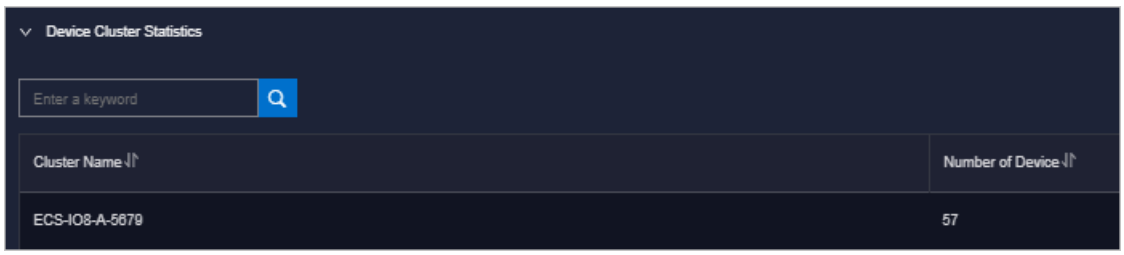

1.1.6.5.15. Slow IO analysis

The Slow IO Analysis page allows you to view the slow IO list, top ten NCs, cluster statistics, top five cluster statistics, and reasons.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > SLOW IO**.

   By default, the system shows the slow IO list, top ten NCs, cluster statistics, top five cluster statistics, and reasons in the last 24 hours.
4. Select the time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or customize the time range) and click **Search**. View the following information:
   - **Slow IO List**

     The **Slow IO List** section shows the slow IO-related cluster name, NC IP address, virtual machine, device ID, storage type, start time, recovery time, number of slow IOs, and causes.

     To view the information of a cluster, an NC, or a 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 data by Cluster Name, NC IP, Virtual Machine, Device ID, Storage Type, Start Time, Recovery Time, Number of Slow IO, or Reason.
   - **Top10 NC**

     The system shows the information of the top ten NCs on a graph and table.

     Notes:

     - The **Graphical Analysis** section shows the proportion for the number of slow IO in each cluster of the top ten NCs by using a pie chart.
     - The **Top10 NC** section shows the NC IP address, cluster name, number of slow IOs, percentage, and primary cause of slow IOs on the top ten NCs.

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

       You can also sort data by NC IP, Cluster Name, Slow IO, and Major Reason.
   - **Cluster Statistics**

     The **Cluster Statistics** section shows the cluster name, number of devices, number of slow IOs, percentage, and primary cause of slow IOs on clusters.

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

     You can also sort data by Cluster Name, Number of Device, Number of Slow IO, and Major Reason.
   - **Top Five Cluster Statistics**

     The system shows the statistics of top five clusters by using a graph and a table.

     Notes:

     - The **Graphical Analysis** section shows the proportion for the number of slow IOs on each of the top five clusters on a pie chart.

- The **Top Five Cluster Statistics** section shows the cluster name, number of devices, number of slow IOs, percentage, and primary cause of slow IOs on the top five clusters on a table.

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

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

- **Reason**

  The system shows the primary cause on a graph and table.

  Notes:

  - The **Graphical Analysis** section shows the proportion of reasons by using a pie chart.
  - The **Reason** section shows the number of slow IO from the dimension of reasons.

    To query the information of a reason, enter the reason information in the search box to perform a fuzzy search.

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

1.1.6.5.16. Product settings

The Product Settings module allows you to view the sales status of a cluster, configure the overcommit ratio of a cluster, and specify whether a cluster is available for sale.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Product Settings**.

   By default, the system shows the data of each cluster within the current environment, including the cluster name, overcommit ratio, and sales status.

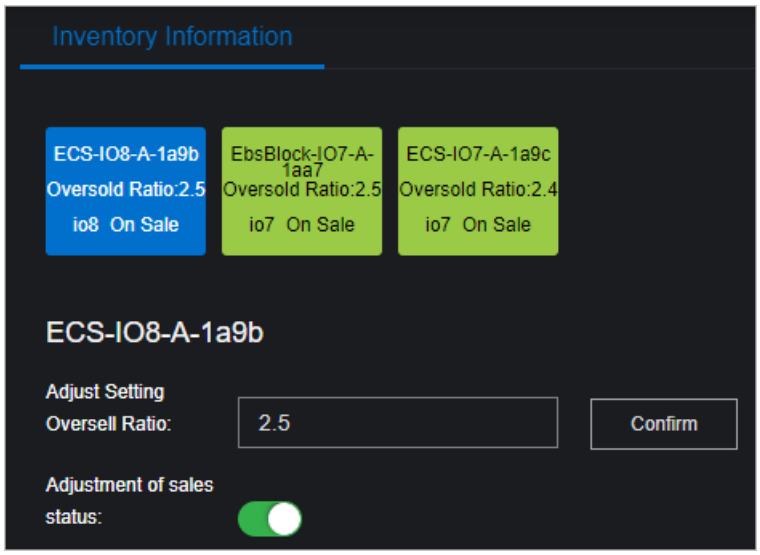

> **Note** An overcommit ratio is the ratio of the marketable capacity of a storage device to the physical capacity. For example, if the physical storage capacity of a storage device is 1 TB and marketabe capacity is 2.5 TB, and the overcommit ratio is 2.5.

4. Perform the following operations:
   - Select a cluster, enter a number in the **Adjust Setting Oversell Ratio** field, and then click **Confirm** to set the oversold ratio of the cluster.
   - Select a cluster and turn on or off **Adjustment of sales status** to enable or disable the cluster for sale.

1.1.6.5.17. View ECS disk size rankings

The ECS Disk Size Ranking module allows you to view the amount of space occupied by all disks in the elastic block storage attached to an Elastic Compute Service (ECS) cluster in Apsara Distributed File System.

Context

When an ECS cluster occupies a large amount of space in Apsara Distributed File System, the on-site O&M personnel must check the space occupied by each disk in the elastic block storage attached to the ECS cluster. Then, they must contact the business side to migrate data and release disks. The ECS disk size ranking feature helps O&M personnel identify which disks occupy a large amount of space in Apsara Distributed File System so that they can perform targeted cleanup and quickly lower space usage.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > EBS > ECS Disk Size Ranking**.
4. Select the ECS cluster that you want to query from the **Cluster** drop-down list and click **Search**.

   All disks in the elastic block storage attached to the selected ECS cluster are listed from large to small based on the actual amount of space that the disks occupy in Apsara Distributed File System. You can view the cluster name, cluster ID, and zone of the selected cluster. You can also view the storage type, size, and identifier of each disk.
5. (Optional) You can click **Reset** to clear the preceding search conditions.

1.1.6.6. 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.

1.1.6.6.1. Overview

This topic describes the features of the Task Management module.

The Task Management module has the following features:

- Supports task overview and allows you to create tasks in a quick manner.
- Supports four task execution modes: manual, scheduled, regular, and advanced.
- Supports the breakpoint feature for scripts, which allows a task to be paused between two scripts to wait for manual intervention.
- Allows you to query tasks by name, status, and creation time.
- Allows you to upload scripts by using .tar packages.

1.1.6.6.2. View task overview

The Task Overview page shows the overall running conditions of tasks in the system. You can also create tasks on this page.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Task Management > Overview**.

   The **Task Overview** page appears.

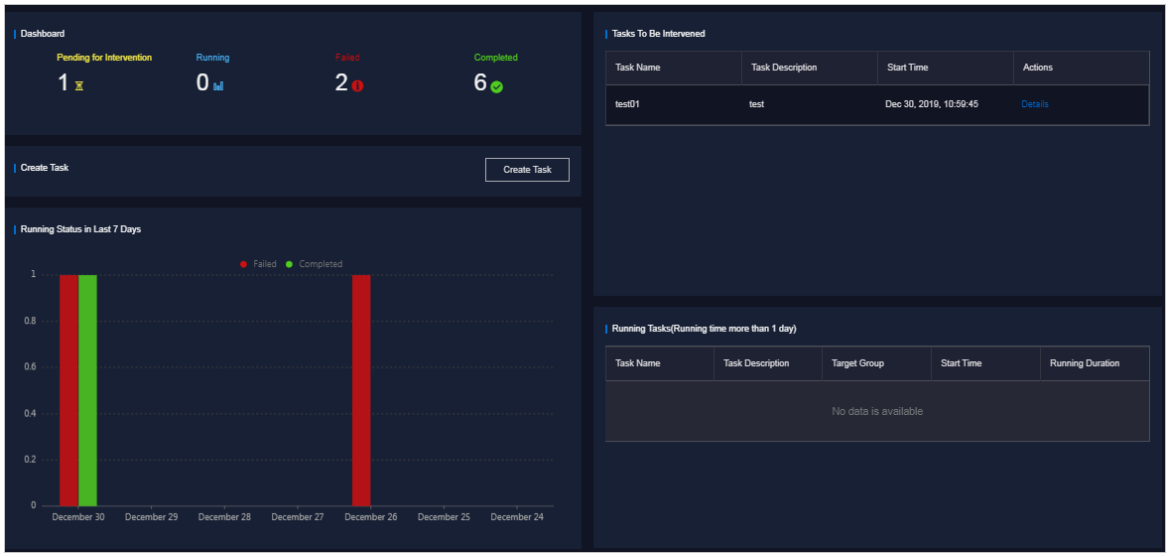

4. Perform the following operations:
   - In the **Dashboard** section, view the number of tasks that are in the **Pending for Intervention**, **Running**, **Failed**, or **Completed** state.

     Click a state or number to view tasks in 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 to proceed. 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 run trend and success information of tasks within the last seven days.
   - In the **Running Tasks (Running time more than 1 day)** section, view the running status of

tasks within the last 24 hours.

1.1.6.6.3. Create a task

You can make regular modifications as tasks and run tasks in the Apsara Uni-manger Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Task Management > Tasks**.
4. Click **Create**.
5. In the dialog box that appears, configure the parameters.

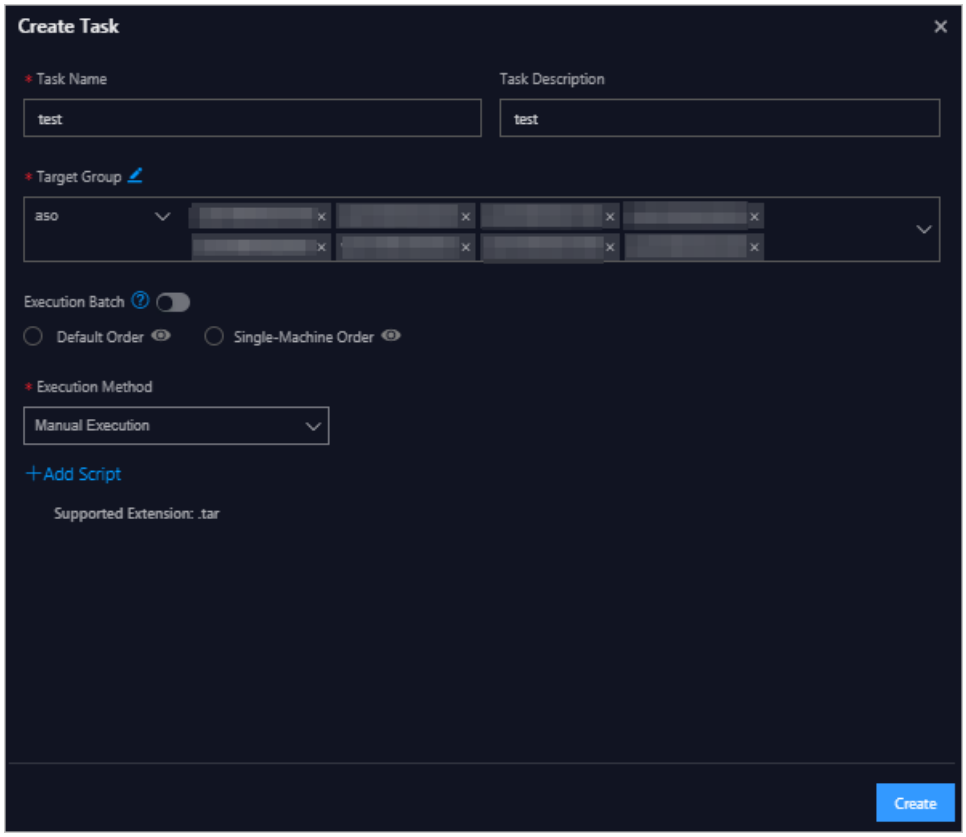

Parameter	Description
Task Name	The name of the O&M task.
Task Description	The description of the O&M task.

Parameter	Description
Target Group	The task target. You can use one of the following methods to configure the target group: ○ Select a product. Enter the VM or physical machine in the field and 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**.
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 > Tasks**. If **Execution Batch** is specified, **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 remaining machines in batch 5. You can adjust the batch for machines based on your needs. 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. You can change the number of machines in each batch. ○ **Single-Machine Order**: By default, each batch has one machine. You can adjust the batch for machines based on your actual needs.
Execution Method	If **Execution Batch** is specified, **Execution Method** can be set only to **Manual Execution**. If **Execution Batch** is not enabled, you can select one of the following execution methods: ○ **Manual Execution**: Start the task manually. When **Manual Execution** is selected, 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. If the execution condition is met, the task starts again. ○ **Advanced**: Configure the command to periodically run the task.

Parameter	Description
Add Script	Click **Add Script**. Select one or more .tar packages to upload the script file. After you upload a script, you can delete and re-upload the script. After you upload a script, if **Execution Method** is set to **Manual Execution**, you must specify whether to enable **Intervention Required**. If manual intervention is enabled, when you run the script, the task is suspended and waits for manual intervention.

6. Click **Create**.

Result

The created task is displayed in the task list.

1.1.6.6.4. View the execution status of a task

After a task starts, you can view its execution status.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Task Management > Tasks**.
4. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Search** to search for tasks.
5. Find the task that you want to view and 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 **Execution Batch** is selected when you create a task, **Batch Execution Policy** is displayed in the **Target Group** column.

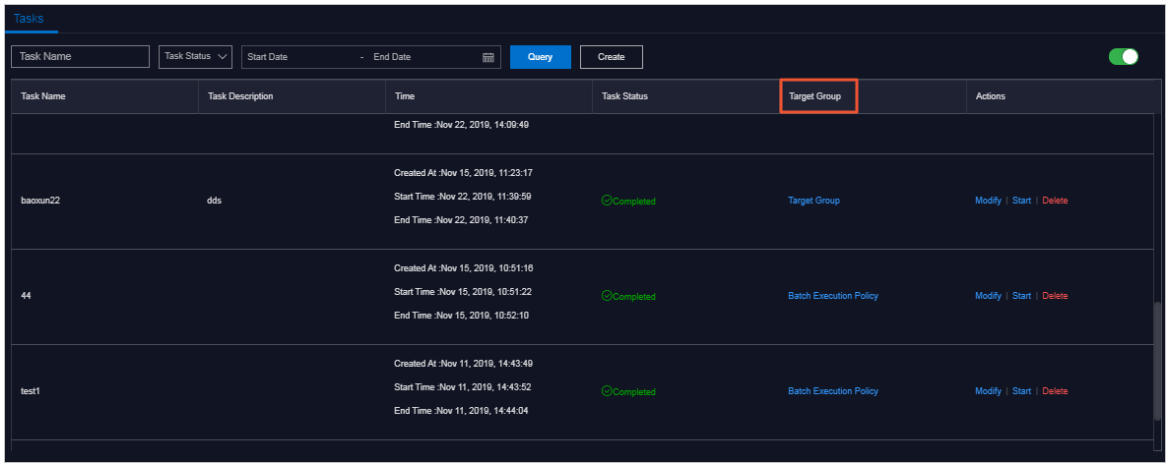

6. In the dialog box that appears, view the task execution status based on the machine color. Click a machine name to view the task execution results on it.

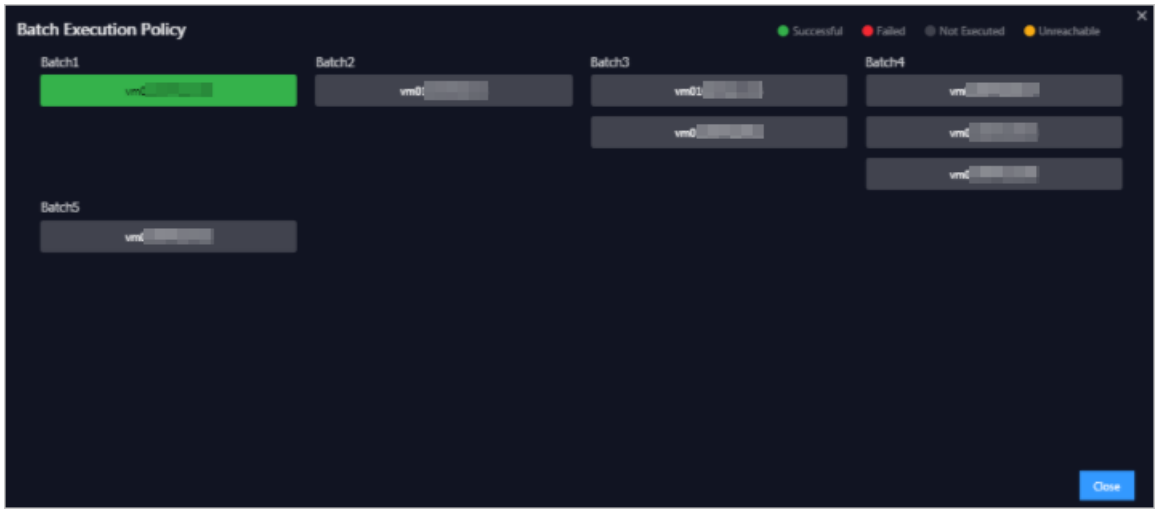

1.1.6.6.5. Start a task

If you select **Manual Execution** when you create a task, you must manually start the task after it is created.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Task Management > Tasks**.
4. (Optional)On the Tasks page, enter the task name, select the task status, start date, and end date, and then click **Query**.
5. Find the task that you want to start and click **Start** in the **Actions** column.
6. In the dialog box that appears, select the batches to start and click **Start**.

   For a new task, after you click **Start** for the first time, the system prompts you that the task is started. The virtual machines (VMs) or physical machines in batch 1 start to run the task. Click **Start** again. 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 value of **Task Status** is changed to **Pending for Intervention**, and the task can be resumed only by clicking **Continue** in the **Actions** column.

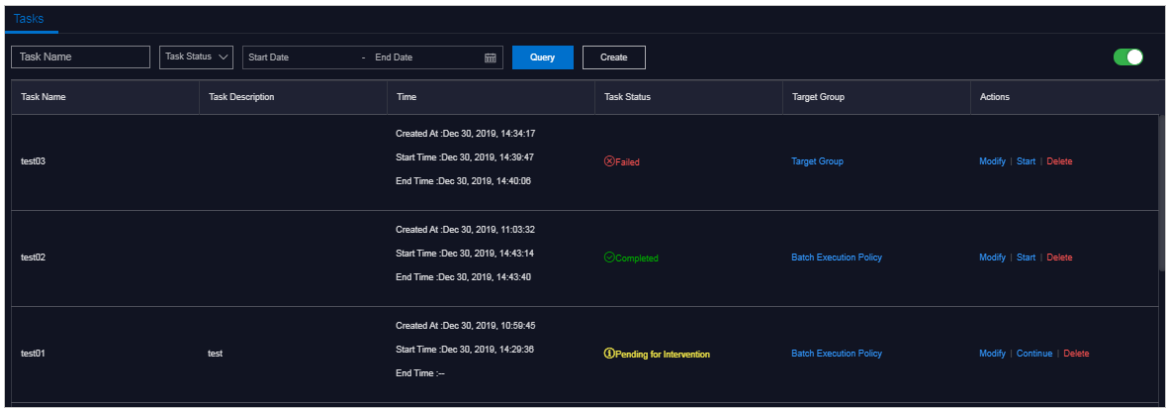

1.1.6.6.6. Delete a task

You can delete tasks that are no longer needed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Task Management > Tasks**.
4. (Optional)Enter a task name, select a task status, start date, and end date, and then click **Query** to search for tasks.
5. Find the task that you want to delete and click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

1.1.6.6.7. Process tasks to be intervened

If a task reaches a breakpoint, the task stops and waits for manual confirmation. The task proceeds only after you provide confirmation.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Task Management > Overview**.
4. In the **Tasks To Be Intervened** section, find the task to be intervened and click **Details** in the **Actions** column.

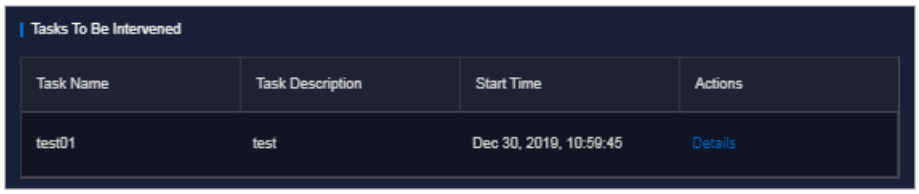

5. On the **Task Details** tab, check the information and click **Continue** for the task to continue.

1.1.6.7. Security operations

The security operations feature provides CLI logon and remote O&M, supports blocking and approval of high-risk operations, and can audit all operations.

The security operations feature provides a web terminal that can be used to:

- Log on to user machines such as virtual machines, hosts, and containers.
- View the environment metadata, OOB information, and cluster configurations of the current user.
- Upload and download files.
- Audit all operations.
- Deliver blocking, secondary verification prompt, and execution after approval for high-risk operations.
- Implement security control configurations for projects to be connected to Apsara Stack Online.

1.1.6.7.1. Fast arrival

You can log on to machines in the Apsara Stack environment such as virtual machines, hosts, and containers and run Linux commands to perform operations. You can also view environment metadata, OOB information, and cluster configurations.

1.1.6.7.1.1. Log on to the host where a server role is deployed

You can log on to the virtual machine, host, or container where a server role is deployed and upload or download files.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Fast Arrival**.

   By default, the **Server Role Logon** tab appears.
4. Select a server role name from the **Server Role** drop-down list.

   > **Note** The drop-down list also supports fuzzy search. You can move the pointer to and click the drop-down list, enter a keyword, and select a name from the fuzzy search results.

5. In the **Host** column, find a host and click Log On to log on to the virtual machine or host where the server role is deployed.

   > **Note** If two links appear in the **Host** column, the first one is used to log on to the virtual machine and the second one is used to log on to the host of the virtual machine.

   i. After you log on to the virtual machine or host where a server role is deployed, enter Linux commands in the CLI window to perform related operations.

   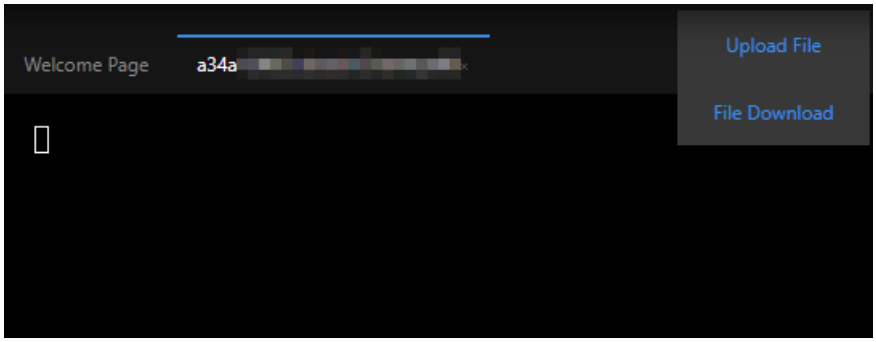

   ii. Click **Upload File** in the CLI window. The **Upload File** dialog box appears. You can upload an object in one of the following ways:

   - Click the dotted box. In the dialog box that appears, select the file to be uploaded, click **Open**. Click **Upload** in the **Upload File** dialog box.
   - Drag the file to the dotted box and then click **Upload**.

iii. Click **File Download**. The **File Download** dialog box appears. Set **File Name** and **File Directory** and then click **Download** to download the file to the default download directory of the local browser.

> **Note** The uploaded and downloaded files cannot exceed 200 MB in size.

6. In the **Docker** column, you can click the relevant links to log on to and restart the container, or view logs and inspection reports of the container.

1.1.6.7.1.2. Log on to the virtual machine where a server role group is deployed

You can log on to the virtual machine where a server role group is deployed and upload or download files.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Fast Arrival**.
4. Click the **Server Role Group Logon** tab.
5. Select a server role group of a product from the **Server Role Group** drop-down list. The server roles included in the group are displayed in the lower part of the page.

> **Note** The drop-down list also supports fuzzy search. You can move the pointer to and click the drop-down list, enter a keyword, and select a name from the fuzzy search results.

6. Find a server role and click **Log On** in the **Machine** column to log on to the virtual machine where the server role is deployed.
   i. After you log on to the virtual machine where a server role is deployed, enter Linux commands in the CLI window to perform related operations.

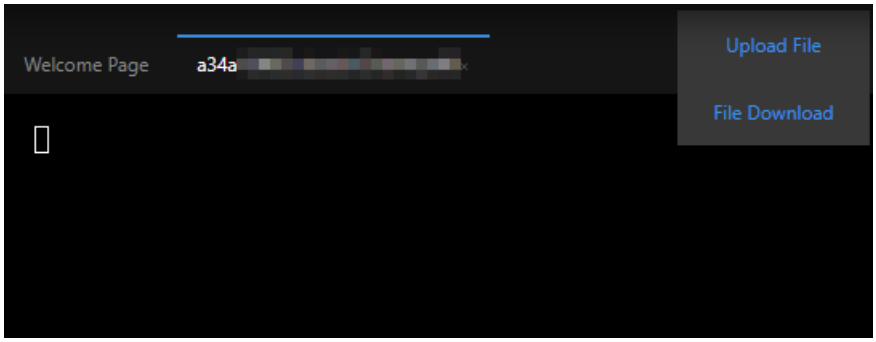

   ii. Click **Upload File** in the CLI window. The **Upload File** dialog box appears. You can upload an object in one of the following ways:
   - Click the dotted box. In the dialog box that appears, select the file to be uploaded, click **Open**. Click **Upload** in the **Upload File** dialog box.
   - Drag the file to the dotted box and then click **Upload**.

iii. Click **File Download**. The **File Download** dialog box appears. Set **File Name** and **File Directory** and then click **Download** to download the file to the default download directory of the local browser.

> **Note** The uploaded and downloaded files cannot exceed 200 MB in size.

1.1.6.7.1.3. Query environment metadata

You can view the metadata of the service registered in Apsara Infrastructure Management Framework.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Fast Arrival**.
4. Click the **Environment Metadata Query** tab.
5. Select the service name from the **Service** drop-down list. The metadata of the service registered in Apsara Infrastructure Management Framework is displayed in the lower part of the page.

> **Note**
>
> - The drop-down list also supports fuzzy search. You can move the pointer to and click the drop-down list, enter a keyword, and select a name from the fuzzy search results.
> - You can also enter a keyword in the box above the displayed metadata to filter metadata.

1.1.6.7.1.4. Query OOB information

You can query the OOB information by specifying an IP or a serial number.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Fast Arrival**.
4. Click the **OOB Information Query** tab.
5. Enter an IP or a serial number in the field. In the **Query** column, click **Query** to view the OOB information.

1.1.6.7.1.5. Query cluster configurations

You can view the configuration files of all services deployed in a cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.

3. In the left-side navigation pane, choose **Security Operations > Fast Arrival**.
4. Click the **Cluster Configuration Query** tab.
5. Select a cluster of a product from the **Cluster Name** drop-down list. The configuration files of all services deployed in the cluster are displayed in the lower part of the page.

   > **Note** The drop-down list also supports fuzzy search. You can move the pointer to and click the drop-down list, enter a keyword, and select a name from the fuzzy search results.

6. Click a configuration file on the left to view its details on the right.

1.1.6.7.1.6. Log on to a metadatabase

You can log on to a metadatabase used by the server role included in the service and upload or download files.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Fast Arrival**.
4. Click the **Metadatabase Logon** tab.
5. Select a service name from the **Service** drop-down list.

   > **Note** The drop-down list also supports fuzzy search. You can move the pointer to and click the drop-down list, enter a keyword, and select a name from the fuzzy search results.

6. The metadatabases used by all server roles included in the service are displayed in the lower part of the page. Find a database and click **Writable Logon** in the **Actions** column.
   i. After you log on to the metadatabase, enter SQL statements in the CLI window to perform related operations.

   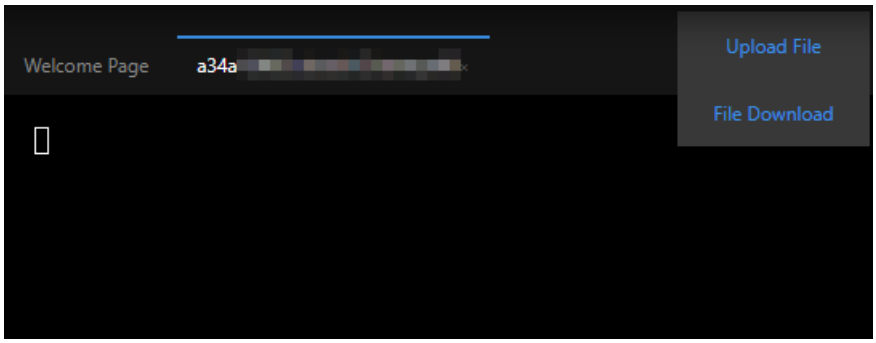

   ii. Click **Upload File** in the CLI window. The **Upload File** dialog box appears. You can upload an object in one of the following ways:
   - Click the dotted box. In the dialog box that appears, select the file to be uploaded, click **Open**. Click **Upload** in the **Upload File** dialog box.
   - Drag the file to the dotted box and then click **Upload**.

iii. Click **File Download**. The **File Download** dialog box appears. Set **File Name** and **File Directory** and then click **Download** to download the file to the default download directory of the local browser.

> **Note** The uploaded and downloaded files cannot exceed 200 MB in size.

1.1.6.7.2. Auditing

You can view command records and videos, file upload and download records, and authorization information related to the security operations feature.

1.1.6.7.2.1. View command records

You can view the command records of the security O&M feature.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Audit**.

   By default, the **Command Records** tab appears.
4. Enter a command in the **Command** field, select a time range, and then click **Search** to view the command records.

   > **Note** The system will audit the command that you enter and give one of the following audit results based on operational risks:
   >
   > - **pass**: The audit is successful.
   > - **fail**: The audit fails.
   > - **multiVerify**: A further verification is required.
   > - **codeVerify**: Authorization is required for use.
5. (Optional)Click **Advanced** in the upper-right corner. The **Machine**, **Service**, **Server Role**, and **Operated By** fields appear. Enter values in the fields and click **Search** to further filter command records.

   > **Note**
   >
   > - Click **Reset** to clear the values that you enter.
   > - Click **Collapse** to hide the preceding fields.
   > - The preceding fields support fuzzy search.

1.1.6.7.2.2. View file upload and download records

You can view file upload and download records of the security operations feature.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Audit**.
4. Click the **Upload and Download Audit** tab to view file upload and download records.

1.1.6.7.2.3. View authorization information

You can view command authorization information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Audit**.
4. Click the **Authorize** tab to view command authorization information.

1.1.6.7.2.4. View command videos

You can view the videos of all commands executed on the machine.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Audit**.
4. Click the **Playback** tab.
5. Select a time range, enter a value in the **Operated By** field, and then click **Search** to view command records.
6. Find a command record and click **View** in the **Actions** column.
7. In the video playback window, click the icon to play back the command video.
8. (Optional)Click **Advanced** in the upper-right corner. The **Machine**, **Service**, and **Server Role** fields appear. Enter values in the fields and click **Search** to further filter command records.

   > **Note**
   > - Click **Reset** to clear the values that you enter.
   > - Click **Collapse** to hide the preceding fields.
   > - The preceding fields support fuzzy search.

1.1.6.7.3. Rules

To control the risks of Linux commands, you can configure blocking rules for Linux commands.

1.1.6.7.3.1. View rules

You can view the information of the command blocking rules that you create or import, such as their details, confirmation and status.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Rules**.
4. Enter a command in the **Command** field and then click **Search** to view the information of its blocking rules.
5. (Optional)Click **Advanced** in the upper-right corner. The **Target Parameter**, **Service**, **Server Role**, **Verification Rule**, and **Status** fields appear. Enter values in the fields and click **Search** to further filter blocking files.

   > Note
   > - Click **Reset** to clear the values that you enter.
   > - Click **Collapse** to hide the preceding fields.

1.1.6.7.3.2. Create a rule

You can create a command blocking rule.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Rules**.
4. In the upper-left corner, click **Create Rule**. In the dialog box that appears, configure the parameters and then click **OK**.

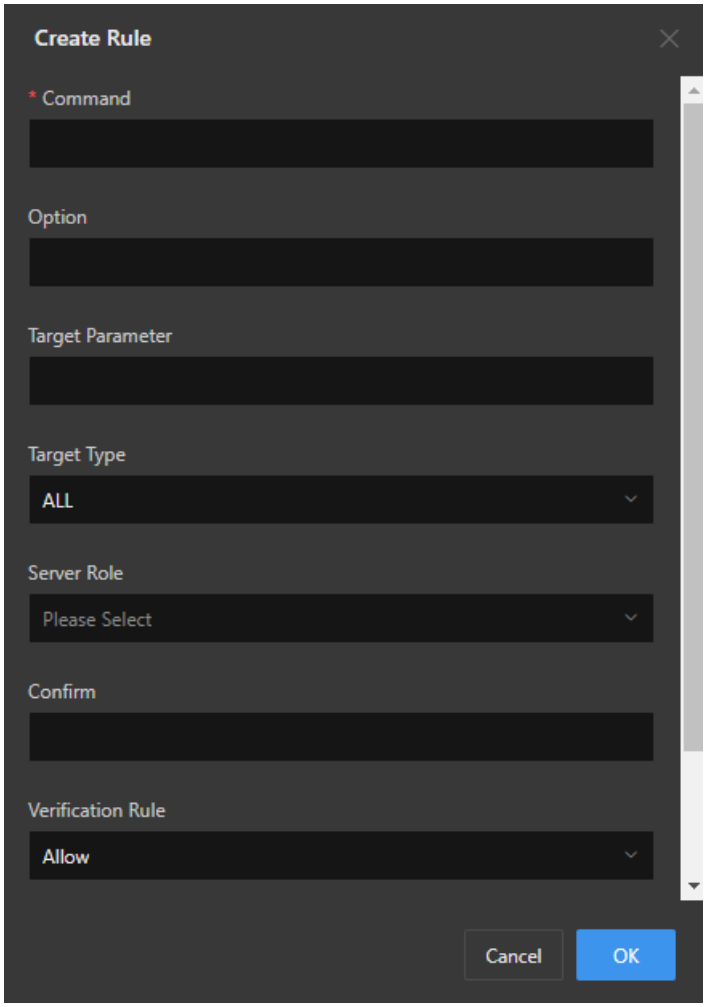

The following table describes the required parameters.

Parameter	Description	Example
Command	The Linux command.	mv
Option	The option of the command. For example, the option of the `rm -rf` command is **rf**.	rf
Target Parameter	The parameter of the command option. For example, in the `find / -name test` command, **name** is the option and **test** is the parameter. If the option does not have a parameter, the value is empty.	test

Parameter	Description	Example
Target Type	The type of the command option. Valid values: ○ ALL ○ FILE ○ DIR ○ OPTION	OPTION
Server Role	The server role of the machine where the command is implemented.	ram-ramService.RamPortalService#
Confirm	The prompt in the CLI window when the command is blocked.	Termination of the process is not allowed.
Verification Rule	The rule to block the command. Valid value: ○ Allow: The command can be run. ○ Block: The command is blocked. ○ Confirm Again: You must confirm again whether the command can be run. ○ Verification Code: The command can be run within a time range after the authorization is approved. If you select this value, a verification code is applied to the system. The verification code is required before the command can be run.	Allow
Status	The status of the command. You can click the button in the **Status** column to modify the status.	The **Status** button is turned on.

1.1.6.7.3.3. Batch import rules

You can batch import command blocking rules.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Rules**.

4. In the upper-left corner, click **Import Rules**. Select the command blocking rule file and then click **Open** to import the file.

   > **Note** The file to be imported must be in the .xlsx format. You can first export the template and then enter information in the template.

1.1.6.7.3.4. Batch export rules

You can batch export command blocking rules that you create.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Rules**.
4. In the upper-left corner, click **Export Rules**. In the dialog box that appears, select the download directory (the Download directory by default), and click **Download** to download the command blocking rule file to your computer.

1.1.6.7.3.5. Modify a rule

You can modify a command blocking rule that you create.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Rules**.
4. Find a rule and click **Modify** in the **Actions** column. In the dialog box that appears, modify the parameters and click **OK**.

1.1.6.7.3.6. Delete a rule

You can delete a command blocking rule that you create.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Rules**.
4. Find a rule and click **Delete** in the **Actions** column. In the message that appears, click **OK**.

1.1.6.7.4. Settings

When a project is connected to Apsara Stack Online, you can configure the IP address and port number of the worker that the Apsara Uni-manager Console can access, and the IP address of Apsara Stack Online that the Apsara Uni-manager Console can access.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Security Operations > Settings**.
4. In the **Configure Worker** section, enter the IP address and port number of the worker, and then click **Save**.
5. In the **Configure Simplified O&M Whitelist** section, enter the IP address of Apsara Stack Online and click **Save**.

   > **Note** Separate multiple IP addresses with commas (,).

   The following figure shows the remote O&M process.

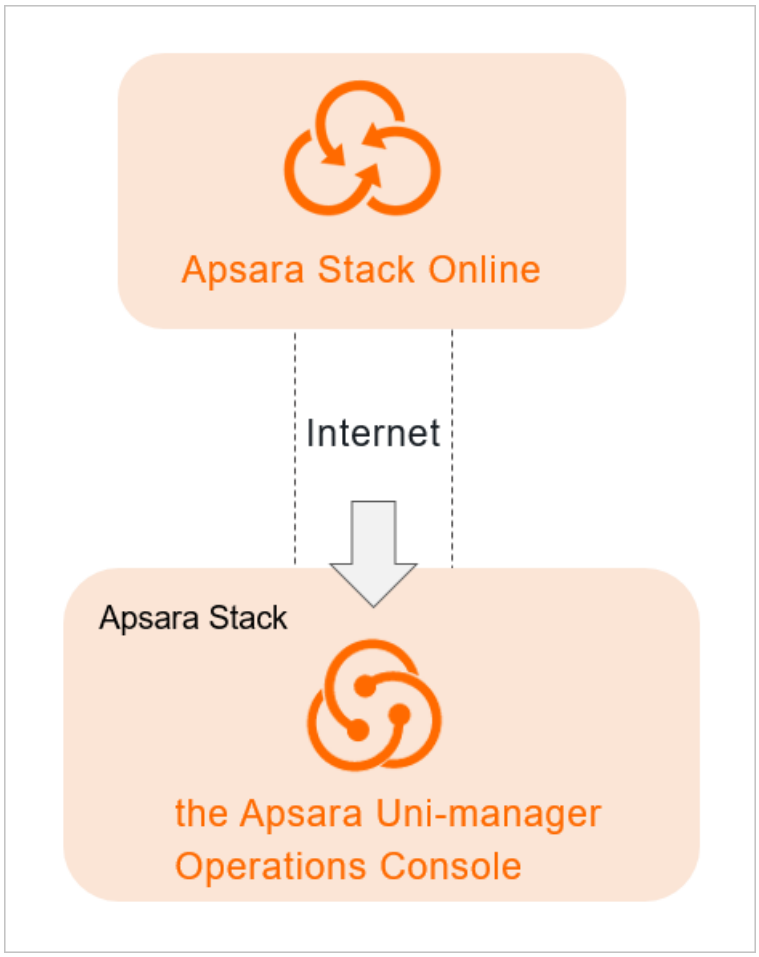

1.1.6.8. Platform encryption

1.1.6.8.1. Disk storage and transmission encryption

1.1.6.8.1.1. SM settings

The SM encryption feature uses the encryption algorithms certified by State Cryptography Administration to encrypt MiniRDS. The feature includes storage encryption for data in tables and transmission encryption such as database connections in SSL mode. You can enable or disable the SM encryption feature, view or update certificates, and add or delete server roles.

Enable SM encryption

After the SM encryption feature is enabled, the encryption algorithms certified by State Cryptography Administration can be used to encrypt MiniRDS.

Prerequisites

The aso-mgr and aso-opr services have reached the desired state.

Context

After the SM encryption feature is enabled, the service may not reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.

   > **Warning** When TDE and SSL are enabled, the Apsara Uni-manager Management Console and Apsara Uni-manager Operations Console cannot be accessed, the AK, SA, and RAM services become unavailable, and the control services are interrupted. Proceed with caution.

4. Find the application and turn on the switch in the **Actions** column. In the dialog box that appears, click **OK**.
5. In the **Task Progress** dialog box, view the task progress. If a step fails, an error message is displayed. The progress is updated in real time.
6. The **Reload** button appears below the step that fails. Click the button to execute the step again.
7. After the SM encryption feature is enabled, the deployment task of changing the kv configuration for the product is triggered. You can query the task details on Apsara Infrastructure Management Framework. Perform the following operations:
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   iii. Find the cluster and then click **Operations** in the **Actions** column.
   iv. On the **Cluster details** page, click the **Operations Logs** tab.
   v. Find the ASO update commit log entry and click **Version Differences** in the **Actions** column. On the **Version Differences** page, view the details.

Disable SM encryption

After the SM encryption feature is disabled, the encryption algorithms certified by State Cryptography Administration cannot be used to encrypt MiniRDS.

Prerequisites

The aso-mgr and aso-opr services have reached the desired state.

Context

After the SM encryption feature is disabled, the service may not reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.

   > **Warning** When SSL is disabled, the Apsara Uni-manager Management Console and Apsara Uni-manager Operations Console cannot be accessed, the AK, SA, and RAM services become unavailable, and the control services are interrupted. Proceed with caution.

4. Find the application and turn off the switch in the **Actions** column. In the dialog box that appears, click **OK**.
5. In the **Task Progress** dialog box, view the task progress. If a step fails, an error message is displayed. The progress is updated in real time.
6. The **Reload** button appears below the step that fails. Click the button to execute the step again.
7. After the SM encryption feature is disabled, the deployment task of changing the kv configuration for the product is triggered. You can query the task details on Apsara Infrastructure Management Framework. Perform the following operations:
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   iii. Find the cluster and then click **Operations** in the **Actions** column.
   iv. On the **Cluster details** page, click the **Operations Logs** tab.
   v. Find the ASO update commit log entry and click **Version Differences** in the **Actions** column. On the **Version Differences** page, view the details.

View execution history

You can view the execution history of product encryption operations.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.
4. Find the application.
5. Click **View** in the **Execution History** column.
6. In the panel that appears, view the execution history.

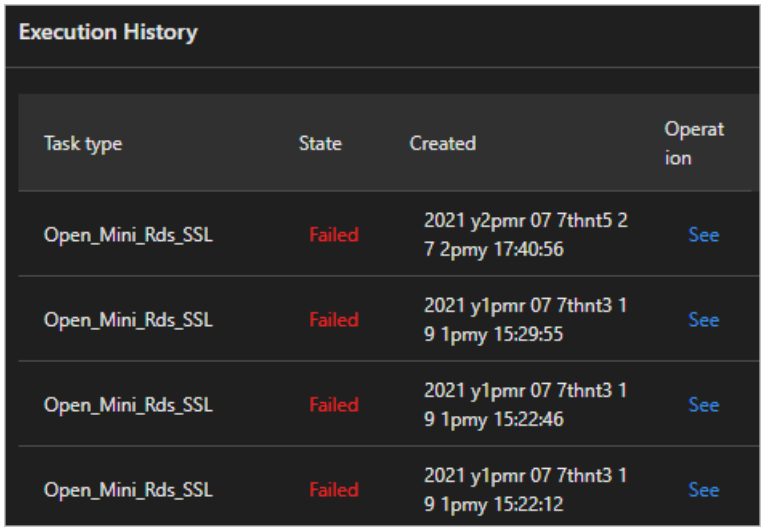

7. Find the task and click **View** in the **Actions** column. In the dialog box that appears, view the task execution details.

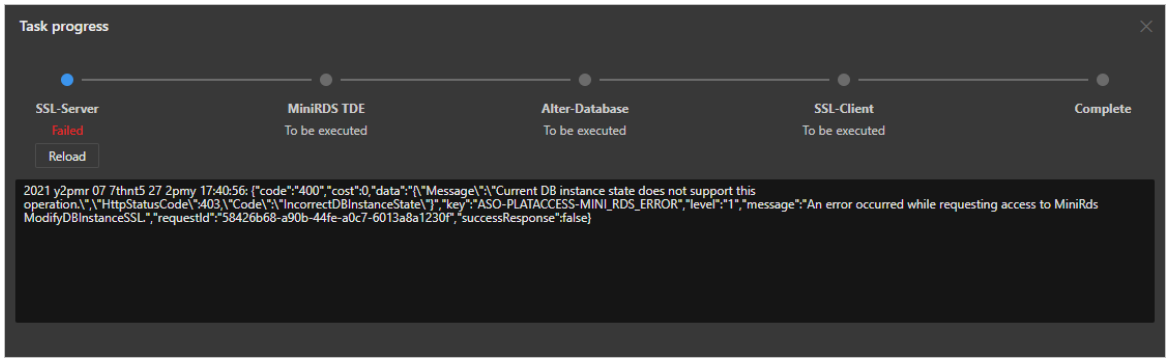

View certificates

After the Data Encryption Service is enabled, certificates are required to connect to the database. You can view application certificates.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.
4. Find the application and click **Certificate** next to it. A dialog box appears, displaying details about the certificate.

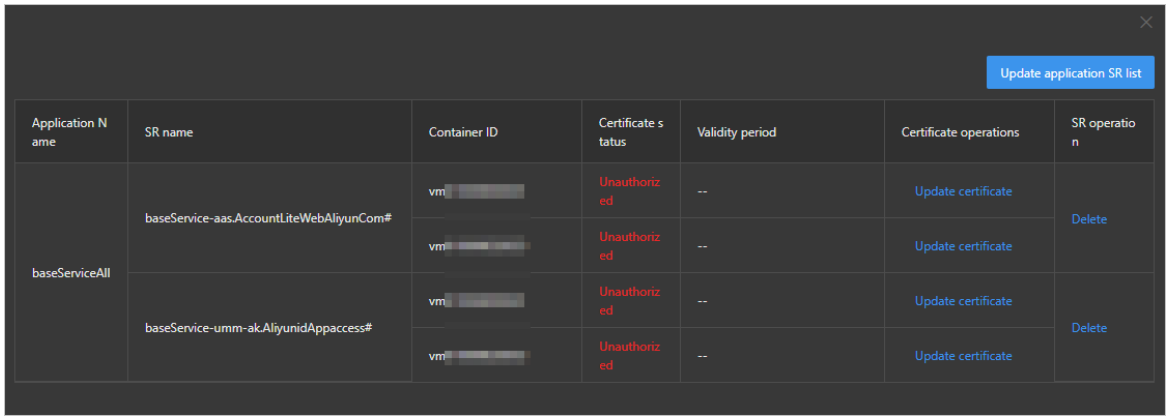

> **Note** The certificate can be in one of the following states: **Normal**, **About to expire**, **Expired**, and **Unauthorized**.

Update a certificate

You can update application certificates.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.
4. Find the application and click **Certificate** next to it. A dialog box appears, displaying details about the certificate.
5. In the dialog box that appears, find the docker and click **Update Certificate** in the **Certificate operations** column. After the certificate is updated, the message **Updated** is displayed.

Add a server role

You can add server roles for product encryption.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.
4. Find the application and click **Certificate** next to it. A dialog box appears displaying details about the certificate.
5. In the upper-right corner of the dialog box, click **Update application SR list**. In the dialog box that appears, enter the SR names. Separate multiple SR names with commas (,). Click **OK**.

> **Note** After the operation is successful, the new server role is displayed in the dialog box.

Delete a server role

You can delete server roles for product encryption.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   The **SM Settings** tab appears.
4. Find the application and click **Certificate** next to it. A dialog box appears displaying details about the certificate.
5. In the dialog box that appears, Find the SR role and then click **Delete** in the **SR operation** column. In the dialog box that appears, click **OK**.

1.1.6.8.1.2. Metadata settings

The metadata encryption feature uses the metadata algorithm to encrypt MiniRDS. The feature only performs disk storage encryption for data in tables. You can only enable metadata encryption.

Enable metadata encryption

After you enable Metadata encryption, the metadata algorithm can be used to encrypt MiniRDS.

Prerequisites

The aso-mgr and aso-opr services have reached the desired state.

Context

After the metadata encryption feature is enabled, the service may not reach the desired state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.

   > **Warning** When TDE is enabled, the Apsara Uni-manager Management Console and Apsara Uni-manager Operations Console cannot be accessed, the AK, SA, and RAM services become unavailable, and the control services are interrupted. Proceed with caution.
4. Click the **Metadata Settings** tab. Find the application and turn on the switch in the **Actions** column. In the dialog box that appears, click **OK**.
5. In the **Task Progress** dialog box, view the task progress. If a step fails, an error message is displayed. The progress is updated in real time.

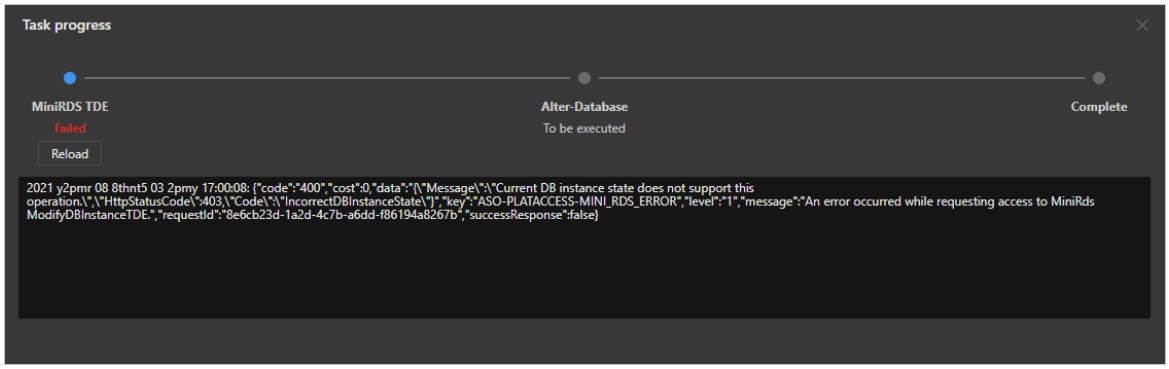

6. The **Reload** button appears below the step that fails. Click the button to execute the step again.

View execution history

You can view the execution history of product encryption operations.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Platform Encryption > Disk and Transmission Encryption**.
4. Click the **Metadata Settings** tab.
5. Find the application.
6. Click **View** in the **Execution History** column.
7. In the panel that appears, view the execution history.

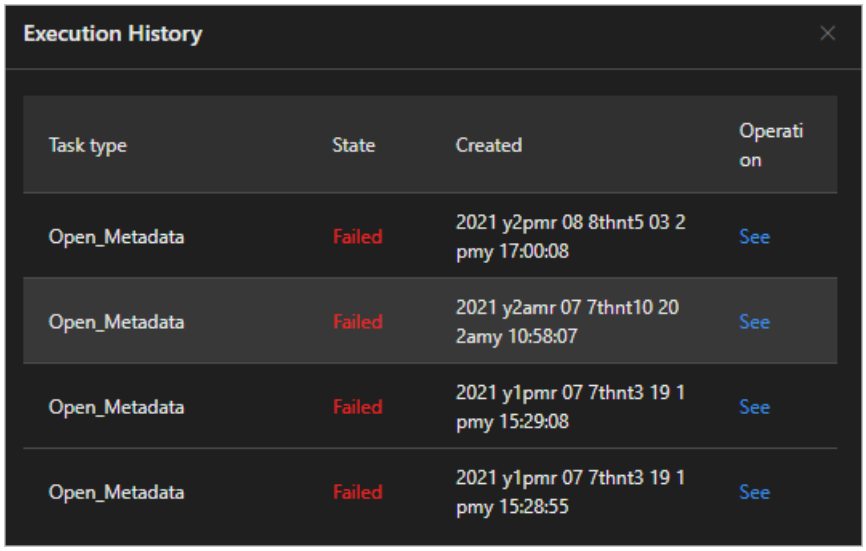

8. Find the task and click **View** in the **Actions** column. In the dialog box that appears, view the task execution details.

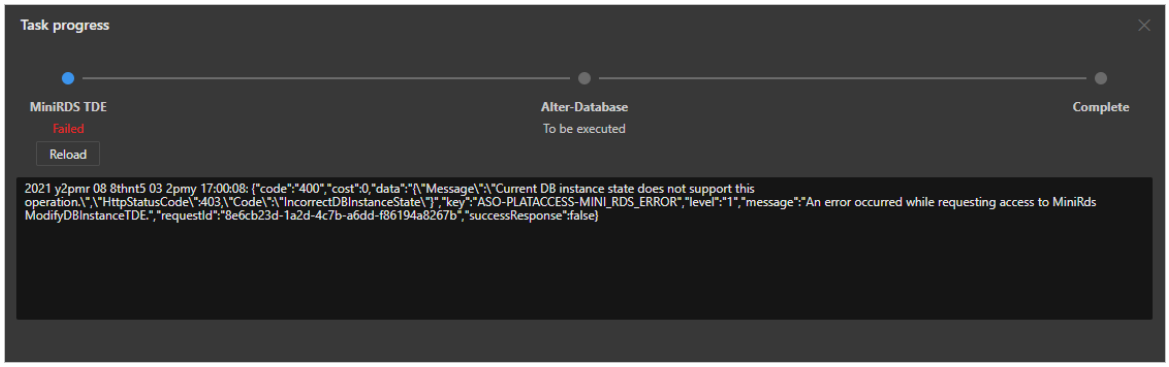

1.1.6.9. Apsara Infrastructure Management Framework O&M

1.1.6.9.1. Old console

1.1.6.9.1.1. Apsara Infrastructure Management Framework

What is Apsara Infrastructure Management Framework?

This topic describes what Apsara Infrastructure Management Framework is from the aspects of core functions and basic concepts.

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

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.

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

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

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

Procedure

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

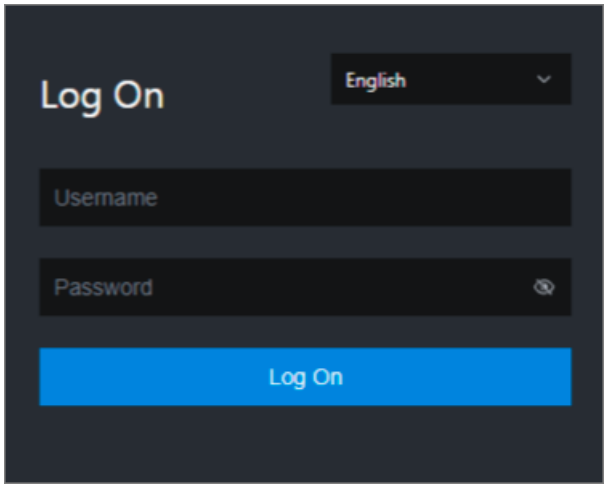

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.
4. Click **Log On**.
5. In the top navigation bar, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Apsara Stack O&M** section, choose **Basic O&M > Apsara Infrastructure Management Framework**.

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.

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

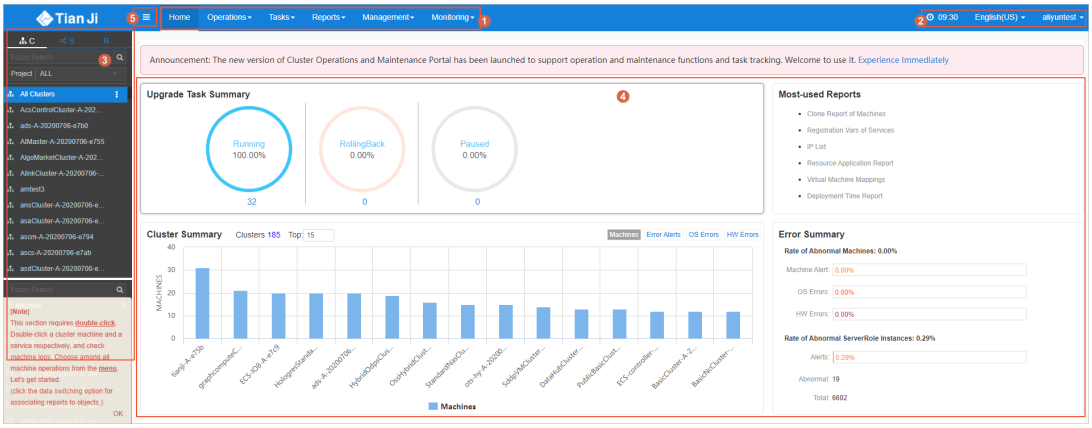

Description of functional sections describes the functional sections on the homepage.

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.
		• 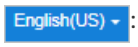 : the current display language of the console. You can select another language from the drop-down list.
		•  : your logon account. You can select **Logout** from the drop-down list to log out of your account.

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.

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.

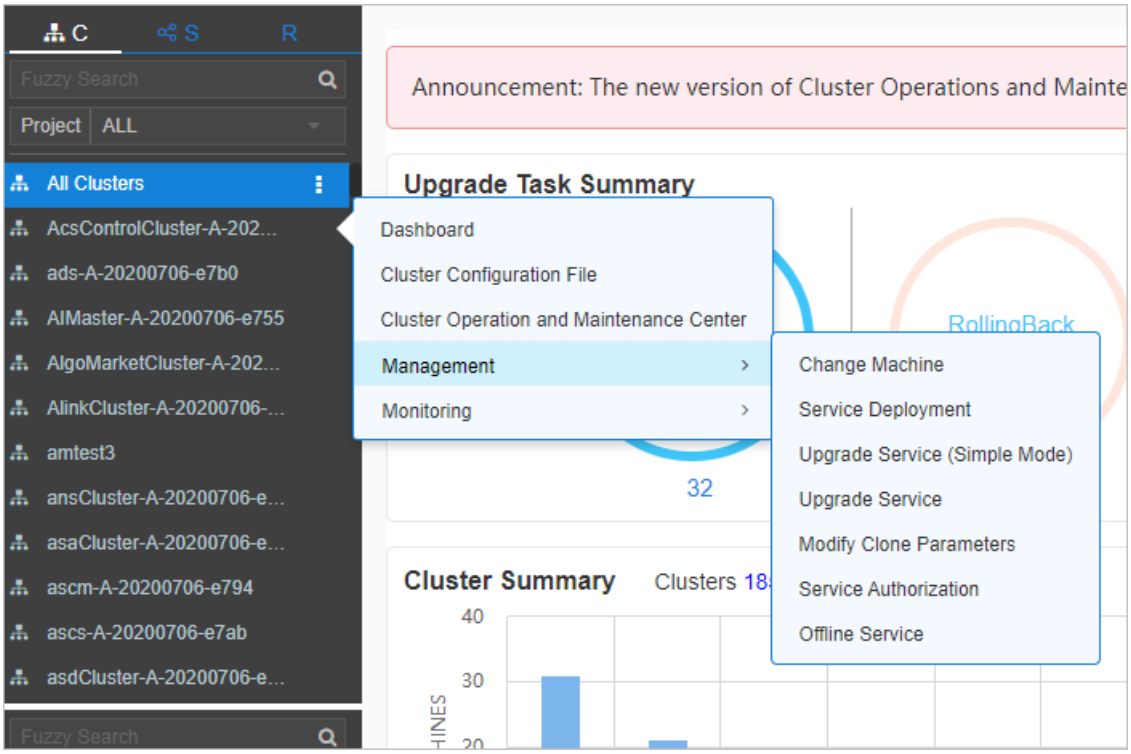

- Click a cluster. All machines and services within the cluster are displayed in the lower part of the left-side navigation pane. Move the pointer over the icon next to a machine or service on the **Machine** or **Service** tab and select menu items to perform corresponding operations on the machine or service.

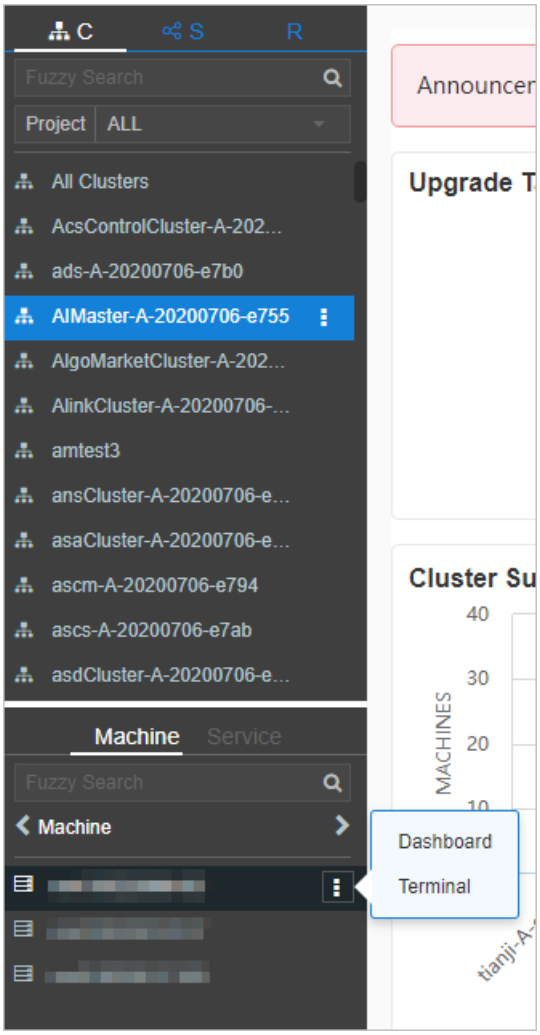

- Click the **Machine** tab. Double-click a machine to view information about all server roles on the machine. Double-click a server role to view applications, and then double-click an application to view log files.
- Click the **Service** tab. Double-click a machine to view information about all server 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.

Cluster operations

This topic describes the actions about cluster operations.

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.
	Monitoring Template	The monitoring template that is used by the service.
	Machine Mappings	The machines where server roles of the service are deployed.

Section	Parameter	Description
Service Information	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.

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.

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.

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.

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.

Service operations

This topic describes the actions about service operations.

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.
Monitoring Templates	The number of monitoring templates.
Service Schemas	The number of service configuration validation templates.

Item	Description
Actions	Click **Management** to view the service instances, service templates, monitoring templates, monitoring instances, service schemas, and detection scripts.

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
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.

Parameter	Description
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.
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.

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.

Item	Description
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.
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.

Item	Description
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.
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.

Machine operations

This topic describes the actions about machine operations.

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.

Item	Description
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.
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.

Monitoring center

You can view the alert status, alert rules, and alert history in the monitoring center.

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.

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.

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.

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.

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.

Tasks and deployment summary

This topic describes how to view rolling tasks, running tasks, history tasks, and deployment summary on Apsara Infrastructure Management Framework.

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.

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.

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.

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.

Reports

The system allows you to search for and view reports based on your business needs, and add commonly used reports to your favorites.

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.
Published At	The published time and created time of the report.

Item	Description
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.

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**.

Appendix

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.

Item	Description
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.

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.

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.

Rolling info report

This report displays the running and completed rolling tasks and the task-related statuses.

Choose a rolling action

This section displays the rolling tasks that are running. If no rolling tasks are running, no data is displayed in this section.

Item	Description
Cluster	The name of the cluster.
Git Version	The version of the 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

When you select a rolling task in the **Choose a rolling action** section, this section displays the rolling statuses of server roles related to the selected task. If no rolling tasks are selected, the statuses of server roles related to all historical rolling tasks are displayed.

Item	Description
Server Role	The name of the server role.
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 for which the rolling task was approved by the decider.
Failure Rate	The proportion of machines on which the rolling task failed.
Success Rate	The proportion of machines on which the rolling task succeeded.

Server Role Rolling Build Information

This section displays the current and desired versions of each application in the server role during 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 of the application before the upgrade.

Item	Description
To Build	The version of the application after the upgrade.

Server Role Statuses on Machines

When you select a server role in the **Server Role in Job** section, this section displays the status of the server role on each machine.

Item	Description
Machine Name	The name of the machine on which the server role is deployed.
Expected Version	The desired version of the server role.
Actual Version	The current version of the server role.
State	The status of the server role.
Action Name	The ongoing action of the server role.
Action Status	The status of the action.

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.

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.

Item	Description
Update Time	The updated time.

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.

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.

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.

Item	Description
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.
APP	The application of the server role.
Name	The resource name.
Type	The resource type.
Status	The resource application status.

Item	Description
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.

Statuses of project components

This report displays the statuses of all abnormal server roles on machines within the current project. This report also displays the alert information of server roles and machines reported to Monitoring System.

Error State Component Table

This section displays the server roles that are not in the GOOD state or that are pending upgrade.

Item	Description
Project	The name of the project.
Cluster	The name of the cluster.
Service	The name of the service.
Server Role	The name of the server role.
Machine Name	The name of the machine.
Need Upgrade	Specifies whether the version has reached the desired state.
Server Role Status	The status of the server role.
Machine Status	The status of the machine.

Server Role Alert Information

When you select a row in the **Error State Component Table** section, this section displays the corresponding information.

Item	Description
Cluster	The name of the cluster.
Service	The name of the service.
Server Role	The name of the server role.
Machine Name	The name of the machine.
Monitored Item	The name of the server role metric.
Level	The severity level of the alert.
Description	The description of the alert.
Updated At	The update time of the alert.

Machine Alert Information

When you select a row in the **Error State Component Table** section, this section displays the corresponding information.

Item	Description
Cluster	The name of the cluster.
Machine Name	The name of the machine.
Monitored Item	The name of the server role metric.
Level	The severity level of the alert.
Description	The description of the alert.
Updated At	The update time of the alert.

Service Inspector Information

When you select a row in the **Error State Component Table** section, this section displays the corresponding information.

Item	Description
Cluster	The name of the cluster.
Service	The name of the service.
Server Role	The name of the server role.
Monitored Item	The name of the server role metric.
Level	The severity level of the alert.

Item	Description
Description	The description of the alert.
Updated At	The update time of the alert.

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.

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.

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.
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.

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.

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.
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.

Item	Description
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.

1.1.6.9.2. New console

1.1.6.9.2.1. Apsara Infrastructure Management Framework 2.0

Introduction to Apsara Infrastructure Management Framework

This topic describes the features and terms of Apsara Infrastructure Management Framework.

What is Apsara Infrastructure Management Framework?

Apsara Infrastructure Management Framework is a distributed data center management system. It can manage applications within clusters that contain multiple machines and provide basic features such as deployment, upgrade, scale-in, scale-out, and configuration change.

Apsara Infrastructure Management Framework also provides data monitoring and report analysis features to facilitate end-to-end operations and maintenance (O&M) and management. In large-scale distributed scenarios, Apsara Infrastructure Management Framework offers automatic O&M to improve O&M efficiency and system availability.

Apsara Infrastructure Management Framework is composed of TianjiMaster and TianjiClient. TianjiClient is installed as an agent on a machine. TianjiMaster delivers the received commands to TianjiClient. Apsara Infrastructure Management Framework uses components to implement different features and provides users with the APIServer and console.

Features

This topic describes the core features of Apsara Infrastructure Management Framework.

Apsara Infrastructure Management Framework provides the following core features:

- Initializes networks within a data center.
- Manages server installation and maintenance processes.
- Deploys, scales, and upgrades cloud services.
- Manages cloud service configurations.
- Applies for cloud service resources.
- Repairs software and hardware faults.
- Monitors software and hardware infrastructure and business processes.

Terms

This topic describes the basic terms related to Apsara Infrastructure Management Framework.

project

A group of clusters. A project provides services for users.

cluster

A group of physical machines. A cluster provides services logically and is used to deploy software of a project.

A cluster can only belong to a single project. Multiple services can be deployed within a cluster.

service

A group of software programs used to provide an independent set of features. A service is composed of one or more server roles. A service can be deployed within multiple clusters to provide service capabilities. For example, pangu, fuxi, and nuwa are all services.

service instance

A service that is deployed within a cluster.

server role

One or more indivisible feature units of a service. A server role is composed of one or more applications. If a service is deployed within a cluster, all server roles of the service must be deployed on machines within the same cluster. Multiple server roles, such as PanguMaster and TianjiClient, can be deployed on the same machine.

service role instance

A service role that is deployed on a machine. A service role can be deployed on multiple machines.

application

A process component contained in a server role. Each application works independently. Applications are the minimum units that can be deployed and upgraded in Apsara Infrastructure Management Framework, and can be deployed on each machine. Typically, an application is an executable software program or Docker container.

If a server role is deployed on a machine, all applications in the server role must be deployed on the machine.

Rolling

A process in which Apsara Infrastructure Management Framework upgrades services and modifies cluster configurations based on the configurations updated by users.

service configuration template

A template that contains the same service configurations. A service configuration template can make it easy to write the same configurations to different clusters, and applies to large-scale deployment and upgrade scenarios.

associated service template

A file named template.conf in service configurations. The file declares that a specific version of a service configuration template is used by a service instance.

service deployment

An action that deploys a service from scratch within a cluster.

desired state

A state in which all hardware and software on each machine of a cluster work normally and all software programs are in the desired versions.

dependency

A dependency relationship between server roles in a service. Tasks are executed or configurations are upgraded based on the dependency relationship. For example, assume that A depends on B. In this case, A is downloaded after B is downloaded and upgraded after B is upgraded. By default, the dependency of configuration upgrade does not take effect.

upgrade

A way to change the current state of a service to the desired state. After a user submits a version change request, Apsara Infrastructure Management Framework can upgrade the service version to the desired version. An upgrade is performed on each server role, and aims to upgrade all machines to the desired version.

Before an upgrade starts, the current and desired states of a cluster are the same. When a user submits a version change request, the current state remains unchanged, but the desired state changes. A rolling task is generated to gradually approximate the current state to the desired state. When the upgrade ends, the current state is exactly the same as the desired state.

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 endpoint of the Apsara Uni-manager Operations Console and the username and password used to log on to the console are obtained from the deployment personnel or an administrator.
- The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id.ops.console.intranet-domain-id*.
- A browser is available. We recommend that you use Google Chrome.

Procedure

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

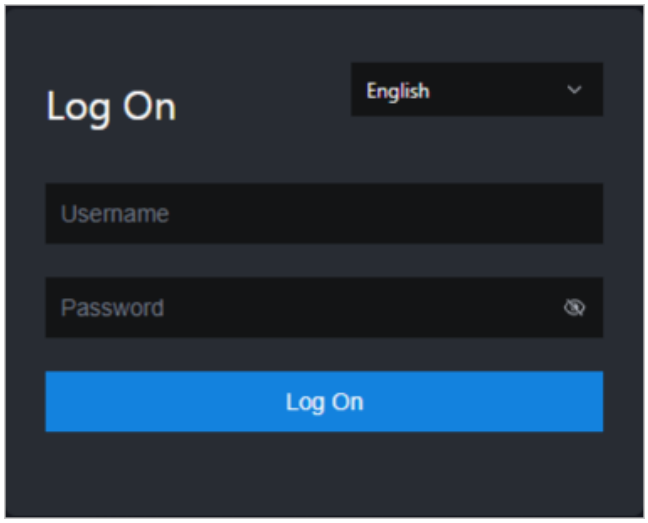

> **Note**
>
> You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

When you log on to the Apsara Uni-manager Operations Console for the first time, you must change the password of your username. For security reasons, your password must meet the following requirements:

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains the following special characters: ! @ # $ %
- The password is 10 to 20 characters in length.

4. Click **Log On**.
5. In the top navigation bar, click **O&M**.

6. In the left-side navigation pane, choose **Product Management** > **Products**.
7. In the **Apsara Stack O&M** section, choose **Basic O&M** > **Apsara Infrastructure Management Framework**.

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 features available on the homepage.

Log on to the Apsara Infrastructure Management Framework console. The homepage appears, as shown in the following figure.

Homepage of the Apsara Infrastructure Management Framework console

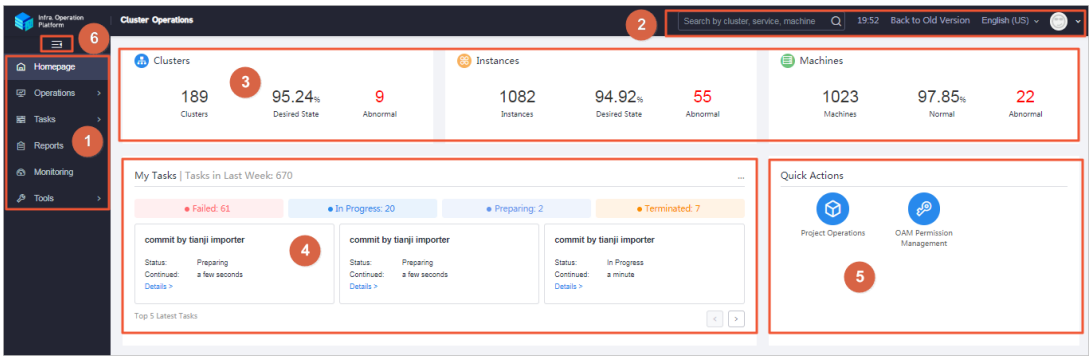

The following table describes the functional sections on the homepage.

Description of functional sections

No.	Section	Description

No.	Section	Description

①	Left-side navigation pane	• **Operations**: the quick entrance to operations & maintenance (O&M) operations, which allows you to find operations and their objects. This menu consists of the following submenus:
		  ◦ **Project Operations**: allows you to manage projects based on your project permissions.
		  ◦ **Cluster Operations**: allows you to perform O&M and management operations on clusters based on your project permissions. For example, you can view the status of clusters.
		  ◦ **Service Operations**: allows you to manage services based on your service permissions. For example, you can view the service list.
		  ◦ **Machine Operations**: allows you to perform O&M and management operations on all machines. For example, you can view the status of machines.
		• **Tasks**: Rolling tasks are generated after you modify 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 situations. This menu allows you to view the alert status, modify alert rules, and search alert history.
		• **Tools**: provides tools such as machine O&M, IDC shutdown, and clone progress.

②	Top navigation bar	• Search box: supports global search. You can enter a keyword in the search box to search for clusters, services, and machines.
		• The following information is displayed when you move the pointer over the time:
		  ◦ **TJDB Sync Time**: the time when the data on the current page is generated.
		  ◦ **Desired State Calc 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 is faulty.
		• **Back to Old Version**: allows you to return to the old version of the Apsara Infrastructure Management Framework console.
		• **English (US)**: the current display language of the console. You can select another language from the drop-down list.
		• Profile picture: allows you to select **Exit** from the drop-down list to log off your account.

No.	Section	Description
③	Status bar of global resources	Displays the overview of global resources. • **Clusters**: displays the total number of clusters, the percentage of clusters that have reached the desired state, and the number of abnormal clusters. • **Service Instances**: displays the total number of instances, the percentage of instances that have reached the desired state, and the number of abnormal instances. • **Machines**: displays the total number of machines, the percentage of normal machines, and the number of abnormal machines. You can move the pointer over each section and then click **Details** to go to the Cluster Operations, Service Operations, or Machine Operations page.
④	Task status bar	Displays the information of tasks submitted within the last week. You can click the number next to a task state to go to the My Tasks page and view the task details. The top 5 latest tasks are displayed in the lower part of the section. You can click **Details** corresponding to each task to view the task details.
⑤	Quick Actions section	Displays links of the following common quick actions: • **Project Operations**: allows you to go to the Project Operations page. • **OAM Permission Management**: allows you to go to the Operation Administrator Manager (OAM) console. OAM is a centralized permission management platform in the Apsara Uni-manager Operations Console.
⑥	Show/hide button	Allows you to expand or collapse the left-side navigation pane to narrow or enlarge the workspace.

Operations

Project operations

This topic describes how to query a project and view its details.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Project Operations**.

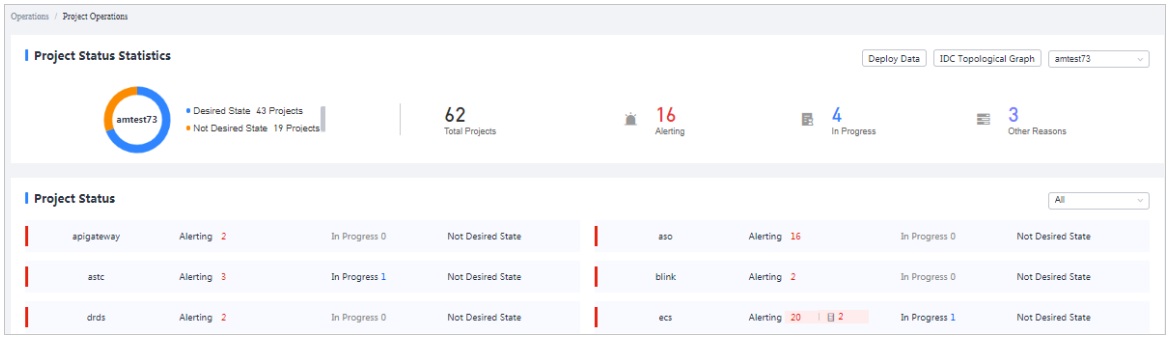

3. On the Project Operations page, perform the following operations:
   - Query a project

     In the upper-right corner of the **Project Status** section, enter the name of a project in the search box to search for the project. The search results include the number of alerts, the number of tasks in progress, and whether the project reaches the desired state.
   - View project details
     - Click the number next to **Alerting** corresponding to a project. In the Alert Information dialog box, view the metric name, metric type, and alert source. Click the alert source to view service details.
     - Click the number next to **In Progress** corresponding to a project. In the Tasks dialog box, view details about service upgrade and machine change.

Cluster operations

This topic describes the actions about cluster operations.

View the cluster list

This topic describes how to view all clusters and their information.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Use one of the following methods to go to the cluster list:
   - On the **Home** page, move the pointer over the **Cluster** section and click Details in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Cluster Operations**.

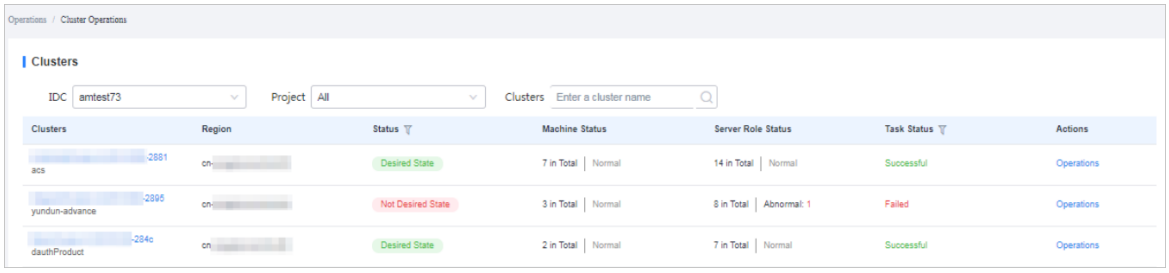

The following table describes the information displayed in the cluster list.

Parameter	Description
Cluster	The name of the cluster. Click the cluster name to view the cluster details.
Region	The region where the cluster is deployed.
Status	Specifies whether the cluster reaches the desired state. Click the icon to filter clusters. ○ Desired State: The cluster has reached the desired state. ○ Not Desired State: The cluster has reached the desired state for the first time but a server role has not reached the desired state due to undefined reasons.

Parameter	Description
Machine Status	The number of machines within the cluster and the machine status. Click the machine status to go to the Machines tab of the Cluster Details page.
Server Role Status	The number of server roles within the cluster and the server role status. Click a server role 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. 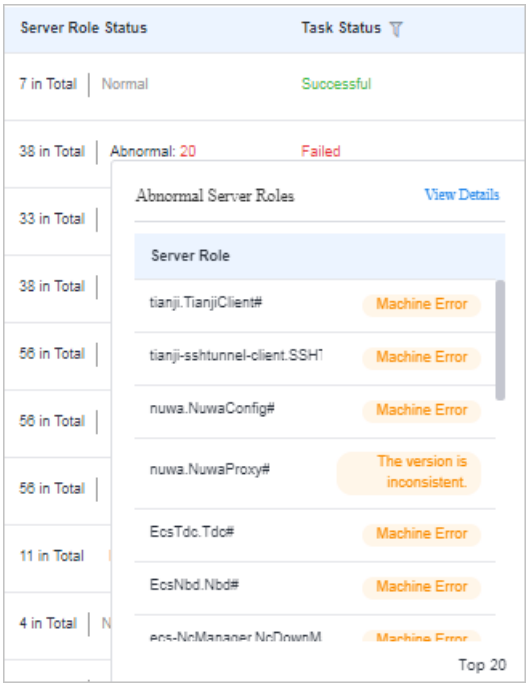
Task Status	The status of the task related to the cluster. Click the icon to filter clusters. Click the task status to view the task details.
Actions	The available operations. Click **Operations** to go to the **Cluster Details** page.

View details of a cluster

This topic describes how to view details of a cluster.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
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 **Operations** in the **Actions** column to go to the **Cluster Details** page.

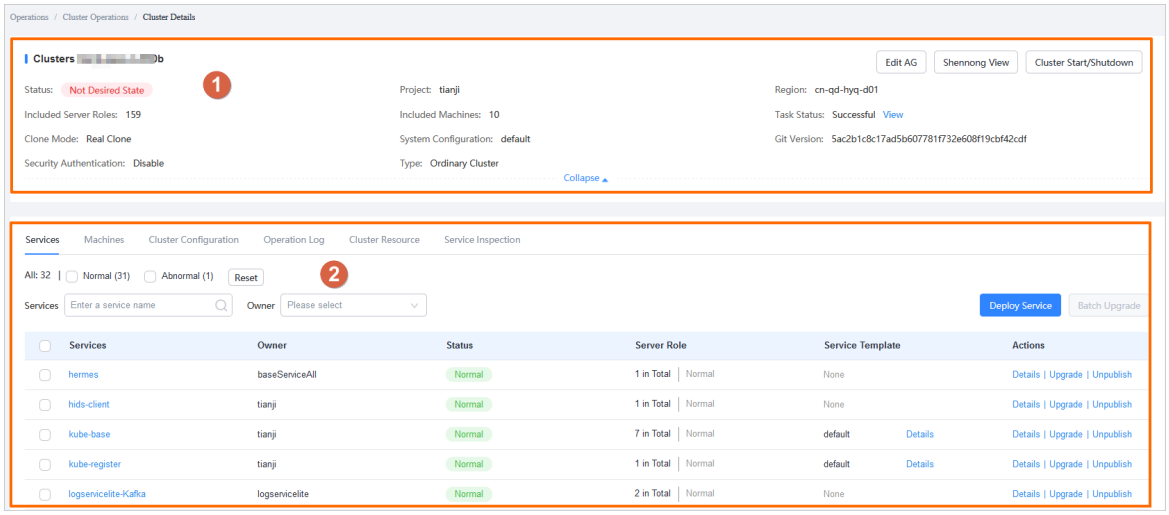

Section	Parameter	Description
①	Status	○ **Desired State**: All clusters in a project have reached the desired state. ○ **Not Desired State**: A project has reached the desired state for the first time but a server role has not reached the desired state due to undefined reasons.
	Project	The project to which the cluster belongs.
	Region	The region where the cluster is deployed.
	Included Server Roles	The number of server roles included in the cluster.
	Included Machines	The number of machines included in the cluster.
	Purpose	The purpose of the cluster. Click the icon. In the dialog box that appears, select a purpose from the drop-down list.
	Task Status	The status of the task. Click **View** to view the task details. ○ **Successful**: The task is successful. ○ **Preparing**: Data is being synchronized and the task is not started. ○ **In Progress**: The cluster has a changing task. ○ **Paused**: The task is paused. ○ **Failed**: This task failed. ○ **Terminated**: The task is manually terminated.

Section	Parameter	Description
	Clone Mode	○ **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.
	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. By default, security authentication is disabled in non-production environments. You can enable or disable security authentication to meet your business requirements.
	Type	○ **Ordinary Cluster**: an operations unit of machine groups, in which multiple services can be deployed. ○ **Virtual Cluster**: an operations unit of services, which can manage versions of software on machines within several physical clusters in a centralized manner. ○ **RDS**: a type of cluster that renders special cgroup configurations based on some rules. ○ **NETFRAME**: a type of cluster that renders special configurations for special scenarios of Server Load Balancer (SLB). ○ **T4**: a type of cluster that renders special configurations for the mixed deployment of e-commerce. Apsara Stack provides only ordinary clusters.
②	**Services**	The status of each service within the cluster. You can also upgrade or unpublish a service. ○ **Normal**: The service works normally. ○ **Not Deployed**: The service is not deployed on machines. ○ **Changing**: Some server roles in the service are changing. ○ **Operating**: No server role is changing, but a server role is performing operations and maintenance (O&M) operations. ○ **Abnormal**: No server role is changing or the machines on which server roles are deployed are not performing O&M operations. However, the server role is in the **not good** state, or the version that the service runs on the machines is different from the desired version.
	Machines	The running status and monitoring status of each machine within the cluster. You can also view details of server roles that are deployed on each machine.

Section	Parameter	Description
	Cluster Configuration	The configuration file used within the cluster.
	Operation Logs	The operation logs. You can also view the version differences.
	Cluster Resource	The details of resources that can be filtered.
	Service Inspection	The inspection information of each service within the cluster.

View configuration information of a cluster

This topic describes how to view configuration files and folders of a cluster.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Use one of the following methods to go to the **Cluster Configuration** tab to view configuration files and folders:
   - Enter a cluster name in the search box in the upper-right corner of the page. Click **Operations** next to the found cluster. On the Cluster Details page, click the **Cluster Configuration** tab.
   - In the left-side navigation pane, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find a cluster and click **Operations** in the **Actions** column. On the Cluster Details page, click the **Cluster Configuration** tab.

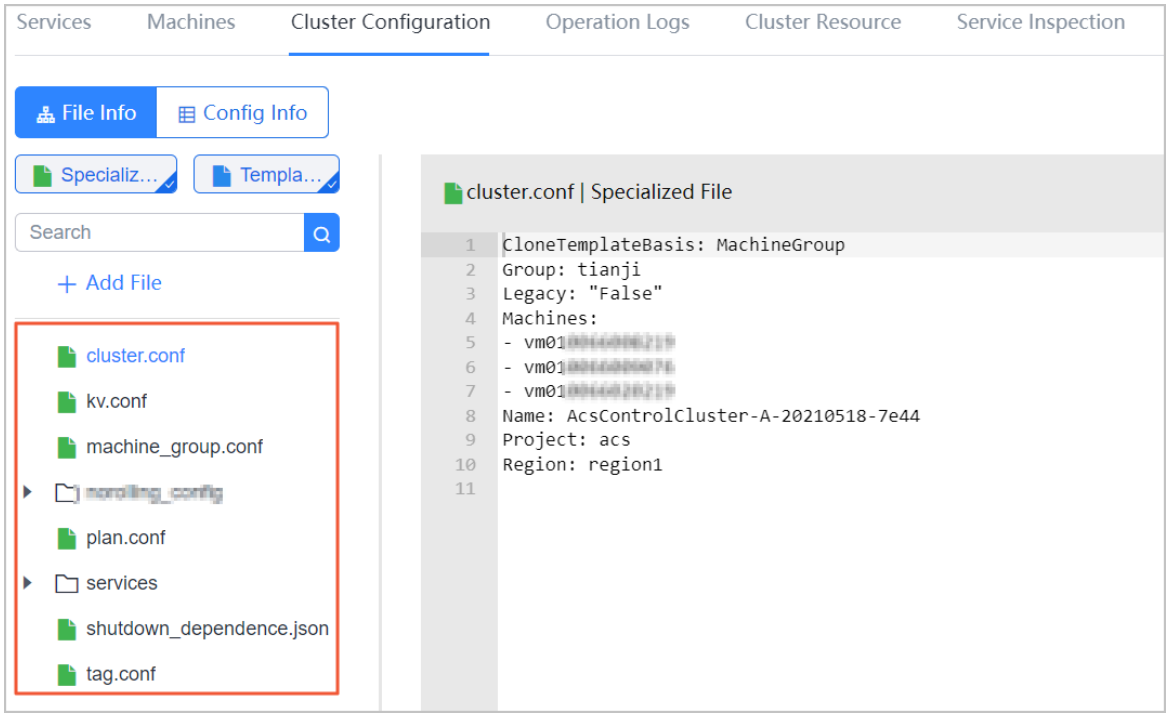

The following table describes configuration files and folders of a cluster.

Parameter	Description
cluster.conf	The configuration file of the cluster, including the cluster name, cluster type, and machines.
kv.conf	The file that stores the values used to replace template placeholders when configurations are rendered.
machine_group.conf	The file that stores information of machine groups within a cluster.
plan.conf	The file that defines dependencies between services and configuration upgrade parameters.
services	The folder where configurations of each service are stored.
shutdown_dependence.json	The shutdown dependency file.
tag.conf	The file that stores the tags used to calculate tag expressions when configurations are rendered.

3. On the **Cluster Configuration** tab, move the pointer over a folder, click the icon next to the folder name, and then select **Add File** to add a configuration file.

   Note You can also click **Add File** below the search box to add a file or folder to the directory.

   i. In the **Add File** dialog box, enter a file or folder name and click **OK**.

   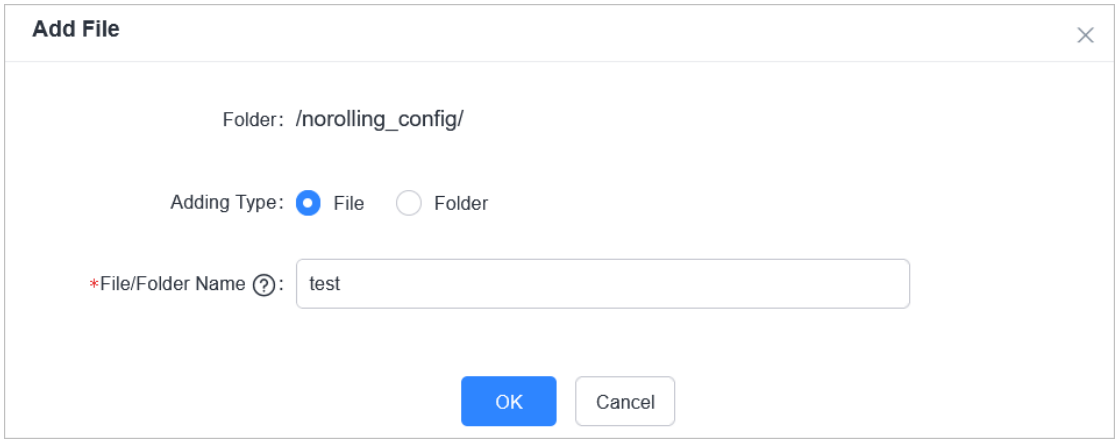

   Note After you enter a folder name and click **OK**, the folder is added.

ii. Enter configuration file information into the **Specialized File** text editor. Click **Preview and Submit**.

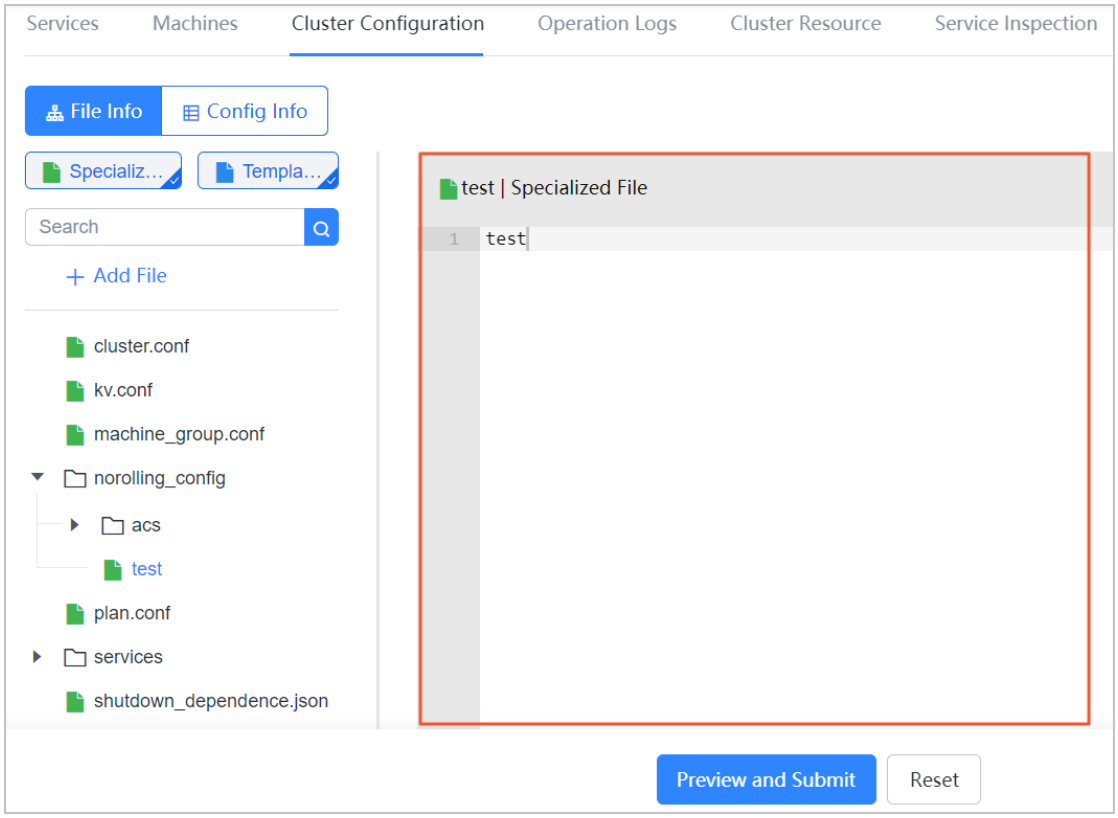

iii. In the **Confirm and Submit** dialog box, enter **Description** and click **Submit**.

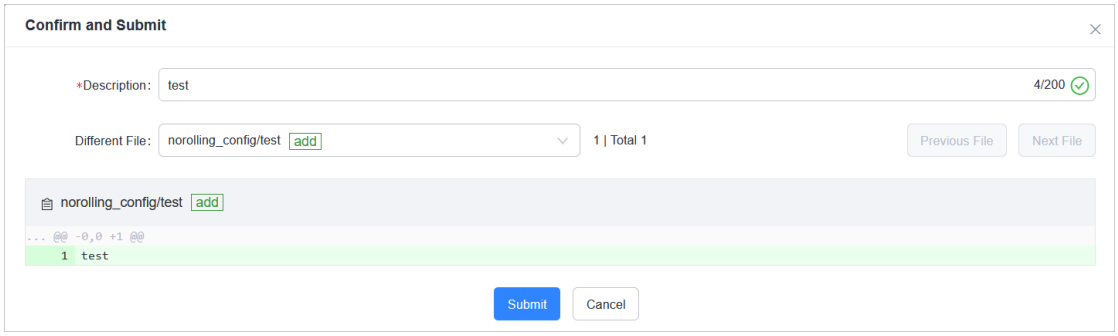

The configuration file is added. You can click the **Operation Logs** tab to view related records.

View operations logs

This topic describes how to view differences between Git versions from operation logs.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Use one of the following methods to go to the operation logs of a cluster:
   - Enter a cluster name in the search box in the upper-right corner of the page. Click **Operations** next to the found cluster. On the Cluster Details page, click the **Operation Logs** tab.
   - In the left-side navigation pane, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find a cluster and click **Operations** in the **Actions** column. On the Cluster Details page, click the **Operation Logs** tab.

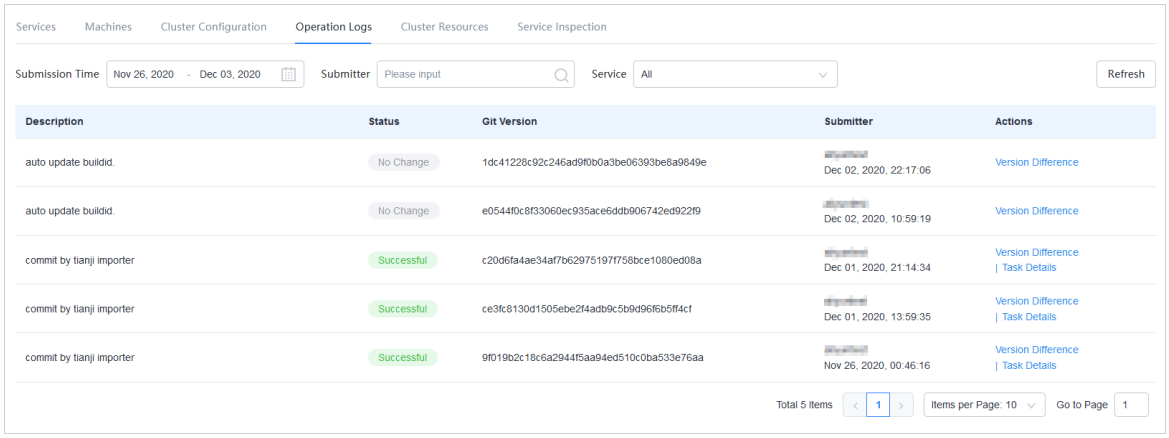

3. View the version differences on the **Operation Logs** tab.
   i. Find the operation log that you want to view and click **Version Difference** in the **Actions** column.
   ii. Set **Configuration Type** to **Show Configuration** or **Cluster Configuration**.
      - **Show Configuration**: displays the cluster configuration merged with the template configuration.
      - **Cluster Configuration**: displays the cluster configuration.
        - Cluster configuration description: Each cluster contains its dedicated configurations, such as the list of machines.
        - Template configuration description: A template that has the same configurations can be used to deploy a service to multiple clusters.
   iii. Select a basic version below **Configuration Type**. Then, a difference file is displayed in the lower part of the page.
   iv. Select a difference file from the **Difference File** drop-down list to view the content of each difference file.

Service operations

View the service list

This topic describes how to view all services and their information.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Use one of the following methods to go to the service list:
   - On the **Home** page, move the pointer over the **Service Instances** section and click Details in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Service Operations**.

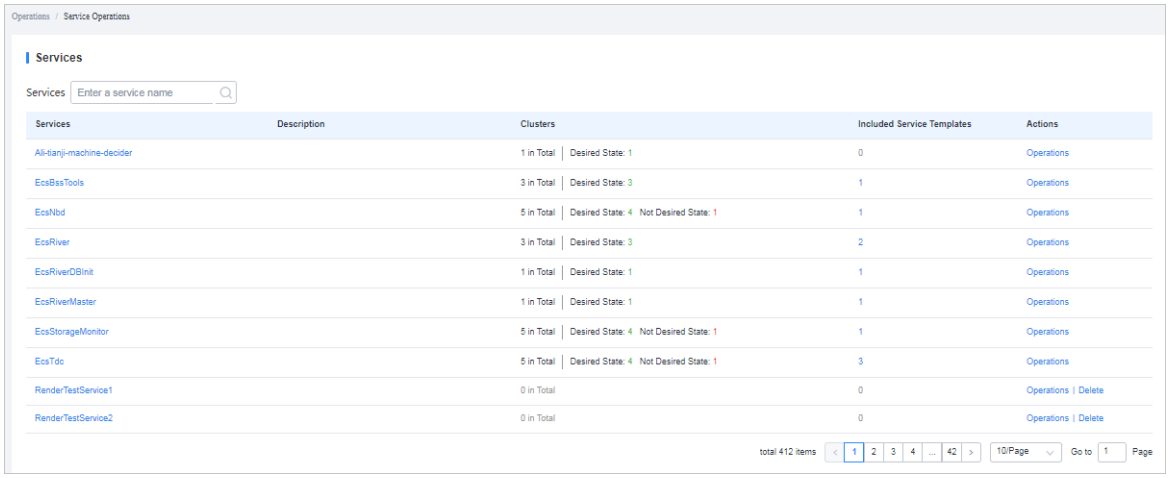

The following table describes the information displayed in the service list.

Parameter	Description
Service	The name of the service. Click the service name to view the service details.
Clusters	The number of clusters in which the service is deployed and the cluster status.
Included Service Templates	The number of service templates that are included in the service.
Actions	○ Click **Operations** to go to the Service Details page. ○ Click **Delete** to delete the service. **Note** A service can be deleted only when the number of clusters in which the service is deployed is 0.

3. (Optional)Enter a service name in the search box to search for the service.

View details of a server role

This topic describes how to view details of a server role.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
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 service.
4. Click the service name or click **Operations** in the **Actions** column.
5. On the **Clusters** tab, click a state in the **Server Role Status** column to view the server roles included in a cluster.

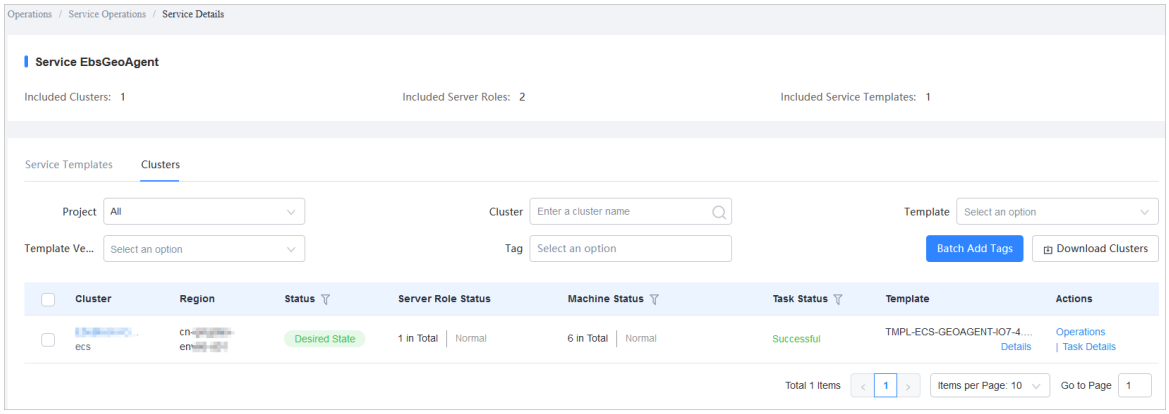

6. Select the server role that you want to view.
   - Click the **Machines** tab to view details of the server role.

Parameter	Description
Machine	The machine where the server role is deployed. Click the machine name to view the machine details.
Actions	■ Click **Metric** to view the server role, machine, and system service metrics. ■ Click **Applications** to view application versions. ■ Click **Terminal** to log on to the machine and perform operations. ■ Click **Restart** to restart the server role.

   - Click the **Upgrade History** tab. Click **Details** in the **Actions** column to view details of a historical task.

Block hardware alerts

This topic describes how to shield hardware monitoring alerts.

Background information

You must block hardware alerts in the following scenarios:

- Alerts are improperly triggered by hardware. In this case, you must block the alerts, and then cancel the block operation after no alerts are reported.
- Upgrades fail to reach the desired state due to hardware faults, and the hardware faults cannot be rectified at this time. In this case, you must block the alerts, and then cancel the block operation after the desired state is reached.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations** > **Service Operations**.
3. In the search box, enter **cruiser** to search for the cruiser service.
4. Find the cruiser service and click **Operations** in the **Actions** column.
5. Click the **Clusters** tab.
6. Click **Operations** in the **Actions** column corresponding to a cluster.

7. Click the **Cluster Configuration** tab. Open the `/user/ignore_monitor_config.json` file in **Cluster File**. Modify the configuration file.

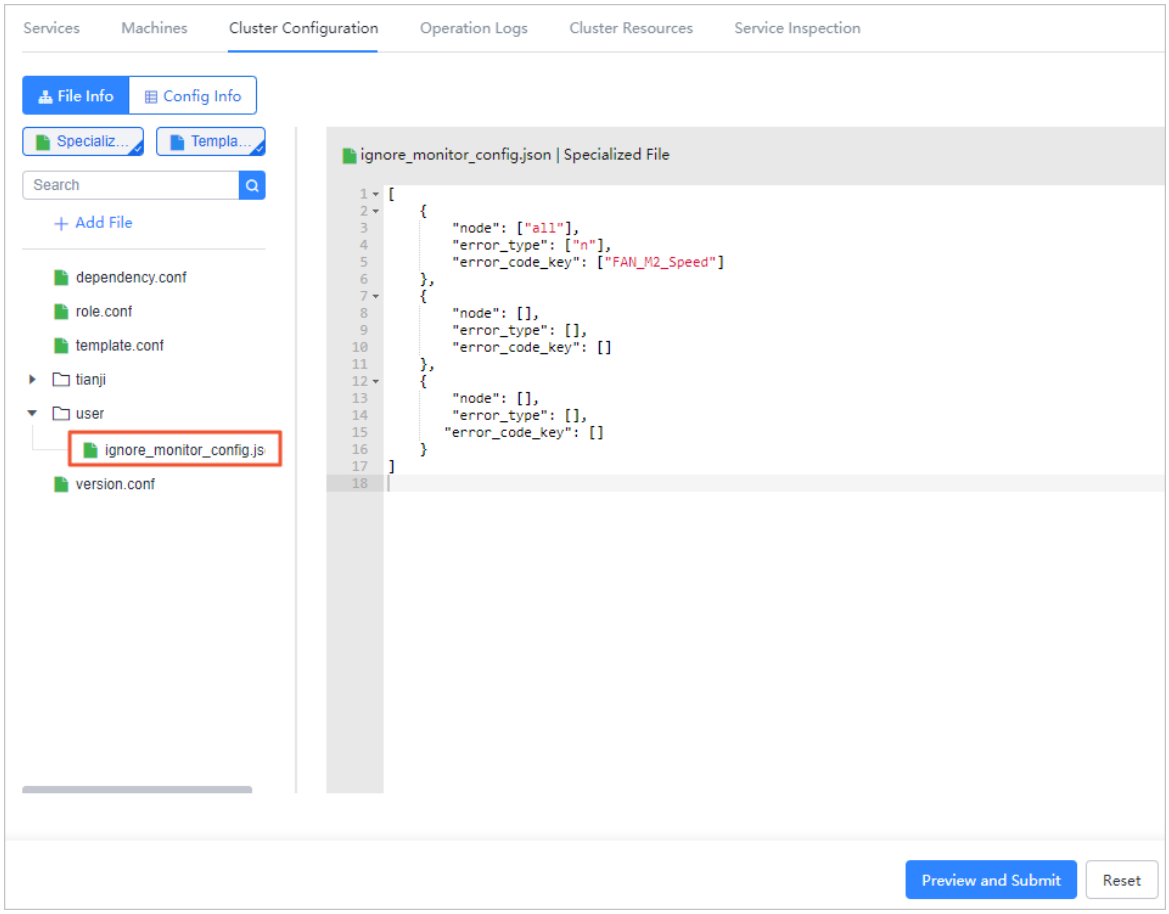

The following table describes parameters in the configuration file.

Parameter	Description	Description
node	The name of the machine where alerts are blocked. If you want to block alerts on all machines, set the parameter to `"all"`.	• All the node, error_type, and error_code_key parameters are in an array format. • The node parameter is required. • At least one of the error_type and error_code_key parameters is required.
error_type	The type of the fault that triggers alerts. Valid values: • 0: LogicDrive fault • 1: hard disk fault • 2: memory fault	

Parameter	Description	Description
error_code_key	The keyword that is used to block alerts. The keyword can be the error code or information.	

Example:

```
{ "node":["vm1243t", "sfas.hostname"], "error_type":["1"] "error_code_key":["BMC", "nic port"] }
```

In the preceding example, the alerts caused by hard disk faults are blocked on the vm1243t and sfas.hostname machines. The error information includes BMC and NIC port.

8. Click **Preview and Submit**.
9. In the **Confirm and Submit** dialog box, enter the description and click **Submit**.

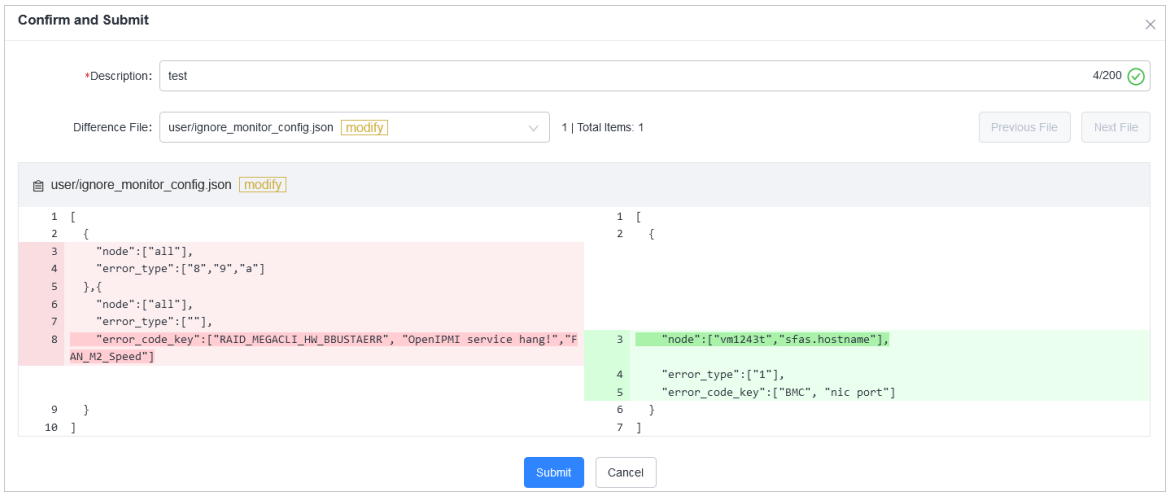

10. Click the **Operation Logs** tab to view related records.

Machine operations

This topic describes how to view the statistics of all machines.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Use one of the following methods to go to the machine list:
   - On the **Home** page, move the pointer over the **Machine** section and click Details in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Machine Operations**.
3. (Optional)Select a project from the drop-down list or enter a cluster or machine name to search for the machine.

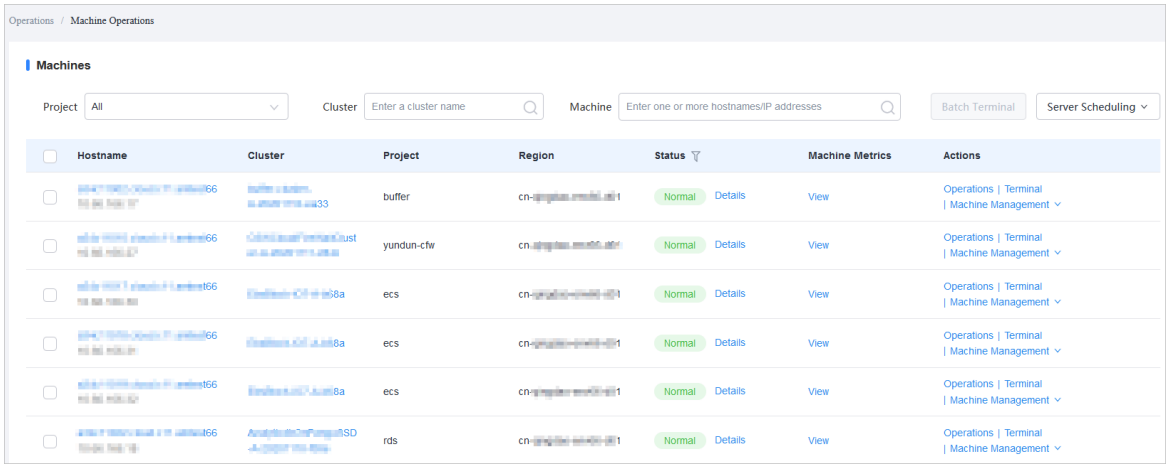

Parameter	Description
Hostname	The hostname of the machine. Click a hostname to go to the Machine Details page.
Cluster	The cluster where the machine is deployed. Click a cluster name to go to the Cluster Details page.
Status	The status of the machine. Click the icon to filter machines. Click **Details**. Then, the **Status Details of Machine** dialog box appears.
Machine Metrics	The metrics of the machine. Click **View**. Then, the **Metrics** dialog box appears. Metrics are displayed on the **Server Role Metric**, **Machine Metrics**, and **System Service Monitor** tabs. You can view the status and update time of each metric. Enter a keyword in one of the search boxes in the upper-right corner to search for a server role or metric. You can also select the status in the upper-left corner 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.

View tasks

This topic describes how to view the submitted tasks and their statuses.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Use one of the following methods to go to the task list:
   - In the left-side navigation pane, choose **Tasks > My Tasks**.

- In the left-side navigation pane, choose **Tasks > Related Tasks**.

3. (Optional) Click the icon in the **Status** column to filter tasks.
4. Find the task that you want to view and click the task name or **Details** in the **Actions** column.
5. View the status and progress of each cluster and server role on the **Task Details** page.

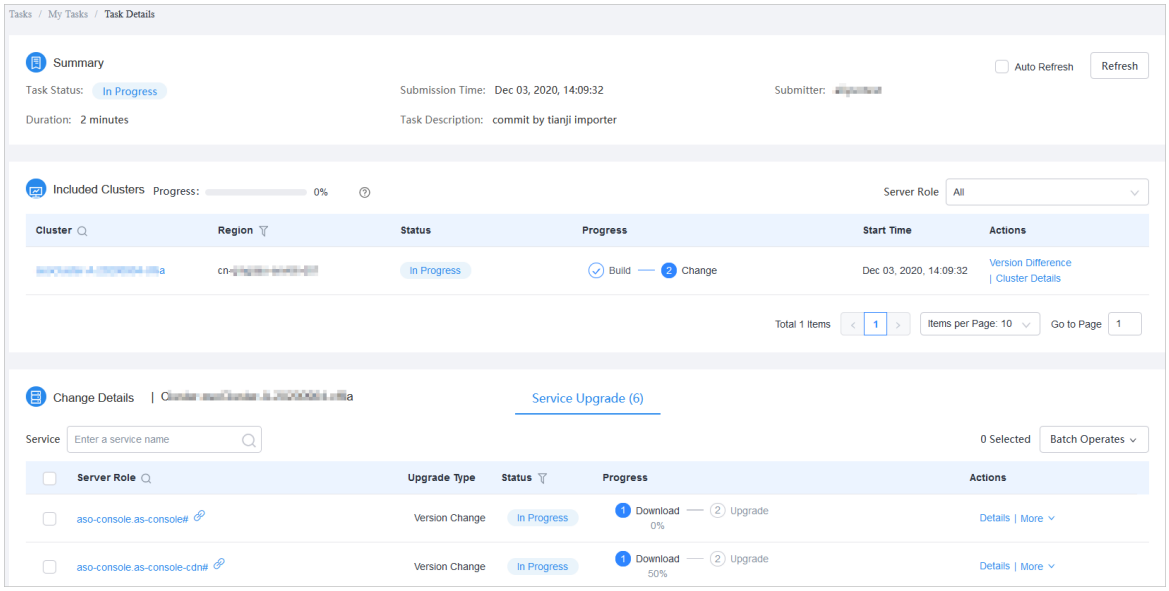

Reports

View reports

This topic describes how to view report data.

Context

The following reports are available in the Apsara Infrastructure Management Framework console:

- System reports: include default and common reports in the system.
- All reports: include system reports and custom reports.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Reports**. On the Reports page, click **Go** to go to the All Reports page.

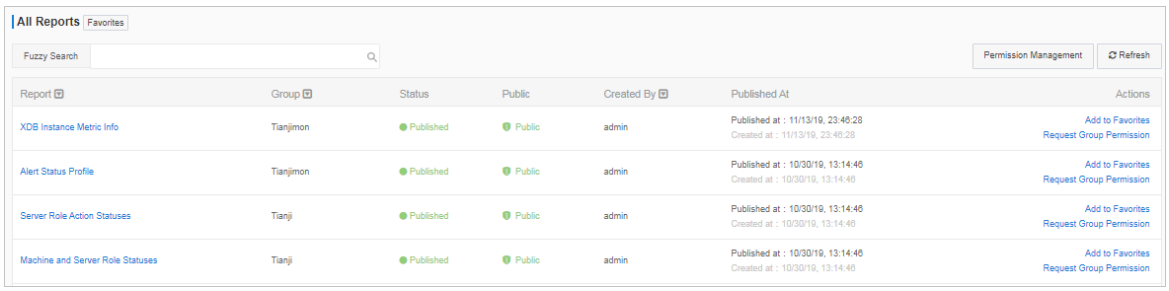

The following table describes information about reports.

Parameter	Description
Report	The name of the report. Move the pointer over the down arrow next to Report and search by report name.
Group	The group to which the report belongs. Move the pointer over the down arrow next to Group and search by group name.
Status	Specifies whether the report is published. ○ Published ○ Not Published
Public	Specifies whether the report is public. ○ Public: visible to all logon users. ○ Private: visible only to the current logon user.
Created By	The person who creates the report.
Published At	The time when the report is created and published.
Actions	○ Click **Add to Favorites** to add the report to your favorites. Then, you can view the report by choosing **Reports > Favorites** in the top navigation bar. ○ Click **Request Group Permission** to go to the Operation Administrator Manager (OAM) console. You can then configure groups and permissions. For more information, see *OAM* in *Operations and Maintenance Guide*.

3. (Optional)Enter a report name in the search box to search for the report.
4. Click the report name to go to the corresponding report details page.

   For more information about reports, see Appendix.

Add a report to favorites

This topic describes how to add frequently used reports to favorites. Then, you can find them on the Home or Favorites page.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Reports**. On the Reports page, click **Go** to go to the All Reports page.
3. (Optional)Search for a report in the search box.
4. Click **Add to Favorites** in the **Actions** column corresponding to the report.
5. In the **Add to Favorites** dialog box, enter tags for the report.
6. Click **Add to Favorites**.

Monitoring center

You can view the alert status, alert rules, and alert history in the monitoring center.

View the status of a metric

This topic describes how to view the status of a metric.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
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 service.
4. Click **Operations** in the **Actions** column corresponding to the service.
5. Click the **Clusters** tab.
6. Find the cluster that you want to view and click **Operations** in the **Actions** column.
7. On the **Services** tab, select a server role and click **Metrics** in the **Actions** column corresponding to a machine to view the server role, machine, and system service metrics.

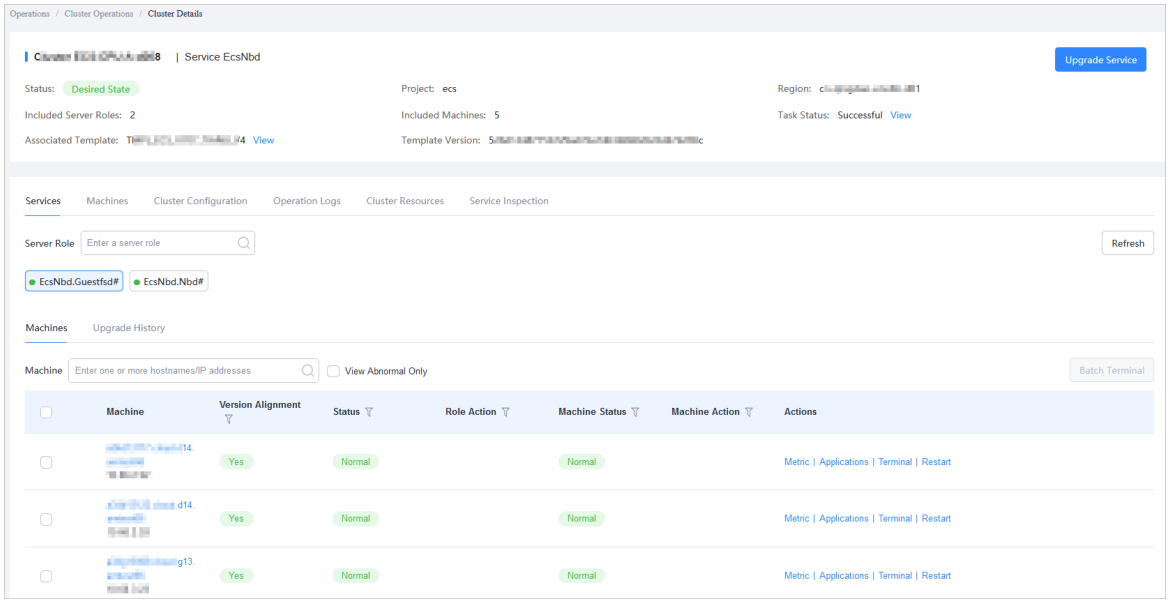

View the alert status

This topic describes how to view the alerts related to different services and the alert details.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Monitoring**. On the Monitoring page, click **Go** to go to the Alert Status page.
3. In the top navigation bar, choose **Monitoring > Alert Status**.

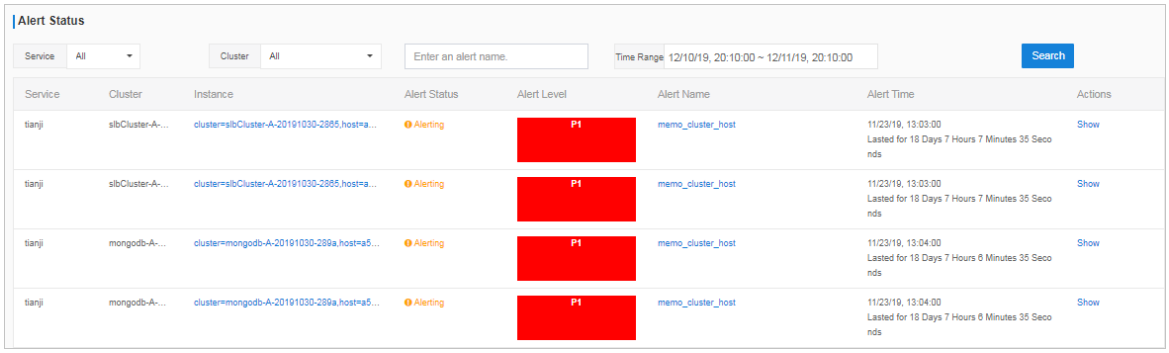

4. (Optional)Search for an alert by service name, cluster name, alert time range, or alert name.
5. View alert details on the **Alert Status** page. The following table describes the related parameters.

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	The state of the alert. Two alert states are available, which are **Normal** and **Alerting**.
Alert Level	The level of the alert. 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	The available operations. Click **Show** to view the data before and after the alert time.

View alert rules

This topic describes how to view the alert rules of a service.

Procedure

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

2. In the left-side navigation pane, click **Monitoring**. On the Monitoring page, click **Go** to go to the Alert Status page.
3. In the top navigation bar, choose **Monitoring > Alert Rules**.

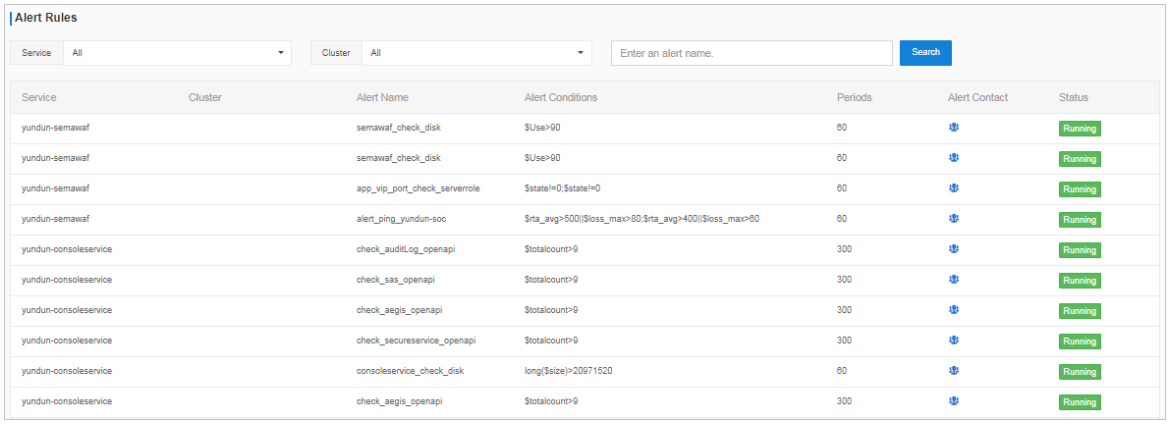

4. (Optional)Search for alert rules by service name, cluster name, or alert name.
5. View alert rules on the **Alert Rules** page. The following table describes the related parameters.

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.

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 the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, click **Monitoring**. On the Monitoring page, click **Go** to go to the Alert Status page.
3. In the top navigation bar, choose **Monitoring > Alert History**.

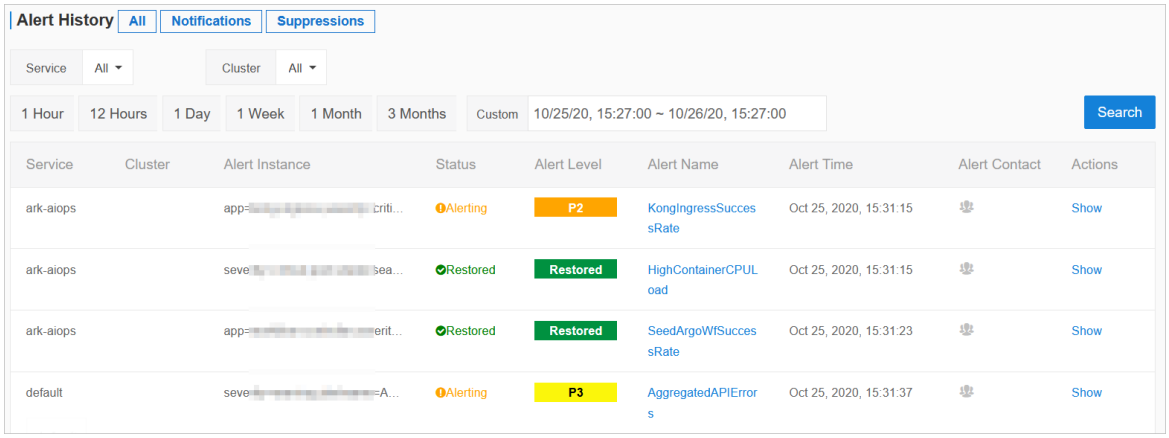

4. (Optional)Search for an alert by service name, cluster name, alert cycle, or alert time range.
5. View the alert history on the **Alert History** page. The following table describes the related parameters.

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	The state of the alert. Two alert states are available, which are **Normal** and **Alerting**.
Alert Level	The level of the alert. 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	The available operations. Click **Show** to view the data before and after the alert time.

Tools

Use machine operations tools

This topic describes how to use machine operations tools in typical scenarios.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Tools > Operation Tools > Machine Tools**. On the Machine Tools page, click **Go** to go to the Operation Tools page.
3. Select a scenario from the Operation Scene drop-down list.

Scenario	Description	Actions
Scene 1. NC Scale-out (with existing machines)	Scales out an SRG of the worker type.	Select a cluster from the Target Cluster drop-down list and an SRG from the Target SRG drop-down list. Select the machines to scale out in the left-side section, click Select> to add them to the right-side section, and then click **Submit**. In the message that appears, click **Confirm**.
Scene 2. Host Scale-out (with existing machines)	Scales out DockerHost#Buffer of a cluster.	Select a cluster from the Target Cluster drop-down list. Select the machines to scale out in the left-side section, click Select> to add them to the right-side section, and then click **Submit**. In the message that appears, click **Confirm**.
Scene 3. NC Scale-in	Scales in an SRG of the worker type.	Select a cluster from the Target Cluster drop-down list and an SRG from the Target SRG drop-down list. Select the machines to scale in in the left-side section, click Select> to add them to the right-side section, and then click **Submit**. In the message that appears, click **Confirm**.
Scene 4. Host Scale-in	Scales in DockerHost#Buffer of a cluster.	Select a cluster from the Target Cluster drop-down list. Select the machines to scale in in the left-side section, click Select> to add them to the right-side section, and then click **Submit**. In the message that appears, click **Confirm**.

Scenario	Description	Actions
Scene 5. VM Migration	Migrates virtual machines (VMs) from a host to another host.	Select a source host from the Source Host drop-down list and a destination host from the Destination Host drop-down list. Select the VMs to migrate in the left-side section, click Select> to add them to the right-side section, and then click **Submit**. In the message that appears, click **Confirm**.
Scene 6. Host Switching	Switches a standby host to the primary host.	Select a source host from the Source Host drop-down list and a destination host from the Destination Host drop-down list. Click **Submit**. In the message that appears, click **Confirm**.

Shut down a data center

This topic describes how to shut down up to 25 machines within all clusters of a data center in scenarios such as vehicle-mounted devices.

Prerequisites

- The total number of machines within all clusters of a data center cannot exceed 25.
- Your browser is connected with the machines on which Apsara Infrastructure Management Framework is deployed over a smooth network. If a proxy is required to log on to the Apsara Infrastructure Management Framework console, the proxy is not configured on a machine that you want to shut down.
- Your browser remains active while the machines are being shut down.
- Data related to operations such as scaling is not retained within the default cluster before the machines are shut down.

Context

When you shut down a data center, business clusters are shut down first, and then the base cluster is shut down.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Tools > IDC Shutdown**. In the right-side workspace, click **Go**.
3. On the **IDC Shutdown** page, click **Start Shutdown**.
4. In the **Confirm Operation** message, enter *SHUTDOWN* and click **Confirm**.

> **Warning**
>
> - The data center shutdown operation shuts down all services and machines and thus cause business interruption.
> - Backend services must communicate with the frontend shutdown page during the data center shutdown process. Do not close the shutdown page until the shutdown is complete.

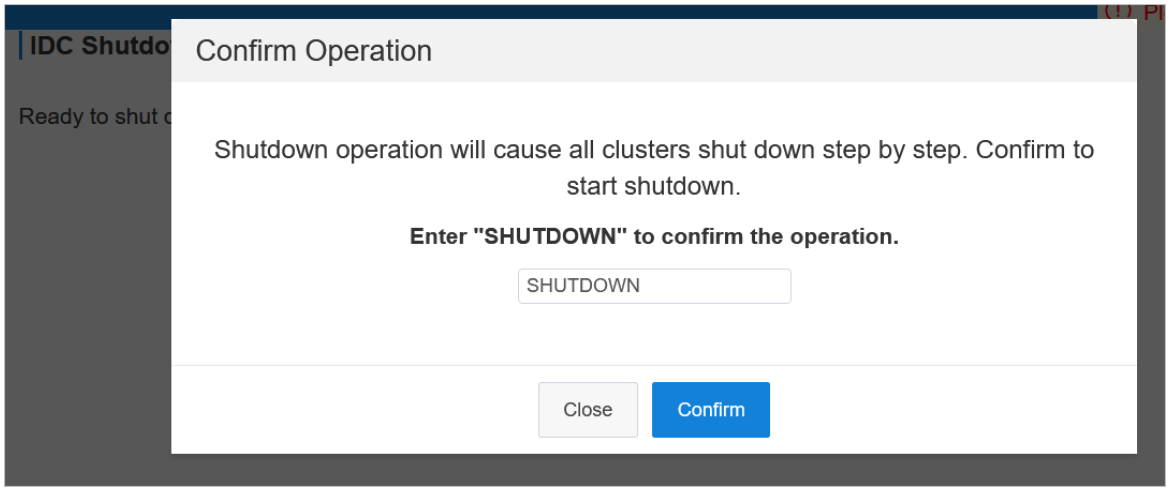

5. View the data center shutdown progress and the statuses of clusters, machines, and server roles.

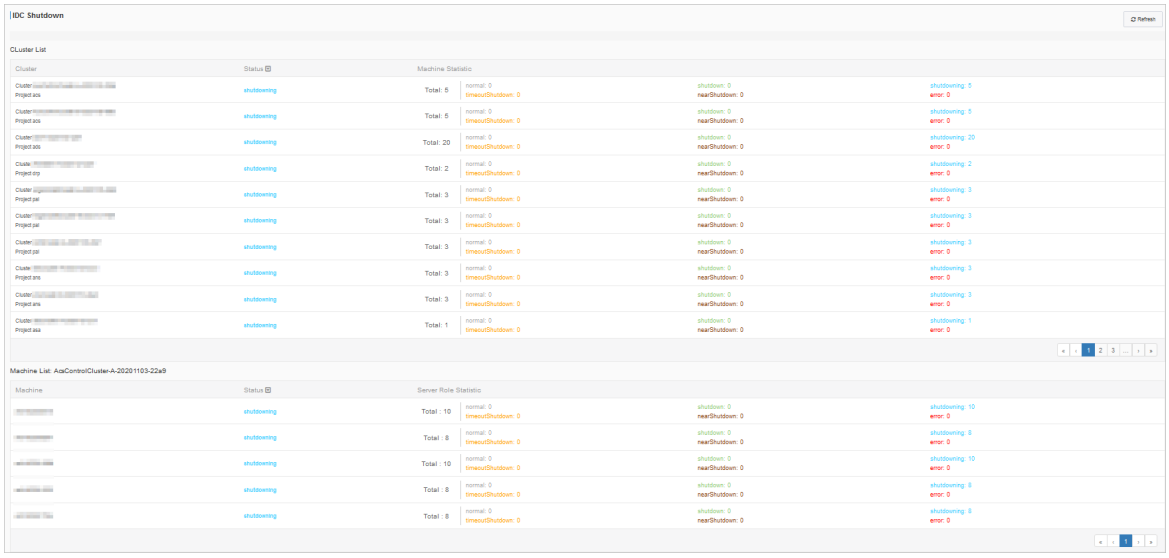

It takes a long time to shut down all clusters and machines within an environment. You can view the shutdown progress on the **IDC Shutdown** page. The following statuses are available for clusters, machines, and server roles:

- **normal**: A cluster, machine, or server role is running normally.
- **shutdown**: A cluster, machine, or server role is shut down.
- **shutdowning**: A cluster, machine, or server role is being shut down.
- **timeoutShutdown**: The shutdown of a cluster, machine, or server role timed out.
- **nearShutdown**: A cluster, machine, or server role is about to be shut down.

- **error**: An error occurred while a cluster, machine, or server role is being shut down.

You can perform the following operations:

- View the data center shutdown progress: In the upper part of the **IDC Shutdown** page, view the data center shutdown progress.
- View the cluster status: In the **Cluster List** section, view the status of each cluster, the total number of machines within each cluster, and the number of machines in each state.
- View the machine status: In the **Cluster List** section, click a state corresponding to a cluster. In the **Machine List** section, view all machines in the corresponding state within the cluster, the total number of server roles on each machine, and the number of server roles in each state.
- View the server role status: In the **Machine List** section, click a state corresponding to a machine. In the **SR List--xxx** message, view all server roles in the corresponding state on the machine.

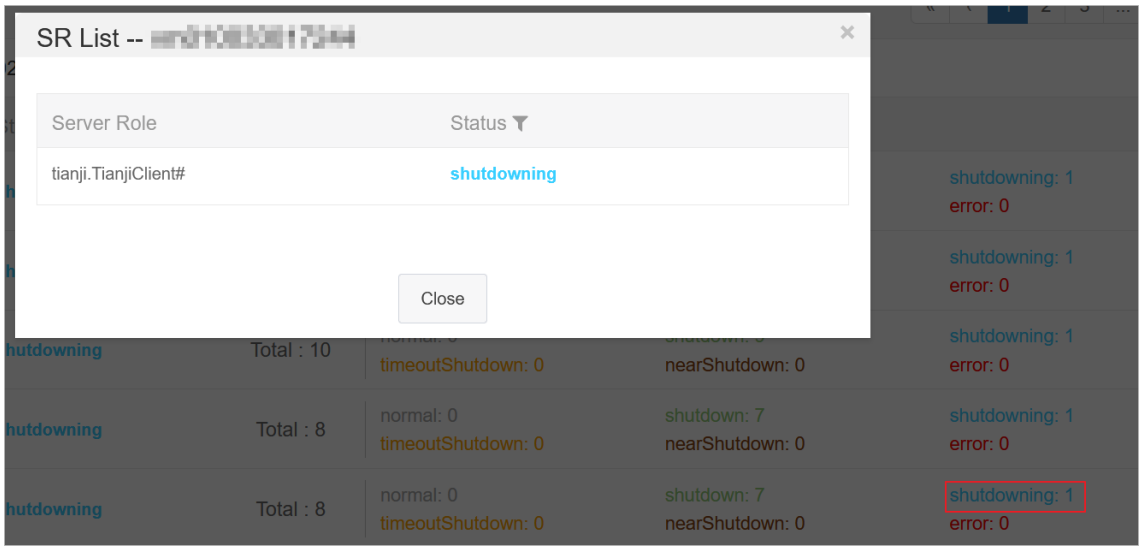

> **Note**
>
> In the left-side navigation pane, click **Go**. On the **All Reports** page, enter the entire or part of **Machine Power On or Off Statuses of Clusters** in the **Fuzzy Search** search box. In the search results, click **Machine Power On or Off Statuses of Clusters** to view the status of each server role.

- Filter clusters or machines: In the **Cluster List** or **Machine List** section, click the filter icon in the **Status** column and select a state to filter all clusters or machines in the corresponding state.
- Refresh data: Click **Refresh** in the upper-right corner to refresh data.

If all clusters in the **Cluster List** section are displayed in the **shutdown** state, the data center shutdown operation succeeds. After the base cluster is shut down, the OPS1 server is also shut down. Then, the Apsara Infrastructure Management Framework console is inaccessible.

6. After all base machines are shut down and inaccessible, go to the data center and confirm that all machines are powered off.

What's next

If you want to use the machines in the future, power on each machine one by one in the data center and wait until all services reach the desired state.

View the clone progress

This topic describes how to go to the OS Provision console (Corner Stone) and check the progress, status, and errors about machine installation.

Prerequisites

The username and password used to log on to the OP Provision console are obtained from delivery personnel.

Context

Apsara Infrastructure Management Framework provides a quick entry to the OS Provision console, which allows you to view details about machine installation. You can then obtain the progress and status about machine installation and then locate the installation faults.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Tools > Clone Progress**.
3. On the logon page of the OS Provision console, enter **Username** and **Password**, and then click **Submit**.

Appendix

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.
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.

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.

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.

Item	Description
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.

Rolling info report

This report displays the running and completed rolling tasks and the task-related statuses.

Choose a rolling action

This section displays the rolling tasks that are running. If no rolling tasks are running, no data is displayed in this section.

Item	Description
Cluster	The name of the cluster.
Git Version	The version of the 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.
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

When you select a rolling task in the **Choose a rolling action** section, this section displays the rolling statuses of server roles related to the selected task. If no rolling tasks are selected, the statuses of server roles related to all historical rolling tasks are displayed.

Item	Description
Server Role	The name of the server role.
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 for which the rolling task was approved by the decider.
Failure Rate	The proportion of machines on which the rolling task failed.
Success Rate	The proportion of machines on which the rolling task succeeded.

Server Role Rolling Build Information

This section displays the current and desired versions of each application in the server role during 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 of the application before the upgrade.
To Build	The version of the application after the upgrade.

Server Role Statuses on Machines

When you select a server role in the **Server Role in Job** section, this section displays the status of the server role on each machine.

Item	Description
Machine Name	The name of the machine on which the server role is deployed.
Expected Version	The desired version of the server role.
Actual Version	The current version of the server role.
State	The status of the server role.

Item	Description
Action Name	The ongoing action of the server role.
Action Status	The status of the action.

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.

Item	Description
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.

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.

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.

Item	Description
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.

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.

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.
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.

Item	Description
Refer Version List	The version that uses the resource.
Error Msg	The exception message.

Statuses of project components

This report displays the statuses of all abnormal server roles on machines within the current project. This report also displays the alert information of server roles and machines reported to Monitoring System.

Error State Component Table

This section displays the server roles that are not in the GOOD state or that are pending upgrade.

Item	Description
Project	The name of the project.
Cluster	The name of the cluster.
Service	The name of the service.
Server Role	The name of the server role.
Machine Name	The name of the machine.
Need Upgrade	Specifies whether the version has reached the desired state.
Server Role Status	The status of the server role.
Machine Status	The status of the machine.

Server Role Alert Information

When you select a row in the **Error State Component Table** section, this section displays the corresponding information.

Item	Description
Cluster	The name of the cluster.
Service	The name of the service.
Server Role	The name of the server role.
Machine Name	The name of the machine.
Monitored Item	The name of the server role metric.
Level	The severity level of the alert.
Description	The description of the alert.

Item	Description
Updated At	The update time of the alert.

Machine Alert Information

When you select a row in the **Error State Component Table** section, this section displays the corresponding information.

Item	Description
Cluster	The name of the cluster.
Machine Name	The name of the machine.
Monitored Item	The name of the server role metric.
Level	The severity level of the alert.
Description	The description of the alert.
Updated At	The update time of the alert.

Service Inspector Information

When you select a row in the **Error State Component Table** section, this section displays the corresponding information.

Item	Description
Cluster	The name of the cluster.
Service	The name of the service.
Server Role	The name of the server role.
Monitored Item	The name of the server role metric.
Level	The severity level of the alert.
Description	The description of the alert.
Updated At	The update time of the alert.

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.

Item	Description
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.

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.

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.
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.

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.

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.

Item	Description
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.
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.

1.1.6.10. Obtain the Prometheus domain name

This topic describes how to obtain the Prometheus domain name from the OPS1 server terminal when Prometheus is used for service monitoring.

Procedure

1. Log on to the OPS1 server terminal as the root user.
2. Run the following command to query the Prometheus domain name:

   ```
   kubectl get ing -n monitoring | grep tianjimon-prometheus-prome-prometheus | awk '{print $2}' | cut -d',' -f1
   ```

   A similar output is displayed:

   ```
   prometheus.tianjimon.cn-xxxx-envxx-d01.intra.envxx.shuguang.com
   ```

3. Replace the `<prometheus-domain>` value in the address `http://<prometheus-domain>/targets` with the domain name `prometheus.tianjimon.cn-xxxx-envxx-d01.intra.envxx.shuguang.com` obtained in the preceding step to obtain the endpoint `http://prometheus.tianjimon.cn-xxxx-envxx-d01.intra.envxx.shuguang.com/targets` of collection services in Prometheus.
4. Access the endpoint `http://prometheus.tianjimon.cn-xxxx-envxx-d01.intra.envxx.shuguang.com/targets` to view the collection services in Prometheus.

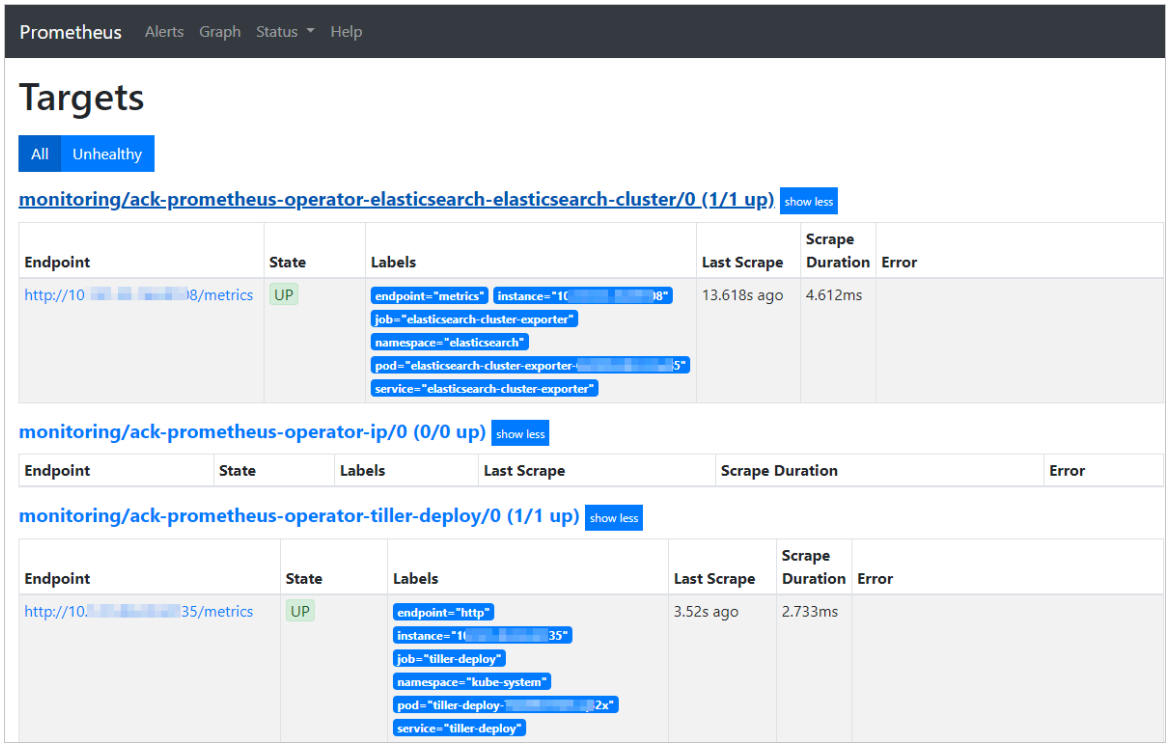

Note If a collection service is displayed in red, the service is abnormal and its metrics cannot be collected. You must troubleshoot the service.

1.1.7. Analysis

1.1.7.1. Inventory analysis

The Inventory Analysis module allows you to predict capacity trends and perform operations based on the available product inventory and usage habits.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. On the **Inventory Analysis** page, view the cloud product inventory.

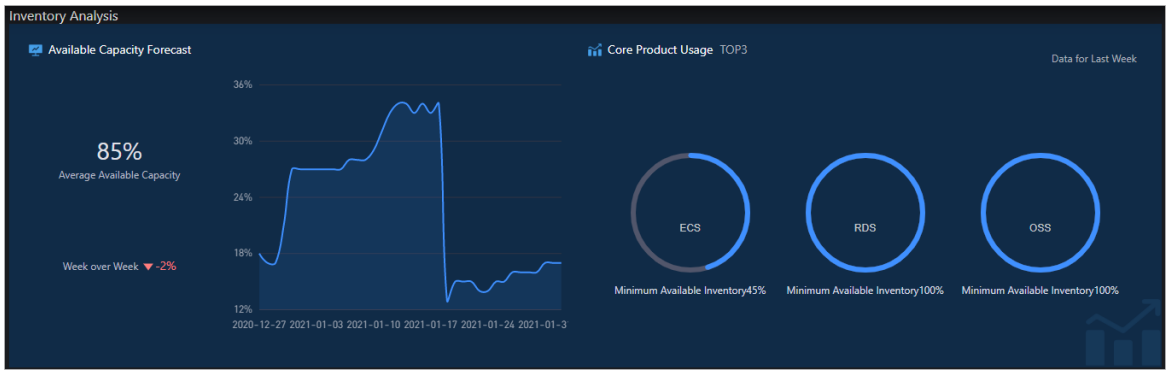

- In the **Inventory Analysis** section, view the average available inventory, changing trends, and

core product usage.

- In the **Product available capacity forecast** section, view the inventory of items related to a single product.
  - Click **Days**, **Weeks**, or **Months** to view the predicted available capacity of the product within the specified time range.
  - Click an inventory item under a product name to view the corresponding product inventory.
  - Move the pointer over a curve. Inventory information at a specific time point is displayed.

1.1.7.2. View the ECS inventory

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. In the left-side navigation pane, click **ECS**.
4. Select a date in the upper part of the page and view the ECS inventory.

> **Note** You can click the icon in the upper-right corner of the page to specify a zone and configure thresholds.

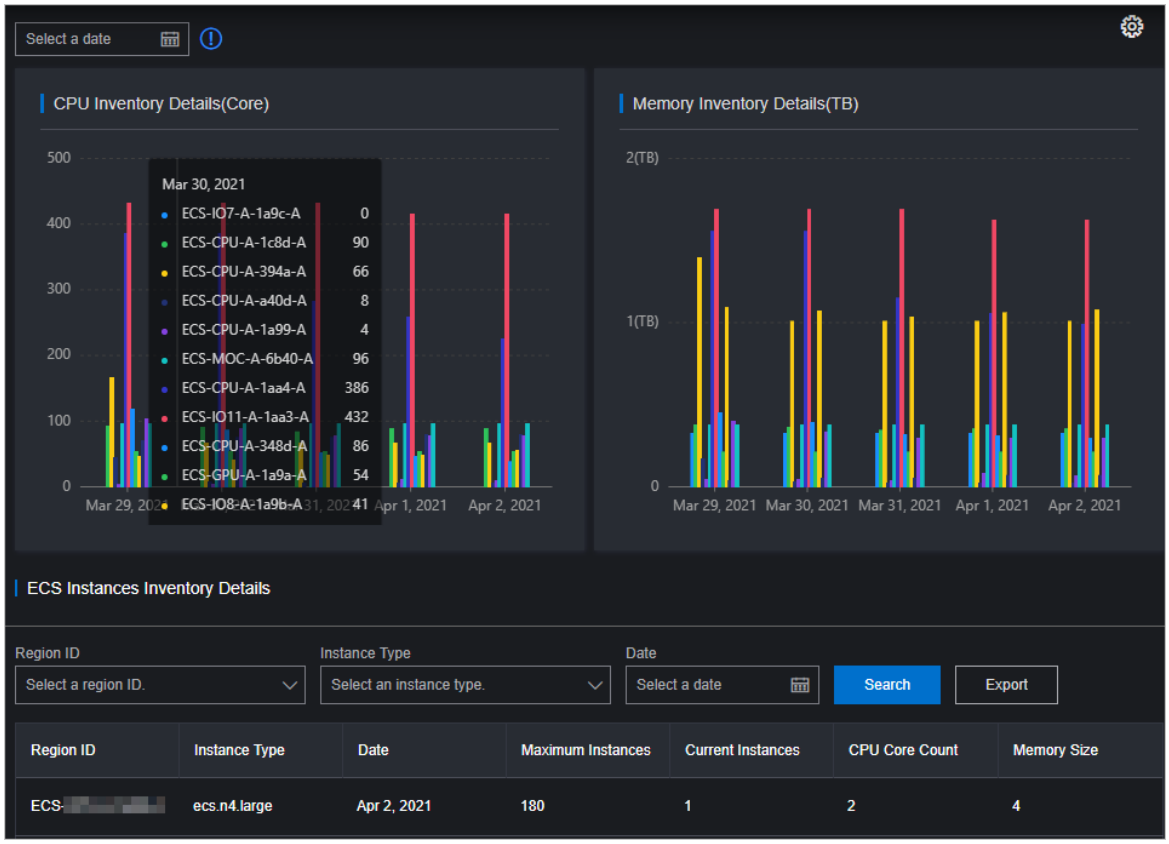

- The **CPU Inventory Details (Core)** and **Memory Inventory Details (TB)** sections show the usage and availability of CPU (core) and memory (TB) of all ECS instance families for the last five days.
- The **ECS Instances Inventory Details** section shows the inventory details of specified ECS instance type at the specified date on multiple pages by **Region ID**, **Instance Type**, and **Date**, as well as the CPU and memory configurations corresponding to each instance of this type.

5. (Optional)Query data by specifying **Region ID**, **Instance Type**, and **Date** in the **ECS Instances Inventory Details** section, and then click **Export** to export the ECS inventory details to your computer.

1.1.7.3. View the SLB inventory

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. In the left-side navigation pane, click **SLB**.

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

4. View the SLB inventory.
   - The Internal VIP Used Inventory and Public VIP Used Inventory sections show the amount and percentage of internal and public VIP inventory that are being used.
   - The Network Card Traffic section shows the inbound and outbound network card traffic.
   - The **SLB Inventory Details** section shows the SLB inventory details on multiple pages by **Type** and **Date**.
5. Select a cluster from the **Cluster** drop-down list in the upper-left corner and click **Search** to view the SLB inventory data of the cluster.

   > **Note**
   > - The clusters with the **slbCluster** keyword in their names are default clusters.
   > - The clusters with the **slbExtraCluster** keyword in their names are expanded clusters.

1.1.7.4. View the RDS inventory

You can view the Relational Database Service (RDS) inventory to query the usage and availability of RDS resources. This way, you can perform O&M operations in an efficient manner.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.

3. In the left-side navigation pane, click **RDS**.

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

4. View the RDS inventory.
   - The **RDS Inventory** section shows the inventories of different types of RDS instances within the last five days. Different colors indicate different types of RDS instances.
   - The **RDS Inventory Details** section shows the RDS inventory details on multiple pages by **Engines** and **Date**.

1.1.7.5. View the OSS inventory

You can view the Object Storage Service (OSS) inventory to learn more about the usage and availability of OSS resources and perform O&M operations more efficiently.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. In the left-side navigation pane, click **OSS**.

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

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

1.1.7.6. View the Tablestore inventory

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. In the left-side navigation pane, click **OTS**.

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

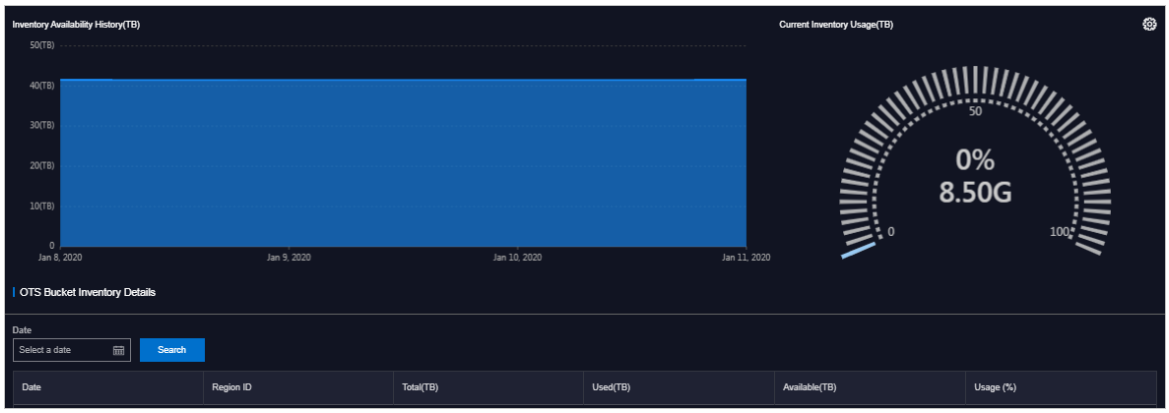

4. View the Tablestore inventory.
   - The **Inventory Availability History (TB)** section shows the available Tablestore inventory for the last five days.

- The **Current Inventory Usage (TB)** section shows the amount and percentage of Tablestore inventory that are in use.
- The **OTS Bucket Inventory Details** section shows the Tablestore inventory details on multiple pages by **Date**.

1.1.7.7. View the Log Service inventory

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Diagnose**.
3. In the left-side navigation pane, click **SLS**.

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

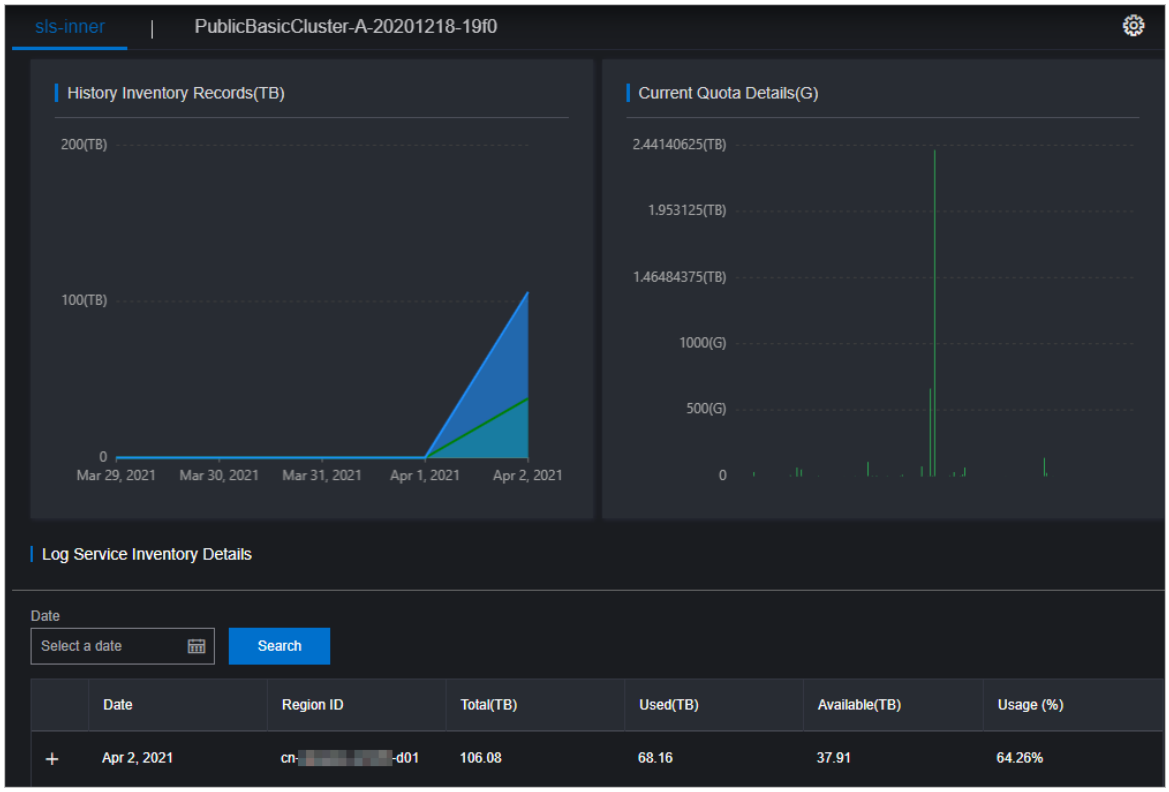

4. On the **sls-inner** tab, view the Log Service inventory details.
   - The **History Inventory Records (TB)** section shows the available and total Log Service inventory for the last five days.
   - The **Current Quota Details (G)** section shows the amount and percentage of Log Service inventory that are currently in use.
   - The **Log Service Inventory Details** section shows the Log Service inventory details on multiple pages by **Date**.

5. Click the **PublicBasicCluster-XXX** tab to view details about the Log Service inventory for which that you have applied.
   - 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 currently in use.
   - The **SLS Bucket Inventory Details** section shows the Log Service inventory details in multiple pages by **Date**.

1.1.7.8. View the EBS inventory

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

Context

> **Note** EBS is the Apsara Distributed File System storage provided for ECS by the base, and ECS IO clusters are used for Apsara Distributed File System storage. Therefore, you can view the EBS inventory in ECS IO clusters.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. In the left-side navigation pane, click **EBS**.

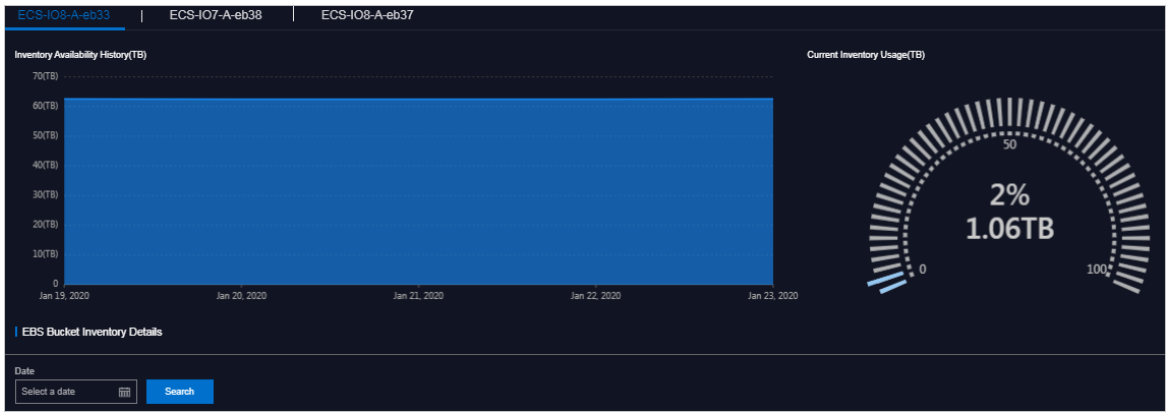

4. If multiple ECS IO clusters exist in the environment, click the tab of each ECS IO cluster to view the EBS inventory.
   - The **Inventory Availability History (TB)** section shows the available EBS inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount 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.

1.1.7.9. View the NAS inventory

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **Analysis**.
3. In the left-side navigation pane, click **NAS**.

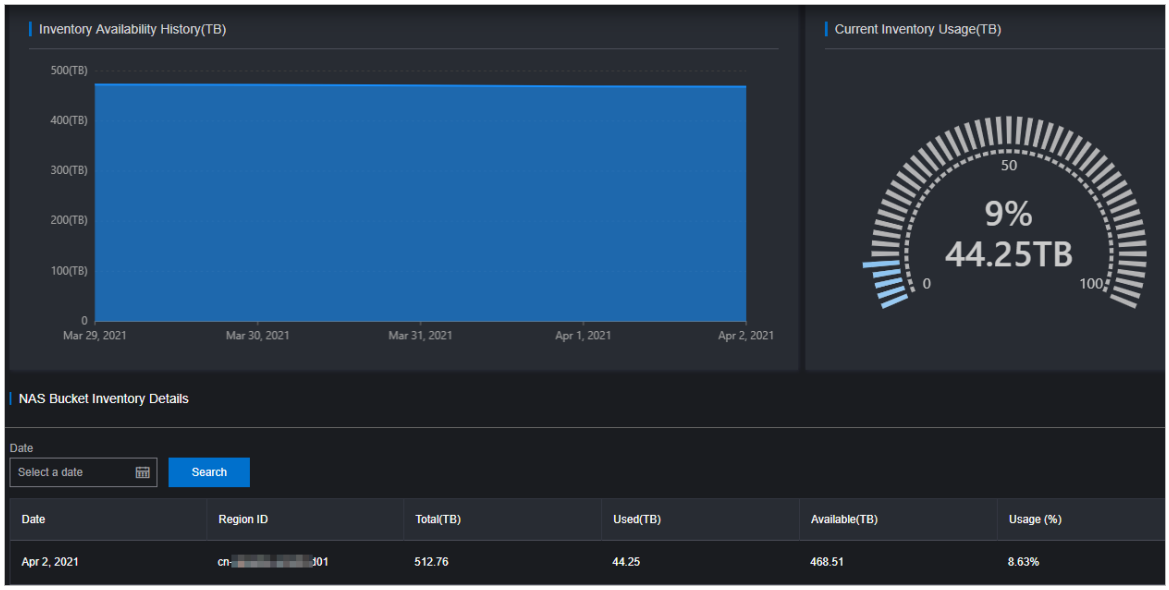

4. View the NAS inventory.
   - The **Inventory Availability History (TB)** section shows the available NAS inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount and percentage of NAS inventory that are being used.
   - The **NAS Bucket Inventory Details**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

1.2. Operations and Maintenance Guide

1.2.1. Overview

This topic describes the purposes, requirements, precautions, and technical support for the Apsara Uni-manager Operations Console.

Purposes

This guide helps O&M personnel perform preventative maintenance on the Apsara Uni-manager Operations Console to make sure that it can stably run for a very long period. O&M personnel can follow the instructions in this guide to handle the system issues that are identified during maintenance. If O&M personnel encounter system issues that are not covered in this guide, they can contact Apsara Stack engineers for technical support.

Requirements

- O&M personnel must acquire IT skills, including knowledge for computer networks, knowledge for computer operations, issue analysis, and troubleshooting.
- In addition, O&M personnel must pass the pre-job training and 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.

> **Notice** Take note that O&M personnel must comply with operation procedures during maintenance 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 from users.

Precautions

To ensure a stable system and avoid unexpected events, you must comply with 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. This prevents system faults that are caused by unauthorized operations.
- System security:
  - Before you perform system operations, you must be aware of their impacts.
  - You must clearly record the details about the issues that you encounter during the operations. This helps you troubleshoot and handle the issues.
- 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 Alibaba Cloud technical support for help.

1.2.2. Architecture

1.2.2.1. System architecture

The Apsara Uni-manager Operations Console provides unified O&M capabilities for Apsara Stack products.

The Apsara Uni-manager Operations Console comprises the following modules:

- Console: The user interface layer. This module provides an interactive user interface.
- Rest API: The console and system capabilities modules interact with each other by using the Rest API module.
- System capabilities:
  - Resource management: allows you to view and manage Apsara Stack resources in product, network, and data center dimensions.

- Unified alerting: provides a unified alert export and alert outbound configuration capabilities for Apsara Stack platforms.
- O&M capabilities: provides vertical O&M capabilities such as network, cloud products, storage, and tasks, as well as horizontal security O&M capabilities.
- System management: Provides system administration capabilities such as security policies, offline backup, log clearing, system configuration, and personal configuration.
- Inventory management: allows you to view product inventories and provides usage threshold alerting capabilities for various products.

The Apsara Uni-manager Operations Console also provides APIs for secondary development by third parties. You can select optional remote O&M capabilities.

The following figure shows the system architecture of the Apsara Uni-manager Operations Console.

System architecture of the Apsara Uni-manager Operations Console

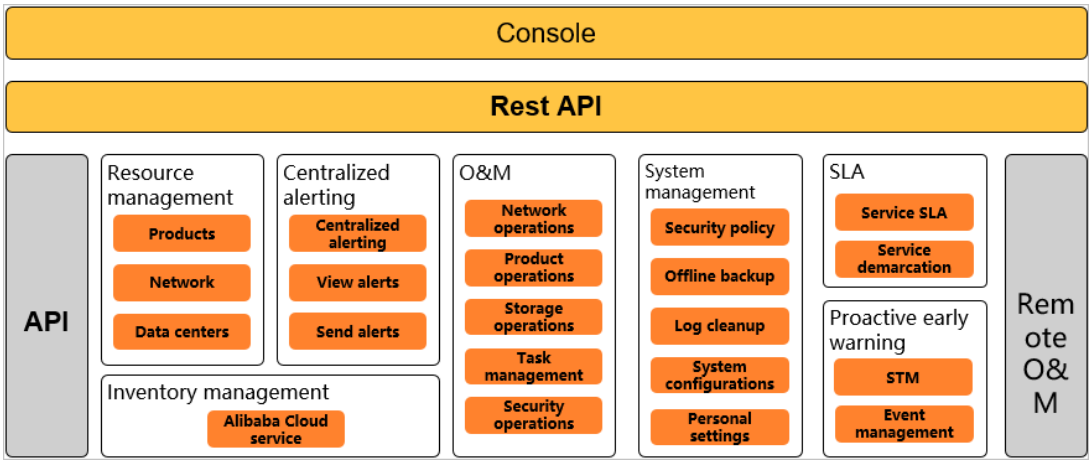

1.2.2.2. Deployment architecture

This topic describes how components are deployed in the Apsara Uni-manager Operations console.

The following table describes the deployment of components in the Apsara Uni-manager Operations Console.

Component	Deployment
aso-console.console	Cluster deployment. console uses the two-node method.
aso-mgr.asmgr	Cluster deployment. asmgr uses the two-node method.
aso-opr.asopr	Cluster deployment. asopr uses the two-node method.
aso-monitor.monitor	Cluster deployment. monitor uses the two-node method. You can change to the single-node or multi-node method based on business needs.

Component	Deployment
aso-monitor.tpcmon	Cluster deployment. tpcmon uses the single-node method.
aso-slalink.diagnosis-api	Cluster deployment. diagnosis-api uses the single-node method.
aso-slalink.full_link	Cluster deployment. full_link uses the two-node method.
aso-slalink.link_mapping	Cluster deployment. link_mapping uses the single-node method.
aso-slalink.noc	Cluster deployment. noc uses the two-node method.
aso-slalink.noc_agent	Cluster deployment. noc_agent uses the two-node method.
aso-slalink.sla-portal	Cluster deployment. sla-portal uses the two-node method.
aso-slalink.sla-thesla	Cluster deployment. sla-portal uses the three-node method.
aso-slalink.sla-topo	Cluster deployment. sla-topo uses the two-node method.
aso-slalink.sla-topo-dumpper	Cluster deployment. sla-topo-dumpper uses the two-node method.
aso-slalink.sla-topo-graphdb	Cluster deployment. sla-topo-graphdb uses the single-node method.
aso-slalink.pgmon	Cluster deployment. pgmon uses the single-node method.
aso-standalone.standalone	Cluster deployment. standalone uses the two-node method.
aso-emergency.stm	Cluster deployment. stm uses the two-node method.
aso-stm.incident	Cluster deployment. incident uses the two-node method.

1.2.2.3. Component architecture

This topic describes the components of the Apsara Uni-manager Operations Console and their features.

The following figure shows the component architecture of the Apsara Uni-manager Operations Console.

Component architecture of the Apsara Uni-manager Operations Console

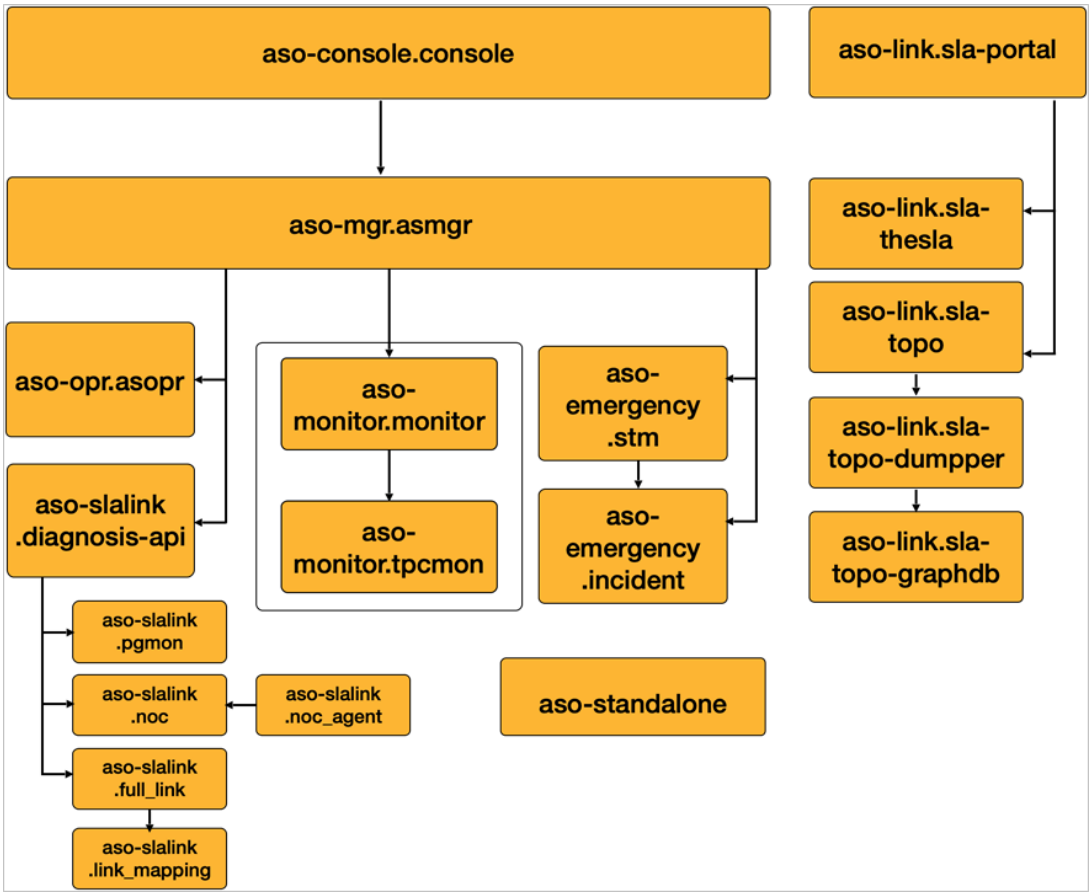

The following table describes the features of the components in the Apsara Uni-manager Operations Console.

Component	Description
aso-console.console	The frontend pages of the Apsara Uni-manager Operations Console, used for data display and analysis
aso-mgr.asmgr	For managing users, organizations, menus, permissions, inventories, physical platforms, product O&M
aso-opr.asopr	For managing physical server passwords, physical server channels, and tasks
aso-monitor.monitor	For viewing, importing, and exporting alert monitoring data and alert templates
aso-monitor.tpcmon	For executing the monitoring script of tpcmon
aso-slalink.diagnosis-api	The gateway of the slalink module, used for distributing requests to slalink subordinate modules

Component	Description
aso-slalink.pgmon	For Storage Operations Center
aso-slalink.full_link	The API module for end-to-end monitoring
aso-slalink.link_mapping	For collecting metadata from ECS, RDS, SLB, and VPC in end-to-end monitoring
aso-slalink.noc	For automated task orchestration in Network Operations Center
aso-slalink.noc_agent	For network data collection in Network Operations Center
aso-slalink.sla-portal	The portal console for SLA and end-to-end diagnostics, used for interaction with the SLA and end-to-end diagnostics modules
aso-slalink.sla-thesla	For processing SLA data and querying API data, analyzing time series data, and electing high availability clients
aso-slalink.sla-topo	For processing end-to-end diagnostic data and querying API data
aso-slalink.sla-topo-dumpper	For collecting end-to-end diagnostic and SLA data
aso-slalink.sla-topo-graphdb	For storing data for graph databases on the end-to-end diagnostic server
aso-standalone.standalone	For standalone output of the installation package for Apsara Uni-manager Operations Console
aso-emergency.stm	For discovering STM events
aso-emergency.incident	For managing STM events

1.2.2.4. Server roles

This topic describes projects, clusters, services, and server roles in the Apsara Uni-manager Operations Console to facilitate O&M personnel to identify and discover issues.

This following table describes projects, clusters, services, and server roles in the Apsara Uni-manager Operations Console.

project	Cluster	Service	Server role	Description

project	Cluster	Service	Server role	Description
aso	asoCluster	aso-console	aso-console.console	The frontend pages of the Apsara Uni-manager Operations Console, used for data display and analysis
		aso-monitor	aso-monitor.monitor	For viewing, importing, and exporting alert monitoring data and alert templates
			aso-monitor.tpcmon	For executing the monitoring script of tpcmon
		aso-slalink	aso-slalink.diagnosis-api	The gateway of the slalink module
			aso-slalink.noc	For Network Operations Center
			aso-slalink.noc_agent	The data collection agent for Network Operations Center
			aso-slalink.full_link	The API module for the slalink module
			aso-slalink.link_mapping	For listing ECS, RDS, SLB, and VPC instances in the slalink module
			aso-slalink.sla-portal	The SLA frontend display module
			aso-slalink.sla-thesla	The SLA backend processing module
			aso-slalink.sla-topo	The topo backend processing module

project	Cluster	Service	Server role	Description
			aso-slalink.sla-topo-dumpper	The data collection agent for topo
			aso-slalink.sla-topo-graphdb	The storage database for topo
			aso-slalink.pgmon	For Storage Operations Center
		aso-mgr	aso-mgr.asmgr	For managing users, organizations, menus, permissions, inventories, physical platforms, product O&M
		aso-opr	aso-opr.asopr	For managing physical server passwords, physical server channels, and tasks
		aso-stanalone	aso-standalone.standalone	For standalone output of the installation package for Apsara Uni-manager Operations Console
		aso-emergency	aso-emergency.stm	For discovering STM events
			aso-emergency.incident	For managing STM events

The following figure shows the call relationship between server roles. Arrows point to the called parties.

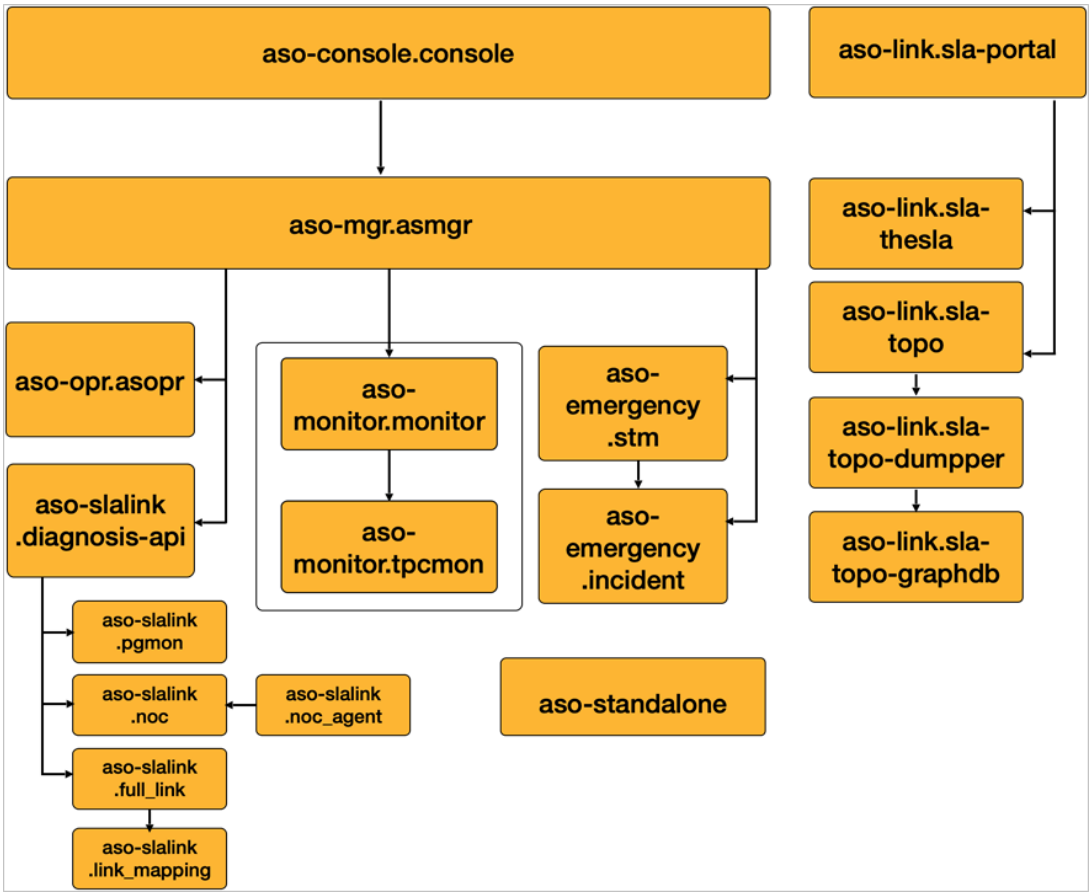

1.2.3. Handle alerts

You can choose **Alerts > Alerts** in the Apsara Uni-manager Operations Console to view and handle alerts.

1.2.4. Security maintenance

1.2.4.1. Maintain account passwords

This topic describes how to change account passwords in the Apsara Uni-manager Operations Console.

Accounts have validity periods. You must change passwords within the specified period of time. Otherwise, your account will be locked and you cannot log on the Apsara Uni-manager Operations Console.

Choose **Settings > Personal Settings > Personal Information**. You can click Change Password on the page to change your account password.

Choose **Settings > Personal Settings > Logon Settings** and click the **Account Validity Period** tab. You can set the validity period of your account on the tab.

1.2.4.2. Log audit

This topic describes how to audit component logs, log file paths, and log query methods.

After you access the container, find the log files listed in the following table to view log details.

Component	File	Description
aso-mgr.asmgr	/logs/asomanage-logger/main_trace.log	The logs and error logs related with asmgr container business
	/logs/asomanage-logger/access_trace.log	asmgr container access logs
	/logs/asomanage-logger/call_trace.log	Other service call logs related with asmgr container business
aso-monitor.monitor	/logs/asomonitor-logger/main_trace.log	The logs and error logs related with monitor container business
	/logs/asomonitor-logger/access_trace.log	monitor container access logs
	/logs/asomonitor-logger/call_trace.log	Other service call logs related with monitor container business
aso-opr.asopr	/logs/asooperator-logger/main_trace.log	The logs and error logs related with asopr container business
	/logs/asooperator-logger/access_trace.log	asopr container access logs
	/logs/asooperator-logger/call_trace.log	Other service call logs related with asopr container business

1.2.5. Troubleshooting

1.2.5.1. Establish emergency response mechanisms

O&M teams must establish emergency response mechanisms, so that services can be recovered in time after errors or security events occur.

1.2.5.2. Designate owners for handling various issues

When monitoring alerts occur on the Apsara Uni-manager Operations Console, onsite O&M personnel must first troubleshoot the issues based on the alerts. If onsite O&M personnel cannot solve the issues, contact Apsara Stack technical support engineers.

1.2.5.3. Troubleshooting

This topic describes troubleshooting in the Apsara Uni-manager Operations Console.

After a system fault is discovered during routine maintenance, O&M personnel can use the O&M feature and system logs in the Apsara Uni-manager Operations Console to view the detailed information of the fault, analyze the causes, and solve the fault. If the faults cannot be fixed, O&M personnel can collect the related information, such as system information and fault symptoms, and contact Apsara Stack engineers for technical support. This way, O&M personnel can fix the faults based on the instructions of the engineers.

After the fault is rectified, verify the solution, review the troubleshooting method, and make improvements.

1.2.5.4. Troubleshooting

1.2.5.4.1. No alert data found in the Apsara Uni-manager Operations Console

The Apsara Uni-manager Operations Console and TianJi portal obtain alert data from TianjiMon. You can first verify from TianJi portal whether TianjiMon generates alert data.

Procedure

1. Check whether alert data can be found on TianJi portal:
   - If yes, contact technical support engineers for the Apsara Uni-manager Operations Console.
   - If no, contact technical support engineers for TianjiMon.

1.2.5.4.2. The Apsara Uni-manager Operations Console prompts that the account has been reclaimed

This topic describes how to handle the issue where you are prompted that the account has been reclaimed when you log on to the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the aso_auth database.
2. Execute the following statement to view the status and validity date of the user:

   ```
   select * from user where name="test";
   ```

3. Check whether the user status (statue) value is 1.
   - If yes, execute the statement.
   - If no, execute the `update user set statue=1 where name="test";` statement to change the user status value to 1.
4. Execute the following statement to change the validity date later than the current day:

   ```
   Update user set validity_date="2019-11-04 15:23:23" where name="test";
   ```

1.2.5.4.3. 403 returned when you log on to SRE Technology Support Platform

This topic describes how to handle issue where that 403 is returned when you log on to SRE Technology Support Platform.

Procedure

1. Log on to the Apsara Uni-manager Operations Console as a security officer.
2. In the top navigation bar, click **Settings**.
3. In the left-side navigation pane, choose **System Settings > Role Management**.
4. Find the role of the user that you use to log on to SRE Technology Support Platform and click **Edit** in the **Actions** column.
5. In the **Edit Role** dialog box, modify the parameters, add the tongque_administrator role permissions to the role, and then click **OK**.

1.2.6. Appendixes

1.2.6.1. Common HTTP status codes

This topic describes the common HTTP status codes and their types in the Apsara Uni-manager Operations Console.

HTTP status code types

HTTP status codes can be divided into five types based on the first digit:

- 1xx: informational response. The request was received, continuing process.
- 2xx: successful. The request was received, understood, and accepted.
- 3xx: redirection. Further action needs to be taken in order to complete the request.
- 4xx: client error. The request contains bad syntax or cannot be fulfilled.
- 5xx: server error. The server failed to fulfil an apparently valid request.

Common HTTP status codes

HTTP status code	Description
200	Successful response
404	Resources not found
400	Bad request, such as syntax errors
401	Request authorization failed
402	Cookie verification failed
403	Forbidden request

HTTP status code	Description
404	Files, queries, or URLs not found
406	Account not activated
409	Resource already exists
418	Verification code invalid
419	Service authorization failed
420	Signature data check failed
421	Certificate check failed
422	Random does not exist
423	SN and account mismatch
424	Username or ID does not match
425	Service does not exist
426	Authentication does not exist
427	Authentication failed
428	Request from Apsara Infrastructure Management Framework failed
500	Internal errors in the server
501	Requested function not supported in the server
510	Database operation errors

2.Apsara Uni-manager Management Console Operations

2.1. Operations and Maintenance Guide

2.1.1. Overview

This topic describes the purposes, requirements, precautions, and technical support for the Apsara Uni-manager Management Console.

Purposes

This guide helps O&M personnel perform routine maintenance and troubleshooting on the Apsara Uni-manager Management Console. O&M personnel can follow the instructions in this guide to handle the system issues that are identified during maintenance. If O&M personnel encounter system issues that are not covered in this guide, they can contact Apsara Stack engineers for technical support.

Requirements

- O&M personnel must acquire IT skills, including knowledge for computer networks, knowledge for computer operations, issue analysis, and troubleshooting. They include common commands in Linux, Docker, and DNS.
- In addition, O&M personnel must pass the pre-job training and learn the knowledge for the Apsara Stack system. The required knowledge for the system includes but is not limited to the usage of maintenance tools such as Apsara Infrastructure Management Framework and Apsara Stack Doctor.

> **Notice**
>
> O&M personnel must comply with operation procedures during maintenance 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 from users.

Precautions

To ensure a stable system and avoid unexpected events, you must comply with 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. This prevents system faults that are caused by unauthorized operations.
- System security
  - Before you perform system operations, you must be aware of their impacts.
  - You must clearly record the details about the issues that you encounter during the operations. This helps you troubleshoot and handle the issues.
- 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 Alibaba Cloud technical support for help.

2.1.2. Architecture

2.1.2.1. System architecture

This topic describes the major components of the Apsara Uni-manager Management Console and their features.

The Apsara Uni-manager Management Console consists of four components: frontend (portal), resource hosting, ASAPI, and backend services. The following figure shows the architecture of the Apsara Uni-manager Operations Console.

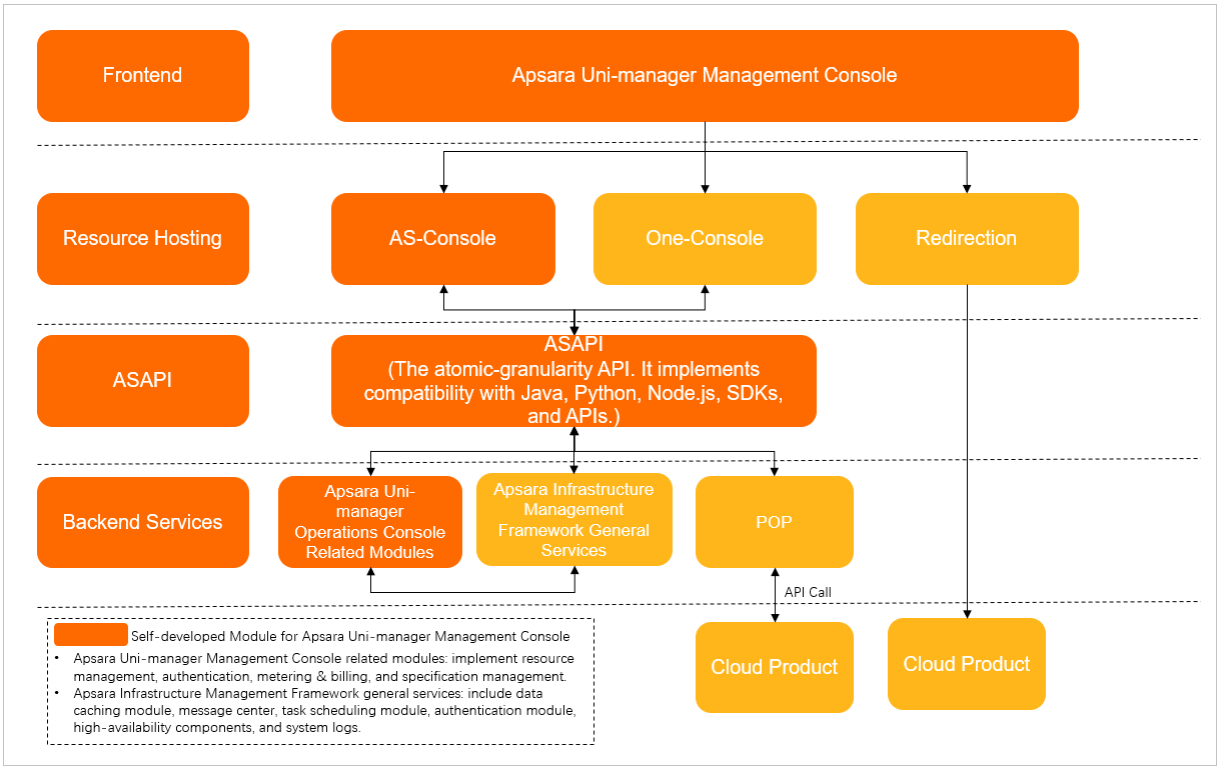

These modules provide the following features:

- Frontend: implements interaction and visualization of cloud resources, and enables you to manage the lifecycle of cloud resources in the console.
- Resource hosting: includes AS-Console and One-Console. They can host static files used at the frontend of different cloud products and provide unified API call and configuration services for the frontend.

- ASAPI: the unified gateway of the Apsara Uni-manager Management Console. It implements compatibility with Java, Python, Node.js, SDKs, and APIs, receives requests from the frontend and SDKs, and manages APIs in the cloud. Specifically, this component provides the following features:
  - Allows you to use SDKs to call APIs of cloud products.
  - Provides account verification, automatic routing, throttling, and log auditing.
  - Calls cloud products, Apsara Uni-manager Management Console related modules, and third-party integrated systems.
  - Adopts the asynchronous and highly concurrent request solution to provide Jetty, NIO, and Servlet 3.0 non-blocking asynchronous requests and implement excellent characteristics such as high concurrency and high throughput.
  - Uses the thread pool isolation technology to boost request processing speed by dividing threads within ASAPI into business processing threads, routing threads, and result processing threads.
- Back-end services: include Apsara Uni-manager Management Console related modules, Apsara Infrastructure Management Framework general services, and POP. They offer the following features:
  - Apsara Uni-manager Management Console related modules: implement resource management, authentication, metering & billing, and specification management.
  - Apsara Infrastructure Management Framework general services: include data caching module, message center, task scheduling module, authentication module, high-availability components, and system logs.
  - POP: the unified cloud computing underlying gateway. After ASAPI completes data processing, requests from APIs of cloud products are forwarded to POP and then back to cloud products for resource processing.

2.1.2.2. Deployment architecture

This topic describes how components are deployed in the Apsara Uni-manager Management Console.

Component	Deployment
Portal	Cluster deployment. Portal uses the two-node method.
OneConsole	Cluster deployment. OneConsole uses the two-node method.
ASAPI	Cluster deployment. ASAPI uses the two-node method.
Auth	Cluster deployment. Auth uses the two-node method.
Common	Cluster deployment. Common uses the two-node method.
Manage	Cluster deployment. Manage uses the two-node method.

Component	Deployment
ResourceMgr	Cluster deployment. ResourceMgr uses the two-node method.
Metering	Cluster deployment. Metering uses the two-node method.
Billing	Cluster deployment. Billing uses the two-node method.
Cosmos	Cluster deployment. Cosmos uses the two-node method.

2.1.2.3. Component architecture

This topic describes the components in the Apsara Uni-manager Management Console.

Component	SR on Apsara Infrastructure Management Framework	Description
Portal	ascm-portal.Portal#	Displays at the frontend pages of self-developed product consoles.
OneConsole	oneService.one-console-aliyun-com#	Hosts cloud product console pages and processes page requests. ASConsole features have been incorporated into OneConsole.
ASAPI	asapi.ApiServer#	The API gateway, used to process all internal and external calls that the Apsara Uni-manger Management Console received. SDKs are used to initiate calls.
Auth	unimanager-base.IdentityService#	Implements authentication for organizations, users, and permissions.
Common	unimanager-base.CommonServer#	Used for basic services such as pulling basic data (endpoints) and menu configuration of cloud products.

Component	SR on Apsara Infrastructure Management Framework	Description
Manage	ascm-brm.Manage#	Processes the core processes for Apsara Stack resource management, such as the processes to create, modify, query, and delete cloud resources.
ResourceMgr	ascm-brm.ResourceMgr#	Synchronizes the resource list of cloud products at a regular basis.
Metering	ascm-metering.metering_server#	Metering features.
Billing	unimanager-billing.billing#	Billing features.
Cosmos	ascm-cosmos.studio#	Dashboard displays.

The following figure shows the key dependencies between components.

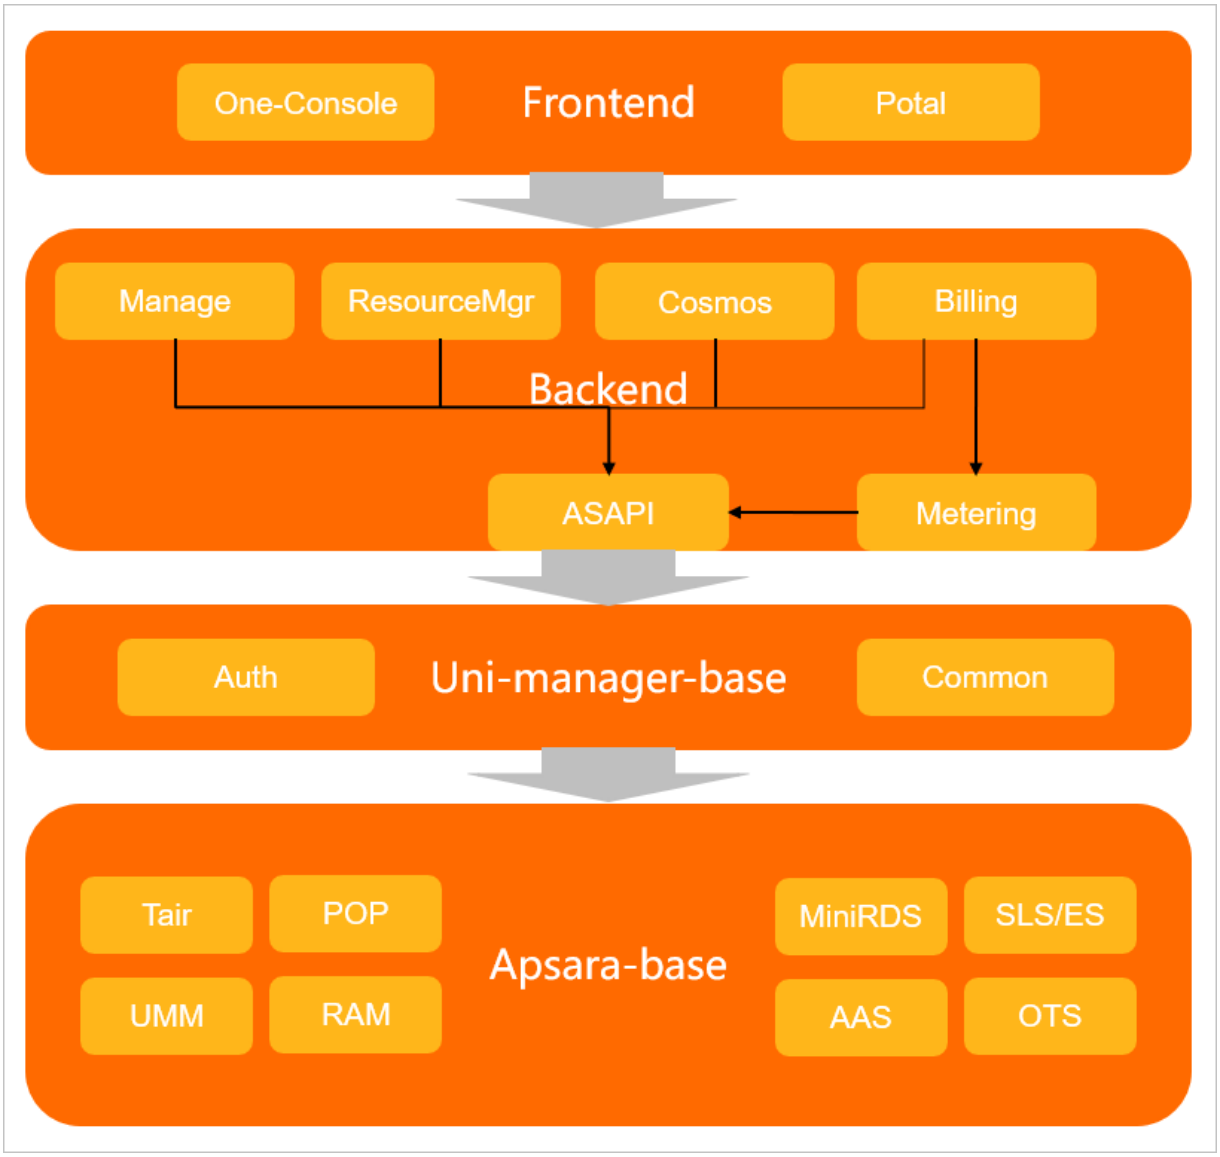

2.1.2.4. Dependent base services

This topic describes the base services that the Apsara Uni-manager Management Console depends on.

Whether base services are normal is critical to the stability of the Apsara Uni-manager Management Console. The Apsara Uni-manager Management Console depends on the following base services:

- MiniRDS: used for business data persistency, such as storing organization and resource set data.
- Tair: caches frequently accessed data, such as the cloud resource list.
- UMM: used to obtain your AccessKey ID and AccessKey secret for authentication.
- POP: the access gateway of cloud products, used to call product API operations.
- RAM: implements user authorization and controls operation permissions on resources.
- AAS: used for account management.
- Tablestore: processes message data and time series data, such as message push and metering.

- Log Service: stores call logs, including backend service logs (such as access, call, and main) and API call audit logs.

2.1.2.5. Server roles

This topic describes the server roles of the Apsara Uni-manager Management Console and their features.

The following table describes the server roles of the Apsara Uni-manager Management Console.

Project	Cluster	Service	Server role	Description
	ascm	Asapi	ApiServer#	A core component of ASAPI. Almost all features related with the Apsara Uni-manager Management Console depend on ASAPI for forwarding API requests. ApiServer# depends on ApiDbInit# and on middleWare-tair, webapp-pop, ram-ramService, baseService-umm-ak, MinRDS, and unimanager-base in base services.
			ApiDbInit#	Used for DB initialization (API definition initialization). It is executed only once during deployment or upgrade.
			ServiceTest#	Checks the service status.
		ascm-brm	Manage#	The core business logic for resource management.
			ResourceMgr#	The scheduled resource call task.
			ServiceTest#	Checks the service status.
		ascm-metering	metering_server#	Meters core services.
			ServiceTest#	Checks the service status.

Project ascm	Cluster	Service	Server role	Description
		ascm-portal	Portal#	A self-developed frontend service (will be integrated into OneConsole).
		unimanager-billing	billing#	Bills core services.
			ServiceTest#	Checks the service status.
		unimanager-base	DbInit#	Used for DB initialization.
			IdentityService#	Used for users, roles, and permissions.
			BaseInit#	Initializes configuration data (cloud product configuration metadata).
			CommonSever#	A public component, used to provide cloud product metadata service and configure menus.
			ServiceTest#	Checks the service status.
	ascmcommander	ascm-cosmos	studio#	The frontend of the dashboard.
			insight#	The backend of the dashboard.
			DbInit#	Used for DB initialization.
			ServiceTest#	Checks the service status.
			cdn-aliyun-com#	The frontend static resource server.

Project	oneConsole Cluster	oneService Service	Server role	Description
oneConsole			one-console-aliyun-com#	Used for front-end service hosting, including the DB initialization script, foreground page rendering and data APIs, DB operation module and VIPER configuration module.
	oneServiceAdminCluster	oneService-admin	one-console2#	Encapsulates API calls for some non-POP products, including Tablestore, OSS, and API Gateway.

2.1.3. Routine maintenance

2.1.3.1. Monitoring metrics

This topic describes monitoring metrics and how to view them.

View monitoring metrics

You can use one of the following methods to view monitoring items:

- Use Apsara Stack Doctor
  i. Log on to Apsara Stack Doctor.
  ii. In the left-side navigation pane, choose **End-to-end Diagnostics** > **Machine Metrics**.
  iii. Select values from the drop-down lists in the upper-left corner of the page, and click **Query Metrics** to view the trend charts of machine metrics.
- Use Linux commands

  The following commands can be used: vmstat, top, tsar, and jstat(/opt/taobao/java/bin). The Apsara Uni-manager Management Console provides common O&M scripts. For more information, see Common O&M commands.

Monitoring metrics

Metric	Description	Criteria
CPU	Docker CPU utilization	High CPU utilization, such as over 70%.
Memory	Container memory usage	High memory usage, such as over 70%, or memory usage continues to grow over a period of time.

Metric	Description	Criteria
Load	Average CPU load	It can be judged based on the daily average baseline. Typically, when the load value is greater than 0.7 * CPU, it can be regarded as high load.
I/O usage (%)	Disk I/O usage	It can be judged based on the daily average baseline. Typically, when the I/O usage value is greater than 70%, it can be regarded as high I/O usage.
TCP retransmission rate (%)	Reflects the network transmission condition	Such as not more than 5%.
Network traffic	The traffic on your network	It can be judged based on the daily average baseline. If a surge occurs, you must focus on network traffic.
GC operations	GC is an important metric for the memory usage by Java applications	Focus on the occurrence of Full GC operations (Concurrent Mark-Sweep).
GC duration	GC is an important metric for the memory usage by Java applications	Typically, the GC duration of tens of milliseconds is normal. If the GC duration last for seconds, focus on this metrics.
JVM memory usage (MB)	Details about JVM memory size	It can be judged based on GC and JVM conditions. View JVM memory usage,
netstat	Details about TCP connections	It can be judged based on the daily average baseline. Check whether the TIME_WAIT value surges.

2.1.3.2. Alert settings

You can use Apsara Stack Doctor (ASD) to configure alert settings based on monitoring metrics and allow you to view real-time alerts and discover system problems.

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **End-to-end Diagnostics** > **End-to-end Diagnostic Settings**.

3. Click the **Alert Settings** tab.

   The page displays the alert rule information, including the server role, status, and alerts.

4. Adds alert rules.

   Click **Add Alert Settings** in the upper-left corner. In the dialog box that appears, configure the parameters and click **OK**.

   The following table describes the parameters.

Parameter	Description
Server Role	The name of the server role.
Metric	The name of the metric that triggers alerts.
Condition	The metric threshold. Alerts can be triggered only when the metric is greater than or less than the threshold value.
Level	Set the parameter to **Critical** or **Normal**.
Duration	The period when the metric threshold is exceeded. When the specified period is reached, alerts are triggered.

5. (Optional) Modify alert rules.

   Find a server role and click **Modify** in the **Actions** column. In the dialog box that appears, modify the parameters and click **OK**.

6. (Optional) Disable or enable alert rules.

   Find a server role and click **Disable** or **Enable** in the **Actions** column to disable or enable the alert rules for the server role.

2.1.3.3. Routine inspections

This topic describes the methods of routine inspections.

Routine inspections by Apsara Infrastructure Management Framework

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations** > **Cluster Operations**.
3. In the **Cluster** search box, enter `ascm`.
4. Find the cluster and then click **Operations** in the Actions column.
5. Click the **Services** tab and check whether the service reaches the desired state.

Manual inspections

The inspection steps are the same for different services. The ASAPI service is used in this topic.

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations** > **Cluster Operations**.
3. In the **Cluster** search box, enter `ascm`.
4. Find the cluster and then click **Operations** in the Actions column.
5. Click the **Service Inspection** tab.
6. Find the **asapi** service and then click **Details** in the **Actions** column.
7. Find the machine and click **Terminal** in the **Actions** column.

You can run the following commands to perform inspections.

Project	Cluster	Service	Server role	Health check
		asapi	ApiServer#	Run the following command in the docker: `curl http://127.0.0.1:7001/api/version` The output for a successful execution: `{"version":"v3.13","code":"200","apiCount":13122,"desc":"update xml pretty"}`
			ApiDbInit#	Typically, Docker will automatically exit after an execution of about 5 to 10 minutes. If Docker has not exited for a long time (Apsara Infrastructure Management Framework has not reached the desired state), you can view the log and check the error information in the log: `docker logs <dockerid>` The output for a successful execution: `Progress is 100%`

Project	Cluster	Service	Server role	Health check

ascm	ascm	ascm-brm	Manage#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

			ResourceMgr#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

		ascm-metering	metering_server#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

		ascm-portal	Portal#	Run the following command in the docker:
				`curl -s http://127.0.0.1/check.node`
				The output for a successful execution:
				`success`

		unimanager-billing	billing#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

Project	Cluster	Service	Server role	Health check
		unimanager-base	DbInit#	Run the following command in the docker: `docker logs <dockerid>` View the execution results.

			IdentityService#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

			BaseInit#	Run the following command in the docker: `docker logs <dockerid>` View the execution results.

			CommonServer#	Run the following command in the docker:
				`curl --location --request GET 'localhost:8008/ascm/basic/welcome'`
				The output for a successful execution:
				`welcome to basic.`

	ascmcommander	ascm-cosmos	studio#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

			insight#	Run the following command in the docker:
				`/alidata/bin/check_health.sh`
				The output for a successful execution:
				`ok`

Project	Cluster	Service	Server role	Health check
			DbInit#	Run the following command in the docker: `docker logs <dockerid>` View the execution results.

oneConsole	oneConsole Cluster	oneService	cdn-aliyun-com#	Run the following command in the docker:
				`curl <IP address of the docker>:80/api/console/check-health`
				The output for a successful execution:
				`ok`

			one-console-aliyun-com#	Run the following command in the docker:
				`curl http://127.0.0.1:7003/preload.htm`
				The output for a successful execution:
				`success`

	oneServiceAdminCluster	oneService-admin	one-console2#	Run the following command in the docker:
				`curl 127.0.0.1:7001/preload.htm`
				The output for a successful execution:
				`success`

2.1.4. Security maintenance

During maintenance of the Apsara Uni-manager Management Console, you must take note of account password maintenance and log audit.

Account password maintenance

Global roles such as operations administrator in the Apsara Uni-manager Management Console have operation permissions on all resources. Therefore, you must keep relevant passwords strictly confidential. We recommend that you adopt complex password policies and update them on a regular basis.

Log audit

The Apsara Uni-manager Management Console provides two levels of log audit:

- Operation log audit

  For more information about how to view operations logs, see *View operations logs* in *Apsara Uni-manager Management Console User Guide*.

- API operation call audit

  i. Log on to the Apsara Uni-manager Management Console as an administrator or a resource auditor.

  ii. In the top navigation bar, choose **Products** > **Others** > **APIs Tool**.

  iii. In the left-side navigation pane of the API Tool page, click **API Audit**.

  iv. (Optional) Select or enter information about the call record at the upper part of the page based on your needs.

  v. On the API Audit page, you can view the ID, product name, operation name, call status, operation duration, access time, and caller.

  vi. Find a call record and click the icon in the **Actions** column to view the request parameters, response parameters, and call logs of the operation.

Due to the limitation of storage space in the base, the preceding logs cannot be stored for a very long period. The maximum storage period is seven days.

If you want to store logs for more days, use Log Service, which can dump logs. Or . For more information, contact the Log Service and Elasticsearch teams.

2.1.5. Troubleshooting

2.1.5.1. Determine fault effects

This topic describes the general idea to determine fault effects.

1. Determine fault effects. Examples:
   - Whether the physical environment is normal (as network, DNS, and firewall).
   - Whether the base environment is normal (whether the base services have reached the desired state on Apsara Infrastructure Management Framework, especially the base services that the Apsara Uni-manager Management Console depends on).
   - Whether the Apsara Uni-manager Management Console can be logged on.
   - Which regions are affected.
   - Which products are affected.
   - Which features are affected.
2. Determine fault types. Examples:
   - Physical environment faults (contact hardware vendors).
   - Base environment faults (contact Alibaba Cloud onsite O&M engineers).
   - Product performance issues.
   - Product feature issues.

2.1.5.2. Collect fault logs

This topic describes how to collect fault logs.

1. Use the following template to collect fault information.

Item	Example
Operator	An operator
Operator role	Organization administrator
Organization (level)	Level-2 organization
Resource set	Default resource set
Product	ECS
Operation	Create an ECS instance
Console exception screenshot	Failed to create an ECS instance (screenshot)
Screenshot for abnormal return values	View the request and response parameters on the browser
Exception condition	An exception occurs if you select Zone B
Reproducible or not	Yes

2. Collect fault logs.

i. Obtain the EagleEye ID when the fault occurred.

- If you can log on to the Apsara Uni-manager Management Console and reproduce the fault, you can obtain the EagleEye ID from the browser by pressing F12. The eagleeye-traceid from the ResponseHeaders is preferred, followed by the eagleEyeTraceId from the response.

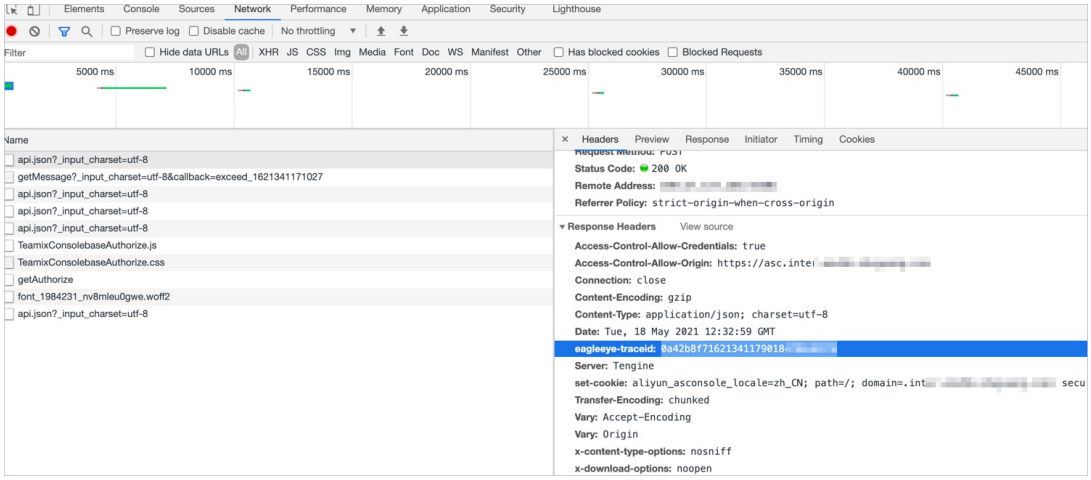

- If you can log on to the Apsara Uni-manager Management Console but cannot reproduce the fault, you can query the API call logs when the fault occurred. For more information, see Security maintenance.
- If you cannot log on to the Apsara Uni-manager Management Console but can query the access and call logs in the ASAPI dockers, you can obtain the EagleEye ID from the container logs by using filters such as API names and other accessible key parameters for further analysis.

ii. Obtain API operation call trace logs.

- Use API audit to obtain API operation call records. For more information, see Security maintenance.
- Use trace details in Apsara Stack Doctor to view logs. For more information, see View logs.
- Use Log Service to query API operation call logs. For more information, see View logs.
- Use the ASAPI containers to query API operation call logs. For more information, see Module logs.

2.1.5.3. Quickly troubleshoot issues

This topic describes how to quickly troubleshoot issues.

1. Troubleshoot issues in the base environment.

   i. Use different test environments to troubleshoot client issues due to browsers, networks, firewalls, and proxies.

   ii. Check Apsara Infrastructure Management Framework and the Apsara Uni-manager Operations Console to see whether services have not reached the desired state or alerts occur.

   iii. Check the network connectivity and whether DNS can resolve addresses.

2. Analyze business exceptions.

   View exceptions displayed in the Apsara Uni-manager Management Console and API call logs to analyze exception causes. For more information, see Common HTTP status codes in Appendix.

Key	Description	Purpose
API	The exceptional API call, including the product name and version	Effectively locate exceptional products or modules
domain	The domain name of an API request	
serverRole	The exceptional server role	
ASAPIErrorCode	The error code	Effectively discover the causes for exceptions
ASAPIErrorMessge	The error message	
ASAPISuccess	Indicates whether the call is successful.	
eagleEyeTraceId	The EagleEye ID of a trace	Continue querying logs of backend products
ASAPIRequestId	The request ID for a backend product	

To further troubleshoot exceptions, you can search the backend logs of the product based on the trace EagleEye ID to obtain more details.

3. Troubleshoot performance issues on Apsara Stack Doctor.

i. View slow API operation statistics.

For gateway-layer applications, Apsara Stack Doctor aggregates and collects all gateway API operations in the following time dimensions: last 5 minutes, last 30 minutes, last 1 hour, last 3 hours, and last 6 hours. Slow API operations are displayed in the descending order of average RTs.

a. Log on to Apsara Stack Doctor.
b. In the left-side navigation pane, choose **End-to-end Diagnostics** > **End-to-end Diagnostic Settings**.
c. Click the **Slow API Operations** tab.
d. Select values from the **Server Role** and **Time** drop-down lists, and then click **Query Slow API Operations**.

The information of all slow API operations is displayed in the descending order of **Average RT**.

ii. If a single API operation is slow, you can enter the trace EagleEye ID and select a time range on the Trace Details tab to view trace details.

a. Log on to Apsara Stack Doctor.
b. In the left-side navigation pane, choose **End-to-end Diagnostics** > **End-to-end Diagnostic Settings**.
c. Click the **Trace Details** tab.
d. Enter the Trace ID in the **TraceId** field, select a time range, and then click **Query Trace** to view the details of the trace.

iii. View machine performance metrics.

For more information about how to view the metrics, see Monitoring metrics.

2.1.6. Appendixes

2.1.6.1. Module logs

This topic describes the features of different log files.

Component	Log file	Description
	/logs/asapi-logger/access_trace.log The log path of Enterprise Edition V3.10.0: /opt/api-server/logs/apiServer-access.log	Logs that ASAPI receives API operation call requests

Component	Log file	Description
asapi.ApiServer#	/logs/asapi-logger/call_trace.log The log path of Enterprise Edition V3.10.0: /opt/api-server/logs/apiServer-call.log	Logs that ASAPI calls backend services and products
	/logs/asapi-logger/main_trace.log The log path of Enterprise Edition V3.10.0: /opt/api-server/logs/server.log /opt/api-server/logs/error.log	ASAPI debugging logs
	/logs/asapi-logger/exception_trace.log The log path of Enterprise Edition V3.10.0: /opt/api-server/logs/apiServer-exception.log	ASAPI exception logs
asapi.ApiDbInit#	docker logs	Docker startup logs
ascm-portal.Portal#	/root/logs/ascm-portal/ascm-error.log	Portal exception logs
	/root/logs/ascm-portal/ascm-access.log	Logs that Portal receives page call requests
	/root/logs/ascm-portal/ascm-agent.log	Portal process logs
	/root/logs/ascm-portal/ascm-web.log	Portal application logs
	/root/logs/ascm-portal/egg-schedule.log	Portal scheduled task logs
	/root/logs/ascm-portal/egg-web.log	Portal framework kernel and plug-in logs
	/logs/ascm-logger/access_trace.log	Logs that Manage receives ASAPI call requests
	/logs/ascm-logger/call_trace.log	Logs that Manage calls ASAPI

Component	Log file	Description
ascm-brm.Manage#	/logs/ascm-logger/exception_trace.log	Manage exception logs
	/logs/ascm-logger/main_trace.log	Manage debugging logs
ascm-brm.ResourceMgr#	/logs/ascm-logger/access_trace.log	Logs that Manage receives requests
	/logs/ascm-logger/call_trace.log	Logs that scheduled tasks call ASAPI
	/logs/ascm-logger/exception_trace.log	Scheduled task exception logs
	/logs/ascm-logger/main_trace.log	Scheduled task debugging logs
ascm-metering.metering_server#	/logs/metering/access_trace.log	Logs that Metering receives requests
	/logs/metering/call_trace.log	Logs that Metering calls other ascm services
	/logs/metering/exception_trace.log	Metering exception logs
	/logs/metering/main_trace.log	Metering debugging logs
unimanager-billing.billing#	/logs/billing/access_trace.log	Logs that Billing receives requests
	/logs/billing/call_trace.log	Logs that Billing calls other ascm services
	/logs/billing/exception_trace.log	Billing exception logs
	/logs/billing/main_trace.log	Billing debugging logs
unimanager-base.DbInit#	docker logs	Docker startup logs
unimanager-base.BaseInit#	docker logs	Docker startup logs

Component	Log file	Description
unimanager-base.IdentityServic e#	/logs/ascm-logger/access_trace.log	Logs that IDENTITYSERVICE receives requests
	/logs/ascm-logger/call_trace.log	Logs that IDENTITYSERVICE calls other base services such as RAM
	/logs/ascm-logger/exception_trace.log	IDENTITYSERVICE exception logs
	/logs/ascm-logger/main_trace.log	IDENTITYSERVICE debugging logs
unimanager-base.CommonServ er#	/ascm/.logs/basic.log	CommonServer application logs
	/ascm/.logs/basic-error.log	CommonServer exception logs
ascm-cosmos.insight#	/logs/data-processor/access_trace.log	Logs that INSIGHT receives requests
	/logs/data-processor/call_trace.log	Logs that INSIGHT calls other ascm services
	/logs/data-processor/exception_trace.log	INSIGHT exception logs
	/logs/data-processor/main_trace.log	INSIGHT debugging logs
ascm-cosmos.DbInit#	docker logs	Docker startup logs
oneService.one-console-aliyun-com#	/home/admin/logs/access_trace.log	Logs that ONE-CONSOLE-ALIYUN-COM receives requests
	/home/admin/logs/call_trace.log	Logs that ONE-CONSOLE-ALIYUN-COM calls ASAPI
	/home/admin/logs/exception_trace.log	ONE-CONSOLE-ALIYUN-COM exception logs

Component	Log file	Description
	/home/admin/logs/main_trace.log	ONE-CONSOLE-ALIYUN-COM debugging logs
oneService-admin.one-console2#	/home/admin/one-console2/logs/ascm-logger/access_trace.log	Logs that One-Console2 receives requests
	/home/admin/one-console2/logs/ascm-logger/call_trace.log	Logs that One-Console2 calls cloud products
	/home/admin/one-console2/logs/ascm-logger/exception_trace.log	One-Console2 exception logs
	/home/admin/one-console2/logs/ascm-logger/main_trace.log	One-Console2 debugging logs

2.1.6.2. Common error codes

This topic describes common error codes and their solutions.

Module	Error code	Description	Solution
	asapi.server.request.unknown.host	ASAPI failed to request the domain name of a backend product.	Perform the following operations: • Check on Apsara Infrastructure Management Framework whether the product is deployed. • Check on Apsara Infrastructure Management Framework whether the domain name exists. • Use the ASAPI container to disable the domain name link.

Module	Error code	Description	Solution
ASCM->asapi	asapi.server.api.notfound	ASAPI could not find the API operation to be called.	Perform the following operations: • Check whether the request parameters are correct. • Check whether the API operation is registered in the ASAPI database.
	asapi.server.endpoint.notfound	ASAPI could not obtain the endpoint of the API operation to be called.	Perform the following operations: 1. Check whether the corresponding product is deployed in the environment and has reached the desired state. 2. Then contact ascm_common O&M engineers.
	asapi.server.internal.error	An internal error occurred in ASAPI.	Contact ascm_asapi O&M engineers.
	asapi.server.endpoint.notmatch	An error occurred when ASAPI obtains the endpoint of the API operation to be called.	Contact ascm_common O&M engineers.
	asapi.server.timeout.connect	A timeout occurred when ASAPI connects to the server.	Contact ascm_asapi O&M engineers.
	asapi.server.request.parameter.accesskeyid.error	The AccessKey pair is invalid or the AccessKey secret could not be obtained.	Perform the following operations: • Check whether the AccessKey ID and AccessKey Secret are correct in the request parameters. • Contact ascm_asapi O&M engineers.

Module	Error code	Description	Solution
	asapi.server.request.parameter.regionid.not.found	A region ID could not be found.	Check whether the `regionId` is empty in the request parameters.
	asapi.server.request.validate.signature.failed	ASAPI failed to verify the signature.	Perform the following operations: • Check whether ASAPI SDK is used to initiate the call. • Contact ascm_asapi O&M engineers.
	asapi.server.request.flow.control	The interface is throttled.	Perform the following operations: • Check whether stress testing is performed in the current environment. • Contact ascm_asapi O&M engineers.
	asapi.server.request.upload.file.error	An error occurred when a file is upload.	Contact ascm_asapi O&M engineers.
	asapi.server.request.download.file.error	An error occurred when a file is downloaded.	Contact ascm_asapi O&M engineers.
	asapi.server.api.no.permission.operate	Insufficient permission for calling an API operation.	Contact ascm_asapi O&M engineers.
	ascm.manage.InvalidParameter.CloudAccount.OnlySupportAttachToTopOrg	An account can be bound to only level-1 organizations.	N/A
	ascm.manage.EntityNotExist.CloudAccount	The account does not exist.	Create the account.
	ascm.manage.DeleteConflict.CloudAccount.Organization	Before you delete an account, you must unbind it from the organization.	N/A

Module	Error code	Description	Solution
ASCM->Manage	ascm.manage.EntityNotExist.AccessKey	The AccessKey pair does not exist.	Create the AccessKey pair.
	ascm.manage.EntityNotExist.Instance	The instance does not exist.	Create the instance.
	ascm.manage.EntityNotExist.Role	The role does not exist.	Create the role.
	ascm.manage.InvalidParameter.Role.NotVisibleToTheUser	The role is not visible to the user.	Create a role that is visible to this user, or change the role attributes to make it visible to the user.
	ascm.manage.EntityNotExist.Organization	The organization does not exist.	Creates the organization.
	ascm.manage.EntityNotExist.ResourceSet	The resource set does not exist.	Creates the resource set.
	ascm.manage.EntityAlreadyExisted.User	The user already exists.	N/A
	ascm.manage.EntityAlreadyExisted.User.UserName	The username already exists.	Modify the username.
	ascm.auth.invalidParameter.token	The token is invalid.	Use a valid token.
	ascm.auth.account.deleted	The user has been deleted.	Recover the user as the administrator and use the user account to log on again.
	ascm.auth.account.inactive	The user has been disabled.	Enable the user as the administrator and use the user account to log on again.

Module	Error code	Description	Solution
	ascm.auth.account.type.notSupported	This account type is not supported.	N/A
	ascm.auth.account.notLoggedIn	The user is not logged on or the logon status is invalid.	Use the user account to log on again.
	ascm.auth.username.null	The username is empty. Enter it again.	Enter the username.
	ascm.auth.password.null	The password is empty. Enter it again.	Enter the password.
	ascm.auth.password.expired	The password has expired.	Reset the password as the administrator and use the new password to log on again.
	ascm.auth.account.usernameOrPasswordError	The username or password is invalid. Logon fails.	Enter the correct username and password.
	ascm.auth.organizationId.null	The organization ID is empty.	Enter the organization ID.
	ascm.auth.organization.rootNotSupported	The root organization is not supported.	N/A
	ascm.auth.loginPolicy.deActive	The user could not log on because the logon policy bound to the user has been disabled.	Enable the logon policy as the administrator and use the user account to log on again.
	ascm.auth.role.deActive	The operation could not be performed because the current role of the user has been disabled. Try again after the administrator enables the role.	Enable the role as the administrator and use the user account to log on again.

Module	Error code	Description	Solution
ASCM->Common	exposed resource not found <region.Identifier.dns>	An error occurred when ASAPI obtains DNS information.	Perform the following operations: • Obtain DNS information again. • Contact ascm_common O&M engineers.
ASCM->ONE-CONSOLE	signature.verification.failed	Signature verification failed.	Perform the following operations: • Obtain the signature again and verify it. • Contact ascm_ONE-CONSOLE O&M engineers.
	console.invalid.parameters	Console system error - An input parameter is invalid.	Contact ascm_ONE-CONSOLE O&M engineers.
	console.need.login	Log on again.	N/A
	console.system.error	Console system error - A system error occurred.	Contact ascm_ONE-CONSOLE O&M engineers.
	console.api.error	Console system error - An error occurred while calling ASAPI.	Contact ascm_ONE-CONSOLE O&M engineers.
	console.token.not.found	Console system error - Failed to obtain x-as-console-token. Refresh the page and try again.	N/A

2.1.6.3. Common O&M commands

This topic describes common O&M commands for the Apsara Uni-manager Management Console.

- top

  The `top` commands can dynamically display the current process and other conditions. You can continuously refresh the current status. If you run the command in the foreground, the command will exclusively occupy the foreground until you terminate the command.

Command syntax: `top <Parameter>`.

Command	Description
top -b	Batch mode.
top -c	Displays the complete command.
top -l	Ignores failed processes.
top -s	Secure mode.
top -S	Cumulative mode.
top -i <Time>	Sets a time interval.
top -u <Username>	Specifies the username.
top -p <Process ID>	Specifies the process.
top -n <Number>	Specifies the maximum number of iterations.

- vmstat

  The `vmstat` command is the most common monitoring tool in Linux and Unix. It can display the monitoring metrics of the server at a given time interval, including the CPU utilization, memory usage, virtual memory swapping, and I/O usage.

  Command syntax: `vmstat [<Parameter>]`.

Command	Description
vmstat -a	Displays active and inactive memory.
vmstat -f	Displays the number of forks since system startup.
vmstat -m	Displays slab statistics.

Command	Description
vmstat -n	Displays the header only once.
vmstat -s	Displays a table of various event counters and memory statistics.
vmstat delay	Specifies the refresh interval. If this parameter is not specified, only one result is displayed.
vmstat count	Specifies the number of refreshes. If the number of refreshes is not specified but the refresh interval is specified, the number of refreshes is infinite.
vmstat -d	Displays disk statistics.
vmstat -p	Detailed partition statistics.
vmstat -S	Specifies output units. Valid values: k, K, m, and M, which represent 1000, 1024, 1000000, and 1048576 bytes. Default value: K (1024 bytes).
vmstat -V	Displays version information.

- tsar

  The `tsar` command is a collection tool independently developed by Taobao. It is mainly used to collect server system information such as CPU, I/O, memory, and TCP, and data of applications such as squid, HAproxy, and NGINX. The collected data is stored on disks, to allow you to query historical information at any time and provide diverse output methods. To run the tsar command to display data, you can specify a module and merge data with multiple pieces of information. When the `-live` parameter is used, real-time data can be displayed in seconds.

  Command syntax: `tsar <Modules enabled> -i 1 -l <option>`

  The following table describes the option parameters.

Parameter	Description
-check	Displays the last collected data.
--check/-C	Displays the last help information of tsar.

Parameter	Description
--cron/-c	Use tsar in crond mode.
--interval/-i	Specifies the display interval of tsar. Unit: minutes. Default value: 5. When the `-live` parameter is used, the unit is seconds and the default value is 5.
--list/-L	Lists enabled modules.
--live/-l	Enables real-time mode. It is similar to iostat and can be used with the -i and module parameters.
--file/-f	Specifies the input file.
--ndays/-n	Displays the historical data of a specified time range. Default value: 1 day.
--date/-d	Specifies the date. It is in the format of YYYYMMDD or n (indicates n days ago).
--detail/-D	Specifies whether to view the main fields or all fields of the module.
--spec/-s	Displays the field after the -s value, such as `tsar -cpu -s sys,util` .
--watch/-w	Displays the historical records of last N minutes.
--merge/-m	Aggregates the displayed data. If three squid applications are run on the machine, you can run the `tsar-squid -m` command to aggregate the data.
--item/-I	Displays data of the specified items.
--help/-h	Displays help information and module information.

The following table describes the modules enabled parameters.

Parameter	Description
--cpu	CPU utilization data.
--mem	Physical memory usage data.
--swap	Virtual memory usage data.
--tcp	TCP IPv4 data.
--udp	UDP IPv4 data.
--traffic	Network traffic data.
--io	Linux I/O data.
--pcsw	Process start and context switch.
--partition	Disk usage data.
--tcpx	TCP connection data.
--load	System load data.

- jstat

  The jstat command can display the usage of each part of the heap memory and the number of loaded classes.

  Command syntax:

  ```
  jstat -<option> [-t] [-h<lines>]<vmid> [<interval> [<count>]]
  ```

  ```
  jstat [ generalOption | outputOptions vmid[interval[s|ms] [count]] ]
  ```

  Parameters:

  - `generalOption`: a single common command line option, such as `-help`, `-options`, or `-version`.
  - `outputOptions`: one or more output options, consisting of individual statOption options. They can be used with options such as `-t`, `-h`, and `-J`.

The following table describes the option parameters.

Parameter	Description
# -t	Adds a Timestamp column to the printed column to show the running time of the system.
# -h	Specifies a header is added after how many rows when periodical data collection is used.
# vmid	Virtual machine ID (PID).
# interval	The execution interval. Unit: milliseconds.
# count	Specifies how many records are displayed. No upper limit is set by default.

The following table describes the jstat parameters.

Parameter	Description
-class	Displays ClassLoad information.
-compiler	Displays JIT compilation information.
-gc	Displays GC-related heap information.
-gccapacity	Displays information about the capacities of the generations and their usage.
-gcmetacapacity	Displays the sizes of metaspace.
-gcnew	Displays information about the new generation.
-gcnewcapacity	Displays information about the capacities of the new generation and their usage.
-gcold	Displays information about the old generation and the permanent generation.

Parameter	Description
-gcoldcapacity	Displays information about the capacities of the old generation.
-gcutil	Displays a summary about GC information.
-gccause	Displays a summary about GC information (same as -gcutil), with the cause of the last and current GC events.
-printcompilation	Displays JIT compilation method information.

2.1.6.4. View logs

This topic describes how to view API call logs on the Trace Details tab in Apsara Stack Doctor, on Log Service.

View logs on the Trace Details tab in Apsara Stack Doctor

1. Log on to Apsara Stack Doctor.
2. In the left-side navigation pane, choose **End-to-end Diagnostics** > **End-to-end Diagnostic Settings**.
3. Click the **Trace Details** tab.
4. Enter the Trace ID in the **TraceId** field, select a time range, and then click **Query Trace** to view the details of the trace.

View logs on Log Service

1. Access the docker for ASAPI (serverRole:asapi.ApiServer#). For more information about how to access the docker, see Routine inspections.
2. Run the following command to obtain the endpoint, AccessKey ID, and AccessKey secret:

   ```
   env | grep sls
   ```

3. Replace `data` in the `domain_dataendpoint` value in the returned code with `portal` to obtain the endpoint of Log Service.
4. Enter the endpoint of Log Service in a browser. The account and password are the AccessKey ID and AccessKey secret obtained in Step 2.
5. After you select the project, you can view logs in the Logstore.

2.1.6.5. Common commands for Apsara Stack Agility PaaS

Apsara Stack Agility PaaS uses the Kubernetes base. This topic describes common Kubernetes commands for accessing ASAPI containers.

The following O&M commands are commonly used in the PaaS environment:

- Obtain the endpoint of the Apsara Uni-manager Management Console.

```
For #ascm
kubectl get ingress -n ascm
For #oneconsole
kubectl get ingress -n oneconsole
```

- Obtain pods.

```
For #ascm
kubectl get pods -n ascm
For #oneconsole
kubectl get pods -n oneconsole
```

- Access pods.

```
For #ascm
kubectl exec -it asapiapiserver-0 bash -n ascm
For #oneconsole
kubectl exec -it oneserviceoneconsole-0 bash -n oneconsole
```

- View docker logs.

```
For #ascm
kubectl logs -f asapiapidbinit-jqzk6 -n ascm
For #oneconsole
kubectl logs -f  oneserviceoneconsole-0 bash -n oneconsole
```

- Restart pods.

```
kubectl get pod asapiapidbinit-6wjvm  -n ascm  -o yaml | kubectl replace --force -f -
// Delete the existing one and create another.
```

- View pod details.

```
kubectl describe pod  asapiapiserver-0 -n ascm
```

- Exit the PaaS environment.

```
exit
```

2.1.7. Troubleshooting

2.1.7.1. Errors when you log on to the Apsara Uni-manager Management Console

This topic describes how to troubleshoot the errors when you log on to the Apsara Uni-manager Management Console.

Problem description

The page does not respond or an internal error is returned after you log on to the Apsara Uni-manager Management Console with the correct account and password.

Cause

Three possible causes:

- Network issues.
- An error is reported when the ascm-auth service calls the ASS service.
- The services in the Apsara Uni-manager Management Console and the base services (ASS, RAM, Tair, and Ummak) that the Apsara Uni-manager Management Console depends on have not reached the desired state.

Solution

1. Check on Apsara Infrastructure Management Framework whether the services in the Apsara Uni-manager Management Console cluster have reached the desired state.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the left-side navigation pane, choose **Operations** > **Cluster Operations**.
   iii. In the **Cluster** field, enter `ascm`.
   iv. Find the cluster and then click **Operations** in the Actions column.
   v. Click the **Services** tab and check the current status of services.

      If a service has not reached the desired state, set the service to the desired state. Log on to the Apsara Uni-manager Management Console again.
2. If all services in the cluster have reached the desired state and the logon is still abnormal, you can check network issues in the order of frontend page > Portal dockers > ASAPI dockers > Auth/Manager dockers".

   Troubleshooting procedure: Run the `curl` command in each container to connect the endpoint of the next service.
   i. Click the **Machines** tab.
   ii. In the **Machine** field, enter the name of the docker SR.
   iii. Find the machine and click **Terminal** in the Actions column.
   iv. In the left-side machine list, click the machine and run the following command.

   ```
   bash -c bash
   ```

   v. Run the following command to view the ID of the docker:

   ```
   docker ps
   ```

   vi. Run the following command to access the docker:

   ```
   sudo docker exec -it <Docker ID> bash
   ```

vii. Run the following command in the docker:

If an endpoint cannot be connected, contact onsite network engineers to solve network issues.

3. If no network issues are found between dockers, you must check logs to discover issues.
   i. Open the Apsara Uni-manager Management Console page, enter the account and password, and then press F12.
   ii. Click **Log On**.
   iii. Click the Headers tab and obtain `eagleeye-traceid`.
   iv. Perform the operations in step 2 to access the **ascm-portal.Portal#** docker.
   v. Run the following command in the docker:

```
grep <eagleeye-traceid>/root/logs/ascm-portal/ascm-error.log*
```

   Find the service that returns errors based on the error logs. Contact onsite O&M engineers to further discover possible causes.
4. If the Auth docker is normal, the username and password are correct, and the logon is still abnormal, the possible cause usually is that an error is reported when the AAS service is called.
   i. Run the following command in the Auth docker:

```
grep {eagleeye-traceid} /logs/ascm-logger/*.log -5
```

   ii. If a return value error is reported in the log, it can be determined that an error is reported when the AAS service is called. Contact AAS O&M engineers to further discover possible causes.

2.1.7.2. Timeouts when VPC, ECS, SLB, and RDS API operations are called

This topic describes how to troubleshoot the timeout issues in the Apsara Uni-manager Management Console when VPC, ECS, SLB, and RDS API operations are called.

Problem description

Timeouts occur in the Apsara Uni-manager Management Console when VPC, ECS, SLB, and RDS API operations are called.

Cause

This may be caused by frequent POP Full GC processes.

Solution

1. Access the `webapp-pop. PopAliyunCom#` docker. For more information about how to access the docker, see Routine inspections.
2. Run the following command in the docker:

```
tail -f/home/admin/logs/eagleeye/stat-eagleeye-jvm.log | grep gc
```

```
#tail -100f /home/admin/logs/eagleeye/stat-eagleeye-jvm.log | grep gc
2021-05-06 14:31:00|3|131,gc,ConcurrentMarkSweep|3,52876,115144,726999682
2021-05-06 14:31:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:33:00|3|131,gc,ConcurrentMarkSweep|3,53555,115147,727053237
2021-05-06 14:33:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:34:00|3|131,gc,ConcurrentMarkSweep|2,21406,115151,727105707
2021-05-06 14:34:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:36:00|3|131,gc,ConcurrentMarkSweep|2,35293,115157,727187937
2021-05-06 14:36:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:37:00|3|131,gc,ConcurrentMarkSweep|2,30999,115159,727218936
2021-05-06 14:37:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:38:00|3|131,gc,ConcurrentMarkSweep|2,32498,115161,727251434
2021-05-06 14:38:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:40:00|3|131,gc,ConcurrentMarkSweep|2,37706,115166,727333517
2021-05-06 14:40:00|3|131,gc,ParNew|0,0,646751,116449762
2021-05-06 14:42:00|3|131,gc,ConcurrentMarkSweep|2,34489,115171,727425227
2021-05-06 14:42:00|3|131,gc,ParNew|0,0,646751,116449762
```

ConcurrentMarkSweep stands for fullgc logs. You can see that a Full GC process occurs every one or two minutes.

3. Restart the dockers in sequence: POP > ASAPI > Manage > ResourceManage > OneConsole.
4. After the restart is complete, view the POP Full GC processes.

   If the issue persists, contact onsite O&M engineers.

2.1.7.3. Timeouts when you create ACK clusters in the Apsara Uni-manager Management Console

This topic describes how to troubleshoot the timeout issue when you create ACK clusters in the Apsara Uni-manager Management Console.

Problem description

An error is reported when you create an ACK cluster in the Apsara Uni-manager Management Console. The error code is "java.net.SocketTimeoutException:Read timed out".

Cause

Two possible causes:

- Network issues.
- High loads in backend services cause performance degradation.

Solution

1. Find the call logs for ASAPI and POP.
   i. Press F12. Click **Create** again.
   ii. Obtain `requestid` from the response data.

iii. Log on the Apsara Infrastructure Management Framework console and access the **asapi.ApiServer#** docker.

iv. Run the following command in the container to obtain the EagleEye ID:

```
grep <requestid> /logs/asapi-logger/*
```

v. Run the following command to query the duration on ASAPI based on the EagleEye ID:

```
grep <EagleEye ID> / logs/asapi-logger/*
```

vi. Access the **webapp-pop. PopAliyunCom#** docker and run the following command to obtain POP logs and the duration on POP:

```
grep <EagleEye ID>/opt/tianji/alidata/www/logs/java/gateway/logs/trace/*
```

2. Compare the duraiton on ASAPI with that on POP.
   - If the duration on POP is longer than that on ASAPI and the difference is small (for example, within 1s to 2s), contact product engineers to further troubleshoot the timeout issue.
   - If the duration on ASAPI is much longer than that on POP, you can check whether POP has a fullgc problem. For more information, see Timeouts when VPC, ECS, SLB, and RDS API operations are called.

2.1.7.4. Timeouts when you obtain ACK clusters in the Apsara Uni-manager Management Console

This topic describes how to troubleshoot the timeout issue when you obtain ACK clusters in the Apsara Uni-manager Management Console.

Problem description

An error is reported when you obtain ACK clusters in the Apsara Uni-manager Management Console. The error code is “java.net.SocketTimeoutExeption:Read timed out”.

Cause

Two possible causes:

- Network issues.
- High loads in backend services cause performance degradation.

Solution

1. Find the call log of the backend service.

i. Press F12. Refresh the page and obtain ACK clusters again.

On the **Headers** tab, find the `_preventCache` field.

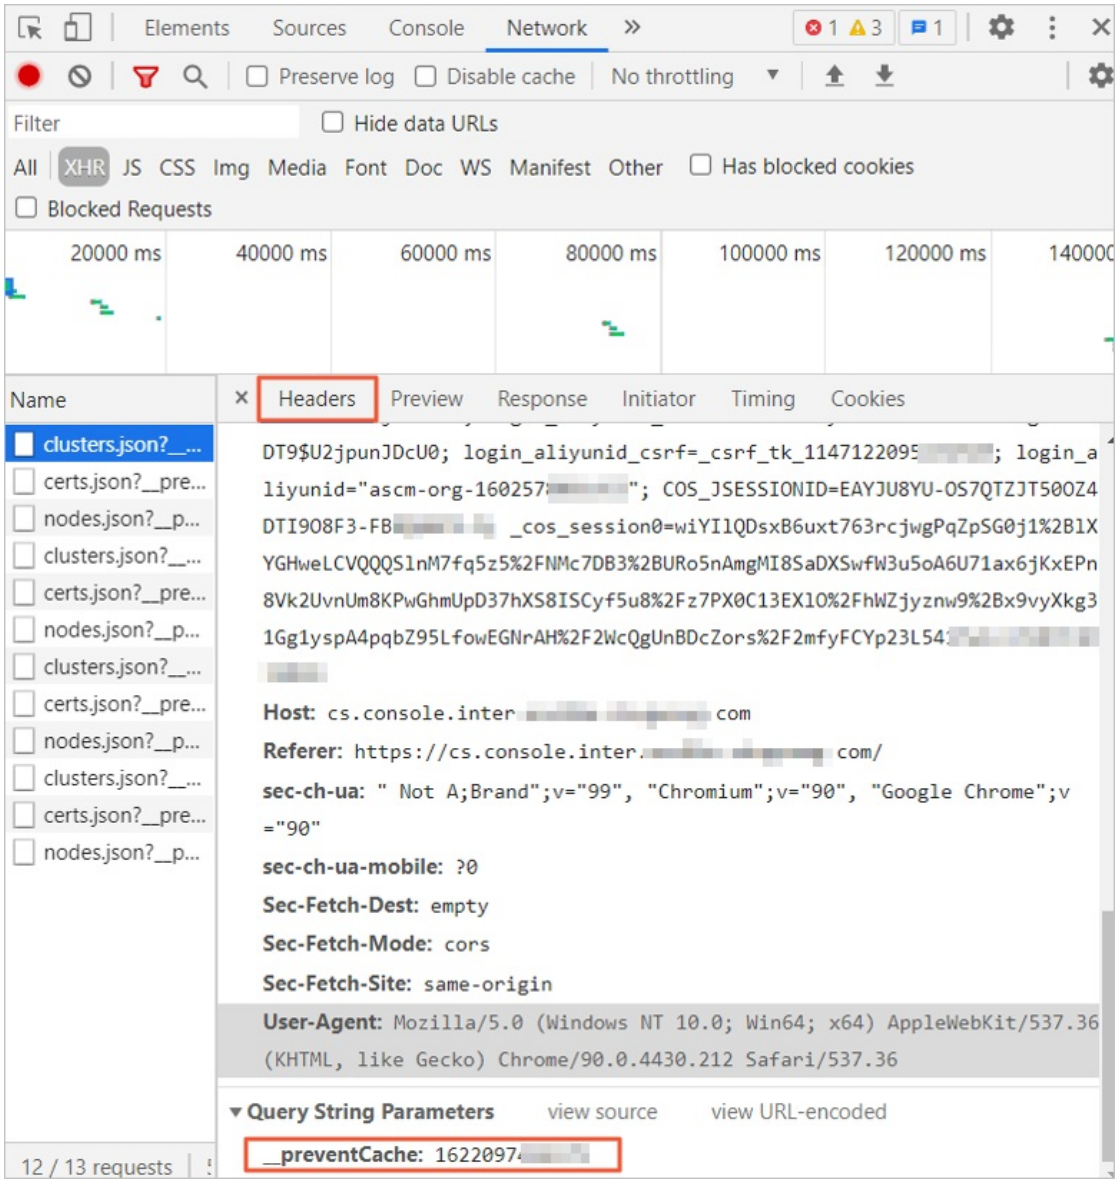

ii. Access the **ascm-brm.Manage#** docker.

iii. Run the following command to obtain the Manage logs based on `_preventCache` :

```
cat/logs/ascm-logger/main_trace* | grep <_preventCache>
```

iv. Obtain the EagleEye ID from the Manage logs. Run the following command to find the log where Manage calls ASAPI.

```
grep /logs/ascm-logger/* | grep <EagleEye ID>
```

2. Access the **asapi.ApiServer#** docker. Run the following command to obtain the log where ASAPI calls POP based on the EagleEye ID.

```
grep <EagleEye ID> logs/asapi-logger/*
```

Obtain the `x-acs-signature-nonce` value.

3. Log on the Log Service console. Select the project for POP. Obtain the log where POP calls ACK based on the `x-acs-signature-nonce` value obtained in Step 3.

   For more information about how to view logs in Log Service, see View logs.

4. Based on the preceding logs, analyze which module consume longer time.

2.1.7.5. Errors when Blink API operations are called

This topic describes how to troubleshoot the errors in the Apsara Uni-manager Management Console when Blink API operations are called.

Problem description

An error is reported in the Apsara Uni-manager Management Console when a Blink API operation is called. The error code is "ascm.manage.EntityNotExist.Instance".

Solution

1. Access the **ascm-brm.Manage#** docker. For more information, see Routine inspections.
2. Run the following command in the docker to access the BRM database:

   ```
   mysql -u${db_user} -h${db_host} -p${db_password} -D${db_name} -P${db_port}
   ```

3. Run the following command in the database to check whether the `resource_instance.no_show_count` value is 1 and whether the `instance_id` value contains uppercase letters.

   ```
   select * from resource_instance where type ='foas_project' and instance_id = '$project_name' ;
   ```

   - If the `instance_id` value contains uppercase letters, run the following command to back up the table data and then modify the uppercase letters of the `instance_id` value to lowercase letters.

     If 10 minutes later the `no_show_count` value automatically changes to 0, you can skip subsequent steps.

     ```
     create tableresource_instance_bak20210204 as select * from resource_instance;
     updateresource_instance set instance_id = lower(instance_id) where type ='foas_project';
     ```

   - If the `instance_id` value does not contain uppercase letters and the `no_show_count` value is 1, run the following command in the **ascm-brm.ResourceMgr#** docker to obtain logs:

```
cat/logs/ascm-logger/call_trace* |grep foas |grep ListProject |grep ECONNREFUSED |tail -5
cat/logs/ascm-logger/call_trace* |grep foas |grep ListProject |grepServiceUnavailable | tail -5
cat/logs/ascm-logger/call_trace* |grep foas |grep ListProject |grep'\"Code\"' | tail -5
cat/logs/ascm-logger/call_trace* |grep foas |grep ListProject |grepResponseTimeoutError | tail -5
cat/logs/ascm-logger/call_trace* |grep foas |grep ListProject |grep InternalError| tail -5
```

If no logs are returned, run the following command to obtain logs:

```
cat /logs/ascm-logger/call_trace.log|grep foas |grep ListProject| Region for grep$project | AccessKey p
air of the organization for grep$project |grep $project_name | tail -5
cat /logs/ascm-logger/call_trace.log |grep foas |grep ListProject| Region for grep$project | AccessKey
pair of the organization for grep$project | tail -5
```

4. Provide the collected logs to onsite O&M engineers for the Apsara Uni-manager Management Console.

2.1.7.6. Insufficient permissions when you perform operations in the Apsara Uni-manager Management Console

This topic describes how to troubleshoot the insufficient permissions issues when you create instances in the Apsara Uni-manager Management Console.

Problem description

When you create instances in the Apsara Uni-manager Management Console, the system prompts insufficient permissions.

> **Note**
>
> This topic describes a general method for determining permissions. Creating an instance is used in the example.

Solution

1. Press F12 and create an instance again. Obtain the `asapiErrorCode` value on the **Preview** tab of the developer tool.
2. Access the **ascm-brm.Manage#** docker.
3. Run the following command to access the BRM database:

   ```
   mysql -u${db_user} -h${db_host} -p${db_password} -D${db_name}-P${db_port}
   ```

4. Run the following command to confirm the permissions:

   ```
   select * from ascm_privilege wherecode='<asapiErrorCode>'\G
   ```

5. Check whether the permissions for the current role are selected.
   i. Log on to the Apsara Uni-manager Management Console as an administrator.

ii. In the top navigation bar, click **Enterprise**.

iii. In the left-side navigation pane, choose **Users** > Users.

iv. Find the user for which an insufficient permissions error is reported and click the username.

v. View the user role next to the user profile.

vi. In the top navigation bar, click **Enterprise**.

vii. In the left-side navigation pane, click **Roles**.

viii. Find the roles and click the role name. Check whether the permissions confirmed in Step 4 are selected.

> **Note**
>
> You can modify only permissions of custom roles. If the preset roles do not meet the requirements, create a custom role and grant permissions. For more information about how to create a role and grant permissions, see *Apsara Uni-manager Management Console User Guide*.

2.1.7.7. Product API operations have not been registered in the Apsara Uni-manager Management Console

This topic describes how to solve issues that product API operations have not been registered in the Apsara Uni-manager Management Console.

Problem description

When you call a product API operation in the Apsara Uni-manager Management Console, you are prompted that the API operation has not been registered.

Solution

1. Obtain the product name and version for the called API operation from the API call logs.

   For more information, see Audit logs.

2. Access the `asapi.ApiServer#` docker.

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

   ii. In the left-side navigation pane, choose **Operations** > **Cluster Operations**.

   iii. In the **Cluster** field, enter `ascm`.

   iv. Find the cluster and then click **Operations** in the Actions column.

   v. Click the **Machines** tab.

   vi. In the **Machine** field, enter the name of the machine.

   vii. Find the machine and click **Terminal** in the Actions column.

   viii. In the left-side machine list, click the machine and run the following command.

   ```
   bash -c bash
   ```

ix. Run the following command to view the ID of the docker:

```
docker ps
```

x. Run the following command to access the docker:

```
sudo docker exec -it <Docker ID> bash
```

3. Run the `cd /alidata/tools/` command to access the registration tools directory.

   Registration tools are stored in this directory. The tools used in the Enterprise edition are asapiTools, and the tools used in the Chinese version are asapiTools-arm64.

4. Run the following command to synchronize all API operations for the current product version that are not registered from POP to ASAPI.

   > **Notice**
   >
   > The product name in this command is case-sensitive. You must enter a correct product name.

   ```
   ./asapiTools pop -p <Product name> <Version>
   ```

   After the API operations are obtained, a list of the API operations is generated. You must confirm whether the list contains the API operation involved in the preceding issue. If not, you must confirm whether the product name and version are correct. If the API operation involved in the preceding issue still cannot be obtained, contact onsite O&M engineers.

5. After the synchronization is successful, call the product API operation again.

2.1.7.8. Timeouts when ASAPI calls API operations of cloud products

This topic describes how to troubleshoot the timeout issues in the Apsara Uni-manager Management Console when ASAPI calls API operations of cloud products.

Problem description

Timeouts occur in the Apsara Uni-manager Management Console when ASAPI calls API operations of cloud products.

Solution

1. Obtain the EagleEye ID when the timeout occurred.
   i. Press F12. Refresh the page and call the API operation again.
   ii. Click the **Headers** tab and obtain the `eagleeye-traceid`.
2. Obtain the product name and version for the called API operation from the API call logs based on the `eagleeye-traceid`.

   For more information, see Audit logs.

3. Access the `asapi.ApiServer#` docker.
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the left-side navigation pane, choose **Operations** > **Cluster Operations**.
   iii. In the **Cluster** field, enter `ascm` .
   iv. Find the cluster and then click **Operations** in the Actions column.
   v. Click the **Machines** tab.
   vi. In the **Machine** field, enter the name of the machine.
   vii. Find the machine and click **Terminal** in the Actions column.
   viii. In the left-side machine list, click the machine and run the following command.

   ```
   bash -c bash
   ```

   ix. Run the following command to view the ID of the docker:

   ```
   docker ps
   ```

   x. Run the following command to access the docker:

   ```
   sudo docker exec -it <Docker ID> bash
   ```

4. Run the following command to access the ASAPI database:

   ```
   mysql -h <db_asapi_db_host> -P<db_asapi_port> -u<db_asapi_user> -p<db_asapi_password> asapi
   ```

5. Run the following command in the ASAPI database:

   ```
   # The PolarDB-X API operation is called in the following example. Modify the timeout period based on th
   e actual time for the backend API operation.
   update ops_api set timeout = 30000 where namespace = "Drds" and name ="DescribeDrdsRdsInstances
   " and api_version = "2019-01-23";
   ```

   ```
   insert into api_timeout_unlimited(namespace,name,api_version,category_id) values ("Drds "," Describ
   eDrdsRdsInstances ","2019-01-23","privateCloud");
   ```

   > **Note**
   > If the table api_timeout_unlimited is not provided in the current version, you can skip the second part of the command.

6. After the execution is successful, wait for 5 minutes. ASAPI calls the API operation again and check the call results.

   If the problem persists, contact Alibaba Cloud O&M engineers.

3.Network operations

3.1. Apsara Network Intelligence

3.1.1. 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.

3.1.2. Log on to the Apsara Network Intelligence console

This topic describes how to log on to the Apsara Network Intelligence console.

Prerequisites

- You must log on to the the Apsara Uni-manager Operations Console to access Apsara Network Intelligence. Before you start, you must obtain the URL, username, and password of the Apsara Uni-manager Operations Console from the engineer that deploys the service.

  The URL of the Apsara Uni-manager Operations Console is in the `ops.asconsole.intranet-domain-id.com` format.
- We recommend that you use the Google Chrome browser.

Procedure

1. Open your browser.
2. In the address bar, enter the URL of the Apsara Uni-manager Operations Console: `ops.asconsole.intranet-domain-id.com` and press enter.
3. Enter your username and password.

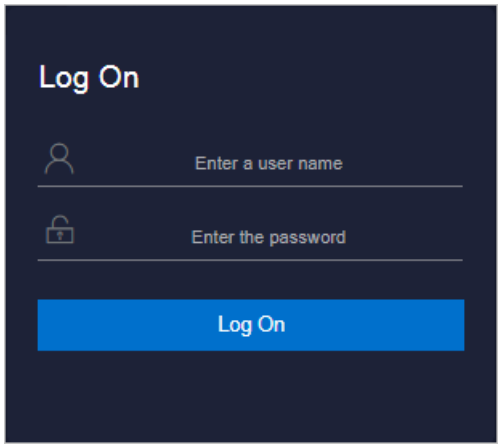

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

If this is the first time you log on to the Apsara Uni-manager Operations Console, you must change your password as prompted.

For higher security, make sure that the password meets the following requirements:

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- The password is 10 to 20 characters in length.

4. Click **Log On** to go to the **Apsara Uni-manager Operations Console** homepage.
5. In the top navigation bar, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Product List**.
7. In the **Infrastructure as a Service (IaaS)** section, click **Apsara Network Intelligence**.

3.1.3. Query information about a network instance

You can query details about a network instance by specifying its ID.

Procedure

1. Log on to the Apsara Network Intelligence console.
2. Enter the ID of a virtual private cloud (VPC) or a Server Load Balancer (SLB) instance to query details.
   - Enter the ID of a VPC to query details about the VPC, vRouters, and vSwitches.

- VPC details

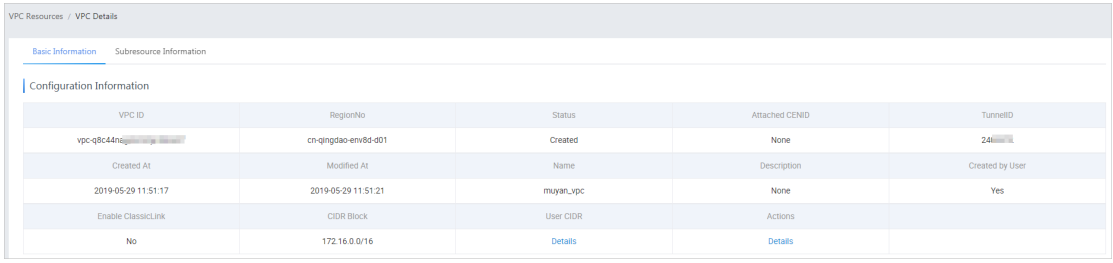

- Information about vRouters, route tables, router interfaces, and vSwitches

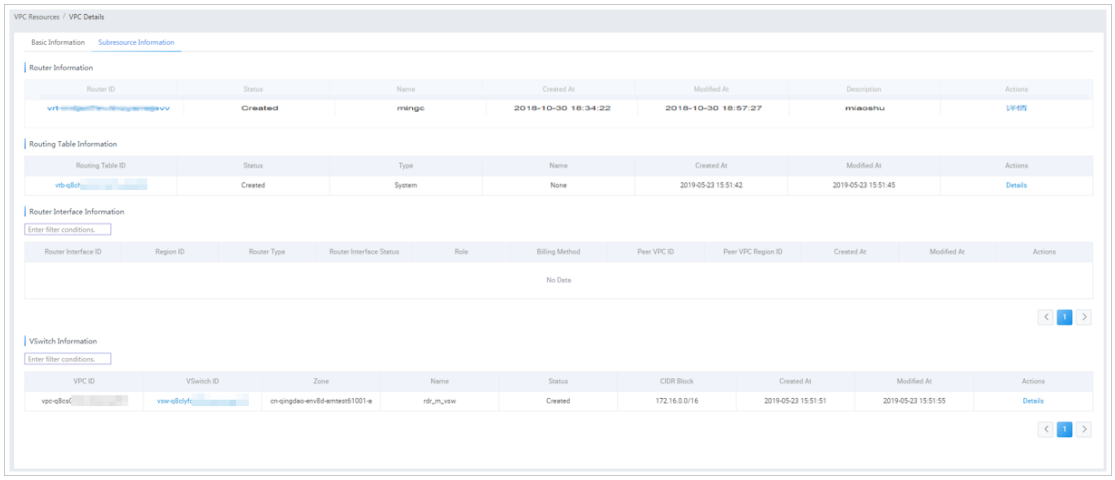

- Enter the ID of an SLB instance to query instance details.
  - Information about SLB instance configurations, virtual IP addresses (VIPs), specifications, and users

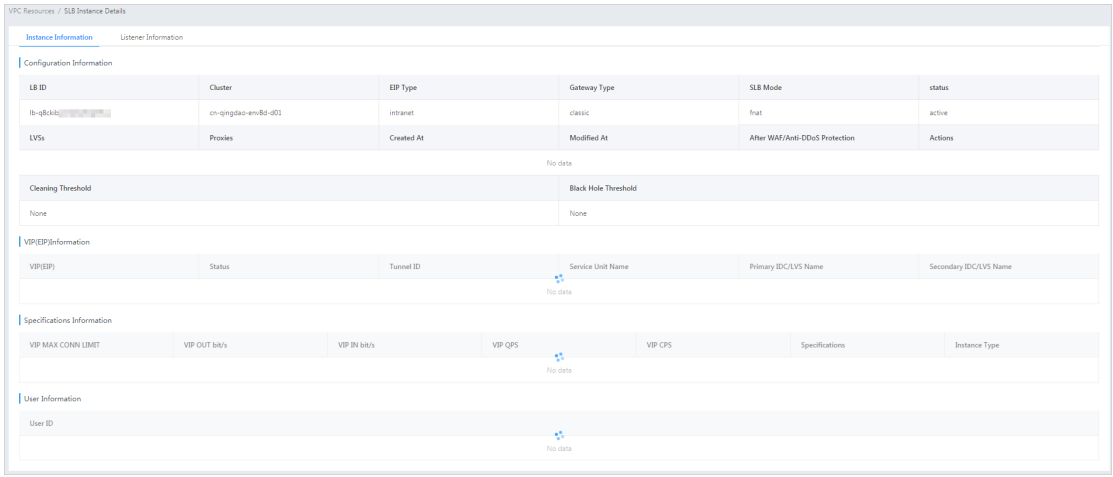

  - Information about listeners

    Click **Show** in the **Back-end Server/Health Check** column to view details about backend servers.

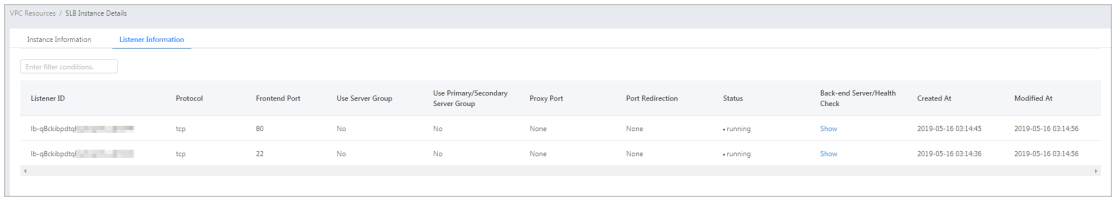

3.1.4. Manage cloud service instances

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.

3.1.5. Tunnel VIP

3.1.5.1. 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.

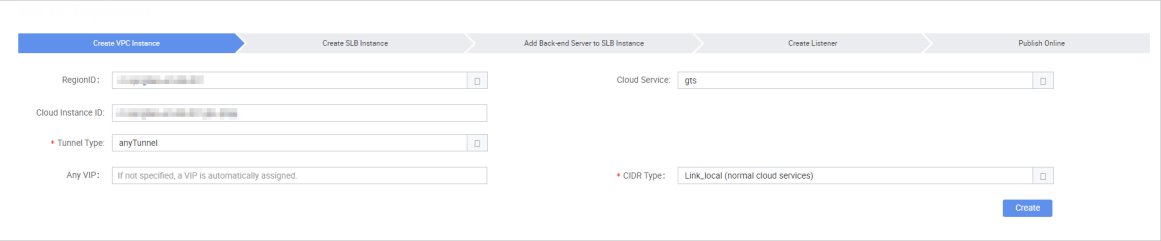

The following tunnel types are available:

- **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.

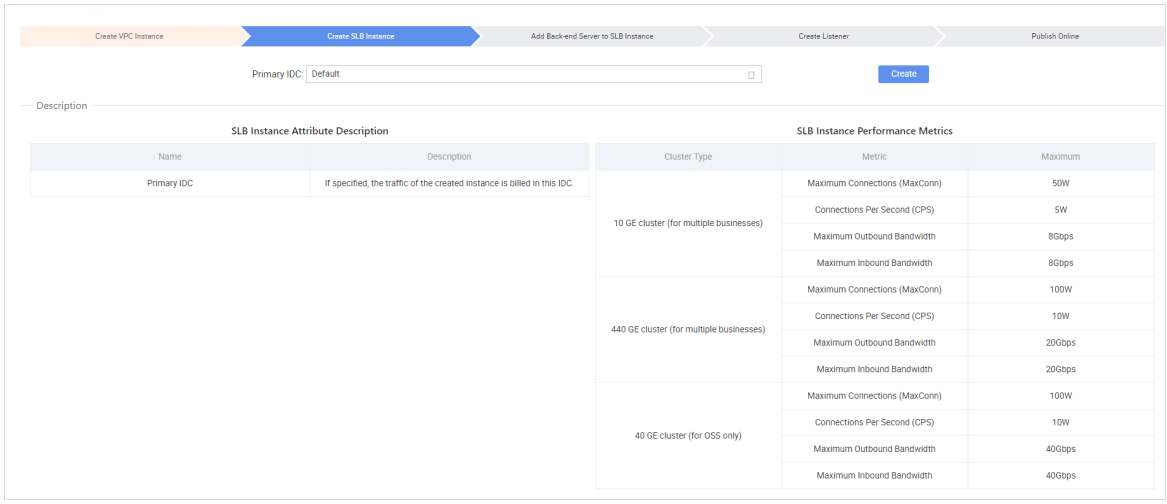

7. Click **Create**.
8. In the **Add Back-end Server to SLB Instance** step, specify the following information:
   - **VPC ID**: Enter the ID of the VPC to which target ECS instances belong. This parameter must be set if the target ECS instances are deployed in a VPC.
   - **Back-end Servers**: Specify 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.

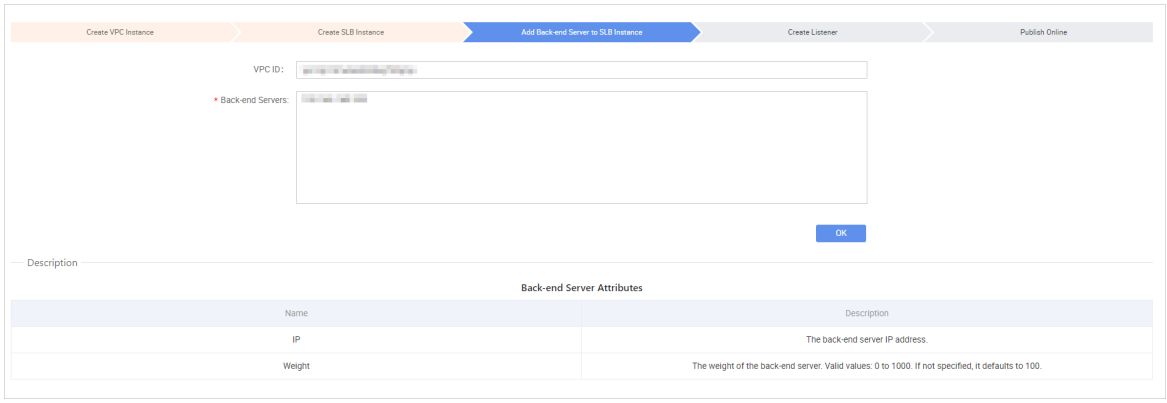

9. Click **OK**.
10. In the **Create Listener** step, click **Add** to configure a User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) listener.

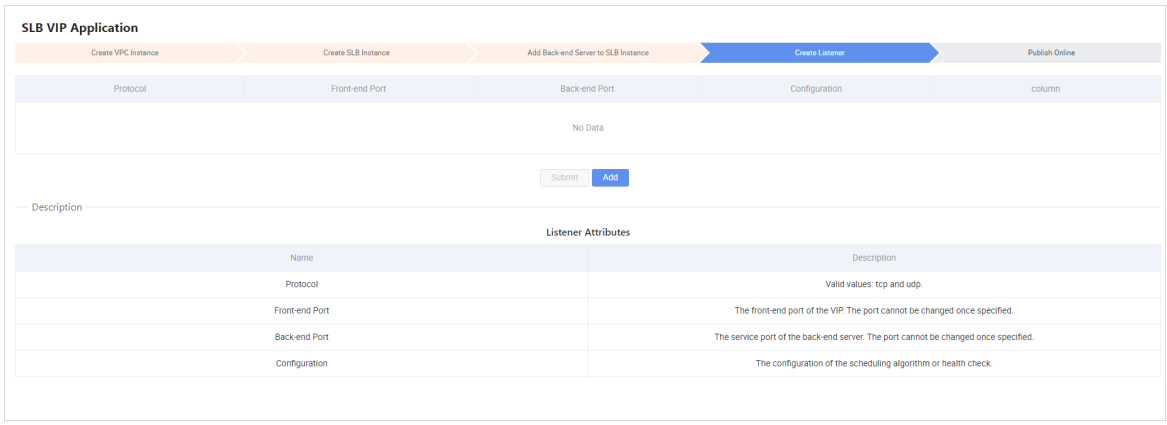

11. Click **Submit**.
12. In the **Publish Online** step, click **Yes** and then click **OK**.

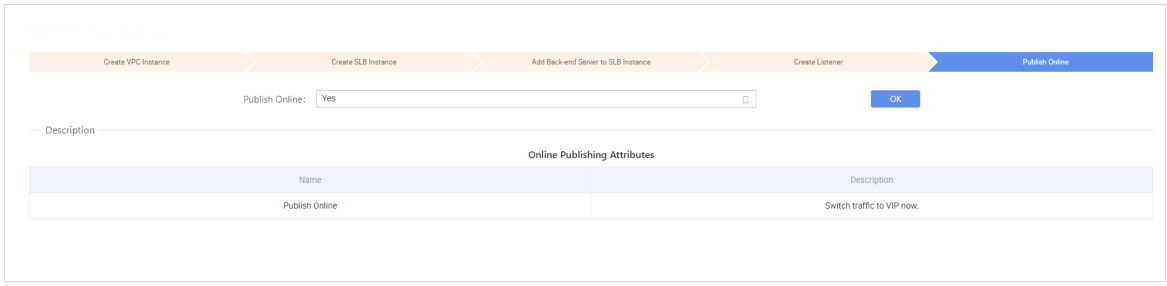

Result

The cloud services for which you have created the VIP can forward traffic through the created Layer-4 listener.

3.1.5.2. 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**.

3.1.6. 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.

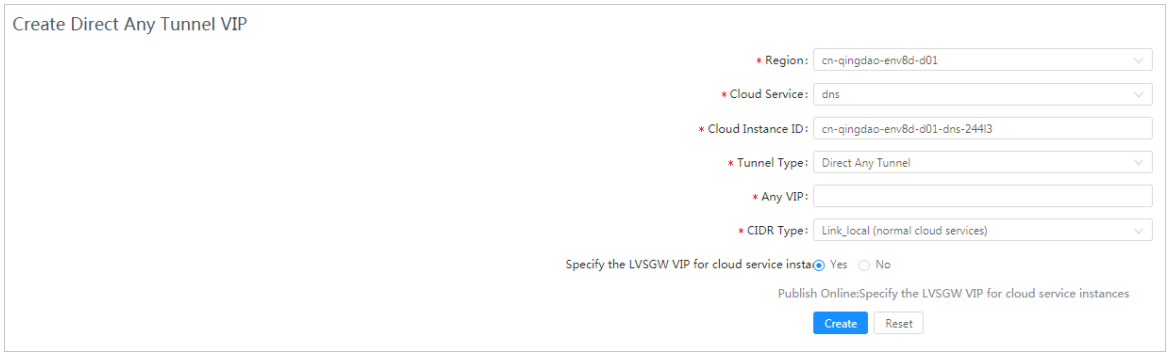

5. Click **Create**. Cloud service instances that have Direct Any Tunnel VIPs can forward traffic through XGW.

3.1.7. Leased line connection

3.1.7.1. 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.

3.1.7.2. 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**.

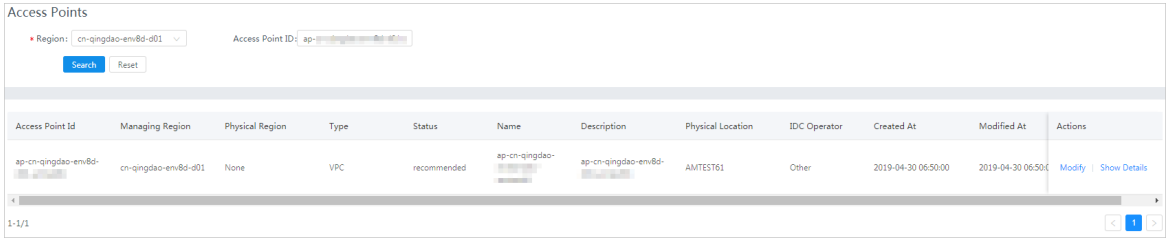

Modify access point information

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.

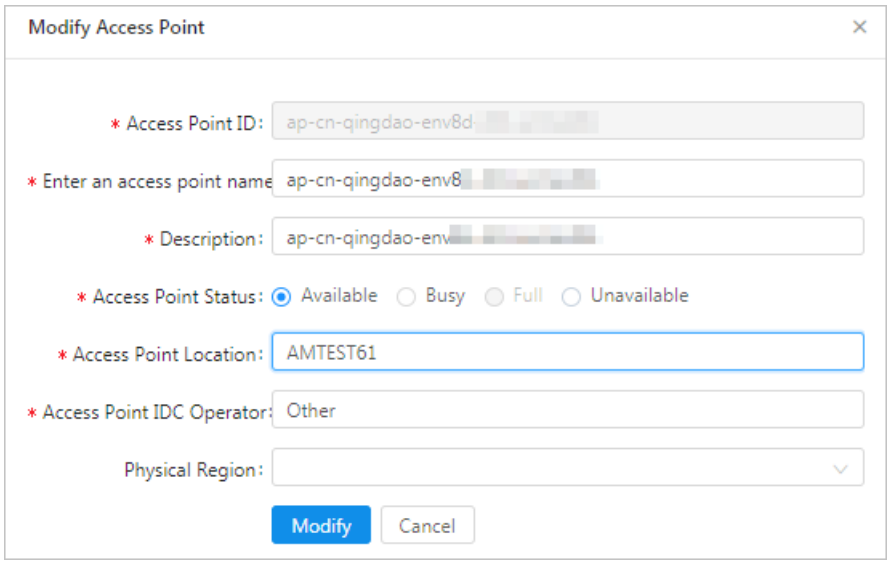

3.1.7.3. Manage access devices

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**.

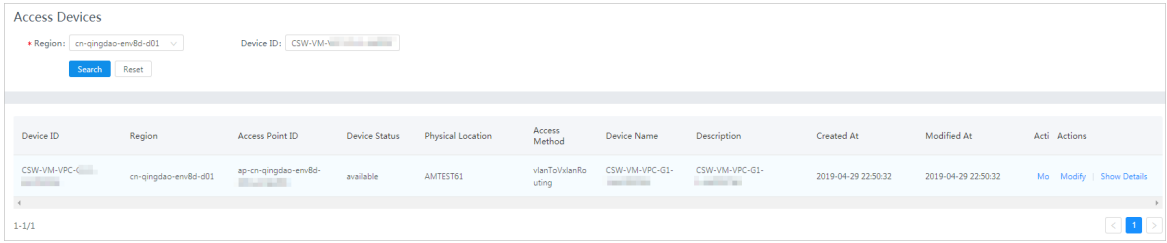

6. Find the target access device and click **Show Details** in the **Actions** column to view details of the access device.

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.

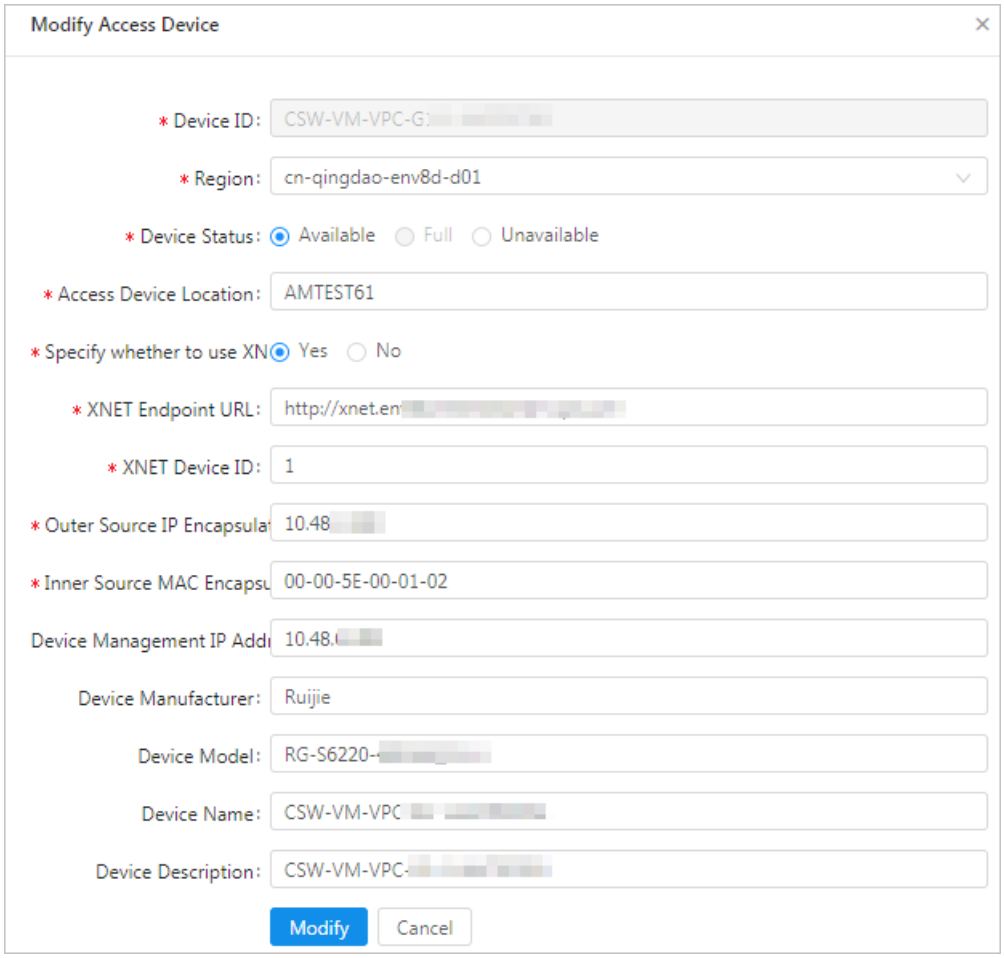

3. Click **Modify**.

3.1.7.4. 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.

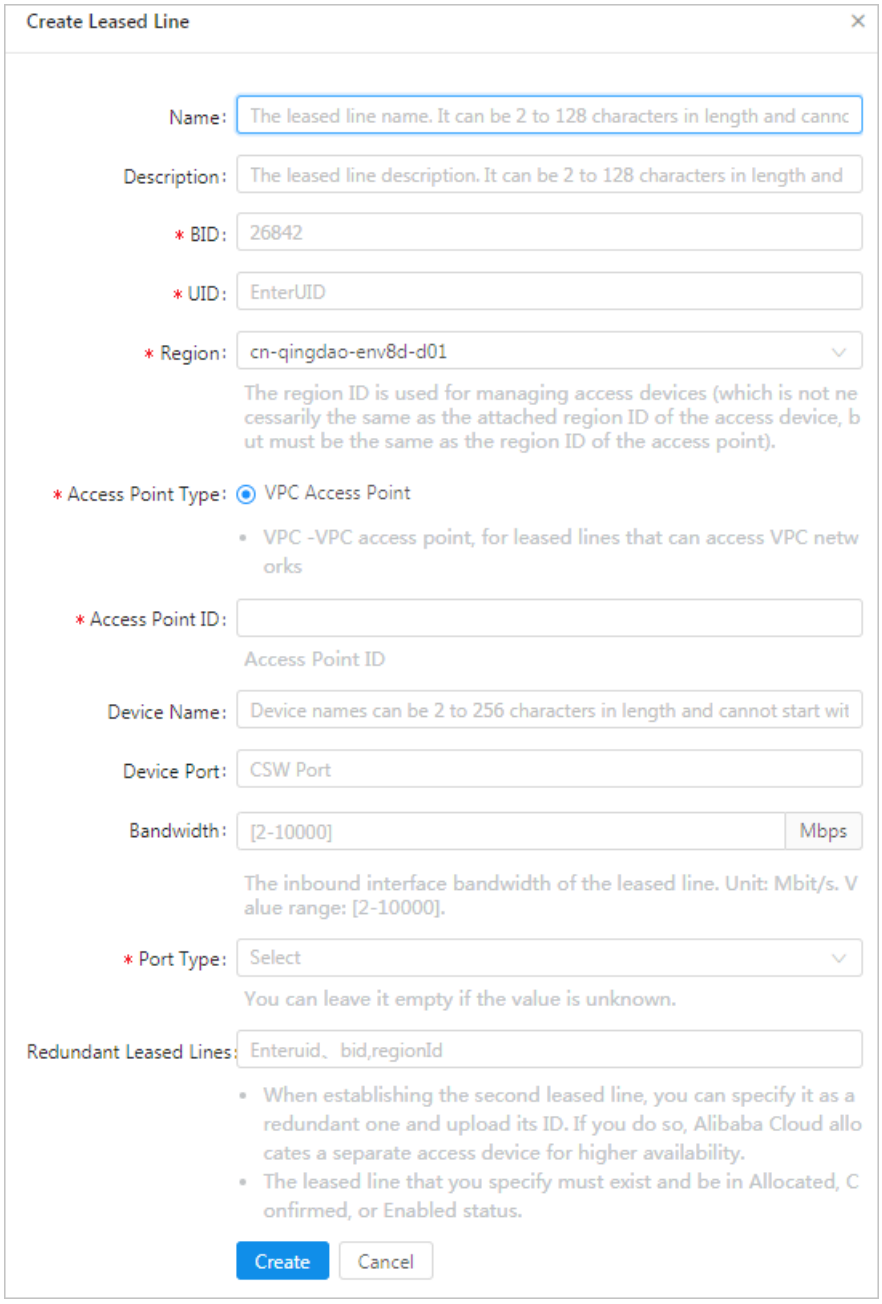

When the leased line state is **Confirmed**, the line is created.

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.

3.1.7.5. 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**.

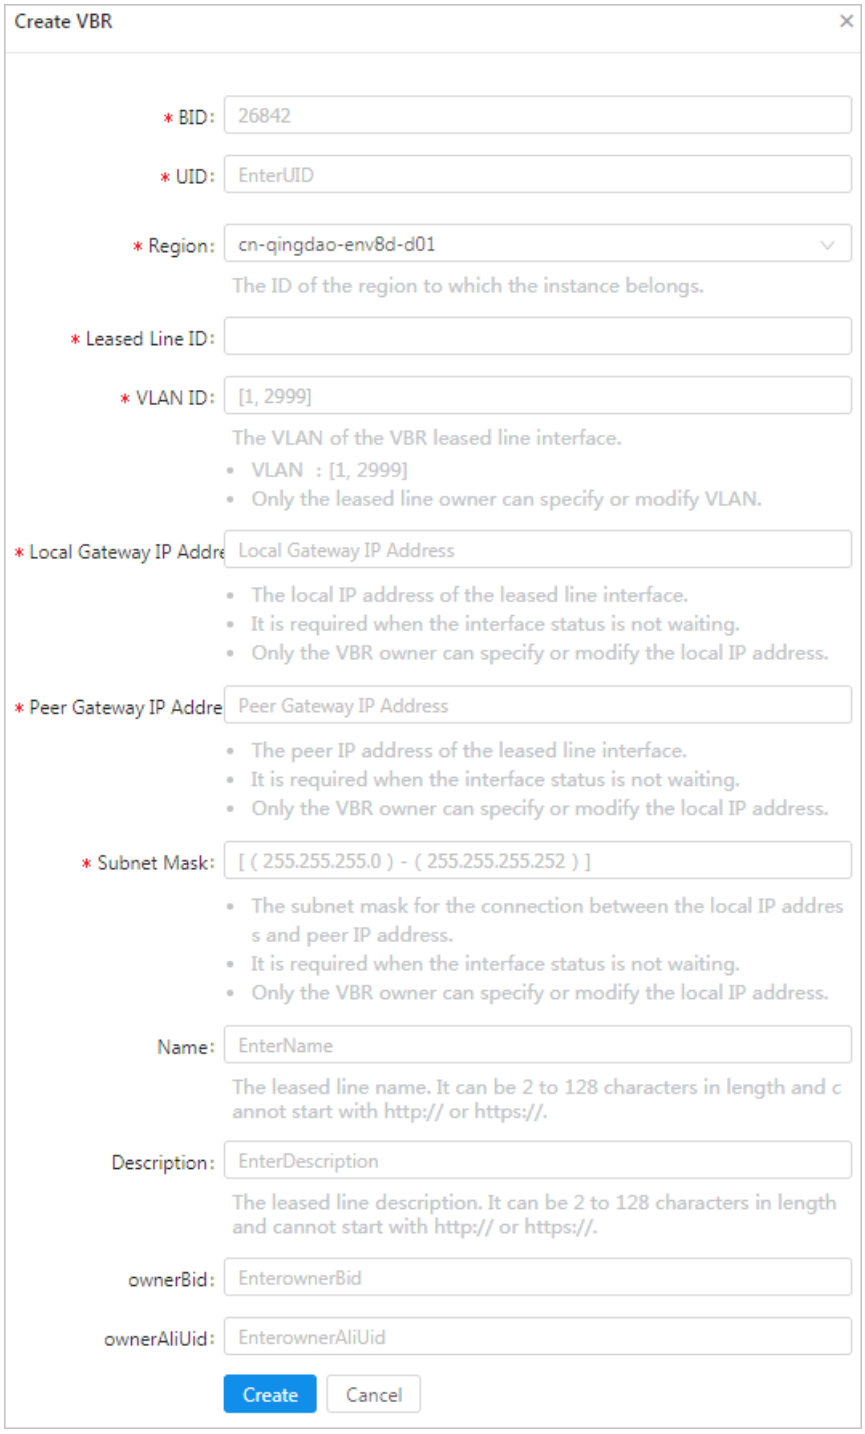

5. Follow the on-screen prompts to configure the VBR parameters.

   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.

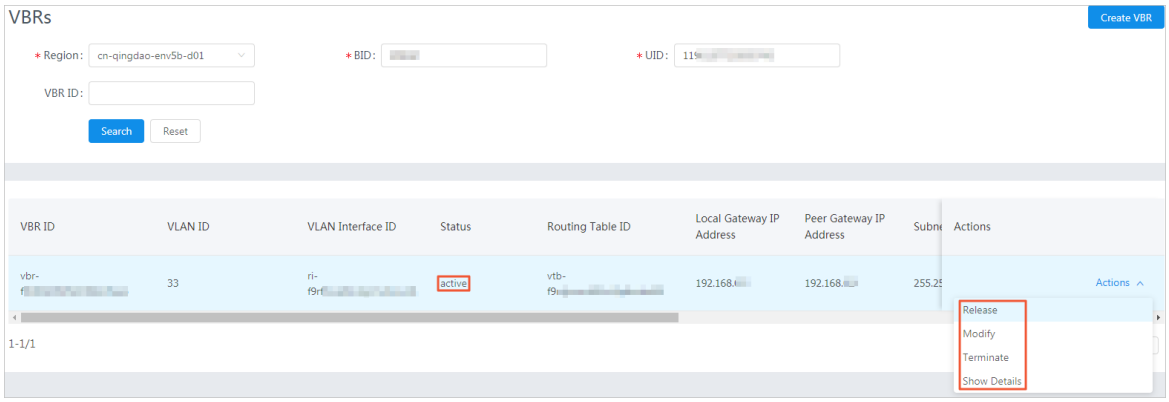

   You can click **Release**, **Modify**, **Terminate**, or **Show Details** in the **Actions** column to manage a VBR.

3.1.7.6. 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.

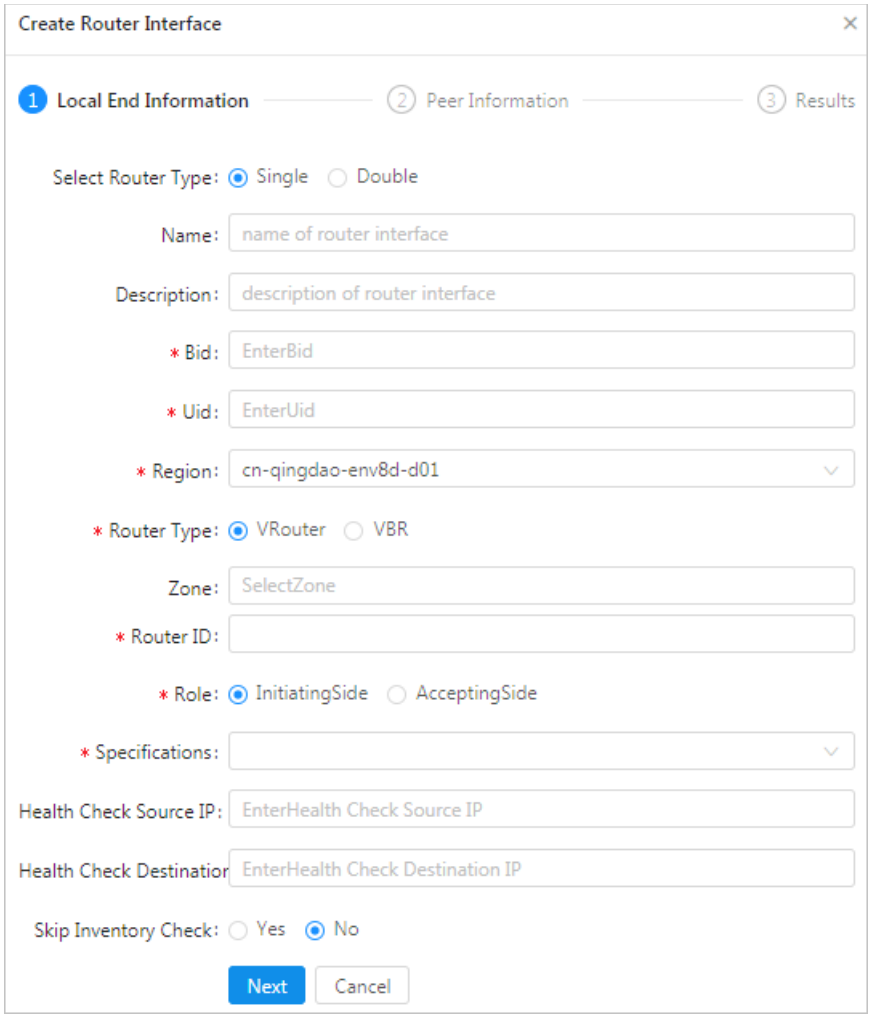

When the router interface state is **Active**, the interface is created.

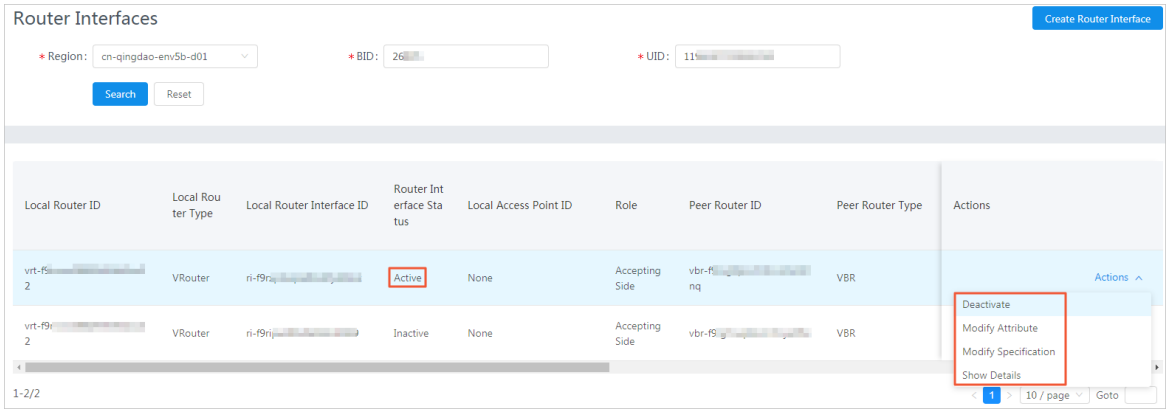

3.1.7.7. Create a routing table

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

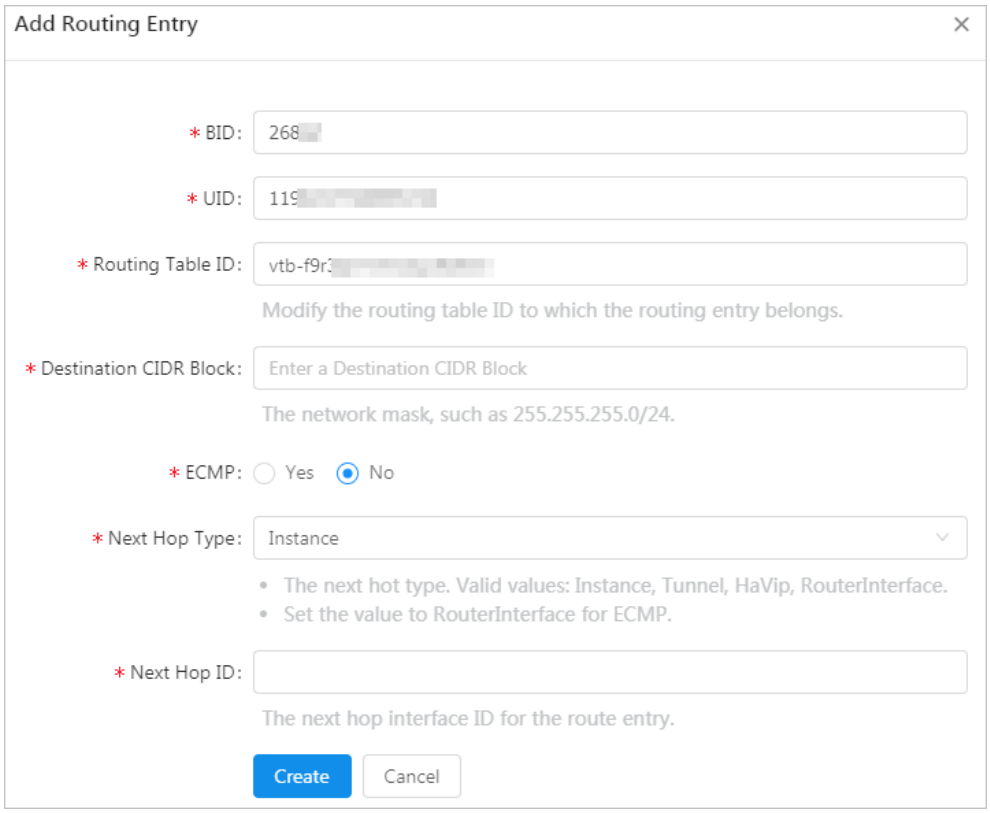

vii. Repeat the preceding steps to add a route destined for a target IDC.

> **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.

3.1.8. 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

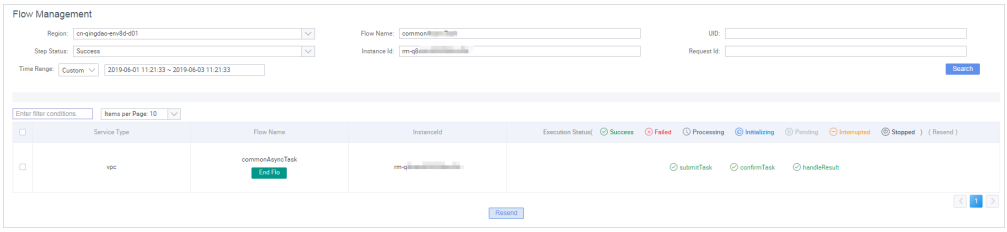

3.1.9. Configure reverse access to cloud services

Cloud services cannot be directly accessed from external networks. You must configure reverse access to allow external networks to access cloud services.

Prerequisites

Log on to the Apsara Uni-manager Operations Console, and obtain the **AccessKey ID** and **AccessKey secret** on the **Personal Information** page.

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. Enter the **AccessKey ID** and **AccessKey secret** and click **OK**.
4. On the **Cloud Service Reverse Access** page, click **Create Cloud Service Reverse Access**.
5. On the **Allocate Cloud Service ID** wizard page, select a region and enter a name and description.
6. Click **Continue**. The following information is automatically created and displayed on the **Create Address Pool** wizard page: the application ID of the cloud service that allows reverse access and the address pool that is used for reverse access to the cloud service.
7. Click **Continue**. On the **Add Server Address** wizard page, configure the Elastic Compute Service (ECS) instance to be used for reverse access.
   - **VPC ID**: Enter the ID of a virtual private cloud (VPC) and 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**. On the **Create Mapping IP** wizard page, enter the ID of a vSwitch and the mapping IP address of the ECS instance. The IP address is used to allow access from external networks.
9. Click **Continue**. On the **Complete Authorization** wizard page, specify VPC ID, Server IP, and Instance Port for reverse access.

   Separate multiple port numbers with commas (,). For example, `10,20,30`. You can specify up to 10 ports.

4.Operations of basic cloud products

4.1. Elastic Compute Service (ECS)

4.1.1. Operations and Maintenance Guide

4.1.1.1. 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

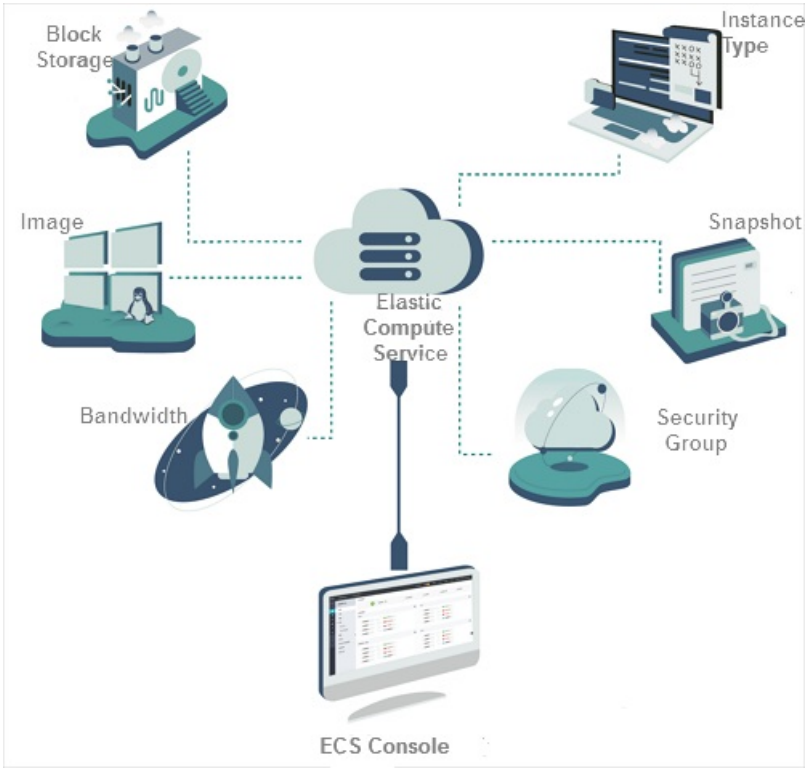

4.1.1.2. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager Operations Console.

Prerequisites

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

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

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

Procedure

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

   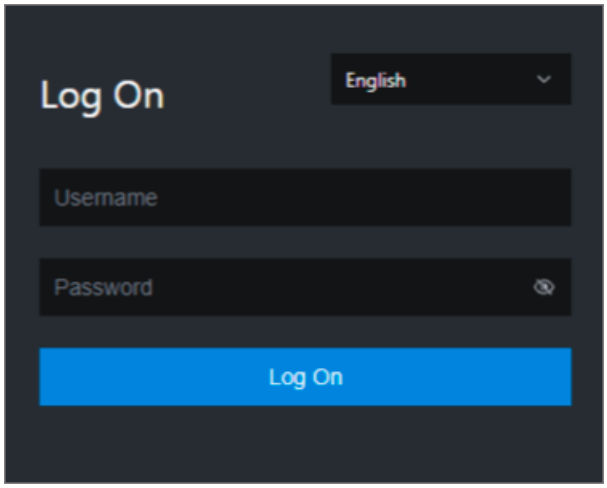

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.

4. Click **Log On**.

4.1.1.3. ECS operations and maintenance

4.1.1.3.1. 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.

4.1.1.3.2. VM

4.1.1.3.2.1. 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.

4.1.1.3.2.2. Query VMs

In the ECS Operations and Maintenance Platform, you can view the list of existing VMs and their information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.
6. In the VM list, click a VM ID. You can view the information of the VM in the **VM Details** panel.

4.1.1.3.2.3. Start a VM

In the ECS Operations and Maintenance Platform, you can start a VM in the same manner as you start a real server.

Prerequisites

The VM to start is in the **Stopped** state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.

6. In the VM list, select the VM to start. 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**.

4.1.1.3.2.4. Stop a VM

In the ECS Operations and Maintenance Platform, you can stop a VM in the same manner as you stop a real server.

Prerequisites

The VM to stop is in the **Running** state.

Context

This operation may interrupt the programs running on the VM. Perform this operation during off-peak hours to minimize the impact on services.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.
6. In the VM list, select the VM to stop. 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**.

4.1.1.3.2.5. Restart a VM

In the ECS Operations and Maintenance Platform, you can restart a VM in the same manner as you restart a real server.

Prerequisites

The VM to restart is in the **Running** state.

Context

This operation may interrupt the programs running on the VM. Perform this operation during off-peak hours to minimize the impact on services.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.
6. In the VM list, select the VM to restart. 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**.

4.1.1.3.2.6. Cold migration

In the ECS Operations and Maintenance Platform, 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 cannot be performed.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.
6. In the VM list, select the VM to migrate. 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.

Parameter	Description
Switchover Policy	The switchover policy. Valid values: ○ **Force Migrate** ○ **Active Migrate**
Start	The startup mode. Valid values: ○ **Normal** ○ **Repair**
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**.

4.1.1.3.2.7. Hot migration

In the ECS Operations and Maintenance Platform, 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, hot migration cannot be performed and you must perform cold migration instead. For more information, see Cold migration.
- Security risks may arise if you perform hot migration. Exercise caution when you perform hot migration.
- Hot migration does not interrupt services running on the VM.
- Hot migration can be performed only within the same zone. Cross-zone hot migration cannot be performed.

Prerequisites

You can perform hot migration only on VMs in the **Running** state.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.

6. In the VM list, select the VM to migrate. Choose **More > Online Migrate** above the list.
7. Set Throughput Limit.

   The value of Throughput Limit can range from 1 to 1000. Unit: MByte/s. Default value: 20.
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.

4.1.1.3.2.8. Reset a disk

In the ECS Operations and Maintenance Platform, you can reset disks to restore them to their initial status.

Prerequisites

- When you reset a disk, installed applications are cleared from the disk. 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 the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **VMs** tab.
5. On the VMs tab, set the filter conditions and click **View**.
6. In the VM list, select the VM to which the disk that you want to reset is attached. Choose **More > Reset Disk** above the list.
7. In the Reset Disk dialog box, select the disk that you want to reset and set Operation Reason. Click **OK**.

4.1.1.3.3. Disks

4.1.1.3.3.1. Overview

In an ECS instance, cloud disks can be considered as physical disks. You can mount, detach, and create snapshots for disks.

4.1.1.3.3.2. Query disks

In the ECS Operations and Maintenance Platform, you can view the list of existing disks and their information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set the filter conditions and click **View**.

4.1.1.3.3.3. View snapshots

In the ECS Operations and Maintenance Platform, you can view the list of snapshots created for a disk and their information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set the filter 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.

4.1.1.3.3.4. Attach a disk

After a disk is created, you can attach the disk to a VM.

Context

Only disks in the **Available** state can be attached.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set the filter conditions and click **View**.
6. Find the disk to attach and choose > **Mount**.
7. In the Mount Disk dialog box, set VM ID and Operation Reason. Click **OK**.

4.1.1.3.3.5. Detach a disk

In the ECS Operations and Maintenance Platform, only data disks can be detached. System disks and local disks cannot be detached.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set the filter conditions and click **View**.
6. Find the disk to detach and choose  **> Detach**.
7. In the Detach Disk dialog box, set Operation Reason. Click **OK**.

4.1.1.3.3.6. Create a snapshot

In the ECS Operations and Maintenance Platform, you can manually create snapshots for disks.

Context

Snapshots can be created only for system disks.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Disks** tab.
5. On the Disks tab, set the filter 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**.

4.1.1.3.4. Snapshots

4.1.1.3.4.1. 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.

4.1.1.3.4.2. Query snapshots

In the ECS Operations and Maintenance Platform, 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 the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Snapshots** tab.
5. On the Snapshots tab, set the filter conditions and click **View**.

   AliUid is a required filter condition.

4.1.1.3.4.3. Delete a snapshot

In the ECS Operations and Maintenance Platform, 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 the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Snapshots** tab.
5. On the Snapshots tab, set the filter conditions and click **View**.

   AliUid is a required filter condition.
6. Find the snapshot that you want to delete and choose > **Delete**.
7. Set Operation Reason. Click **OK**.

4.1.1.3.4.4. Create an image

In the ECS Operations and Maintenance Platform, you can create a custom image from a snapshot. The image contains 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 the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Snapshots** tab.
5. On the Snapshots tab, set the filter conditions and click **View**.

   AliUid is a required filter 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 uses a public image or a custom image. Click **OK**.

4.1.1.3.5. Images

4.1.1.3.5.1. 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.

4.1.1.3.5.2. Query images

In the ECS Operations and Maintenance Platform, you can view the list of existing images and their information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Images** tab.
5. On the Images tab, set the filter conditions and click **View**.

> **Note** If you set Image Type to Custom Image, you must also set AliUid.

4.1.1.3.6. Security groups

4.1.1.3.6.1. 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.

4.1.1.3.6.2. Query security groups

In the ECS Operations and Maintenance Platform, you can view the list of existing 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 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 the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Security Groups** tab.
5. On the Security Groups tab, set the filter conditions and click **View**.

4.1.1.3.6.3. Add security group rules

You can add rules to security groups to control access to or from instances in the security groups.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Security Groups** tab.
5. On the Security Groups tab, set the filter 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 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.
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**.

4.1.1.3.7. Custom instance types

4.1.1.3.7.1. Add custom instance types

When existing instance types do not meet your business requirements, you can add custom instance types in the ECS Operations and Maintenance Platform and create instances of the custom instance types.

Context

Custom instance types are assigned to the ecs.anyshare instance family. When you add a custom instance type, you can configure the Instance Type, vCPUs, and Mem (GiB) parameters. The values of other parameters such as Base Bandwidth, Packet Forwarding Rate (10,000 PPS), and NIC Queue are automatically generated.

> **Note**
>
> - All custom instance types are shared instance types.
> - Custom instance types support only the Intel x86 chip architecture.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Custom Instance Type** tab.
5. Click **Add**.
6. In the **Add Instance Type** panel, configure the Instance Type, vCPUs, and Mem (GiB) parameters.
7. 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 select ecs.anyshare as the instance family and create instances of the instance type. For more information, see the "Create an instance" topic in *ECS User Guide*.

4.1.1.3.7.2. Query custom instance types

In the ECS Operations and Maintenance Platform, you can view the custom instance types that you have added and their information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > ECS Operations and Maintenance Platform**.
4. Click the **Custom Instance Type** tab.
5. View the information about the custom instance types. If the custom instance type list does not

automatically refresh, click **Search**.

4.1.1.3.7.3. Modify custom instance types

If you want to retain a custom instance type but the specifications of this instance type do not meet your requirements, you can modify the number of vCPUs and memory size of the instance type.

Prerequisites

A custom instance type is added.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > 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** panel, set the vCPUs and Mem (GiB) parameters.
7. Click **OK**.

4.1.1.3.7.4. Delete custom instance types

In the ECS Operations and Maintenance Platform, 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 the Apsara Uni-manager Operations Console.
2. In the top navigation bar, select an environment version and a region.
3. In the top navigation bar, click **O&M**. Then, the in the left-side navigation pane, choose **Product Management > 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.

4.1.1.4. Apsara Distributed File System Management

4.1.1.4.1. View ECS disk size rankings

The ECS Disk Size Ranking module allows you to view the amount of space occupied by all disks in the elastic block storage attached to an Elastic Compute Service (ECS) cluster in Apsara Distributed File System.

Context

When an ECS cluster occupies a large amount of space in Apsara Distributed File System, the on-site O&M personnel must check the space occupied by each disk in the elastic block storage attached to the ECS cluster. Then, they must contact the business side to migrate data and release disks. The ECS disk size ranking feature helps O&M personnel identify which disks occupy a large amount of space in Apsara Distributed File System so that they can perform targeted cleanup and quickly lower space usage.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > EBS > ECS Disk Size Ranking**.
4. Select the ECS cluster that you want to query from the **Cluster** drop-down list and click **Search**.

   All disks in the elastic block storage attached to the selected ECS cluster are listed from large to small based on the actual amount of space that the disks occupy in Apsara Distributed File System. You can view the cluster name, cluster ID, and zone of the selected cluster. You can also view the storage type, size, and identifier of each disk.
5. (Optional) You can click **Reset** to clear the preceding search conditions.

4.1.1.4.2. EBS dashboard

The EBS Dashboard module allows you to view the overview information and cluster usage trend charts of EBS clusters.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > EBS Dashboard**.

   On the page that appears, cluster overview information and cluster usage trend charts of all EBS clusters are displayed.
4. Select a cluster from the **Cluster Name** drop-down list.
5. 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 Disks**, **Abnormal Masters**, **Deleting Status**, **Abnormal Block GcWorkers**, or **Abnormal Block Servers** is greater than 0, the corresponding value is displayed in red.
   - The **Trend Chart of Cluster Usage** section shows the storage usage curve of the cluster for the last 30 days.

4.1.1.4.3. Block master operations

The Block Master Operations module shows the block master node information of Elastic Block Storage (EBS) clusters, including the IP addresses and roles. The module also allows you to switch the role of a node to leader as well as query and configure flags.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > 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.
4. Select a cluster from the **Cluster Name** drop-down list.
5. In the **Master List** section, 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.

     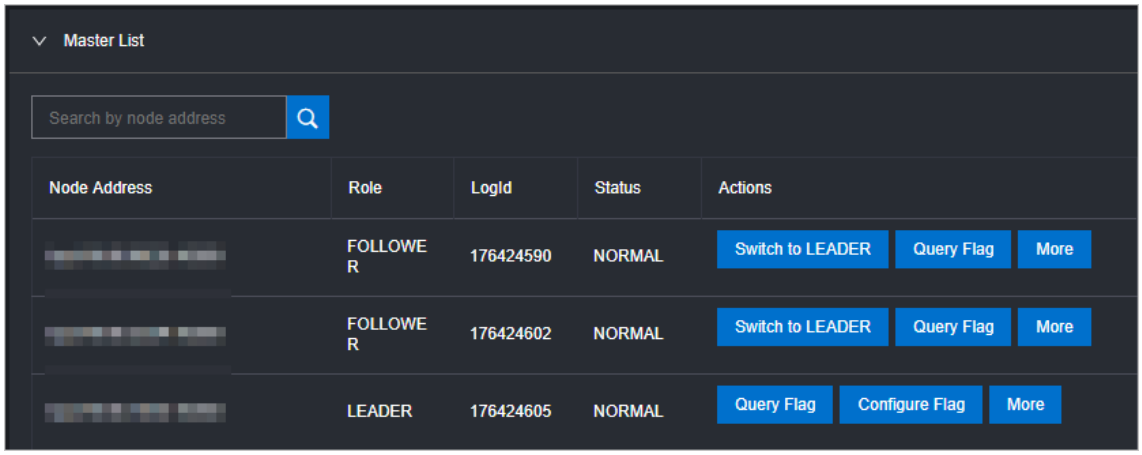

   - Switch to leader

     A leader role for a master node has the same features 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 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 list, find a node that assumes the leader role and click **Configure Flag** in the **Actions** column. In the dialog box that appears, configure the parameters and click **OK**.

The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_blockmaster_flag.json* file.
flag_value	The custom flag value.
flag_type	The 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.

6. In the **Cluster Overview** section, you can query the disk size, number of segments, total storage size, and storage usage of the cluster.

4.1.1.4.4. Block server operations

The Block Server Operations module shows the block server node information of Elastic Block Storage (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. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.

3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Block Master Operations**.

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

4. Select a cluster from the **Cluster Name** drop-down list.
5. In the **Server List** section, perform the following operations:
   - View the server node list

     You can view the 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 list, find a node and click **Query Flag** in the **Actions** column. In the dialog box that appears, set flag_key and 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. Click the **Configure** tab.
     e. 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 list, find a node and click **Configure Flag** in the **Actions** column. In the dialog box that appears, set flag_key and flag_value, select flag_type, and then click **OK**.

     The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_blockserver_flag.json* file.
flag_value	The custom flag value.
flag_type	The flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the server node status

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

  The following table describes the server node status.

Status	Description
NORMAL	The node is running normally.
DISCONNECTED	The node is disconnected.
OFFLOADING	The node is being disabled.
OFFLOADED	The node is disabled.
UPGRADE	The node is upgraded.
RECOVERY	The node is restored.

- Query the version information

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

6. In the **Block Server Blacklist** section, perform the following operations:
   - 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 click **OK**.

     The block server node that is added to the blacklist is disabled and no longer provides services.
   - View the block server blacklist

     In the **Block Server Blacklist** section, you can view all block server nodes that are added to the blacklist.
   - 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 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 services.

4.1.1.4.5. Snapshot server operations

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.

3. In the left-side navigation pane, choose **Apsara Distributed File System Management > SnapShotServer**.

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

4. Select a cluster from the **Cluster Name** drop-down list.
5. Perform the following operations:
   - View the snapshot server node list

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

     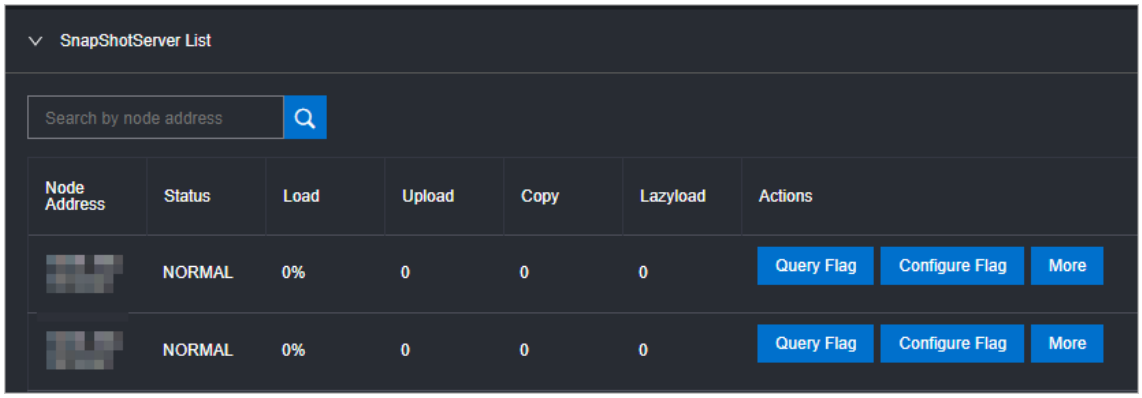

   - Query a flag

     In the snapshot server node list, find a node and click **Query Flag** in the **Actions** column. In the dialog box that appears, set flag_key and 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. Click the **Configure** tab.

     e. 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, find a node and click **Configure Flag** in the **Actions** column. In the dialog box that appears, set flag_key, flag_value, and flag_type, and click **OK**.

     The following table describes the parameters.

Parameter	Description

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_snapshotserver_flag.json* file.
flag_value	The custom flag value.
flag_type	The flag type. Valid values: ■ **int** ■ **bool** ■ **string** ■ **double**

- Configure the snapshot server node status

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

  The following table describes the snapshot server node status.

Status	Description
NORMAL	The node is running normally.
DISCONNECTED	The node is disconnected.
OFFLOADING	The node is being disabled.
OFFLOADED	The node is disabled.

- Query the version information

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

4.1.1.4.6. Block gcworker operations

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Block Gcworker Operations**.

   On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list

is displayed.

4. Select a cluster from the **Cluster Name** drop-down list.
5. Perform the following operations:
   - View the gcworker node list

     You can view the IP addresses and status of the block gcworker nodes in the selected cluster.

     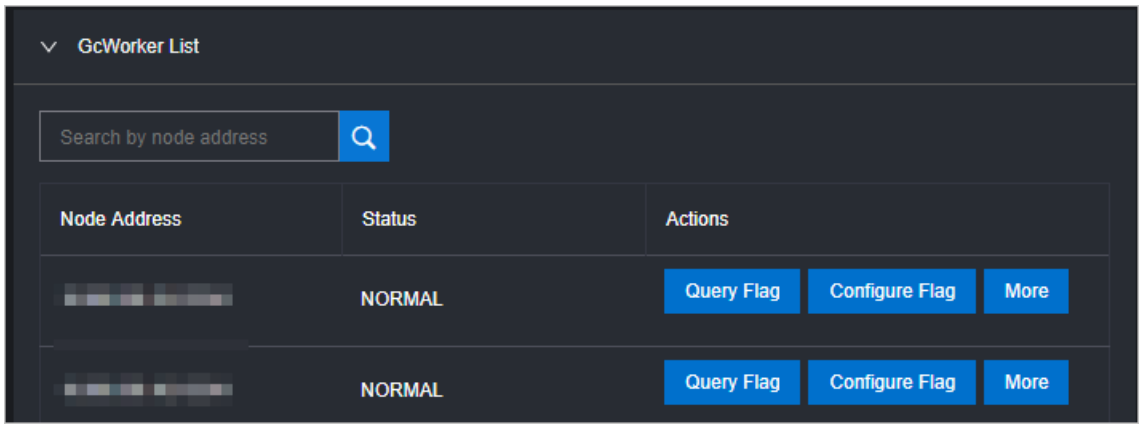

   - Query a flag

     In the gcworker node list, find a node and click **Query Flag** in the **Actions** column. In the dialog box that appears, set flag_key and 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, find a node and click **Configure Flag** in the **Actions** column. In the dialog box that appears, set flag_key, flag_value, and flag_type, and click **OK**.

     The following table describes the parameters.

Parameter	Description
flag_key	The flag key. The value of this parameter is obtained from the service template of the EBS cluster that is stored in the *pangu_blockgcworker_flag.json* file.
flag_value	The custom flag value.

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

- Configure the gcworker node status

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

  The following table describes the gcworker status.

Status	Description
NORMAL	The node is running normally.
DISCONNECTED	The node is disconnected.
OFFLOADING	The node is being disabled.
OFFLOADED	The node is disabled.

- Query the version information

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

4.1.1.4.7. Device operations

The Device Operations module allows you to view disk information in Elastic Block Storage (EBS) clusters such as the disk ID, state, capacity, and category. You can also perform flush operations, modify disk configurations, query segment information, and enable, disable, delete, or restore devices.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > EBS > Device Operations**.

   On the page that appears, the information of the first cluster in the **Cluster Name** drop-down list is displayed.
4. Select a cluster from the **Cluster Name** drop-down list.
5. You can perform the following operations:
   - View the device list

You can view the total number of devices, the total amount of logical space of devices, and information of each device in the cluster, including the device ID, state, logical capacity, number of segments, mode, and flags.

- Check all segments in a cluster

  In the upper-right corner of the **Device List** section, click **Global Check Segment** to view all the segments in the selected cluster and their indexes and states.

- Check disk status

  In the upper-right corner of the **Device List** section, click **Check Cloud Disk Status** to 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 disk information such as the disk ID, state, and capacity.

- Delete a device

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

  After the device is deleted, its state changes to **DELETING** and the device becomes unavailable. While the device is in this state, operations such as enabling the device and modifying configurations cannot be performed on the device.

- Restore a device

  In the device list, find a deleted device that is in the **DELETING** state and click **Restore** in the **Actions** column. In the message that appears, click **OK** to restore the deleted device to a normal state.

  After the device is restored, it becomes available and operations such as enabling the device and modifying configurations can be performed on the device.

- Enable a device

  In the device list, find a device and choose **More > Turn On** in the **Actions** column. In the dialog box that appears, configure 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 used to enable a device.

Parameter	Description
client_ip	Optional. The IP address of the client on which to enable the disk. The client IP address is the IP address of the block server. If the client IP address is not specified, the IP address of your computer is used.
token	The string used as a token to disable the device.

Parameter	Description
mode	The disk mode. Valid values: ■ **ro**: read-only ■ **rw**: read and 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 when you disable a disk.

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

The following table describes the parameters used to disable a device.

Parameter	Description
client_ip	The IP address of the client on which to disable the disk. If the client IP address is not specified, the IP address of your computer is used.
token	The token to use to disable the device. This token was configured when the device is enabled. You can run the **dev -query** command on a server in the EBS cluster to query the token
open_ver	The current openversion of the device when the client IP address is not specified. If a client IP address is specified, you do not need to specify the openversion. You can run the **dev -query** command on a server in the EBS cluster to query the openversion.

- Flush a device

In the device list, find a device and choose **More > Flush** in the **Actions**. In the dialog box that appears, configure parameters and click **Submit** to flush the transaction logs of the disk or its segments.

The following table describes the parameters.

Parameter	Description

Parameter	Description
segment	The segment that you want to flush. If you do not specify this parameter, all segments are flushed.
ifnsw	Specifies whether to flush the index file. Valid values: ■ **0**: The index file is flushed. ■ **1**: The index file is not flushed.
dfnsw	Specifies whether to flush the data files. Valid values: ■ **0**: The data files are flushed. ■ **1**: The data files are not flushed.

- Perform a global flush operation for a cluster

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

  On the right of the **Device List** section, click **Global Flush**. In the dialog box that appears, select ifnsw and dfnsw and click **OK** to flush the transaction logs of all disks or segments in the current cluster.

- Query configuration status

  In the device list, find a device and choose **More > Query Configuration Status** in the **Actions** column. In the dialog box that appears, enter config_ver and click **OK**. You can determine whether the disk is configurable based on the check result.

  config_ver is the config_version parameter among the queried device information.

- Modify device configurations

  You can modify the configurations of a disk, such as the compression algorithm, storage mode, and whether data compression is enabled.

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

  The following table describes the parameters.

Parameter	Description
compress	Specifies whether to enable data compression. Valid values: ■ **enable** ■ **disable**

Parameter	Description
algorithm	The data compression algorithm. Valid values: ■ 0: No data compression algorithms are used. ■ 1: The snappy data compression algorithm is used. ■ 2: The lz4 data compression algorithm is used.
ec	Specifies whether to enable the EC storage mode. Default value: disable. Valid values: ■ enable ■ disable
data_chunks	The number of data chunks. Default value: 8.
parity_chunks	The number of parity chunks. Default value: 3.
packet_bits	The size of a single data block in EC mode. Default value: 15.
copy	The number of data replicas. Default value: 3.
storage_mode	The storage mode of the disk.
cache	Specifies whether to enable the cache mode. Default value: 0. Valid values: ■ 0: disables the cache mode. ■ 1: enables the cache mode.
storage_app_name	The data storage name.
simsuppress	Specifies whether to enable the delay simulation feature. Default value: disable. Valid values: ■ enable ■ disable
baselatency	The basic latency. Default value: 300.
consumespeed	The processing speed. Default value: 256 B/μs.
lat80th	The 80th percentile latency. The default value is 110% of the specified basic latency.
lat90th	The 90th percentile latency. The default value is 150% of the specified basic latency.
lat99th	The 99th percentile latency. The default value is 500% of the specified basic latency.

- Query segment information

  In the device list, find a device and choose **More > Segment Information** in the **Actions** column. In the dialog box that appears, view the information of the segments, such as the indexes and states.

- Check a segment

  In the device list, find a device and choose **More > Check Segment** in the **Actions** column. In the dialog box that appears, select a segment and click **Submit** to view the information of the segment such as the index and state.

4.1.1.4.8. Enable or disable Rebalance

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Rebalance**.
4. 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**.

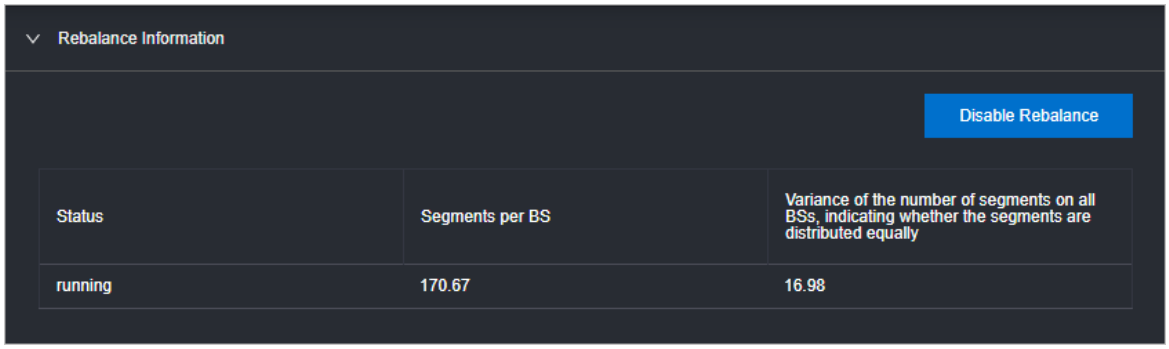

4.1.1.4.9. I/O hang fault analysis

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

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > IO HANG**.

   By default, the system shows the affected VM list, VM cluster statistics, and device cluster statistics for the last 24 hours.

4. Select a time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or a customized time range) and click **Search**. View the following information:
   - **Affected VM List**

     The **Affected VM List** section shows the I/O hang start time and recovery time of all the VMs, as well as the cluster name and user ID of the cluster to which these VMs belong.

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

     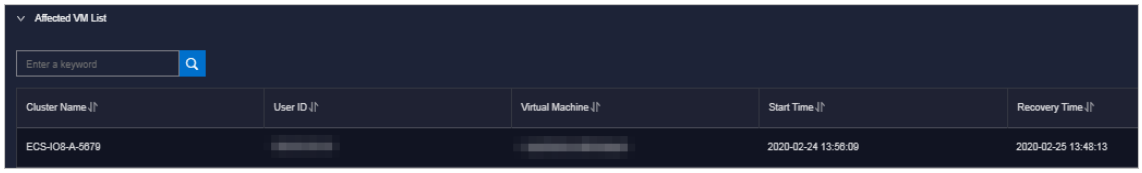

   - **VM Cluster Statistics**

     The **VM Cluster Statistics** section shows 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.

     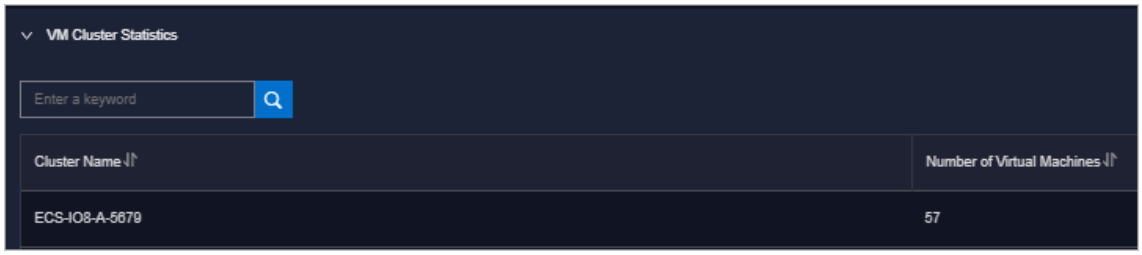

   - **Device Cluster Statistics**

     The **Device Cluster Statistics** section shows 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.

     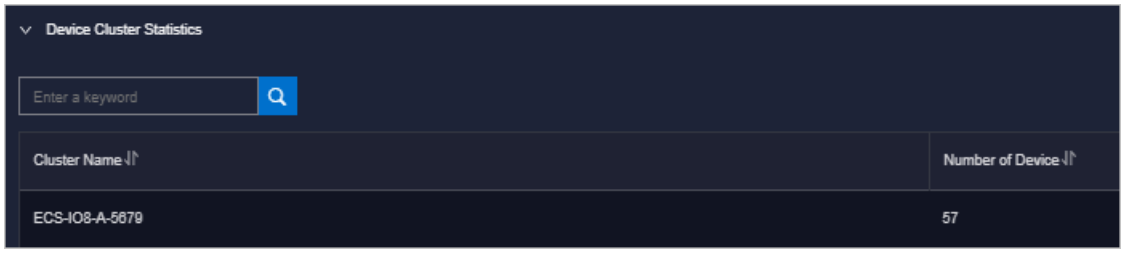

4.1.1.4.10. Slow IO analysis

The Slow IO Analysis page allows you to view the slow IO list, top ten NCs, cluster statistics, top five cluster statistics, and reasons.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > SLOW IO**.

By default, the system shows the slow IO list, top ten NCs, cluster statistics, top five cluster statistics, and reasons in the last 24 hours.

4. Select the time range (**One Hour**, **Three Hours**, **Six Hours**, **One Day**, or customize the time range) and click **Search**. View the following information:
   - **Slow IO List**

     The **Slow IO List** section shows the slow IO-related cluster name, NC IP address, virtual machine, device ID, storage type, start time, recovery time, number of slow IOs, and causes.

     To view the information of a cluster, an NC, or a 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 data by Cluster Name, NC IP, Virtual Machine, Device ID, Storage Type, Start Time, Recovery Time, Number of Slow IO, or Reason.

   - **Top10 NC**

     The system shows the information of the top ten NCs on a graph and table.

     Notes:

     - The **Graphical Analysis** section shows the proportion for the number of slow IO in each cluster of the top ten NCs by using a pie chart.
     - The **Top10 NC** section shows the NC IP address, cluster name, number of slow IOs, percentage, and primary cause of slow IOs on the top ten NCs.

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

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

   - **Cluster Statistics**

     The **Cluster Statistics** section shows the cluster name, number of devices, number of slow IOs, percentage, and primary cause of slow IOs on clusters.

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

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

   - **Top Five Cluster Statistics**

     The system shows the statistics of top five clusters by using a graph and a table.

     Notes:

     - The **Graphical Analysis** section shows the proportion for the number of slow IOs on each of the top five clusters on a pie chart.
     - The **Top Five Cluster Statistics** section shows the cluster name, number of devices, number of slow IOs, percentage, and primary cause of slow IOs on the top five clusters on a table.

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

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

   - **Reason**

     The system shows the primary cause on a graph and table.

Notes:

- The **Graphical Analysis** section shows the proportion of reasons by using a pie chart.
- The **Reason** section shows the number of slow IO from the dimension of reasons.

  To query the information of a reason, enter the reason information in the search box to perform a fuzzy search.

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

4.1.1.4.11. Product settings

The Product Settings module allows you to view the sales status of a cluster, configure the overcommit ratio of a cluster, and specify whether a cluster is available for sale.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Apsara Distributed File System Management > Product Settings**.

   By default, the system shows the data of each cluster within the current environment, including the cluster name, overcommit ratio, and sales status.

> **Note** An overcommit ratio is the ratio of the marketable capacity of a storage device to the physical capacity. For example, if the physical storage capacity of a storage device is 1 TB and marketabe capacity is 2.5 TB, and the overcommit ratio is 2.5.

4. Perform the following operations:
   - Select a cluster, enter a number in the **Adjust Setting Oversell Ratio** field, and then click **Confirm** to set the oversold ratio of the cluster.
   - Select a cluster and turn on or off **Adjustment of sales status** to enable or disable the cluster for sale.

4.1.1.5. ECS Diagnose

4.1.1.5.1. Overview

ECS Diagnose is an O&M service provided in Alibaba Cloud Apsara Stack that allows you to diagnose exceptions on Elastic Compute Service (ECS) instances or hosts and perform proactive O&M to automate O&M.

ECS Diagnose provides the features described in the following table.

Feature	Description
End-to-end demarcation	Diagnoses current and historical exceptions that occur on ECS instances and hosts to identify the exception causes. • For ECS instances: This feature associates all components involved in the link that spans from the ECS console to the data, and factors in exception information about infrastructure devices such as hosts to identify exception causes. • For hosts: This feature associates all ECS instances that reside on a host and factors in exception information about infrastructure devices such as hosts to identify exception causes.
Proactive O&M	Executes the following proactive O&M workflows: • **0x410A01-[NE] Lock NC->Complete** This workflow can lock hosts and put them into the **mlock** state so that ECS instances cannot be created on the hosts. For example, if the hardware of a host (such as CPUs or memory) fails, you must execute this workflow to lock the host, and then shut down the host after ECS instances are failed over from the host. • **0x400012-Force Restart NC** This workflow can forcibly restart hosts on which exceptions have occurred **Notice** regardless of whether ECS instances reside on the hosts. • **0x400011-Bring NC Online Again** After a host is restarted, it enters the **nc_down** state and ECS instances cannot be created on it. If you have already performed troubleshooting on the host and confirmed that the host is normal, you can execute this workflow to put the host into the **free** state and bring the host back online. • **0x10001000008002-Clear VM Residues from NC** After an ECS instance is hot migrated from a source host to a destination host, you can execute this workflow to clear the residual of the instance that may be left on the source host.

4.1.1.5.2. Diagnose exceptions on a host or an ECS instance

If exceptions occur on an Elastic Compute Service (ECS) instance (VM) or a host (NC), you can perform a diagnosis on the End-to-end Demarcation page in the ECS Diagnosis console and identify the causes of the exceptions based on the diagnostic results.

Context

After the exceptions on the ECS instance or host are diagnosed and their causes are identified, you can perform a proactive O&M event to troubleshoot the exceptions. For more information, see Proactive O&M on hosts.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. Go to the **ECS Diagnose** console.
   i. In the top navigation bar, select an environment version and a region.
   ii. In the top navigation bar, click **O&M**.
   iii. In the left-side navigation pane, choose **Product Management > Products**.
   iv. Find the Computing Services section in the **Infrastructure as a Service (IaaS) > >** column. Click ECS Diagnose in the section.
3. On the **End-to-end Demarcation** page, configure search conditions and click the icon.

Search condition	Description
Start and End Time	Specify the beginning and end of the time range to query. By default, exception information generated in the last two days is displayed. You can specify a time range to query historical records. The beginning of the time range can be up to 10 days earlier than the end of the time range.
Machine ID	Specify the IP address of an ECS instance or a host. You can query the IP addresses of ECS instances or hosts on the **ECS Operations and Maintenance Platform** page in the Apsara Uni-manager Operations Console. For more information, see Search for VMs.

4. View the diagnostic information about the ECS instance or host in the search results.

Section	Information contained in the section	Description
①	Resource tags	Shows the information of the diagnostic object, such as the state. You can move the pointer on a tag to view details.
②	Diagnostic overview	Shows diagnostic information such as the number of diagnostic items and the number of exceptions.
③	Diagnostic results	Shows the types and diagnostic information of diagnosed exceptions. Types of exceptions: ○ **Console Exception**: affects user operations, such as starting an instance, stopping an instance, migrating an instance, and creating a snapshot. ○ **NC Exception**: affects host availability, such as service failures. ○ **NC Performance Exception**: affects host performance, such as L3 cache miss, memory bandwidth contention, and a large number of tdc retries. ○ **VM Availability Exception**: affects the availability of ECS instances, such as read/write failures to disks and service failures. ○ **VM Performance Exception**: affects the performance of ECS instances, such as packet loss due to throttling and packet loss due security group settings. Diagnostic information includes the host that hosts the ECS instance, exception information, and exception causes. You can click links to view details.
④	Basic information	Shows the basic information of the diagnosed ECS instance or host, such as the model, cluster, and instance family. You can click links to view details.
⑤	Diagnostic details	Shows a list of exceptions related to the diagnosed ECS instance or host. This section consists of two parts: ○ Diagnostic information sorted by exception type ○ A list of exceptions that shows detailed exception information You can click links to view details.

4.1.1.5.3. View historical cold migration records

You can view cold migration records about Elastic Compute Service (ECS) instances (VMs) on the End-to-end Demarcation page in the ECS Diagnose console for backtracking.

Context

During cold migration, an ECS instance must be stopped before it can be migrated to another host.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. Go to the **ECS Diagnose** console.
   i. In the top navigation bar, select an environment version and a region.
   ii. In the top navigation bar, click **O&M**.
   iii. In the left-side navigation pane, choose **Product Management > Products**.
   iv. Find the Computing Services section in the **Infrastructure as a Service (IaaS) > >** column. Click ECS Diagnose in the section.
3. On the **End-to-end Demarcation** page, configure search conditions and click the icon.

Search condition	Description
Start and End Time	Specify the beginning and end of the time range to query. By default, exception information generated in the last two days is displayed. You can specify a time range to query historical records. The beginning of the time range can be up to 10 days earlier than the end of the time range.
Machine ID	Specify the ID of an ECS instance. You can view the IP addresses of ECS instances or hosts on the **ECS Operations and Maintenance Platform** page in the Apsara Uni-manager Operations Console. For more information, see Search for VMs.

4. After the specified ECS instance is found, click **Query Cold Migration History**.
5. In the **Cold Migration History** list, view the cold migration records about the ECS instance.

   You can view the information described in the following table:

Item	Description
Instance ID	The ID of the ECS instance.
Creation Time	The time when the cold migration of the ECS instance starts.
Modification Time	The time when the cold migration of the ECS instance ends.
Migration Status	The cold migration state.
Migration Destination IP	The IP address of the destination host to which the ECS instance is migrated.

Item	Description
Migration Source IP	The IP address of the source host from which the ECS instance is migrated.
Post-migration Configurations	The post-migration configurations such as the number of CPUs and memory size.
Pre-migration Configurations	The pre-migration configurations such as the number of CPUs and memory size.
Migration Reason	The reason of the cold migration.
Migration Performance Details	The details about migration performance.

4.1.1.5.4. Proactive O&M on hosts

After you diagnose exceptions on Elastic Compute Service (ECS) instances or hosts on the End-to-end Demarcation page in the ECS Diagnose console, you can use the proactive O&M feature to perform troubleshooting.

Context

When exceptions occur on an ECS instance or a host, you can perform the following steps to proactively operate and maintain the host.

1. Step 1: Create a proactive O&M event
2. Step 2: View and execute the O&M workflow
3. Step 3: (Optional) Batch manage proactive O&M events

Step 1: Create a proactive O&M event

1. Log on to the Apsara Uni-manager Operations Console.
2. Go to the **ECS Diagnose** console.
   i. In the top navigation bar, select an environment version and a region.
   ii. In the top navigation bar, click **O&M**.
   iii. In the left-side navigation pane, choose **Product Management > Products**.
   iv. Find the Computing Services section in the **Infrastructure as a Service (IaaS) > >** column. Click ECS Diagnose in the section.
3. In the upper-right corner, click **Proactive O&M**.
4. Click **Create O&M Event**.
5. Configure the parameters described in the following table.

Parameter	Description
Server	Specify the host IP addresses. You can specify multiple IP addresses and separate them with commas (,).
O&M Type	Only **Automated O&M** is available.

Parameter	Description
Associated opcode	Select an O&M workflow to execute. ○ **0x410A01-[NE] Lock NC->Complete** This workflow can lock hosts and put them into the **mlock** state to prevent ECS instances from being created on the hosts. ○ **0x400012-Force Restart NC** This workflow can forcibly restart abnormal hosts, **Notice** regardless of whether ECS instances exist on the hosts. ○ **0x400011-Bring NC Online Again** After a host is restarted, it enters the **nc_down** state and ECS instances cannot be created on it. If you have already performed troubleshooting on the host and confirmed that the host is normal, you can execute this workflow to put the host into the **free** state and bring the host back online. ○ **0x10001000008002-Clear VM Residues from NC** After an ECS instance is hot migrated from a source host to a destination host, you can execute this workflow to clear the residual of the instance that may be left on the source host. After you select a workflow, you must configure parameters for the workflow. For more information, see Parameters of O&M workflows.
Send Notification	Only **No** is available.
Cause Category	Select a cause for O&M based on the exception type.
Supplement	Provide more information about the exception.
Ignore Check	By default, **No** is selected.

6. Click **Submit**.
   After the proactive O&M event is created, you can view it in the proactive O&M event list.

Step 2: View and execute the O&M workflow

After the proactive O&M event is created, the system performs O&M based on the specified workflow. You can view the workflow in the event list and perform operations.

1. On the **Proactive O&M** page, view the execution information of the proactive O&M event that you created in the previous step.

   In the proactive O&M event list, you can view the proactive O&M event and information such as the host, O&M state, O&M workflow.

2. In the **Workflow** column corresponding to the proactive O&M event, click Details to view the details of the O&M workflow.

On the details page of the O&M workflow, you can view the execution state of each task node in the workflow.

3. If a task node in the O&M workflow is in an interrupted state such as the **Pending** state, you can click a button in the **Actions** column to perform an operation.

When all task nodes in the O&M workflow are completed, the system returns a final execution result for the workflow.

Step 3: (Optional) Batch manage proactive O&M events

On the **Proactive O&M** page, you can batch manage multiple proactive O&M events.

1. On the **Proactive O&M** page, select multiple proactive O&M events.
2. Click a button in the lower part of the page to batch manage the selected proactive O&M events.

   You can click one of the following buttons:

   - **Unlock**: batch unlocks hosts. When the hosts are unlocked, their **nc status** changes from **mlock** to **free**.
   - **Cancel O&M**: batch cancels O&M for hosts. When the O&M is canceled for the hosts, the values in the **Actions** column corresponding to the hosts change to **Canceled**.
   - **Complete**: batch completes O&M on hosts. When the O&M is complete on the hosts, the values in the **Actions** column corresponding to the hosts change to **Completed**.

Parameters of O&M workflows

Different O&M workflows can be selected for the **Associated opcode** parameter. The following tables describe the parameters of each O&M workflow.

> **Note** Typically, you can directly use the default values of the parameters.

- **0x410A01-[NE] Lock NC->Complete**

Parameter	Description
alertRuleName	The name of the O&M rule to trigger.
preLock	Specifies whether to lock the host in advance. Valid values: ○ 0: The host is locked in advance. ○ 1: The host is not locked in advance.

- **0x400012-Force Restart NC**

Parameter	Description
alertRuleName	The name of the O&M rule to trigger.
fvt	The environment in which to forcibly restart the host. Default value: 0, which indicates the production environment.

Parameter	Description
needDumpCn	Specifies whether to perform memory dumps for the controller node (CN). Valid values: ○ 0: Memory dumps are not performed for the CN. ○ 1: Memory dumps are performed for the CN. Default value: 0.
needLimit	Specifies whether to enable throttling for O&M traffic. Valid values: ○ 0: Throttling is not enabled. ○ 1: Throttling is enabled. Default value: 1.
preLock	Specifies whether to lock the host in advance. Valid values: ○ 0: The host is locked in advance. ○ 1: The host is not locked in advance.

- **0x400011-Bring NC Online Again**

Parameter	Description
fvt	The environment in which to bring the host back online. Default value: 0, which indicates the production environment.
NcReOnline	The interval at which to restart the host.
parallel	The number of hosts to be concurrently brought online. Default value: 4.
preLock	Specifies whether to lock the host in advance. Valid values: ○ 0: The host is locked in advance. ○ 1: The host is not locked in advance.
ReOnline.ignoreStatus	Specifies whether to ignore the Server Controller state when the host is brought online.

- **0x10001000008002-Clear VM Residues from NC**

Parameter	Description
foreCleanWildVm	Specifies whether to check the process states of the ECS instances on the host. Valid values: ○ true: The process state of an ECS instance is not checked when the instance is in the released state. ○ false: The process state of an ECS instance is checked when the instance is not in the released state.

Parameter	Description
preLock	Specifies whether to lock the host in advance. Valid values: ○ 0: The host is locked in advance. ○ 1: The host is not locked in advance.
vm	The ID of the ECS instance.

4.1.1.6. VM hot migration

4.1.1.6.1. 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.

4.1.1.6.2. 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.

4.1.1.6.3. Perform hot migration on the AG

In the Apsara Uni-manager Operations Console, you can run commands to start or cancel hot migration operations.

Trigger hot migration

After hot migration is triggered for a virtual machine (VM), you can run the `go2which` command or use the ECS Operations and Maintenance Platform to check that the VM enters the Migrating state. When hot migration is complete, the VM returns to the Running state.

To trigger hot migration, run the following `go2which` command:

```
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]
```

Parameters

Parameter	Description	Impact	Value
vm_name	The name of the VM that you want to migrate.	N/A	N/A
nc_id	The destination NC to which to migrate the VM.	If the NC does not support the specifications of the VM, the VM may fail to be migrated.	N/A
rate	The amount of bandwidth to allocate for the migration task.	The migration uses the bandwidth resources of physical machines.	• 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
no_check_image	Forcibly migrates images that are not supported.	The SLA may be violated if this parameter is set to false.	false

Parameter	Description	Impact	Value
no_check_load	Forcibly migrates images even when the load threshold requirements are not met.	Downtime cannot be controlled if this parameter is set to false.	false

Cancel hot migration

To cancel hot migration, run the following command:

```
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>
```

Parameters

Parameter	Description	Impact	Value
vm_name	The name of the VM that you want to migrate.	N/A	N/A
region_id	The region ID of the VM.	N/A	N/A

4.1.1.6.4. 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
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

4.1.1.6.5. 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.

4.1.1.7. Hot migration of disks

4.1.1.7.1. 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.

4.1.1.7.2. Limits

Before you perform hot migration on a disk, you must understand the limits that apply.

Limits

- Only disks of the river type support hot migration.
- The source and destination clusters for a hot migration task must belong to the same Object Storage Service (OSS) domain.
- Hot migration is not supported by shared disks.
- Hot migration is not supported by disks that are larger than 2 TB in size.
- Format and capacity changes are not supported.
- Hot migration can be performed to migrate a disk only within the same zone.
- Hot migration is implemented internally. The names of the source and destination clusters must be less than 15 bytes in length.
- Migration cannot be rolled back.

> **Notice**
>
> - Snapshots must be created to back up disk data before hot migration is performed on disks.
> - When hot migration is complete, data exists on both the source and destination disks. You can use the pu tool to delete data from the source disk only after you confirm that data is complete and accessible on the destination disk. Job recycling is unavailable.

API operations

For more information about the API operations related to disk hot migration, see **Disk hot migration** in *ECS Developer Guide*.

> **Note**
>
> - When a disk is being migrated, an I/O latency of less than 1 second is considered to be normal.
> - Migration consumes network bandwidth. You must take measures to limit the number of concurrent migration tasks.

4.1.1.7.3. 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.

4.1.1.8. Upgrade solution

4.1.1.8.1. Overview

For both hot and cold migration of GPU and FPGA clusters, you must understand the limitations that apply to cluster upgrades.

4.1.1.8.2. 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.

4.1.1.8.3. 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.

4.1.1.9. Handle routine alarms

4.1.1.9.1. 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/servi	Monitors the basic performance of the host and the availability of the services	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.

Metric type	Function	Solution
ce availability metric	on the host. This kind of metrics includes CPU, memory, and handle count.	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.

4.1.1.9.2. 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.
check_apiproxy_proc_new	check_apiproxy_proc_new	Checks the memory usage and CPU utilization for nginx and memcache processes.

4.1.1.9.3. 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.

4.1.1.9.4. 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.
	rms_connectivity	Checks the connectivity to RMS.

Metric	Alert item	Description
check_regionmaster_work	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.	-

4.1.1.9.5. 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.

4.1.1.9.6. 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.

4.1.1.9.7. 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.

4.1.1.10. Inspection

4.1.1.10.1. Overview

ECS inspection includes cluster basic health inspection and cluster resources inspection.

4.1.1.10.2. Cluster basic health inspection

4.1.1.10.2.1. Overview

Cluster basic health inspection includes monitoring inspection, inspection of basic software package versions, and basic public resources inspection.

4.1.1.10.2.2. Monitoring inspection

This topic describes basic monitoring inspections and connectivity monitoring inspections.

4.1.1.10.2.3. Inspection of basic software package versions

This topic describes the version inspections of Server Controller components, Apsara system, virtualization packages, and basic service packages.

4.1.1.10.2.4. 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.

4.1.1.10.3. Cluster resource inspection

4.1.1.10.3.1. Overview

Cluster resource inspection includes cluster inventory inspection and VM inspection.

4.1.1.10.3.2. 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.
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.

4.1.1.10.3.3. 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.

4.2. Container Service for Kubernetes

4.2.1. Operations and Maintenance Guide

4.2.1.1. Architecture

4.2.1.1.1. System architecture

This topic describes the system architecture of Container Service for Kubernetes (ACK).

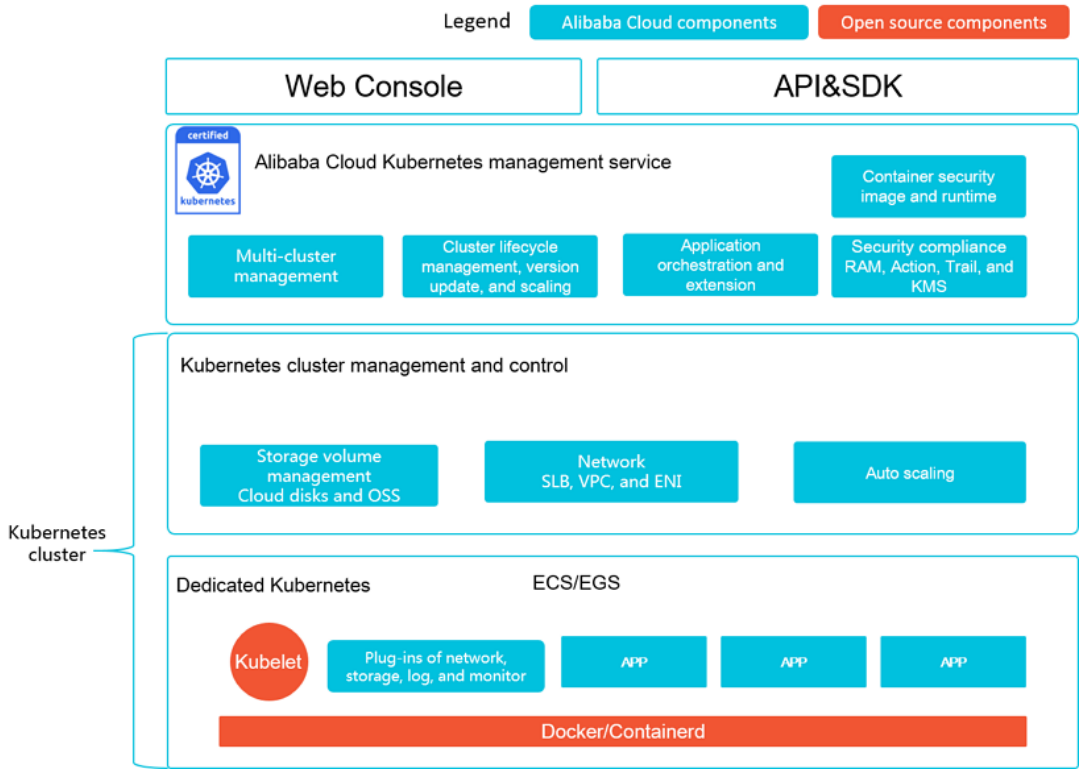

Container Service is adapted and enhanced based on native Kubernetes. This service simplifies cluster creation and scaling, and is integrated with Apsara Stack virtualization, storage, network, and security capabilities. This service provides the optimal environment to run Kubernetes-based containerized applications in the cloud.

Feature	Description
Dedicated Kubernetes mode	Integrated with Apsara Stack virtualization technologies, the service allows you to create dedicated Kubernetes clusters. You can use Elastic Compute Service (ECS), Elastic GPU Service (EGS), and ECS Bare Metal instances as cluster nodes. You can deploy a wide range of plug-ins and configure various specifications for the instances.
Cluster management and control	The service provides powerful network, storage, cluster management, scaling, and application extension features.
Kubernetes management	The service supports secure images and is integrated with Apsara Stack Resource Access Management (RAM), Key Management Service (KMS), and logging and monitoring services to provide a secure and compliant Kubernetes solution.
Convenient and efficient use	ACK provides services through the web console and APIs.

4.2.1.1.2. Component architecture

This topic describes the component architecture of Container Service for Kubernetes (ACK).

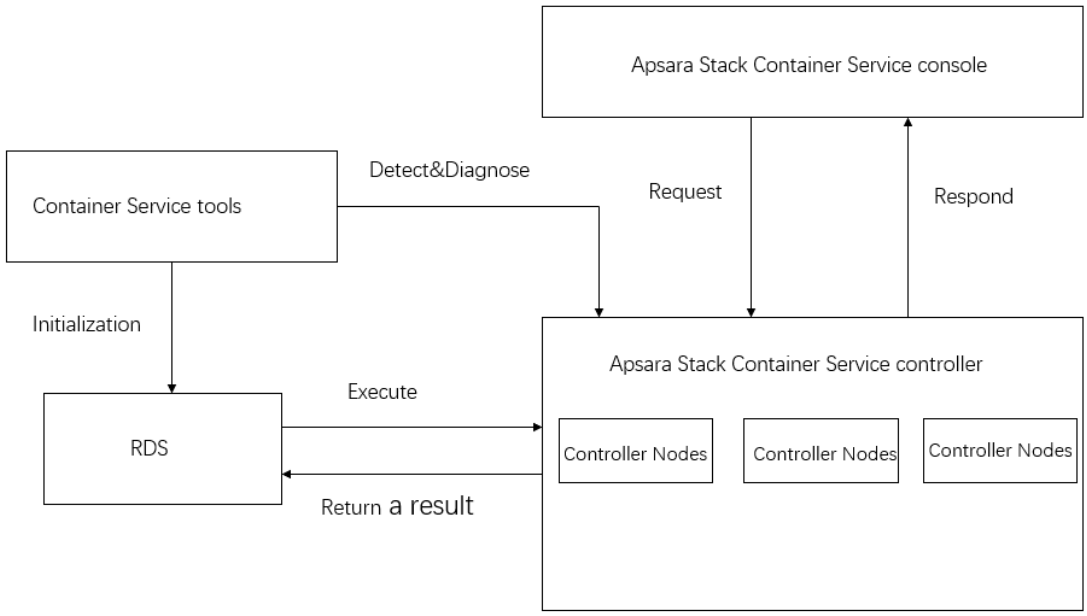

Category	Component	Description
Apsara Stack Container Service console	CosConsoleAliyunCom	The console interface that is used to manage clusters, nodes, and applications.
	Backend	Dashboard.
Apsara Stack Container Service controller	Troopers	Controller backend.
	CloudAccessVpc	Reverse tunneling.
	Cert	Static file service.
	CertControl	Component certificate signing.
	Registry	Image service.
	Mirana	API server traffic forwarding.
	Certificate	Certificate updates.
Container Service tools	ControlInit	Data initialization and API registration.
	ServiceTest	Self-check service.
	DiagnosticTool	Quick diagnostic tool.

4.2.1.1.3. Deployment architecture

This topic describes the deployment architecture of Container Service for Kubernetes (ACK).

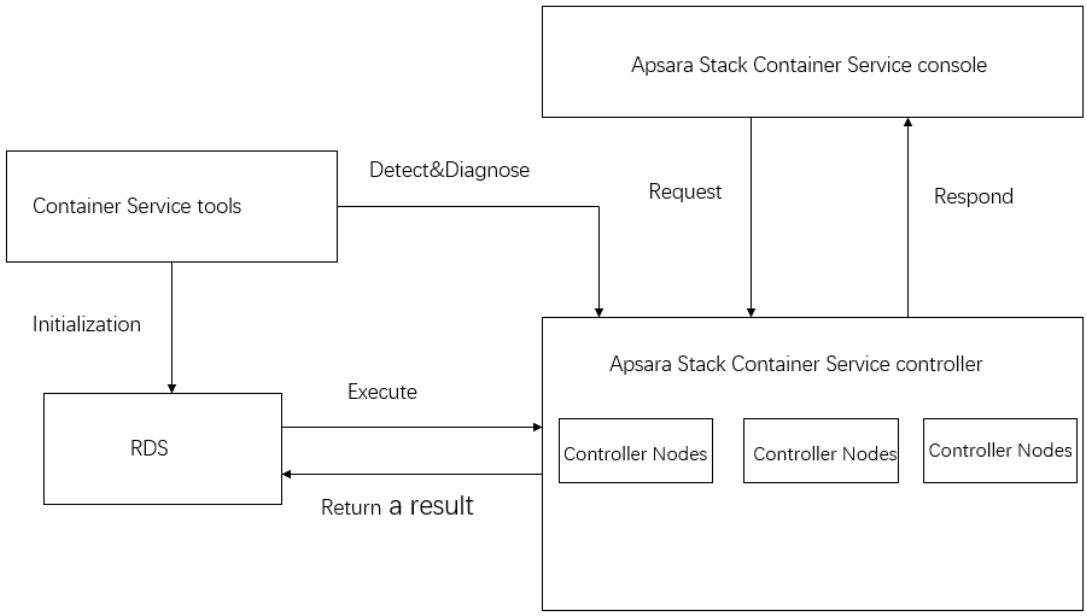

Component	Deployment method
Apsara Stack Container Service console	Containerized deployment
Apsara Stack Container Service controller	Containerized deployment
Container Service tools	Containerized deployment

4.2.1.1.4. Server roles

This topic describes the server roles of Container Service for Kubernetes (ACK).

Project	Cluster	Service	Server role	Description
acs	AcsControlCluster	acs-acs_control	ControlInit	Data initialization and API registration.
			Troopers	Controller backend.
			CloudAccessVpc	Reverse tunneling.
			Cert	Static file service.
			CertControl	Component certificate signing.
			Registry	Image service.

Project	Cluster	Service	Server role	Description
			Mirana	API server traffic forwarding.
			Backend	Dashboard.
			CosConsoleAliyunCom	Console.
			Certificate	Certificate updates.
			ServiceTest	Self-check service.
			DiagnosticTool	Quick diagnostic tool.

4.2.1.2. Components and features

4.2.1.2.1. Console

The Container Service console provides a graphical user interface (GUI) that serves as a portal for all operations on Container Service. The console uses the deployment mode that applies to standard Java applications on Apsara Stack. The Container Service console consists of a Tengine server and a Jetty container.

Log on to the console

1. Log on to the Apsara Uni-manager Operations Console. In the top navigation bar, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. In the **Apsara Stack O&M** section, click **Apsara Infrastructure Management Framework**.
2. You are redirected to the **Infrastructure Operation Platform** console. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. You are redirected to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, find the Container Service cluster that you want to manage.

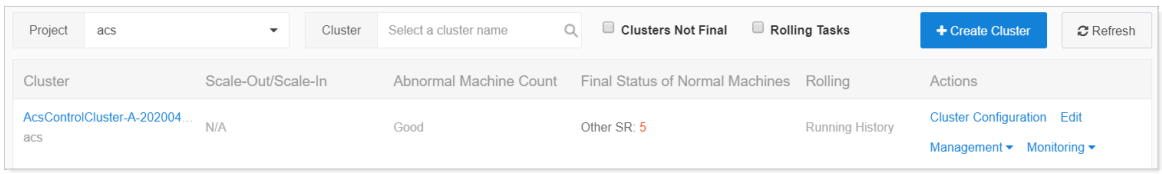

5. Click **Cluster Configuration** in the **Actions** column. On the Cluster Configuration page, find the CosConsoleAliyunCom server role in the **Server Role** section and check the machine of the server role.
6. In the left-side navigation pane, enter the hostname in the Machine search box. Move the 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

query the ID of the cos-console-aliyun-com container.

7. Run the `sudo docker exec -it container_id bin/bash` command to log on to 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*

4.2.1.2.2. Troopers

This topic describes the features of Troopers and how to use them.

The Troopers daemon is used to create clusters and machines. You can also use Troopers to manage clusters and machines in Container Service.

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. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. In the **Apsara Stack O&M** section, click **Apsara Infrastructure Management Framework**.
2. You are redirected to the **Infrastructure Operation Platform** console. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. You are redirected to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, find the Container Service cluster that you want to manage. Click **Cluster Configuration** in the Actions column. On the Cluster Configuration page, find the Troopers server role in the Server Role section and check the machine of the server role.
5. In 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 machine by using a terminal session. On the command line, enter `docker ps` to query the ID of the Troopers container.
6. Run the `sudo docker exec -it container_id bin/bash` command to log on to 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 validate the certificate.

Troopers log files are exported to the stdout stream. No log files are stored in the container. To view log data, run the `docker logs` command outside the container.

4.2.1.2.3. Mirana

This topic describes the deployment modes and features of Mirana.

Server role

1. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. In the **Apsara Stack O&M** section, click **Apsara Infrastructure Management Framework**.
2. You are redirected to the **Infrastructure Operation Platform** console. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. You are redirected to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, find the Container Service cluster that you want to manage. Click **Cluster Configuration** in the Actions column. On the Cluster Configuration page, find the Mirana server role in the Server Role section and check the machine of the server role.

Query log data

In 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 machine by using a terminal session. On the command line, enter `docker ps` to query the ID of the Mirana container. Run the `docker logs contianer_id` command to query 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 modes

- 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 machines 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.
- Allows you to use the Helm client to manage orchestration templates.
- Allows you to call API operations to perform blue-green releases.

- Serves as the proxy for Kubernetes-native API operations.

4.2.1.3. System restart

4.2.1.3.1. Restart a control node

A container control node runs a Docker container where Services such as CosConsoleAliyunCom, Troopers, and etcd are deployed. The following procedure demonstrates how to restart a control node:

Procedure

1. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. In the **Apsara Stack O&M** section, click **Apsara Infrastructure Management Framework**.
2. You are redirected to the **Infrastructure Operation Platform** console. In the left-side navigation pane, click **Reports**.
3. On the **Reports** page, click **Go**.
4. You are redirected to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page, find the Container Service cluster that you want to manage. Click **Cluster Configuration** in the Actions column. On the Cluster Configuration page, find the server role of the control node in the Server Role section and find the machine where the control node is deployed.
5. On the command line, enter `docker ps|grep [app]` to query the container ID.

   `[app]` specifies the name of the application that is deployed in the container. You can query the container ID by using the application name.
6. On the command line, enter `docker restart container_id` to restart the container.

4.2.1.4. Manage certificates for control components

If you do not renew expired certificates for control components, you cannot use the control components. You must renew a certificate within one week after the certificate expires. This topic describes how to query the certificate expiration time and renew a certificate for a control component.

Prerequisites

The permissions to access control components are granted.

Query the certificate expiration time

Method 1: Query the component creation time to determine the certificate expiration time

Log on to the control component and query the creation time of the component. The certificate of the component expires four years after the component is created. In the following steps, the Troopers component is used as an example.

1. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
2. On the **Clusters** tab, move the pointer over the icon to the right of **AcsControlCluster** and click **Dashboard**.

3. On the **Cluster Dashboard** page, find acs-acs-control in the **Service Instances** section and click **Details** in the **Actions** column.
4. On the **Service Instance Information Dashboard** page, find Troopers in the **Server Role List** section and click **Details** in the **Actions** column.
5. In the **Machine Information** section, select the machine that you want to manage and click **Terminal** in the **Actions** column.
6. Run the following command to query the container start time:

   The container start time is the same as the component creation time.

   ```
   docker ps | grep troopers
   ```

   Add four years to the component creation time to obtain the certificate expiration time.

Method 2: Log on to the component to query the certificate expiration time

> **Note** We recommend that you use this method to query the certificate expiration time because this method produces more accurate results.

1. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
2. On the **Clusters** tab, move the pointer over the icon to the right of **AcsControlCluster** and click **Dashboard**.
3. On the **Cluster Dashboard** page, find acs-acs-control in the **Service Instances** section and click **Details** in the **Actions** column.
4. On the **Service Instance Information Dashboard** page, find Troopers in the **Server Role List** section and click **Details** in the **Actions** column.
5. In the **Machine Information** section, select the machine that you want to manage and click **Terminal** in the **Actions** column.
6. Run the following command to go to the directory of the certificate:

   ```
   cd /usr/aliyun/acs/certs/control
   ```

7. Copy the *cert.pem* file in the certificate directory to your on-premises Linux environment.

   > **Note** You must install OpenSSL in your on-premises Linux environment before you copy the file.

8. Run the following command to query the certificate expiration time:

   ```
   openssl x509 -in cert.pem -noout -text
   ```

   The time that follows `Not After` in the `Validity` field is the expiration time of the certificate.

   ```
   Validity
       Not Before: Dec 23 16:02:00 2020 GMT
       Not After : Dec 22 16:07:40 2024 GMT
   ```

Renew the certificate for a control component

1. Log on to the control database and delete all data entries about signed certificates for components such as Console, Mirana, and Troopers. These data entries start with the string `manager`.
   i. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
   ii. On the **Clusters** tab, move the pointer over the icon to the right of **AcsControlCluster** and click **Dashboard**.
   iii. On the **Cluster Dashboard** page, find the resource whose **Type** is db and **Name** is tptemp in the **Cluster Resource** section. Right-click the entry in the **Result** column and click **Show More**.
   iv. In the **Details** dialog box, view database logon information.
   v. Log on to an OPS machine or a VM in the cluster.
   vi. Run the following command to log on to the database:

   Replace the parameters in the following command with the logon information that you obtained in Step d.

   ```
   mysql -h${db_host} -P${db_port} -u${db_user}  -D${dbName} -p
   ```

   Enter the database password when prompted.
   vii. Run the following command to delete the certificates of control components:

   ```
   delete from cert_info where path like "manager/%";
   ```

   viii. Run the following command to check whether the certificates are deleted:

   ```
   select path, created_at  from cert_info where path like "manager/%" \G;
   ```

   The output is empty, which indicates that the certificates are deleted.
2. Restart the CertControl component.

   Chest is used to sign the certificates for the control components used in Apsara Stack Container Service. It is deployed as an HTTP server and provides interfaces that are used to sign certificates for the components.
   i. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
   ii. On the **Clusters** tab, move the pointer over the icon to the right of **AcsControlCluster** and click **Dashboard**.
   iii. On the **Cluster Dashboard** page, find acs-acs-control in the **Service Instances** section and click **Details** in the **Actions** column.
   iv. On the **Service Instance Information Dashboard** page, find CertControl in the **Server Role List** section and click **Details** in the **Actions** column.

v. In the **Machine Information** section, select the machine that you want to manage and click **Restart** in the **Actions** column. In the Confirm message, click **OK**.

> **Note** You must restart both machines in the Machine Information section to restart the server role.

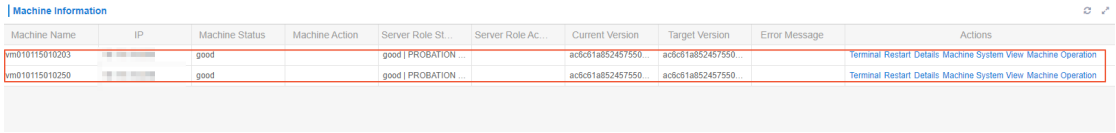

Verify that the entries in the **Machine Action** column change to **restart|pending** in the **Machine Information** section. This indicates that the machines are being restarted. If the entries in the **Machine Action** column are empty, it indicates that the restart is completed.

3. Check whether data entries about component certificates are regenerated in the cert_info table. These data entries start with the string `manager/` .

   i. Run the following command to log on to the database:

   Replace the parameters in the following command with the logon information that you obtained in Step .

   ```
   mysql -h${db_host} -P${db_port} -u${db_user} -D${dbName} -p
   ```

   Enter the database password when prompted.

   ii. Run the following command to view the certificates:

   ```
   select path, created_at from cert_info where path like "manager/%" \G;
   ```

   The result indicates that the certificates are regenerated.

4. Restart control components.

   Restart the following control components: Certificate, Console, Mirana, Static, and Troopers. The startup script of each component image sends a request to the download-cert interface provided by Chest to download a new certificate. The new certificate is valid for four years.

   On the **Service Instance Information Dashboard** page, find the component that you want to restart in the **Server Role List** section and click **Details** in the **Actions** column. The steps to restart a control component are similar to the steps to restart the CertControl component. For more information, see Step .

5. Check whether the components are restarted.

   In the following steps, the Troopers component is used as an example.

   i. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.

   ii. On the **Clusters** tab, move the pointer over the icon to the right of **AcsControlCluster** and click **Dashboard**.

   iii. On the **Cluster Dashboard** page, find acs-acs-control in the **Service Instances** section and click **Details** in the **Actions** column.

   iv. On the **Service Instance Information Dashboard** page, find Troopers in the **Server Role List** section and click **Details** in the **Actions** column.

v. In the **Machine Information** section, select the machine that you want to manage and click **Terminal** in the **Actions** column.

> **Note** To ensure that you obtain accurate results, we recommend that you log on to both machines in the **Machine Information** section to query the certificate expiration time.

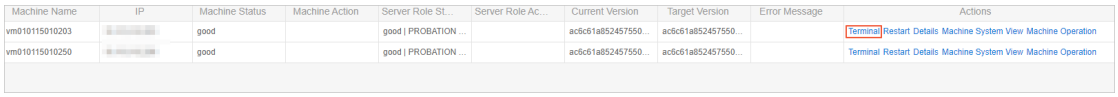

vi. Check whether the certificate exists.

a. Run the following command to query the container ID:

```
docker ps | grep troopers
```

b. Run the following command to log on to the container:

```
docker exec -it <certificate ID> bash
```

c. Run the following command to go to the directory of the certificate:

```
cd /usr/aliyun/ace/certs/control
```

d. Run the following command to view the certificate:

```
ls
```

vii. Check whether the certificate expiration time is renewed.

a. Copy the *cert.pem* file in the certificate directory to your on-premises Linux environment.

b. Run the following command to query the certificate expiration time:

```
openssl x509 -in cert.pem -noout -text
```

The time that follows `Not After` in the `Validity` field is the expiration time of the certificate.

```
Validity
        Not Before: Dec 23 16:02:00 2020 GMT
        Not After : Dec 22 16:07:40 2024 GMT
```

4.3. Auto Scaling (ESS)

4.3.1. Operations and Maintenance Guide

4.3.1.1. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager Operations Console.

Prerequisites

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

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

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

Procedure

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

   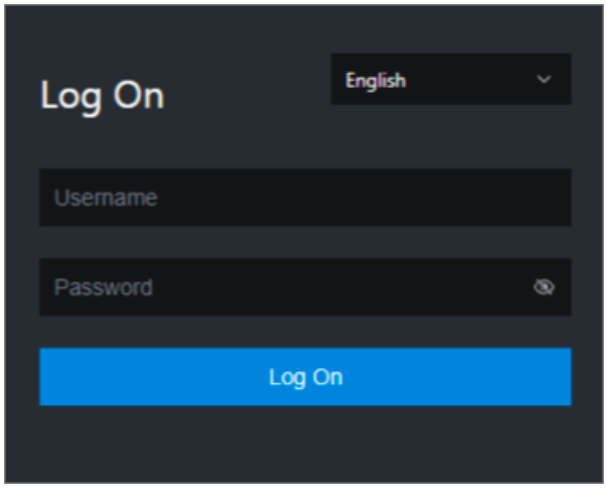

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains 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.

4. Click **Log On**.

4.3.1.2. Product resources and services

4.3.1.2.1. Application deployment

All applications at the Auto Scaling business logic layer are stateless. You must restart the applications by running the docker restart command.

- ess-init

  ess-init first initializes the database service and then pushes all API configuration files of Auto Scaling to the pop configuration center to initialize OpenAPI.

- Trigger (dependent on ess-init)
  - Trigger executes tasks such as performing checks on the health, minimum number, and maximum number of instances as well as deleting scaling groups.
  - Trigger triggers scheduled tasks and event-triggered 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 tasks, such as creating ECS instances, adding instances to SLB backend server groups and RDS whitelists, and synchronizing Cloud Monitor group information.
  - It retries failed tasks and provides the rollback mechanism.

- service-test

  service-test is used for regression tests on the overall application running status. It contains more than 60 regression test cases to test the integrity of features.

4.3.1.2.2. Troubleshooting

This topic describes how to handle issues related to the business logic layer.

Procedure

1. Go to the Alibaba Cloud Support Platform to submit a ticket.
2. Check the status of services that depend on the business logic layer in the Apsara Infrastructure Management Framework console.

   If a service cannot be executed, the Auto Scaling business logic layer is affected. The following table describes the services and their impacts.

   Services and their impacts

Service name	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.

Service name	Key impact
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 RDS association requests become invalid.
tianjimon (Monitoring System)	Some services are unavailable.

4.3.1.3. Inspection

4.3.1.3.1. Overview

Auto Scaling inspection monitors the basic health conditions of clusters.

The inspected basic health conditions include the following aspects:

- Monitoring inspection
- Basic software package version inspection

4.3.1.3.2. Monitoring inspection

The monitoring inspection includes the basic monitoring and connectivity monitoring inspection.

4.3.1.3.3. Basic software package version inspection

The basic software package version inspection includes the version inspection for trigger, coordinator, worker, and base services.

4.4. Resource Orchestration Service (ROS)

4.4.1. Operations and Maintenance Guide

4.4.1.1. ROS component O&M

4.4.1.1.1. 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.

4.4.1.1.2. 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.

4.4.1.1.3. 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`

    ```
    [root@ ~]
    #/usr/local/sbin/rabbitmqctl list_vhosts
    Listing vhosts ...
    /
    /ros_0112
    ```

    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.

4.4.1.1.4. 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.

4.5. Object Storage Service (OSS)

4.5.1. Operations and Maintenance Guide

4.5.1.1. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager 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.

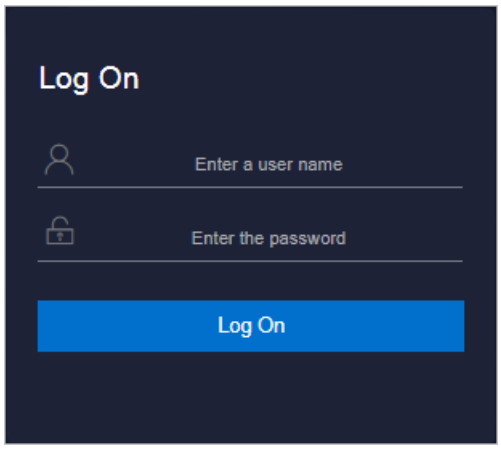

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

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.

4.5.1.2. O&M overview

4.5.1.2.1. System architecture

OSS is a storage solution that is built on the Apsara system. It is based on the infrastructure such as Apsara Distributed File System and SchedulerX. This infrastructure provides OSS and other Alibaba Cloud services with important features such as distributed scheduling, high-speed networks, and distributed storage.

The following figure shows the system architecture of OSS.

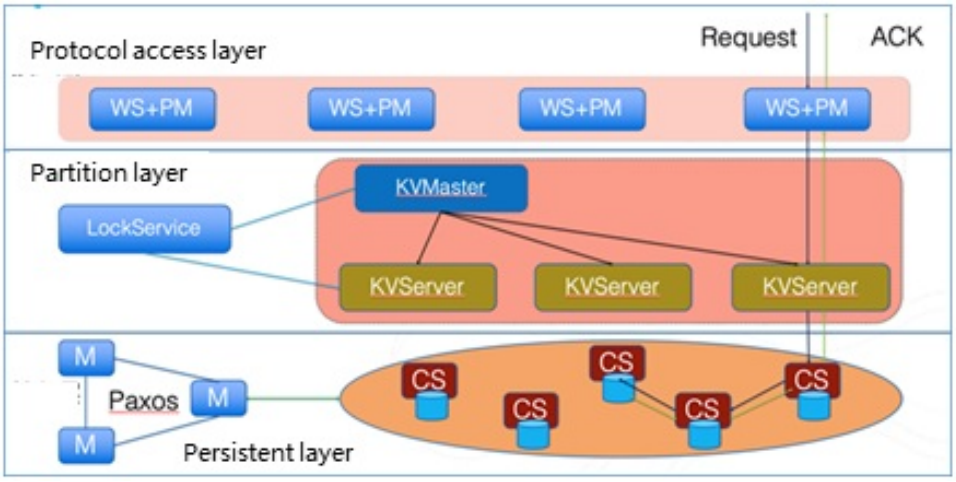

The OSS architecture is composed of three layers: protocol access layer, partition layer, and persistent storage layer.

- Protocol access layer
  - WS: uses the open-source Tengine component, and provides HTTP and HTTPS for external services.
  - PM: parses the HTTP request as the read/write operation on the back-end KV or another module. PM also receives and authenticates the user request sent through a RESTful protocol. If the authentication succeeds, the request is forwarded to KV Engine for further processing. If the request fails the authentication, an error message is returned.

- Partition layer

  The partition layer uses a highly scalable storage system with high performance and strong consistency to manage the index of a large amount of data. The index of objects stored in OSS is managed by range partitions. OSS divides the index of a large amount of objects into range partitions based on the loads and assigns partition servers to manage the range partitions. In addition, the partition layer supports a large number of concurrent requests and provides the automatic load balancing and garbage collection features.

  The index system uses the LSM tree structure that consists of KVMasters and KVServers. KVMaster manages and schedules partitions. KVServer stores indexes and actual data of partitions.

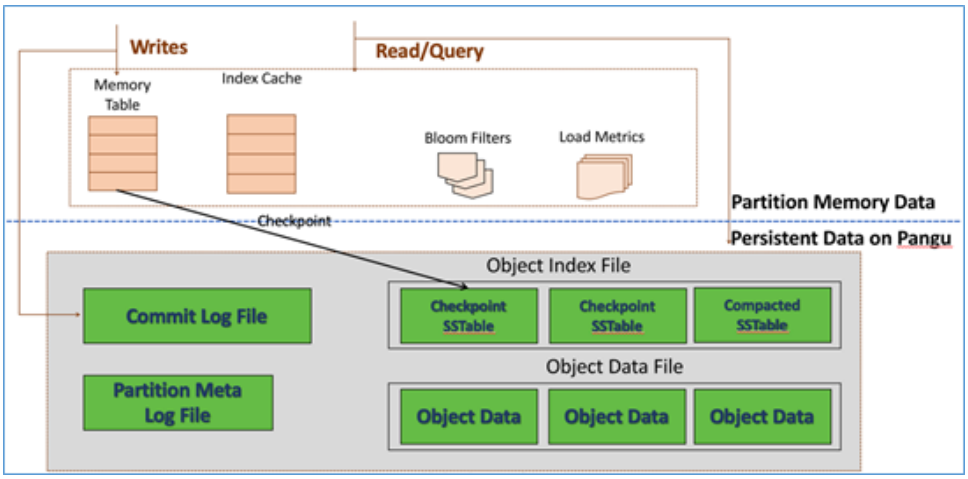

- Persistent layer

  The persistent layer provides a Paxos-based distributed file system that can store exabytes of data and provide high reliability and availability. Masters in this layer use the Paxos protocol to ensure that the metadata stored on the masters is consistent. This way, objects can be effectively stored and accessed in a distributed manner. In addition, data stored in the file system is backed up for redundancy and can be recovered when software or hardware errors occur.

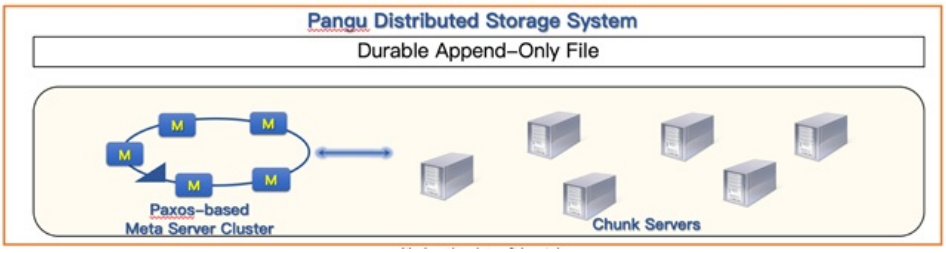

4.5.1.2.2. Component architecture

This topic describes the components of which OSS consists. OSS consists of the following components: front-end, authentication service, storage back-end, management and control, and billing quota.

The following figure shows the components of OSS.

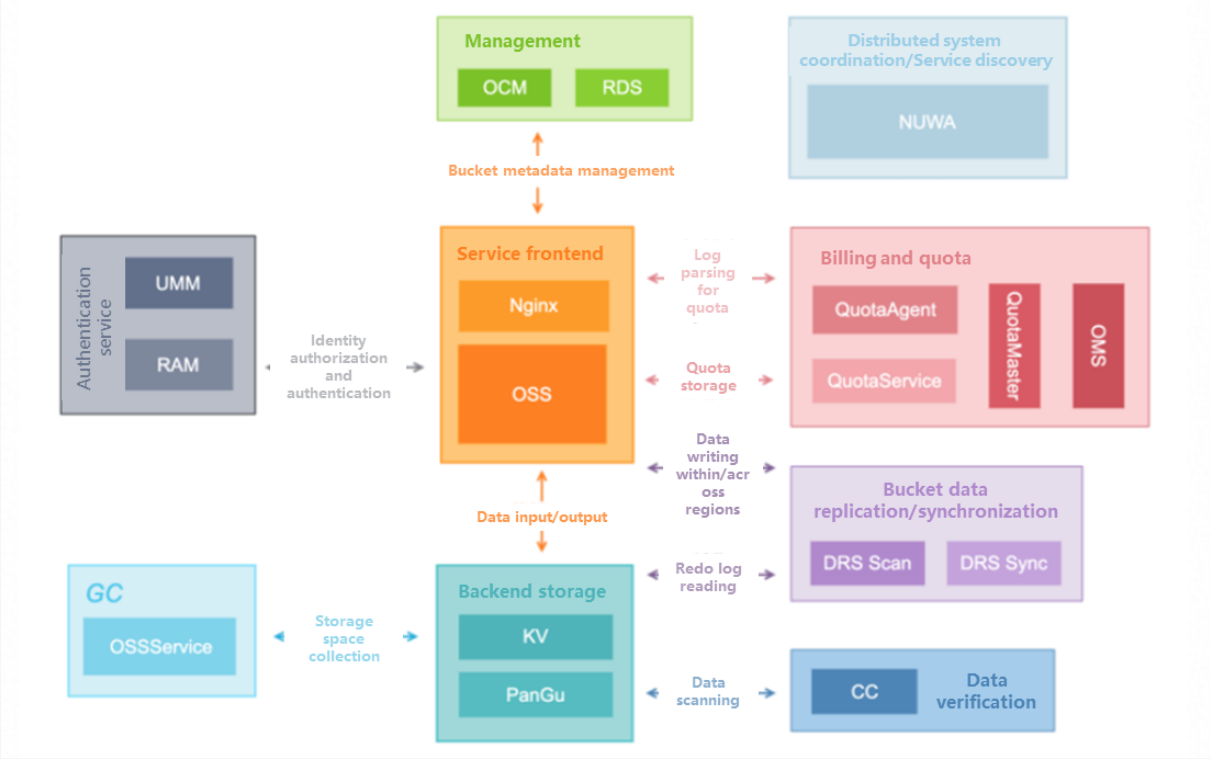

The following table describes the components of OSS.

Category	Component	Description
Front-end	nginx	The reverse proxy web server that redirects your requests to services such as OSS, IMG, or OSS select.
	OSS	The entry for requests that are sent to manage OSS buckets and objects.
Authentication service	UMM	The service that is used to identify users.
	RAM	The service used to authenticate access.
Storage back-end	KV	The data partition layer that provides interfaces based on key-value pairs for ordered access by keys. This component supports versioning and partition splitting.
	PanGu	The distributed file system that is used to permanently store data. • This component provides high-availability file storage. • This component provides different file models to accommodate different requirements.
Management	OCM	The component that is used to manage system information, including buckets, users, and clusters.
	RDS	The database service used to permanently store the management data of OCM.

Category	Component	Description
Billing quota	QuotaAgent	The agent deployed on each front-end server to analyze access logs, generate quota information about each front-end server, and send the information to QuotaService.
	QuotaService	The service that is used to collect and aggregate the quota information about front-end servers. This service also collects the storage statistics of KV and saves the statistics in a specific bucket.
	QuotaMaster	The service that is used obtain the quota information saved by QuotaService and send the information to OMS on a regular basis based on specific requirements.
	OMS	The external service that is used to generate detail reports of fees for users.
Data replication and synchronization across buckets	DRSscan	The component that is used to scan object-related operations from KV redo logs and create synchronization tasks.
	DRSsync	The component that is used to perform the operations described in synchronization tasks in the destination bucket.
Data verification	CC	The component that is used to scan data in KV and verify the integrity of the data.
OSSService		The component that provides data features, such as GC and lifecyle management. Asynchronous tasks for buckets or data are performed by OSS Service or in a manner that similar to OSS Service. For example, DRS and CC tasks are performed in the master/worker mode that is similar to OSS Service.
Distributed storage coordination and service discovery	NUWA	The service that is depended on by OSS Service, DRS, and CC. This service is used to provide the following features: • Elect a master among multiple control nodes to ensure that a service has only one master instance at the same time. • Store the IP addresses of service providers. • Elect the KVMaster in KV.

4.5.1.2.3. Deployment architecture

This topic describes the deployment architecture of Object Storage Service (OSS) in Apsara Stack. OSS is deployed in groups, including the agent groups, Apsara Distributed File System master groups, and independent front-end host groups.

The following figure shows the deployment architecture of OSS.

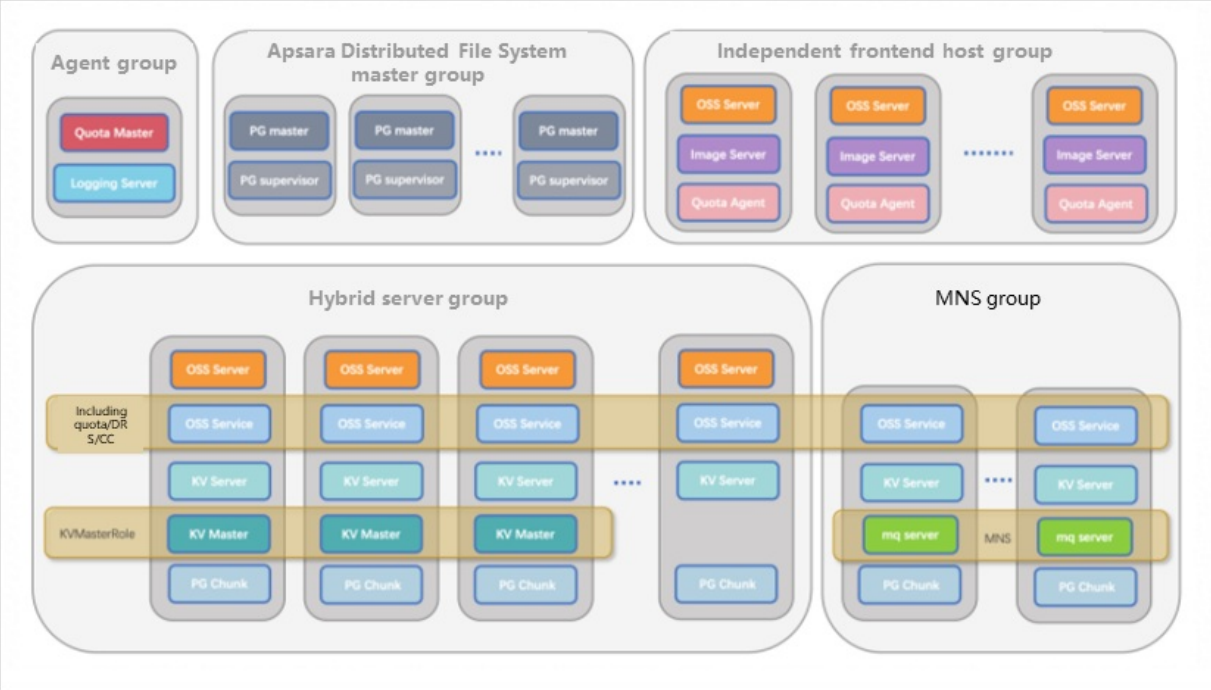

The following table describes the components in the deployment architecture of OSS.

Component	Description
Agent group	One or more agents are deployed in a cluster. The Quota Master service is deployed on an agent that is deployed in one of the clusters deployed in a region.
Apsara Distributed File System master group	Only the master and supervisor services of Apsara Distributed File System are deployed in each group.
Independent front-end host group	Only OSS servers with quota agents and image servers are deployed in each group.
Hybrid server group	Major services are deployed in each group, including major loads such as the access layer, storage layer, and asynchronous tasks. A quota agent is deployed together with the OSS server, which is not displayed in the figure.
MNS group	MNS is the messaging component used in DRS 1.0. The deployment of services in MNS groups is the same as that in hybrid server groups with the exception that OSS servers are not deployed.

4.5.1.2.4. O&M architecture

This topic describes the O&M architecture of OSS.

The following figure shows the O&M architecture of OSS.

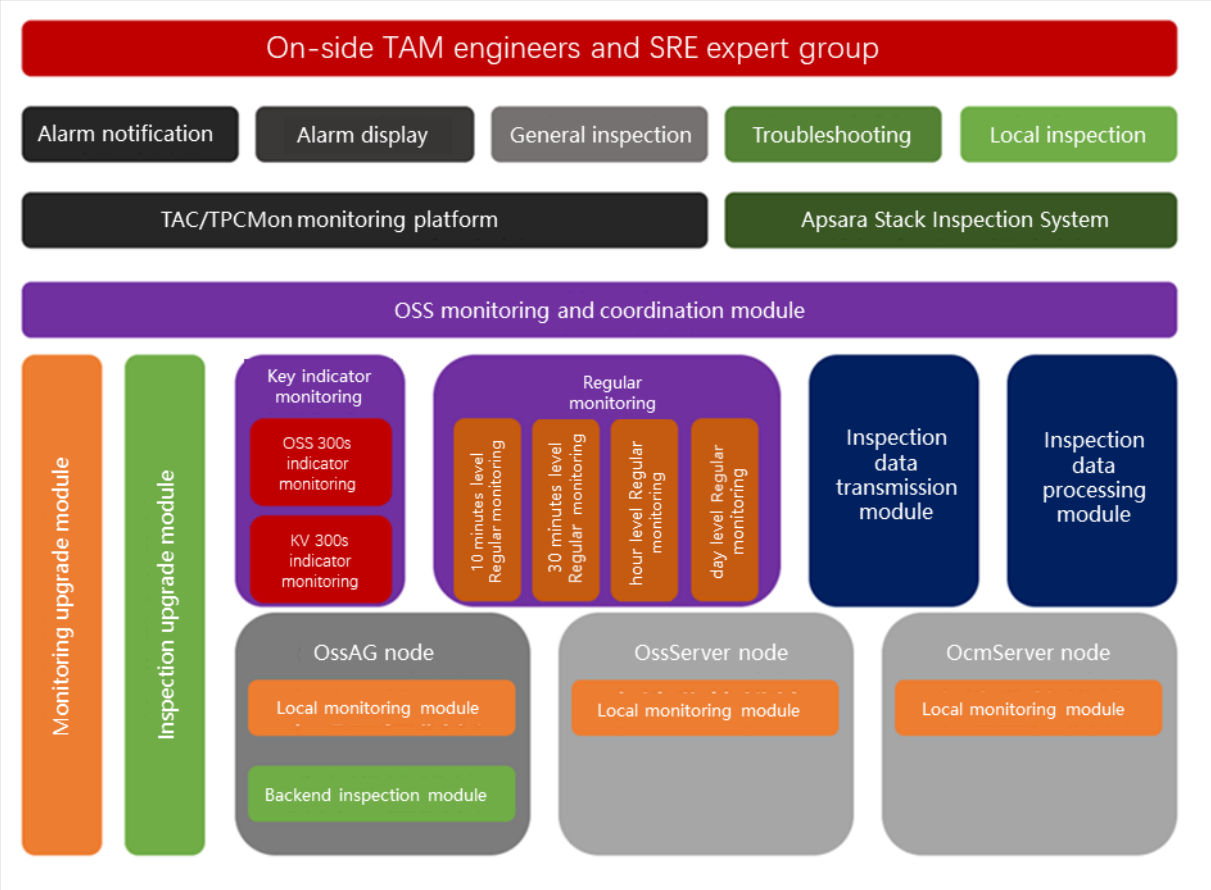

O&M category	Description
Routine maintenance	Includes local inspection, troubleshooting, general inspection, alert viewing, and alert handling.
Fault handling	Performed by on-site TAM engineers and SRE expert groups.

4.5.1.3. OSS operations and maintenance

4.5.1.3.1. User data

4.5.1.3.1.1. Basic bucket information

This topic describes how to query the basic information about a bucket, such as the cluster deployment location, configurations, current capacity, and the number of objects stored in the bucket.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Users**.
4. On the page that appears, select **Bucket Name** from the drop-down list, and then enter the name of the bucket that you want to query.

5. Click **View** to view the basic information about the bucket.

   You can also click the **Data Monitoring** tab to view the SLA, traffic, and QPS of the bucket.

4.5.1.3.1.2. User data overview

OSS allows you to view the statistical data of a user, such as the resource usage and the attributes of resources. You can measure the total resource usage of all buckets owned by a user on a specified day.

Context

The overview of user data is displayed only when you search the data by using UIDs or Alibaba Cloud accounts. You can view the following user data: total storage capacity, total inbound and outbound traffics (including traffics transferred over public networks, internal networks, and CDN), and total charged requests.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Users**.
4. On the page that appears, select **Alibaba Cloud Account** or **UID** from the drop-down list, and then enter the Alibaba Cloud account or UID whose data you want to view.
5. Click **View**. Then, click the **User Data Overview** tab.

   You can select a date and then click **View** to view the user data on a specified day.

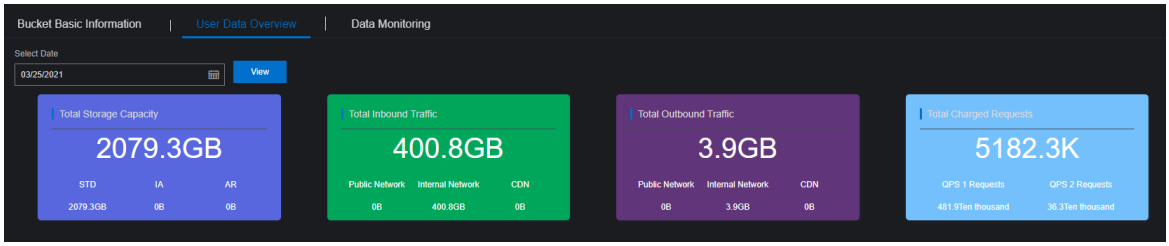

4.5.1.3.1.3. Data monitoring

This topic describes how to monitor the data of OSS in the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Users**.
4. On the page that appears, select **Alibaba Cloud Account** or **UID** from the drop-down list, and then enter the Alibaba Cloud account or UID whose data you want to monitor. Then, click **View**.
5. Click the **Data Monitoring** tab and configure the parameters described in the following table.

Parameter	Description
Bucket Name	Select the bucket you want to monitor from the drop-down list.

Parameter	Description
Specify Time Range	Select the time range that you want to monitor.
Monitoring Items	Select one or more of the following items that you want to monitor. ○ **SLA**: The service level agreement provided by OSS. The SLA of OSS can be calculated based on the following formula: **Number of non-5XX requests per 10 seconds or an hour/Number of valid requests x 100%** . ○ **HTTP Status**: The number of the returned 5XX, 403, 404, 499, 4XX_other, 2XX, and 3XX status codes and the percentage of the requests to which these codes are returned. ○ **Latency**: The latencies of API requests and the maximum value of the latencies. ○ **Storage Capacity**: The storage capacity of the bucket. ○ **Image Processing Capacity**: The number of images processed by using OSS. ○ **Traffic**: The traffic generated by the bucket. ○ **QPS**: The number of API requests sent to the bucket.

6. Click **View**.

4.5.1.3.2. Cluster data

4.5.1.3.2.1. Inventory monitoring

You can monitor the following metrics: total capacity, unused capacity, used capacity, and storage utilization.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Clusters**.
4. Click the **Inventory Monitoring** tab. You can select the data that you want to monitor from the Data Dimension drop-down list.

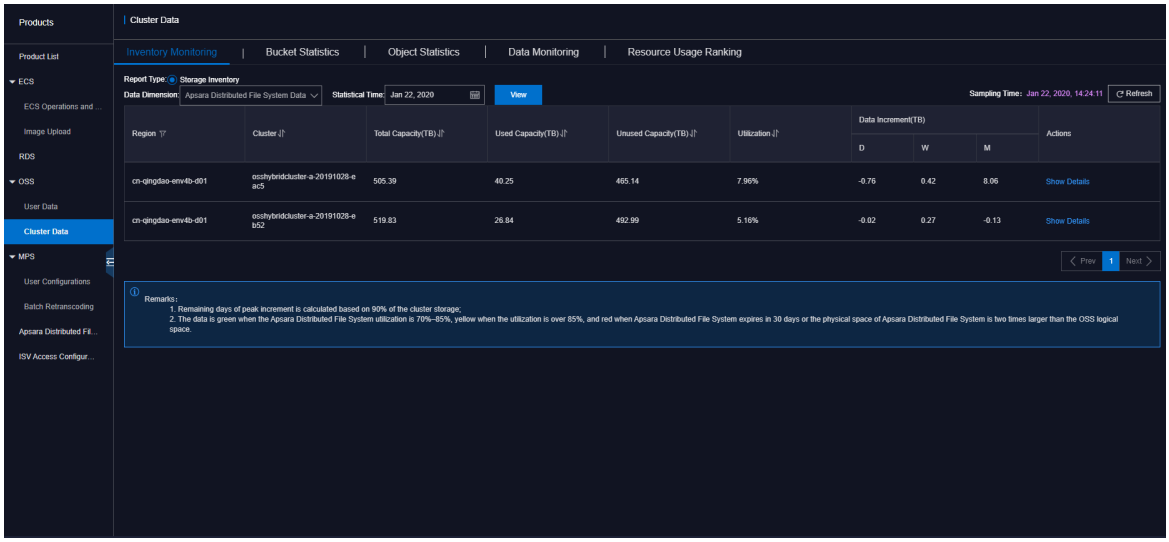

You can select to view the following data:

- Apsara Distributed File System Data: includes the total capacity, used capacity, unused capacity, utilization, and data increment of the cluster.
- Metric Data: includes the statistical capacity that shares and does not share ECS quota.
- KV Data: includes the logic KV data, KV data in the recycle bin, and data increment (by day, week, or month).

4.5.1.3.2.2. Bucket statistics

This topic describes how to measure the number of buckets created in each cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Clusters**.
4. Click the **Bucket Statistics** tab. You can select **Report Statistics**, **Current Overall Statistics**, or **Growth Trend** from the Display Method drop-down list to view the statistics in different methods.

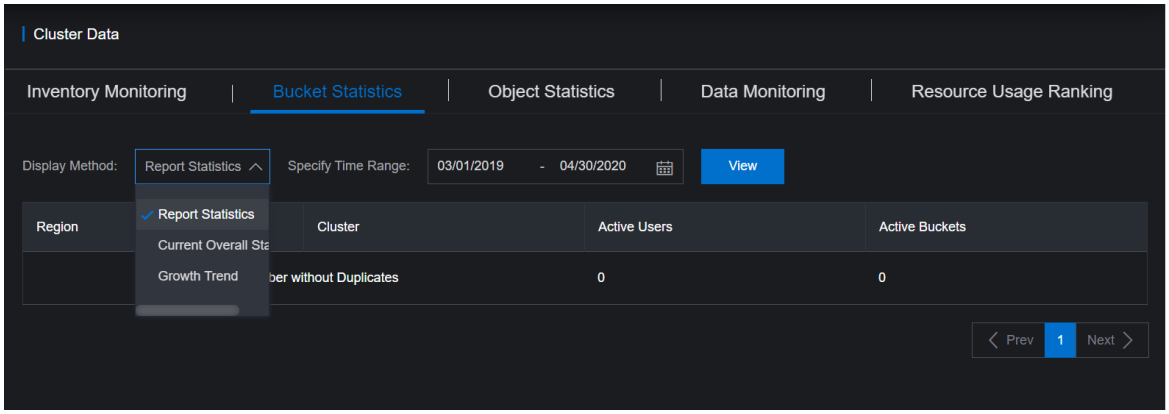

- If you select **Report Statistics**, specify a time range.
- If you select **Current Overall Statistics**, by default, the statistics that you query is generated based on the data in the last hour.

- If you select **Growth Trend**, you can select the following values from the Recent drop-down list to view the statistics within a specific range: *7 Days*, *30 Days*, *3 Months*, *6 Months*, and *1 Year*.

5. Click **View**.

4.5.1.3.2.3. Object statistics

This topic describes how to measure the number of objects stored in each cluster and the trend of the number.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Clusters**.
4. Click the **Object Statistics** tab. You can select **Current Overall Statistics** or **Growth Trend** from the Display Method drop-down list to view the statistics in different methods.

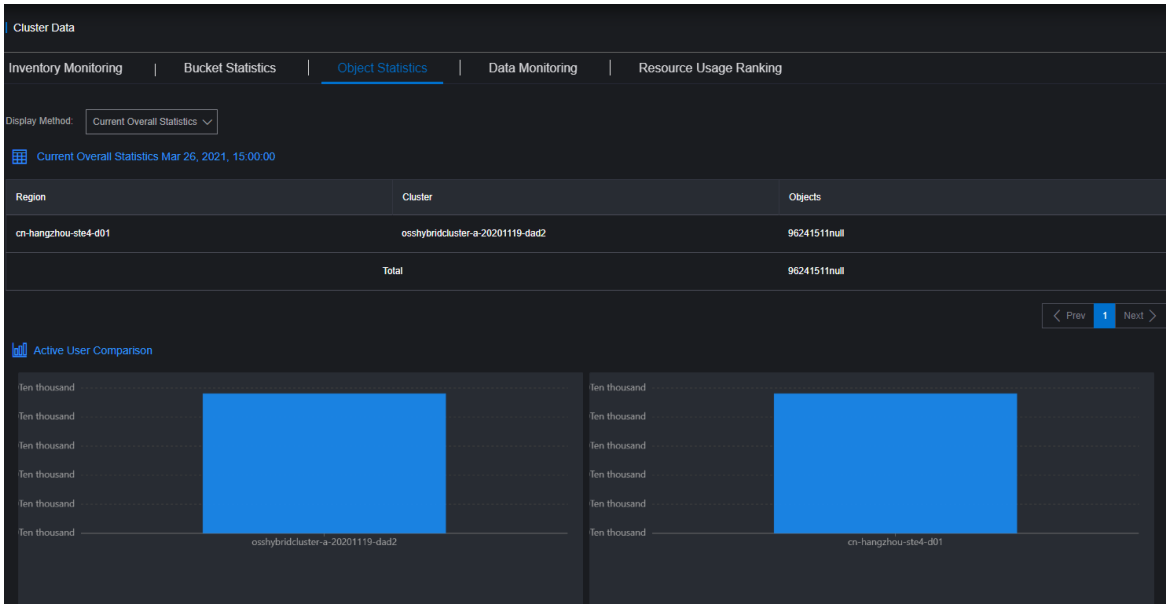

- If you select **Current Overall Statistics**, by default, the statistics that you query is generated based on the data in the last hour.
- If you select **Growth Trend**, you can select the following values from the Recent drop-down list to view the statistics within a specific range: *7 Days*, *30 Days*, *3 Months*, *6 Months*, and *1 Year*.

5. Click **View**.

4.5.1.3.2.4. Data monitoring

This topic describes how to view the metrics used to monitor data.

Context

Metrics used to monitor cluster data are the same as those used to monitor user data except that the monitored data are collected by cluster.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Clusters**.
4. Click the **Data Monitoring** tab, select the metrics from the **Monitoring Items** drop-down list and set a time range in the **Specify Time Range** field. Then, click **View**.

   Note You can select to view the following monitoring metrics:

   - SLA: The service level agreement provided by OSS. The SLA of OSS can be calculated based on the following formula: Number of non-5XX requests per 10 seconds or an hour/Number of valid requests x 100%.
   - Traffic: The traffic generated by the bucket, including the inbound and outbound traffics transferred over public networks, internal networks, and CDN and the inbound and outbound traffic for synchronization.
   - QPS: The number of the following charged requests sent to the bucket: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.
   - Latency: The latencies of API requests such as PutObject, GetObject, and UploadPart and the maximum value of the latencies.
   - HTTP Status: The number of the returned 5XX, 403, 404, 499, 4XX_other, 2XX, and 3XX status codes and the percentage of the requests to which these codes are returned.
   - Storage Capacity: The storage capacity of the bucket, including the capacity and increment of the Standard, IA, and Archive storage.

5. Move the pointer over the trend chart to display data at a specific point in time.

   Data monitoring 1

   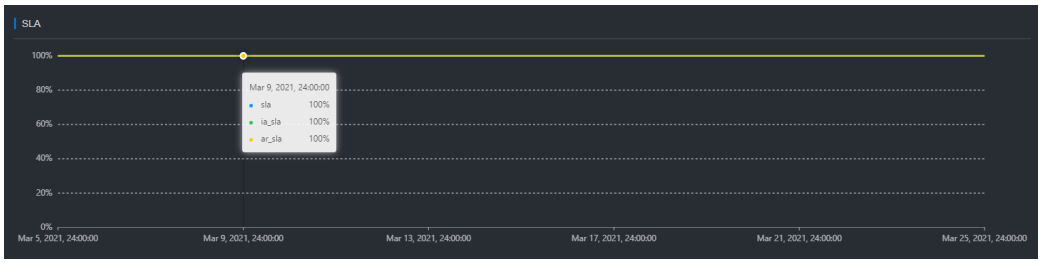

   You can view the following monitoring metrics in the chart:

   - 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 on the ratio of the counts of 5xx, 403, 404, 499, 4xx_others, 2xx, and 3xx status codes to total requests.
   - Latency: collects statistics on the latency of APIs such as put_object, get_object, and upload_part as well as the maximum latency.
   - Storage capacity: collects statistics on the storage capacity of standard, infrequent access (IA), and archive storage and their increments.
   - Image Processing Capacity: The number of images processed by using OSS.

> **Note** By default, this metric is not displayed. You can select this metric from the **Monitoring Items** drop-down list.

- Traffic: The traffic generated by the bucket, including the inbound and outbound traffics transferred over public networks, internal networks, and CDN and the inbound and outbound traffic for synchronization.
- QPS: The number of the following charged requests sent to the bucket: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.

The following examples show the common operations that you can perform on the data monitoring trend chart:

- If you want to query the monitored data by user, you can click the bucket name in the trend chart to show or hide the curve.

  Data monitoring 2

  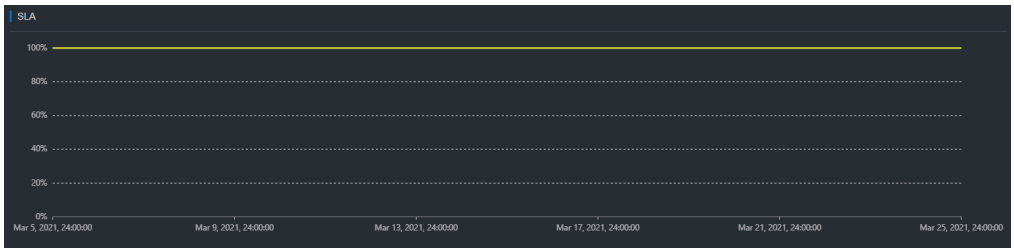

- Move the pointer over the trend chart to display data at a specific point in time.

  Data monitoring 2

  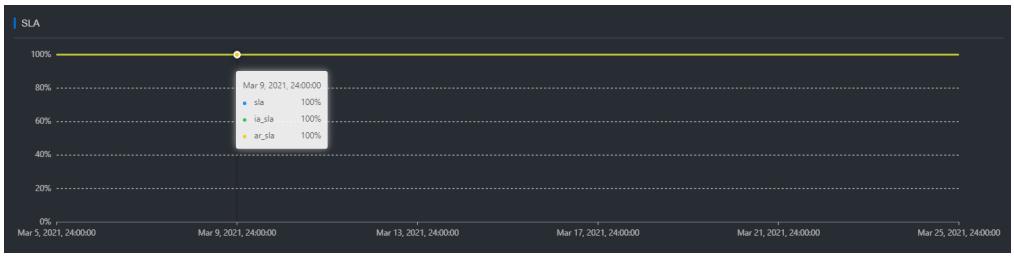

4.5.1.3.2.5. Resource usage rankings

This topic describes how to measure the 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 top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > OSS Clusters**.
4. Click the **Resource Usage Ranking** tab and then select **Report** or **Trend** from the Display Method drop-down list. Select a number from the **Top** drop-down list. Select the metrics from the **Monitoring Items** drop-down list and set a time range in the **Specify Time Range** field to view resource usage.

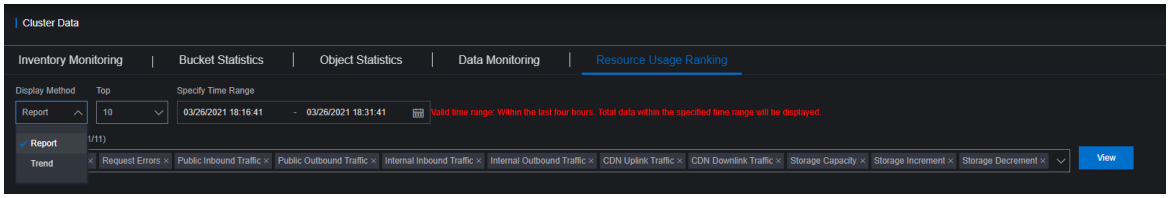

   - If you select the **Report** method, you can view the resource usage of the top **10**, **30**, or **50** buckets in reports.
   - If you select the **Trend** method, you can view the resource usage of the top **10** buckets in trend charts.
5. Click **View**.

4.5.1.4. OSS Operations and Maintenance System

4.5.1.4.1. User and bucket O&M

4.5.1.4.1.1. User data query

User data overview

Object Storage Service (OSS) allows you to view the statistical data of a user, such as the resource usage and the attributes of resources. You can measure the total resource usage of all buckets owned by a user on a specified day.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **User and Bucket O&M > User Data Query**.
5. On the page that appears, select **Apsara Stack Tenant Account** or **UID** from the drop-down list, and then enter the Apsara Stack tenant account or the UID whose data you want to monitor. Then, click **Query**.
6. Click **User Data Overview**. In the **Select Date** field, select a specific date. Then, click **View**.

   On the current page, you can view resource usage, including the total storage capacity, total inbound traffic, total outbound traffic, and total billed API operation calling.

Data monitoring

This topic describes how to monitor Object Storage Service (OSS) data in the OSS Operations and Maintenance System.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **User and Bucket O&M > User Data Query**.
5. On the page that appears, select **Apsara Stack Tenant Account** or **UID** from the drop-down list, and then enter the Apsara Stack tenant account or the UID whose data you want to monitor. Then, click **Query**.
6. Click the **Data Monitoring** tab and configure the parameters. The following table describes the parameters.

Parameter	Description
Bucket	Select the bucket you want to monitor from the drop-down list. You can select multiple buckets. When you select multiple buckets, you can view the total usage of different resources of these buckets.
View Number of Users	You can view the all resource usage except the following items: total storage usage, incremental storage usage, Standard storage usage, incremental Standard storage usage, Infrequent Access (IA) storage usage, incremental IA storage usage, Archive storage usage, incremental Archive storage usage, and resources for Image Processing (IMG). If you select View Number of Users, **Bucket** does not appear.
Select Time Range	Select the time range during which statistics you want to view are generated.
Sampling Interval	Select the interval at which to sample data.

Parameter	Description
Metric	Select one or more of the following metrics that you want to monitor. ○ **SLA**: the service-level agreement (SLA) provided by OSS. The SLA of OSS can be calculated based on the following formula: **Number of non-5xx requests per 10 seconds or per hour/Number of valid requests × 100%** . ○ **HTTP Status**: the number of the returned 5xx, 403, 404, 499, other 4xx, 2xx, and 3xx status codes and the percentage of the requests to which these codes are returned. ○ **Latency**: the latencies of API requests and the maximum value of the latencies. ○ **Storage Capacity**: the storage capacity of the bucket. ○ **Traffic**: the traffic generated by the bucket. ○ **QPS**: the number of API requests sent to the bucket.

7. Click **View**.

Basic bucket information

This topic describes how to query the basic information about a bucket, such as the cluster deployment location, configurations, current capacity, and the number of objects stored in the bucket.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane, choose **User and Bucket O&M > User Data Query**.
5. On the page that appears, select **Bucket Name** from the drop-down list, and then enter the name of the bucket that you want to query.
6. Click **Query** to view the basic information about the bucket.

   You can also click the **Data Monitoring** tab to view the service-level agreement (SLA), traffic, and queries per second (QPS) of the bucket.

4.5.1.4.1.2. Bucket management

You can use the bucket management feature to manage your users and buckets. For example, you can modify the status of users, status of buckets, and access control lists (ACLs) of buckets.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the

Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.

4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **User and Bucket O&M > Buckets**.
5. On the page that appears, select **Apsara Stack Tenant Account**, **UID**, or **Bucket Name** from the drop-down list, and then enter the Apsara Stack tenant account, UID, or bucket name whose data you want to monitor. Then, click **View**.
6. Manage your users or buckets.
   - **User Basic Information**

     In the **User Basic Information** section, you can modify the status of users and the maximum number of buckets that can be created. You can also view the change history of user status.

     - Click **Modify user status**. You can modify the status of the current user as prompted.

       > **Warning** Security risks may arise if you modify user status. We recommend that you do not modify user status if unnecessary.

     - Click **Modify the number of buckets that can be created**. You can modify the maximum number of buckets that the current user can create. Valid values: 0 to 10000.
     - Click **View user status change history**. You can view the change history of user status.
   - **Bucket list information**

     In the **Bucket list information** section, you can manage your buckets.

     - Click **View Details** in the Actions column that corresponds to the required bucket. You can view the basic information about the bucket
     - Click **Manage** in the Actions column that corresponds to the required bucket. You can modify the status and the ACL of the bucket.
     - Click **Delete** in the Actions column that corresponds to the required bucket. You can delete the bucket.

       > **Warning** If you delete a bucket, the bucket and all objects in the bucket are deleted. Deleted buckets and objects cannot be recovered. Exercise caution when you delete a bucket.

4.5.1.4.1.3. Unbind a bucket

This topic describes how to unbind a bucket from a virtual private cloud (VPC) in the Object Storage Service (OSS) Operations and Maintenance System.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.

4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **User and Bucket O&M > Remove Domain Name Mapping from Bucket**.
5. In the **Search Bucket** search bar, enter the name of the required bucket. Then, click **View**.
   - If the bucket is not bound to a VPC, the system displays a message, which indicates that the bucket is not bound to a VPC.
   - If the bucket is bound to a VPC, you can unbind the bucket to the VPC on the unbind bucket page.

4.5.1.4.1.4. QoS configurations

You can modify the quality of service (QoS) configurations for users and buckets in the Object Storage Service (OSS) Operations and Maintenance System.

Configure QoS for users

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **User and Bucket O&M > QoS Configurations**.
5. Click the **User QoS** tab. Then, click **Add Configuration**.
6. In the **Add Configuration** dialog box, configure the QoS parameters. The following table describes the QoS parameters.

Parameter	Description
uid	Enter the UID of the specified user. An asterisk (*) indicates all resources.
cluster	Select the required cluster from the drop-down list.
total_upload_flow(Gb/s)	Select the total upload bandwidth from the drop-down list.
total_download_flow(Gb/s)	Select the total download bandwidth from the drop-down list.
total_qps	Select the total queries per second (QPS) from the drop-down list.
total_active_req	Select the total concurrent responses.

7. Click **More Configuration**. You can configure other parameters.

Parameter	Description

Parameter	Description
extranet_upload_flow(Gb/s)	Select the bandwidth for uploads over the Internet from the drop-down list.
intranet_upload_flow(Gb/s)	Select the bandwidth for uploads over the internal network from the drop-down list.
extranet_download_flow(Gb/s)	Select the bandwidth for downloads over the Internet from the drop-down list.
intranet_download_flow(Gb/s)	Select the bandwidth for downloads over the internal network from the drop-down list.
extranet_qps	Select the QPS over the Internet from the drop-down list.
intranet_qps	Select the QPS over the internal network from the drop-down list.
Remark	Add the remarks.

8. Click the + icon next to **More Configurations**. You can configure parameters related to CPU, Select, and Mirror.
9. Click **OK**.

Configure QoS for buckets

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **User and Bucket O&M > QoS Configurations**.
5. Click the **Bucket QoS** tab. Then, click **Add Configuration**.
6. In the **Add Configuration** dialog box, configure QoS parameters.

   Set the **Bucket** parameter to the name of the destination bucket or an asterisk (*). Configure other parameters in the same manner you configure QoS parameters for users.
7. Click **OK**.

4.5.1.4.2. Cluster O&M

4.5.1.4.2.1. Cluster data

Cluster overview

In the Object Storage Service (OSS) Operations and Maintenance System, you can view the running status of a cluster, such as data monitoring, bucket statistics, and object statistics.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane, choose **Cluster O&M > Cluster Monitoring > Cluster Overview**.
5. On the **Cluster Overview** page, view information about a cluster such as the storage class, service-level agreement (SLA) this year, total requests with 5xx today, total requests today, total Internet inbound traffic today, and total internal network outbound traffic today.

Inventory monitoring

You can monitor the following metrics: total capacity, unused capacity, used capacity, backup ratio, and storage usage.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Cluster O&M > Cluster Monitoring > Inventory Monitoring**.
5. View Inventory Monitoring. Inventory Monitoring includes **Storage Overview** and **Storage Inventory**.
   - Storage Overview

     Storage Overview provides the capacity and usage of Apsara Distributed File System, logical space used by KV data of the recycle bin, logical space used to synchronize data within OSS buckets, and physical space that is not cleared by garbage collection (GC) tasks after users delete OSS data. Click **Export** to export storage overview information as an Excel.
   - Storage Inventory

     Storage Inventory provides usage of clusters. Click **Export**. You can export storage inventory information as an Excel.

> **Note**
>
> - Green indicates the Apsara Distributed File System usage of 70% to 85%. Yellow indicates the Apsara Distributed File System usage of 85% to 90%. Red indicates the Apsara Distributed File System usage of greater than 90%.
> - If the Apsara Distributed File System is available for less than 30 days or the physical space of the Apsara Distributed File System is twice greater than the logical space of OSS, the information is displayed in red.
> - The number of remaining days for peak increments is calculated based on the 90% storage of a cluster by using the following formula: `(Remaining capacity of the Apsara Distributed File System - 0.1× Total capacity of the Apsara Distributed File System)/Average top 10 daily increments in a month)` .

Aside from basic cluster information such as the cluster name and the region, you can also view the following statistics:

- 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 Elastic Compute Service (ECS) instances and other instances.
- KV Data: includes the logical space of KV data in the recycle bin, physical space of KV data in the recycle bin, percentage of the recycle bin space, percentage of the recycle bin space in three days, redo logs, actual backup ratio, ratio of actual backups in three days, redundant backup, and data increment (by day, week, or month).

Click **View Details** in the **Actions** column that corresponds to a region.

On the **View storage capacity details** page, check **Storage that is not physically cleared by OSS GC tasks**, **Logical storage that is not merged by KV**, and **KV layer has been merged but GC quantity has not been completed**. You can check disk space issues.

Data monitoring

This topic describes how to view the metrics used to monitor data by cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Cluster O&M > Cluster Monitoring > Data Monitoring**.
5. On the **Data Monitoring** page, set **Select Time Range** and **Metric**. Click **View**.

   You can choose to view the following monitoring metrics:

   - **SLA**: the service-level agreement (SLA) provided by OSS. The SLA of OSS can be calculated based on the following formula: `Number of non-5xx requests per 10 seconds or per hour/Number of valid requests × 100%` .

- **HTTP Status**: the number of the returned 5xx, 403, 404, 499, other 4xx, 2xx, and 3xx status codes and the percentage of the requests to which these codes are returned.
- **Latency**: the latencies of API requests and the maximum value of the latencies.
- **Storage Capacity**: the storage capacity of the bucket.
- **Traffic**: the traffic generated by the bucket.
- **QPS**: the number of API requests sent to the bucket.

Resource usage rankings

This topic describes how to measure the usage of resources by cluster. This way, administrators can monitor users who consume more resources.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Cluster O&M > Cluster Monitoring > Resource Usage Ranking**.
5. Click the **Resource Usage Ranking** tab and then select **Report** or **Trend** from the **Display Method** drop-down list. Select a number from the **Top** drop-down list. Select the metrics from the **Metric** drop-down list and set a time range in the **Specify Time Range** field to view resource usage.

   Data resource ranking items include:

   - Total Requests
   - Request Errors (5xx, 403, and 499)
   - Internet Inbound Traffic and Internet Outbound Traffic
   - Internal Network Inbound Traffic and Internal Network Outbound Traffic
   - CDN Upstream Traffic and CDN Downstream Traffic
   - Inbound VPC Traffic and Outbound VPC Traffic
   - Storage Capacity, Storage Increment, and Storage Decrement
6. Click **View**.

Bucket statistics

This topic describes how to measure the number of buckets created in each cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance

System, choose **Cluster O&M > Cluster Monitoring > Bucket Statistics**.

5. On the **Bucket Statistics** tab, set **Display Method** to **Report Statistics**, **Current Overall Statistics**, or **Growth Trend**.
   - If you select **Report Statistics**, you must configure **Select Time Range**.
   - If you select **Current Overall Statistics**, the statistics that you query are generated based on the data in the last hour by default.
   - If you select **Growth Trend**, you can configure **Select Time Range** or select **7 Days**, **30 Days**, **3 Months**, **6 Months**, or **1 Year** from the **Recently** drop-down list.
6. Click **View**.

Object statistics

This topic describes how to measure the number of objects stored in each cluster and the trend of the number.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Cluster O&M > Cluster Monitoring > Object Statistics**.
5. On the **Object Statistics** tab, set **Display Method** to **Current Overall Statistics** or **Growth Trend**.
   - If you select **Current Overall Statistics**, the statistics that you query are generated based on the data in the last hour by default.
   - If you select **Growth Trend**, you can configure **Select Time Range** or select **7 Days**, **30 Days**, **3 Months**, **6 Months**, or **1 Year** from the **Recently** drop-down list.
6. Click **View**.

4.5.1.4.2.2. Server management

View the information about servers

In the Object Storage Service (OSS) Operations and Maintenance System, you can view the information about servers in a cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Cluster O&M > Server Management**.

5. Select the required cluster from the **Cluster** drop-down list. Then, click **Query**.

   In the list of servers, all servers in the cluster are listed. The hostnames, IP addresses, and status of servers appear. To view more information about servers, you can also click the following buttons:

   - **Monitoring alarm**

     Click **Monitoring and alarm** next to the required host. You can view information such as the system load, CPU utilization, and memory usage of the host.

   - **Details**

     - Click **Machine information**. You can view the basic information of the host, such as the hostname, IP address, and SN.
     - Click **Disk information**. The **Disk Management** page appears. You can view the status of the disk.

Publish servers

In the Object Storage Service (OSS) Operations and Maintenance System, you can publish servers added to a cluster.

Prerequisites

New servers are added to the required cluster.

Before you publish servers, you must add the servers to the required cluster. For more information about how to add servers to a cluster, contact technical support.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Cluster O&M > Server Management**.
5. Click **Add Machine**.
6. On the **Create task** tab, configure the following parameters. The following table describes the parameters.

Parameter	Description
Cluster	Select a cluster where the required server is deployed.
IP	Specify the IP address of the required server.
Host	Specify the hostname of the server.

Parameter	Description
Timing execution	Set the time when the task is scheduled to execute. ○ **Do not set timing execution**: The task is executed immediately after it is submitted. ○ **Execute after the specified time point**: The task is executed after the specified point in time. ○ **Each day's specified interval execution**: The task is executed within the specified time range each day. ○ **Specify interval execution**: The task is executed within the specified time range.
Execution priority	Set the priority based on which to execute the task. A smaller value indicates a higher priority.
Remarks	Set the remarks of the task.

7. Click **Submit**.

Unpublish servers

In the Object Storage Service (OSS) Operations and Maintenance System, you can unpublish servers you no longer use in a cluster.

Procedure

Warning Security risks may arise if you unpublish servers. Exercise servers when you unpublish servers.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Cluster O&M > Server Management**.
5. Select the required cluster from the **Cluster** drop-down list. Then, click **Query**.
6. Click **Machine offline** in the Operation column that corresponds to the required server.
7. On the **Create Task** tab, configure the following parameters. The following table describes the parameters.

Parameter	Description
Cluster	Specify the cluster where the required server resides.
IP	Specify the IP address of the required server.
Host	Specify the hostname of the required server.

Parameter	Description
Timing execution	Set the time when the task is scheduled to execute. ○ **Do not set timing execution**: The task is executed immediately after it is submitted. ○ **Execute after the specified time point**: The task is executed after the specified point in time. ○ **Each day's specified interval execution**: The task is executed within the specified time range each day. ○ **Specify interval execution**: The task is executed within the specified time range.
Execution priority	Set the priority based on which to execute the task. A smaller value indicates a higher priority.
Remarks	Set the remarks of the task.

8. Click **Submit**.

4.5.1.4.2.3. Disk management

In the Object Storage Service (OSS) Operations and Maintenance System, you can view the disk usage of all servers in a cluster.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Cluster O&M > Disk Management**.
5. Select the required cluster from the **Specify a Cluster** drop-down list. Then, click Query.

   In the list of disks, the types, IDs, total capacity, and used storage space of the disks appear.

4.5.1.4.2.4. Same endpoint used to access different regions

In the Object Storage Service (OSS) Operations and Maintenance System, you can use the same endpoint to access different regions in a cluster.

Enable Same Endpoint Across Regions

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the

Products page. Then, in the Storage Service section, click **OSS Storage Operations and Maintenance System**.

4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Cluster O&M > Network Management > Same Endpoint Across Regions**.
5. Click **All on**.

View the endpoint

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. In the Apsara Stack O&M section, click **Apsara Infrastructure Management Framework**.
4. Click **Reports**.
5. In the top navigation bar, choose **Reports > System Reports**.
6. In the list, click **Registration Vars of Services**.
7. On the **Registration Vars of Services** page, find the **oss-server** row. Right-click the **Service Registration** column and choose **Show More** from the shortcut menu.
8. Click **Formatted Value**.

   The value of **oss-same-endpoint** is the endpoint for access across regions.

4.5.1.4.3. Service O&M of OSS

4.5.1.4.3.1. OSS version query

In the Object Storage Service (OSS) Operations and Maintenance System, you can query the version of OSS.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > OSS Version Query**.
5. Select a region, cluster, and version. Click **View**.

   The list contains information such as the region, cluster, server, and version.

4.5.1.4.3.2. Location query

In the Object Storage Service (OSS) Operations and Maintenance System, you can query the location information of OSS.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > Location Query**.
5. In the **Search for Location** search bar, enter the name of the location. Click **Query**.

4.5.1.4.3.3. OSS service restart

In the Object Storage Service (OSS) Operations and Maintenance System, you can restart OSS.

Procedure

Warning Security risks arise when you restart OSS. Exercise caution when you restart OSS.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > OSS Service Restart**.
5. On the **OSS Service Restart** page, configure the following parameters. The following table describes the parameters.

Parameter	Description
Cluster	Select the required cluster from the drop-down list.
Server Role	Select the required server role from the drop-down list.
IP Address	Select the required IP address from the drop-down list.

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

4.5.1.4.3.4. OSS backend task management

Backend task monitoring

In the Object Storage Service (OSS) Operations and Maintenance System, you can view the running status of backend tasks, including lifecycle and garbage collection (GC) tasks.

View the running status of lifecycle rules

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.

3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > OSS Background Task > Task Monitoring > Lifecycle Management**.
5. On the **Lifecycle Monitoring Result** page, specify the date. Click **View**.

   The list displays information such as the running status of tasks within the specified time, running duration, and the number of deleted objects.

View the running status of GC tasks

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > OSS Background Task > Task Monitoring > Garbage Collection**.
5. On the **Garbage Collection Monitoring Result** page, specify the date. Click **View**.

   The list displays information such as the running status of tasks within the specified time, running duration, and the number of deleted objects.

Backend task configurations

In the Object Storage Service (OSS) Operations and Maintenance System, you can configure lifecycle rules and the time when garbage collector (GC) tasks are run.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > OSS Background Task > Task Configuration**.
5. On the **Task Configuration** page, select the required cluster.

   You can perform the following operations in OSS:

   - Click **Query** in the **Lifecycle Configurations** list. You can query the lifecycle configurations. Click Query in the **GC Configurations** list. You can query the GC task configurations.
   - Click **Modify** in the **Lifecycle Configurations** list. You can modify lifecyle rule configurations. Click Modify in the **GC Configurations** list. You can configure the period during which GC tasks are run.
   - Click **Run Lifecycle**. You can run lifecycle rules. Click **Run GC**. You can run GC tasks.

4.5.1.4.3.5. Synchronization management

In the Object Storage Service (OSS) Operations and Maintenance System, you can create and manage synchronization tasks.

Configure cross-cloud synchronization

For more information about how to configure cross-cloud synchronization, see *Configure cross-cloud replication* in *User Guide*.

Synchronization link management

You can view the cross-cloud synchronization tasks configured for the required cluster and the task running status.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - OSS > Synchronization Management > Source and Destination**.
5. On the **Source and Destination** page, select the required cluster. Click **Query**.

   The list displays the synchronization tasks configured for the cluster, source and destination of the tasks, and synchronization status.

   - Click **Details** next to the required task. You can view the detailed information of the synchronization task.
   - Click **State** next to the task. You can view the detailed running status of the synchronization task.

Distribution of DRS tasks

On the DRS Task Distribution Monitoring page, you can view the times cross-cloud synchronization tasks are performed.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Service O&M - OSS > Synchronization Management > DRS Task Distribution**.
5. On the **DRS Task Distribution Monitoring** page, select the required cluster.

   You can view information such as the synchronization tasks configured for the cluster, source and destination of the tasks, number of historical synchronization tasks, and number of incremental data synchronization tasks.

Cluster synchronization monitoring

You can view the running status of synchronization tasks between two clusters.

1. Log on to the Apsara Uni-manager Operations Console.

2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Service O&M - OSS > Synchronization Management > Cluster Synchronization**.
5. On the **Double-Cluster Synchronization(Sync) Monitoring Result** page, set Select Time Range. Click **View**.

   The list displays the times errors occur when synchronization tasks are run, the number of synchronized objects, and the number of parts.

4.5.1.4.4. Service O&M of KV

4.5.1.4.4.1. KV CheckReady management

In the Object Storage Service (OSS) Operations and Maintenance System, you can check whether KV resources are ready.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV CheckReady Management**.
5. On the **KV CheckReady Management** page, select the required cluster and the app. Click **Execute CheckReady**.

   If `1.checkready:PASS` appears, KV resources are ready.

4.5.1.4.4.2. KV Master management

You can manage KV Masters in the Object Storage Service (OSS) Operations and Maintenance System. For example, you can view configuration details, create snapshots, and switch KV Masters.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV Master Management**.
5. On the **KV Master Management** page, select a required cluster. Click **Confirm**.

   In the list, information such as the IP addresses, status, and snapshot IDs of servers in the required

cluster appears. Perform the following operations based on your business requirements:

- Click **buildinfo** in the Version Information column that corresponds to the required server. You can view the version information of the server.
- Click **Configuration Details** in the Actions column that corresponds to the server. You can view the configurations of the server.
- Click **Create Snapshots** in the Actions column that corresponds to the server. In the dialog box that appears, click **Yes**. Create a snapshot for the server.
- The **state** column indicates the status (Following or Leading) of the server in Master. If the state of the server is Following, you can switch the state of the server.

  Click **Switch to Leader** in the Actions column that corresponds to the server. In the dialog box that appears, click **Yes**.

4.5.1.4.4.3. KV server management

In the Object Storage Service (OSS) Operations and Maintenance System, you can add KV server records to or remove KV server records from a blacklist.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV Server Management**.
5. On the **KV Server Management** page, select the required server from the **Specify a Server** drop-down list or select the required cluster from the **Specify a Cluster** drop-down list.
   - If you set **Specify a Server**:
     a. Enter the IP address or hostname of the server. Click **View**.
     b. Click **Add to Blacklist**.
     c. In the **Add the server to the blacklist** dialog box, select a reason from the **Cause** drop-down list. Click **OK**.
     d. In the **OK** dialog box, click **yes**.
   - If you set **Specify a Cluster**:
     a. Set **Region** and **Cluster**. Click **View**.
     b. In the **KV Server Online Services** and **Abnormal KV Servers** sections, you can also click **Add to Blacklist** in the Actions column that corresponds to the required server.
     c. In the **Add the server to the blacklist** dialog box, select a reason from the **Cause** drop-down list. Click **OK**.
     d. In the **OK** dialog box, click **yes**.

In the **KV Server Blacklist** section, you can also click **Remove from Blacklist** in the Actions column that corresponds to the required server to remove the server from the blacklist or click **Add to Blacklist** in the Actions column that corresponds to the server to add the server to the blacklist.

4.5.1.4.4.4. KV service restart

In the Object Storage Service (OSS) Operations and Maintenance System, you can restart the KV service.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV Service Restart**.
5. On the **KV Service Restart** page, configure the following parameters. The following table describes the parameters.

Parameter	Description
Cluster	Select the required cluster from the drop-down list.
Server Role	Select the required server role from the drop-down list.
IP Address	Select the required IP address from the drop-down list.

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

4.5.1.4.4.5. KV application management

In the Object Storage Service (OSS) Operations and Maintenance System, you can manage KV applications.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV App Management**.
5. On the **KV App Management** tab, select the required cluster from the **Cluster** drop-down list. Then, click **Query**.

   In the list of apps, you can view the details of apps. You can also click an item in the **Actions** column to perform operations including viewing **Partition Details**, **View Configurations**, and **Modify Configurations**.

4.5.1.4.4.6. KV partition management

In the Object Storage Service (OSS) Operations and Maintenance System, you can split and merge KV partitions.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV Partition Management**.
5. On the **KV Partition Management** page, configure the following parameters:
   - **Specify Cluster**: Select the required cluster from the drop-down list.
   - **Specify App**: Select the required app from the drop-down list.
   - **Specify Target**: Select the required item from the drop-down list.
     - If you select **Specify a partition**, **Specify a key**, or **Specify a server**, you must enter the corresponding value.
     - If you select **Information about All Partitions**, you can view the information about all partitions. If you select **View information about all Servers**, you can view the information about all servers.
6. In the list of partitions, split or merge the partition.
   - Split the partition
     a. Click **Split** in the Actions column that corresponds to the partition.
     b. In the **Configure Splitting Method** dialog box, configure the following parameters:
        - **Split By**: Select **Traffic** or **Size**.
        - **Split Into**: Set the number of split subpartitions.
     c. Click **OK**.
     d. In the **OK** dialog box, click **Yes**.
   - Merge
     a. Click **Merge** in the Actions column that corresponds to the partition.
     b. In the **Merge KV Partitions** dialog box, configure the following parameters:
        - **Merge Direction**: You can select **Up** or **Down**.
        - **Number of Merges**: Set the number of partitions to merge.
        - **Partitions to Merge**: The partitions to merge appear.
     c. Click **Confirm**.
     d. In the **OK** dialog box, click **Yes**.

4.5.1.4.4.7. KV management

In the Object Storage Service (OSS) Operations and Maintenance System, you can manage KM, including enabling the auto-split feature, managing KM configurations, and viewing the KM scheduling history.

Enable or disable auto-split

To enable or disable auto-split, perform the following steps:

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KM Management**.
5. On the **Automatic Splitting** tab, select the required cluster from the **Cluster** drop-down list. Then, click **Query**.
6. Turn on or turn off auto-split in the Actions column that corresponds to the required Master.

KV configuration management

To manage KV configurations, perform the following steps:

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KM Management**.
5. On the **KM Management** tab, select the required cluster from the **Cluster** drop-down list. Then, click **Query**.
6. Modify the content of the configuration file. After you modify the content, click **Refresh**.
7. In the message that appears, click **OK**.

KV scheduling history

To view the KV scheduling history, perform the following steps:

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KM Management**.
5. On the **KM Scheduling History** tab, configure the following parameters:
   - Cluster: Select the required cluster from the **Cluster** drop-down list.
   - Time: Select the required time from the **Time** picker.
6. Click **Query**.

   You can view the KV scheduling history within the specified period of time.

4.5.1.4.4.8. KV global flag configurations

In the Object Storage Service (OSS) Operations and Maintenance System, you can query or modify the global flag configurations of KV or KM.

Query global flag configurations

To query or modify global flag configurations, perform the following steps:

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV GlobalFlag Configurations**.
5. On the **Pro** tab, configure the following parameters:
   - **cluster**: Select the required cluster from the drop-down list.
   - **ServerRole**: Select the required server role from the drop-down list.
   - **GlobalFlag**: Select the required global flag from the drop-down list.
   - **server**: When you need to view the specified server, enter the address of the specified server.
6. Click **Query**.

   In the list, you can view all servers and corresponding global flag values.

Modify global flag configurations

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > KV GlobalFlag Configurations**.
5. On the **Pro** tab, configure the following parameters:
   - **cluster**: Select the required cluster from the drop-down list.
   - **ServerRole**: Select the required server role from the drop-down list.
   - **GlobalFlag**: Select the required global flag from the drop-down list.
   - **server**: When you need to view the specified server, enter the address of the server.
6. Click **Query**.

   In the list, you can view all servers and corresponding global flag values. To change the global flag value, you can enter a new global flag value in the **Modify** field and click **Modify**.
7. Change the global flag value.
   - Change the global flag values of all servers
     a. In the **Modify** field, enter a new global flag value and click **Modify**.

b. In the **Submit a Workflow to Modify KV Global Flag** dialog box, configure the following parameters:
   - **comment**: Enter your remarks.
   - **Scheduled Execution**: Set the time when the task is scheduled to execute.
c. Click **Submit**.

- Change the global flag value of the specified server
  a. In the **server** field, enter the value of the required server.
  b. In the **Modify** field, enter the new global flag value. Then, click **Modify**.
  c. In the **Confirm** message, click **OK**.

4.5.1.4.4.9. EC configurations

In the Object Storage Service (OSS) Operations and Maintenance System, you can enable or disable the erasure coding (EC) mode.

Enable the EC mode

For more information about the EC mode, see *OSS storage in EC mode* in *Technical White Paper*.

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Service O&M - KV > EC Management**.
5. Select the required cluster from the **Cluster** drop-down list.
6. Turn on **EC**.
7. Select the EC configuration mode from the **Stripe Width** drop-down list.
   - **2+2**: Select this mode if the number of servers is between 6 to 13.
   - **8+3**: Select this mode if the number of servers is at least 14.
8. Click **Configure**. In the message that appears, click **OK**.

   After you enable the EC mode, OSS uses the EC mode to store data.

4.5.1.4.5. Log monitoring

You can use the log monitoring feature to analyze internal 5xx errors, query logs, and view KV error codes.

Internal error (5xx) analysis

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.

4. In the left-side navigation pane of the Object Storage Service (OSS) Operations and Maintenance System, choose **Log Monitoring > 5XX Error Analysis**.
5. Click the **Total** or **Request flow details** tab. Select the required time range.

   You can select 5Minutes, 15Minutes, 30Minutes, 1Hours, Today, or This week.

   - The **Total** tab displays all errors related to 5xx error codes.
   - The **Request flow details** tab displays request forwarding details.

Query logs

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Products page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Log Monitoring > Backend Log Query**.
5. Select **By ReqId**, **By cluster (5xx only)**, **By Bucket**, **By User**, **Classification by application (5xx only)**, **Custom**, and **By AccessLog** to query logs.

   To analyze the logs you query, you can click **Parse access_log** in the upper-right corner.

View KV error codes

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > Products** to go to the Product List page. Then, in the Storage Service section, click **OSS Operations and Maintenance System**.
4. In the left-side navigation pane of the OSS Operations and Maintenance System, choose **Log Monitoring > KV Error Codes**.

   You can view the status codes and detailed error codes in the KV error comparison table.

4.5.1.5. Tools and commands

4.5.1.5.1. 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

4.5.1.5.2. 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
```

Ensure themod_nginx item is in the *on* state. The following figure shows that the status of mod_nginx is *on*.

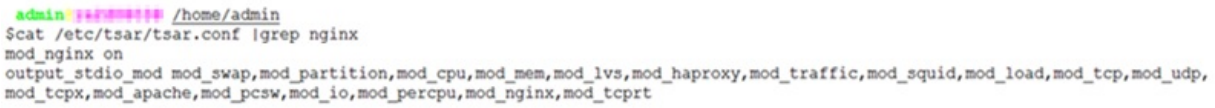

4.6. Tablestore

4.6.1. Operations Guide

4.6.1.1. Tablestore Operations and Maintenance System

4.6.1.1.1. 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.

4.6.1.1.2. User data

4.6.1.1.2.1. 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.

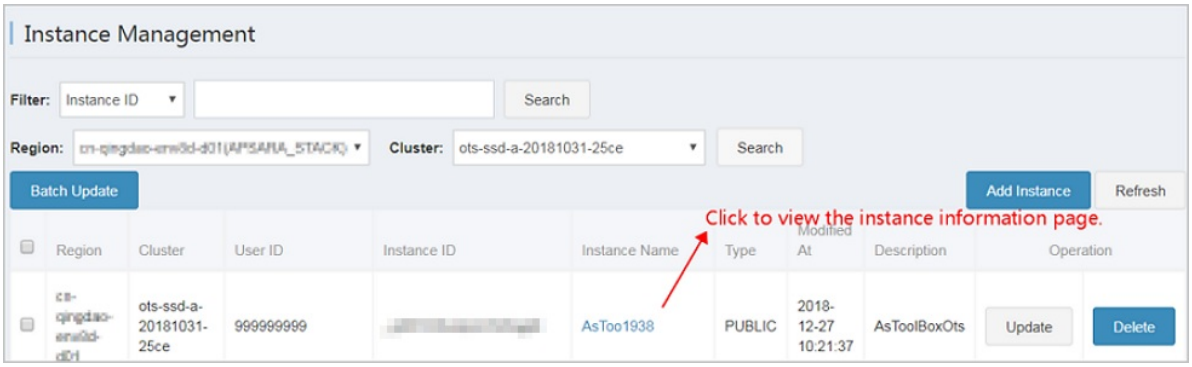

- Search for instances based on specified conditions

  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.

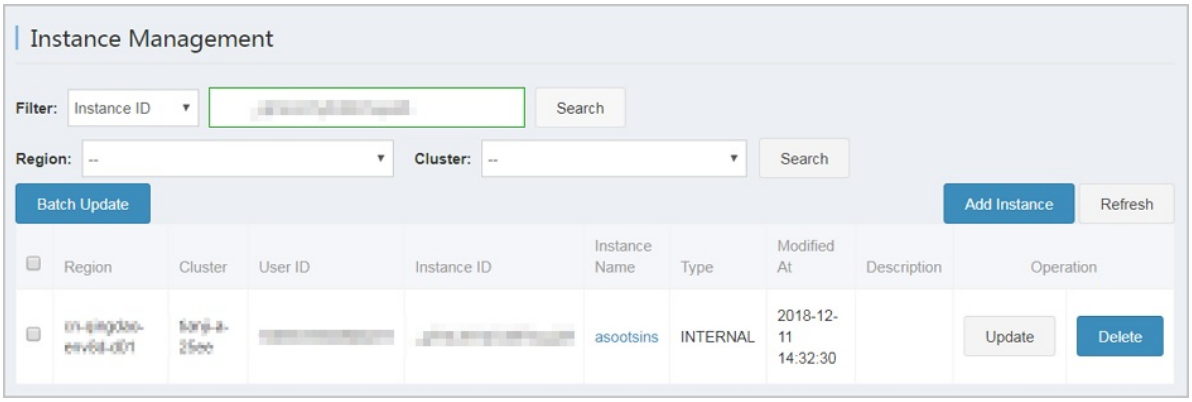

- View instance details

- 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.

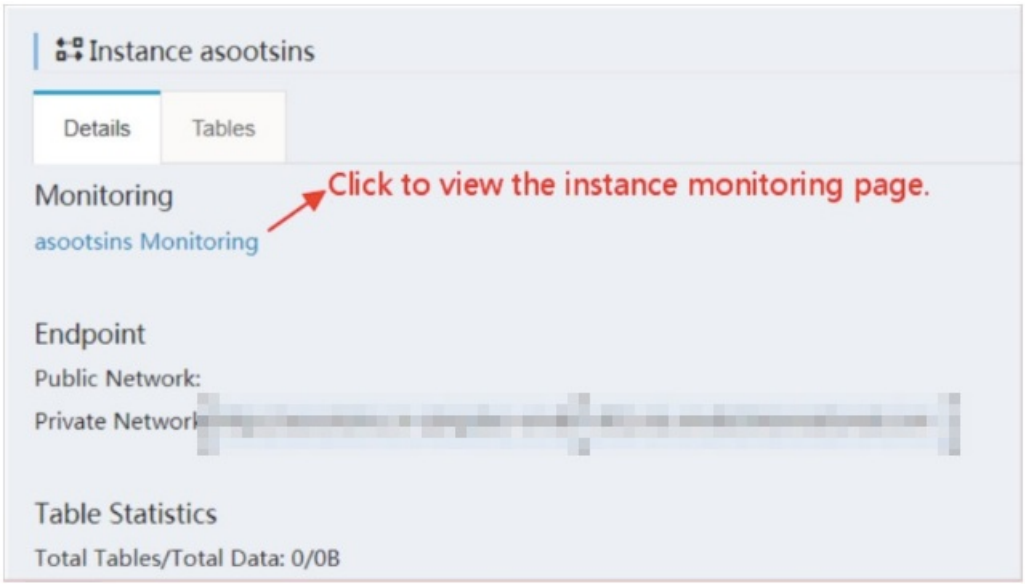

- Tables information

  Click the instance name. On the **Tables** tab, you can view the maxVersion, ttl, readCU, writeCU, and timestamp of tables.

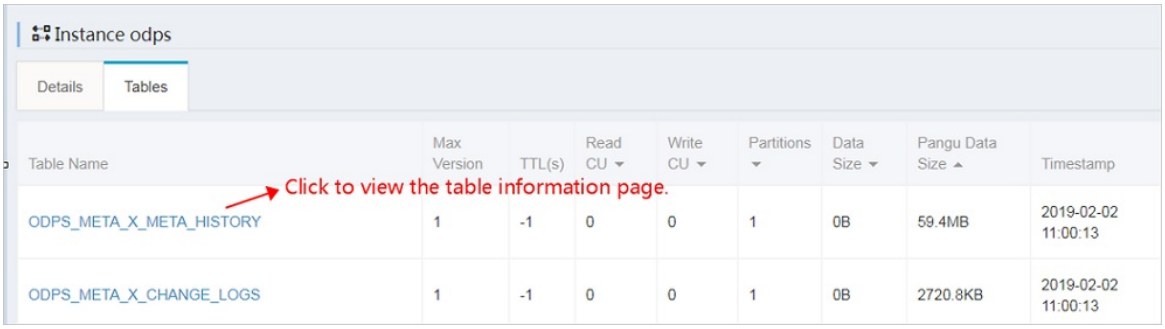

- View table details

- 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.

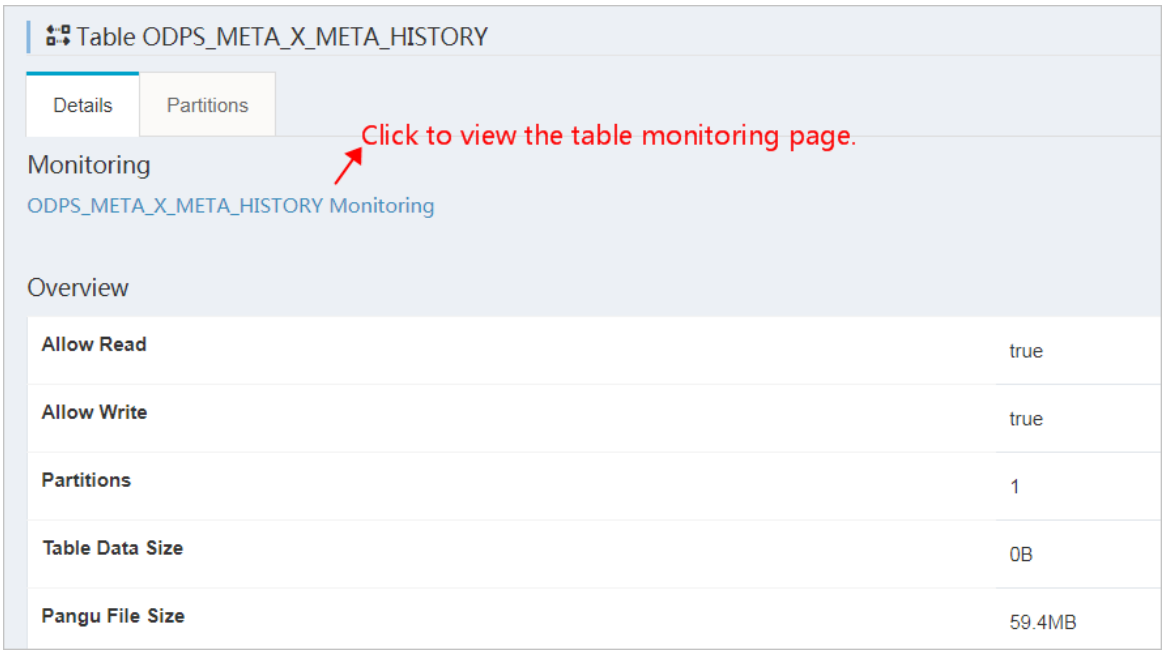

- Partitions

  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.

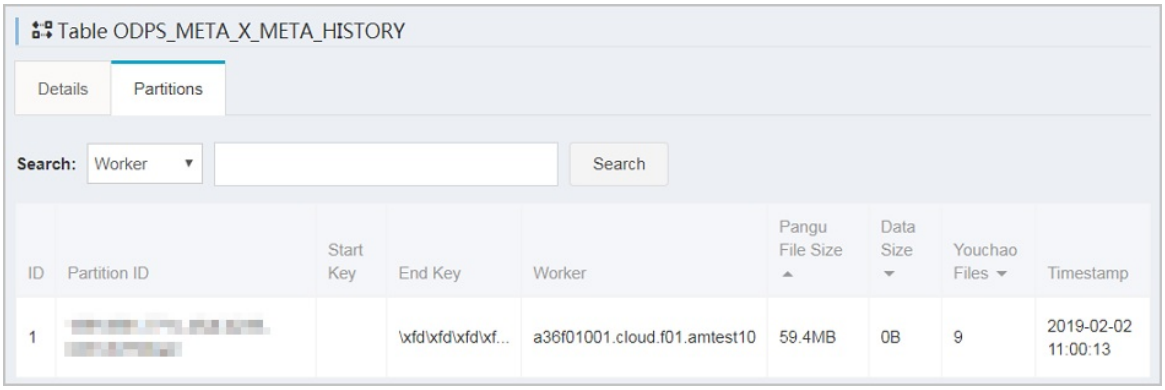

4.6.1.1.3. Cluster management

4.6.1.1.3.1. 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.

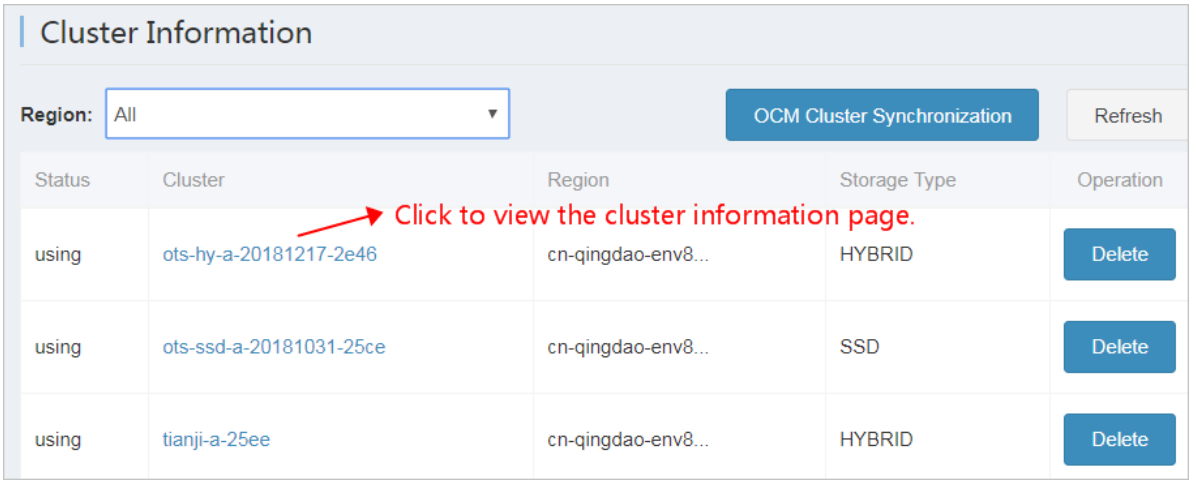

- Cluster details

  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.

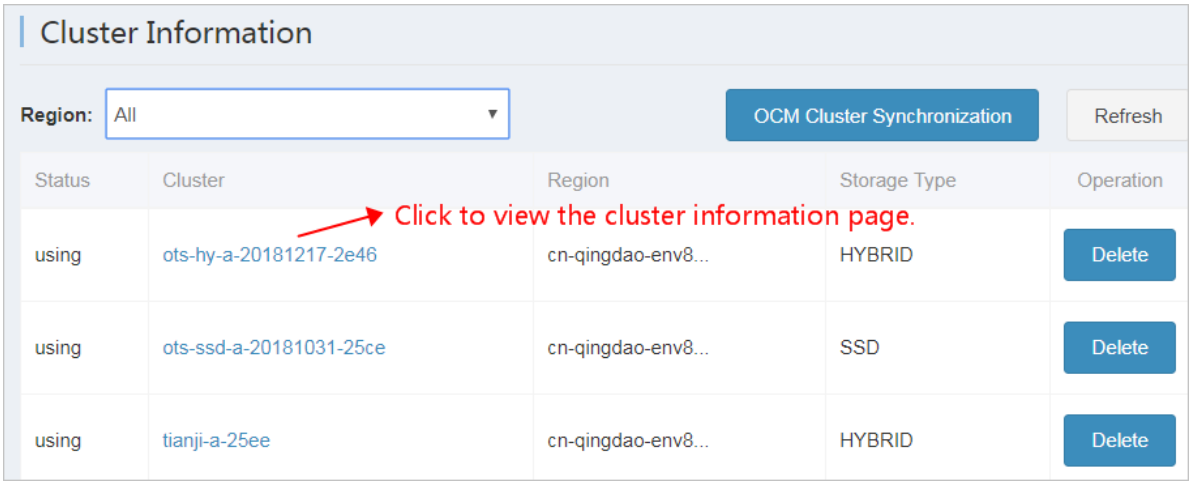

- Overview: provides the basic information of a cluster.

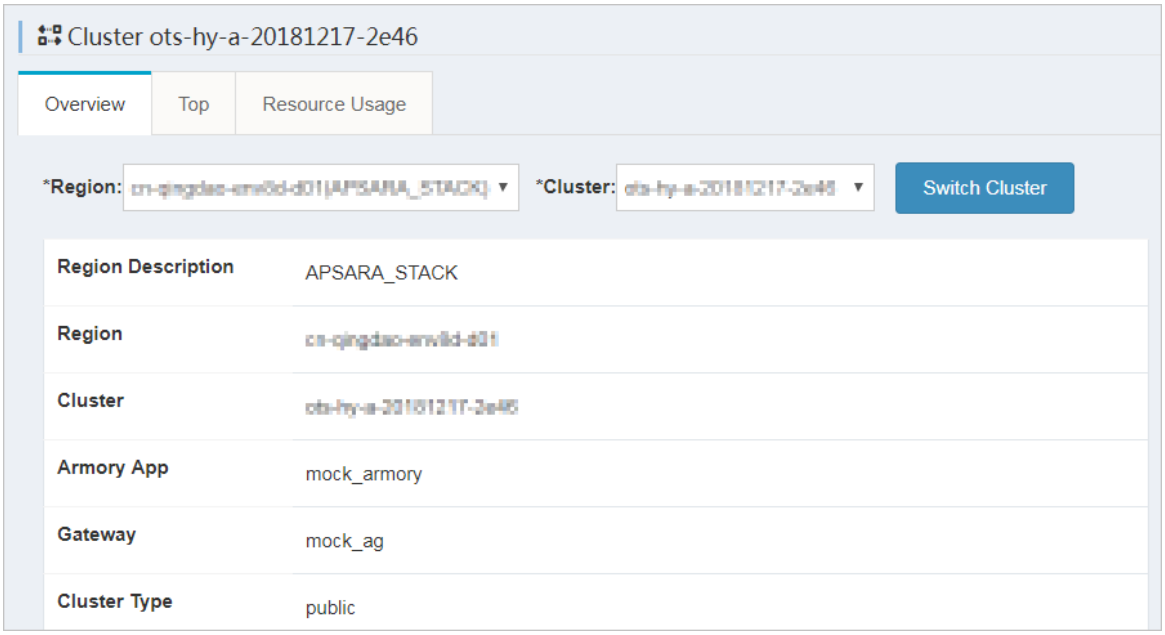

- Top: provides top request information of partitions and tables.

  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.

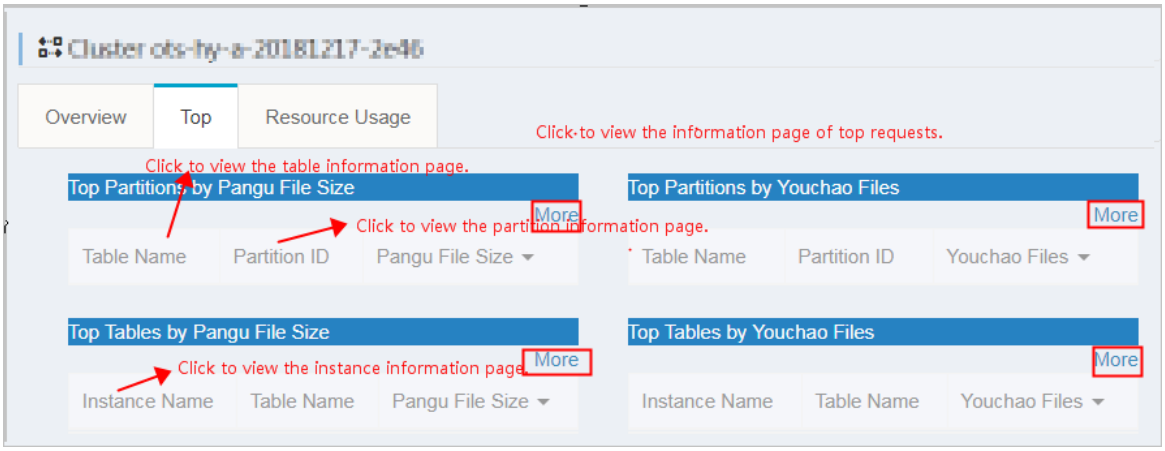

- Resource Usage: provides cluster usage details. The usage statistics collection task is automatically triggered in the background at specific intervals. In special cases, you can click **Collect Data** to manually trigger the usage statistics collection task. After the usage statistics collection task 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.)

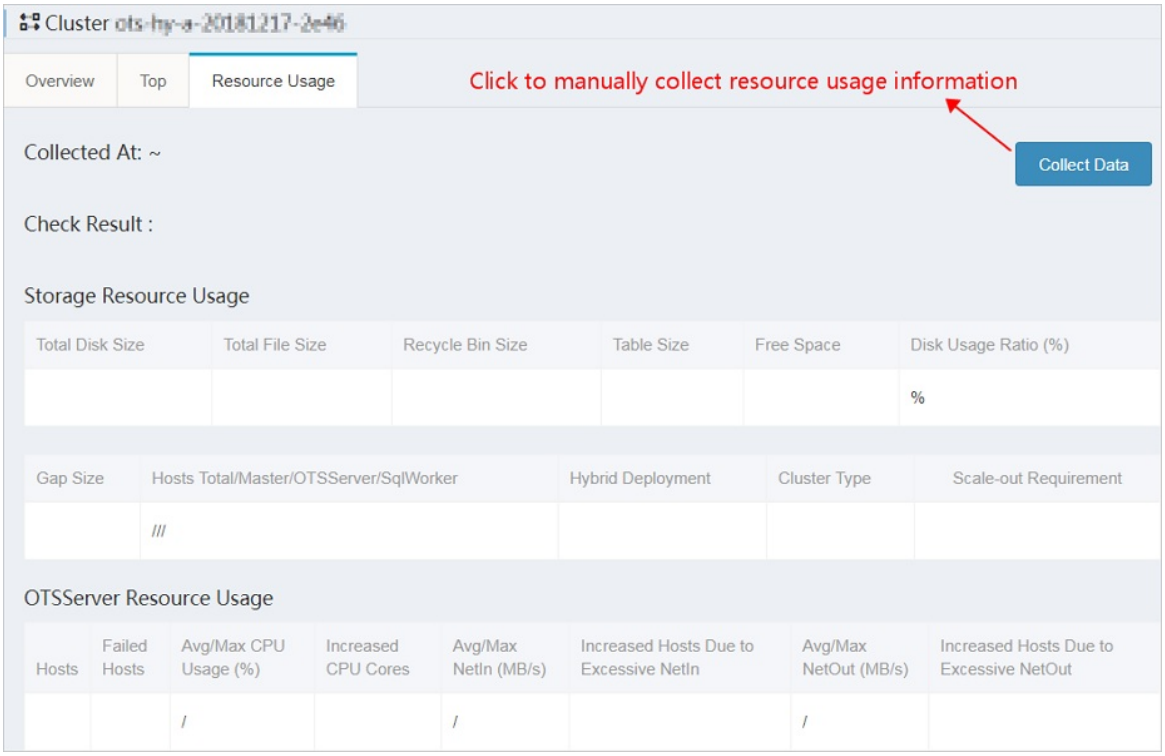

4.6.1.1.4. Inspection center

4.6.1.1.4.1. 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.

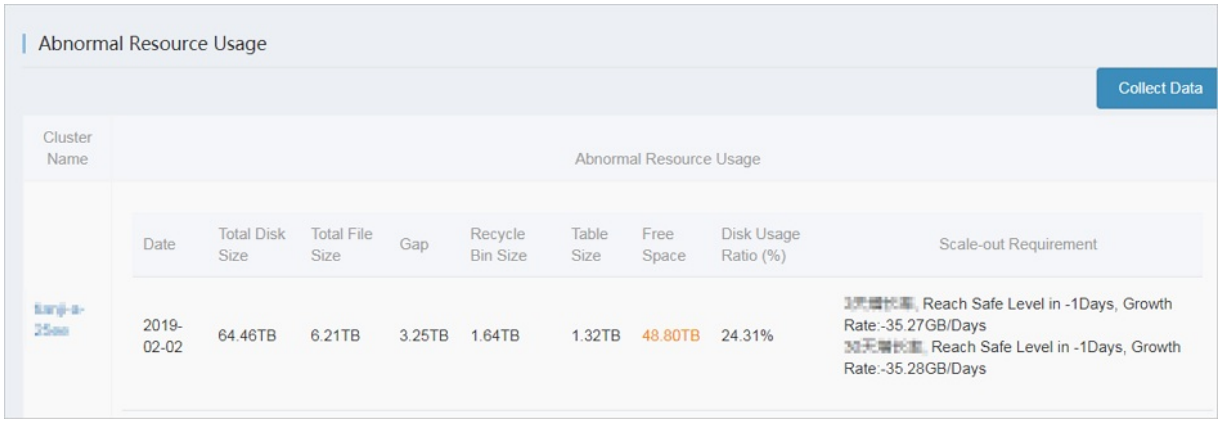

4.6.1.1.5. Monitoring center

4.6.1.1.5.1. 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.

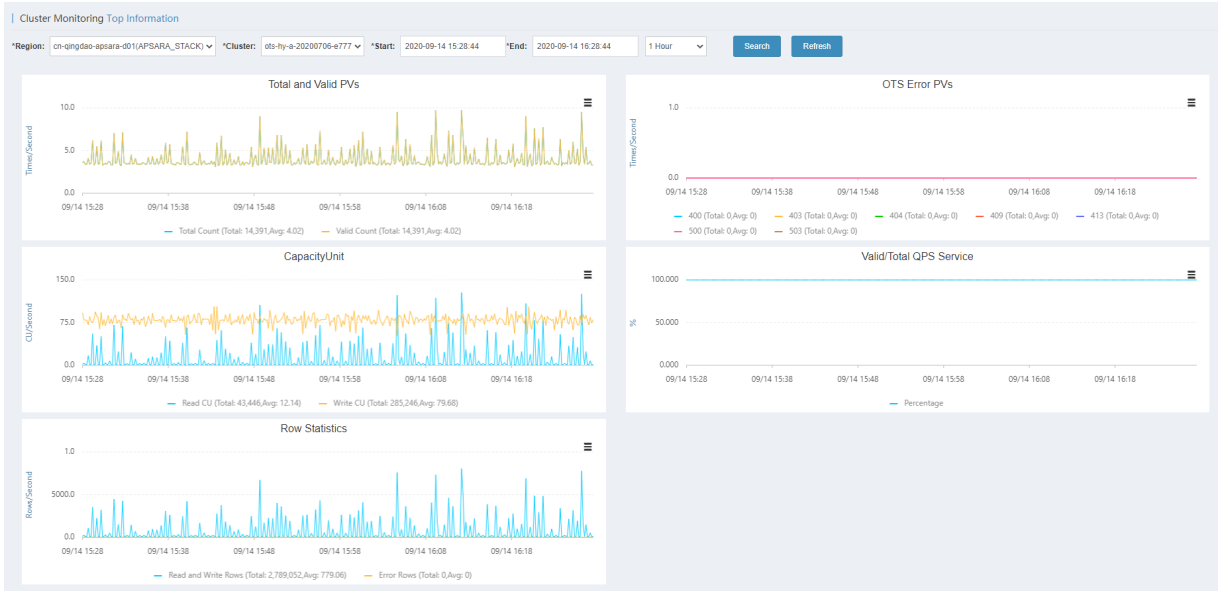

4.6.1.1.5.2. Application monitoring

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.

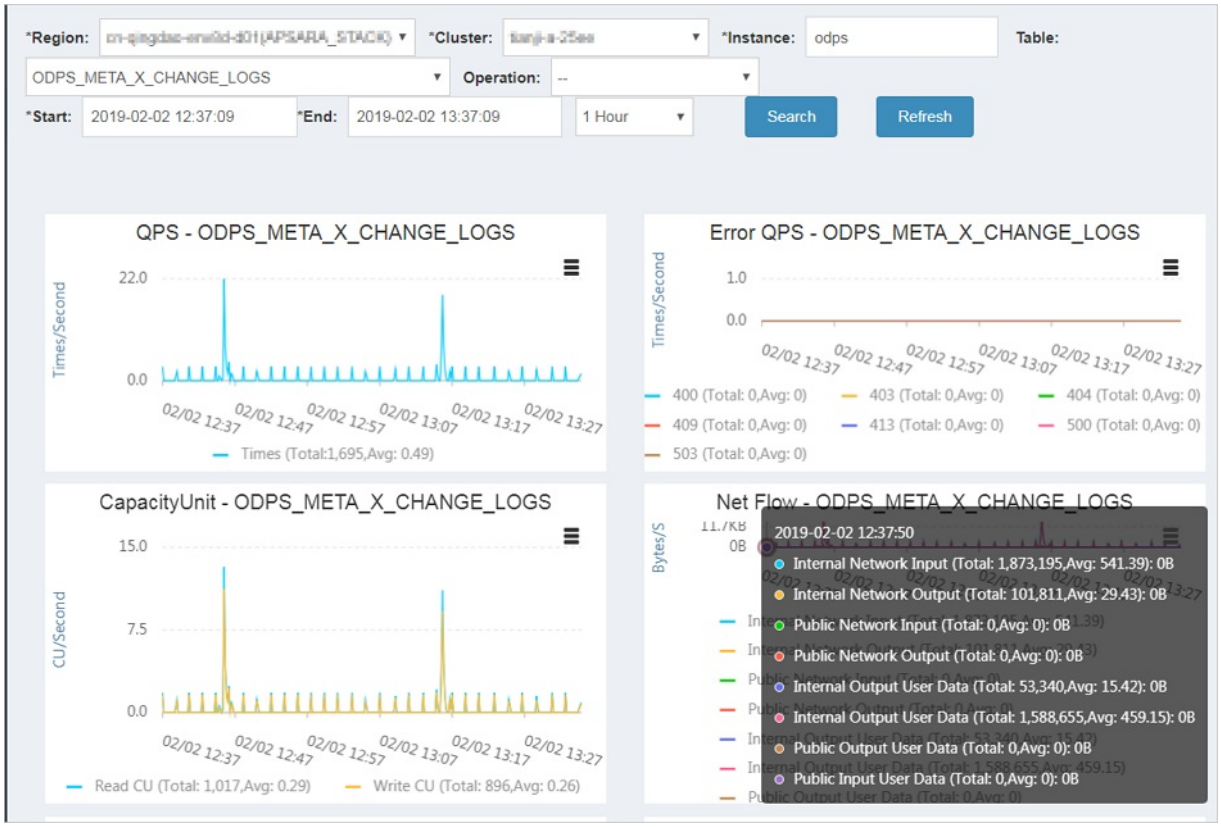

4.6.1.1.5.3. Top requests

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.

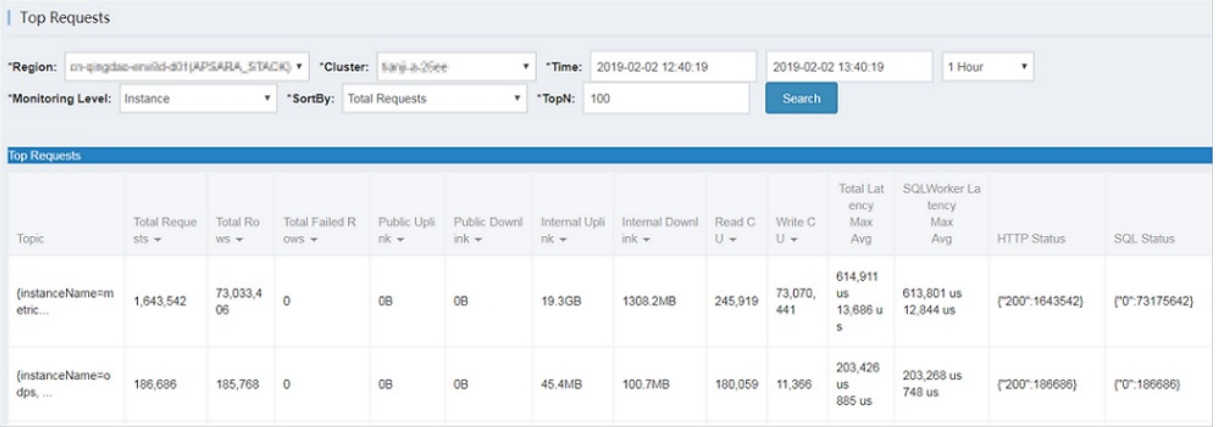

4.6.1.1.5.4. Request log search

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.

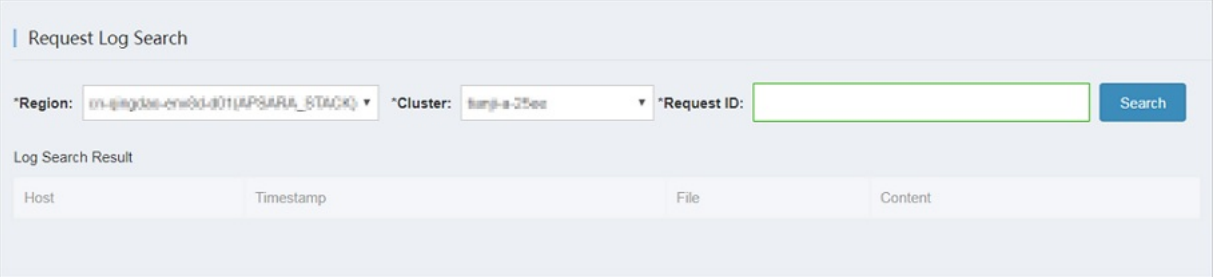

4.6.1.1.6. System management

4.6.1.1.6.1. Manage tasks

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.

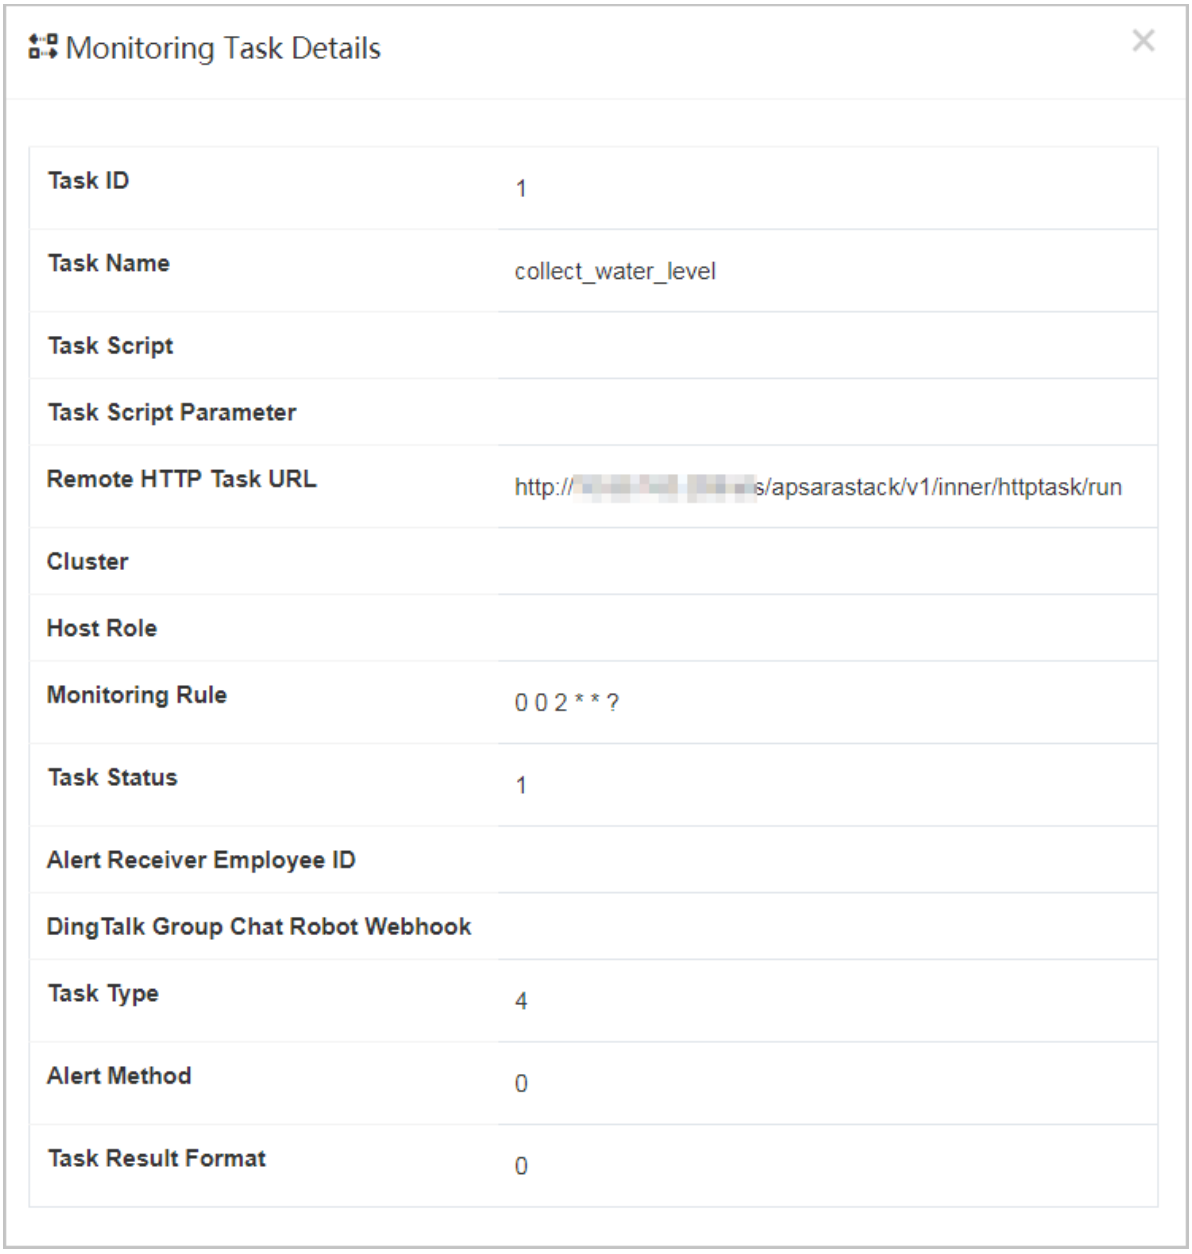

- Enable or disable a task.

  > **Note** Disabled tasks no longer run automatically.

- Run a task immediately.

4.6.1.1.6.2. 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.

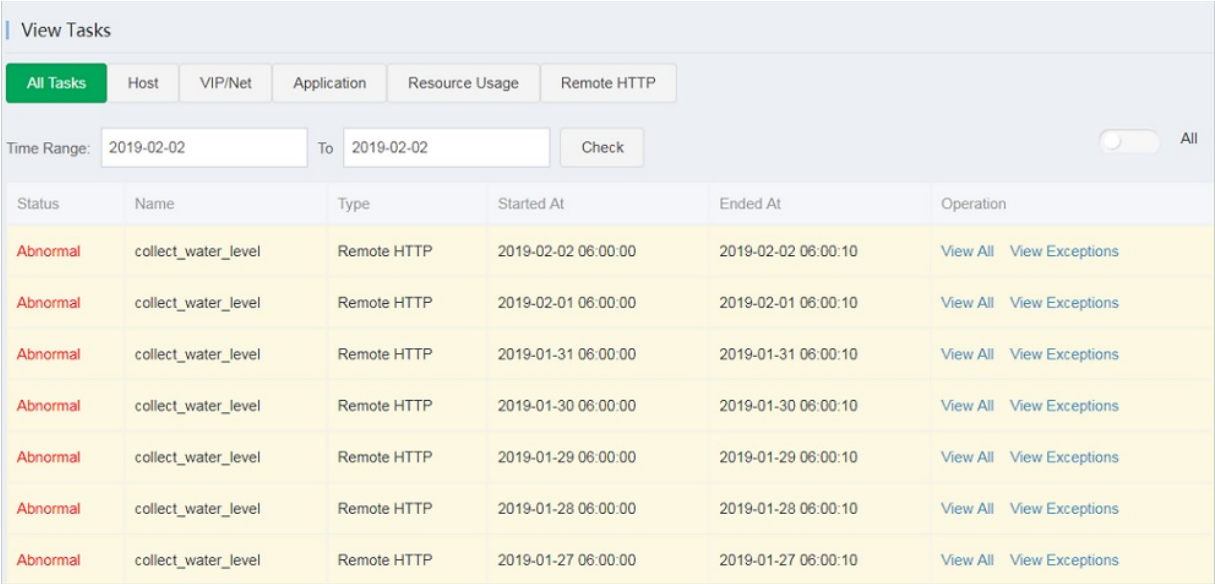

Click **View All** or **View Abnormal** in the Operation column corresponding to the abnormal task to view the specific cause of a task failure, as shown in the following figure.

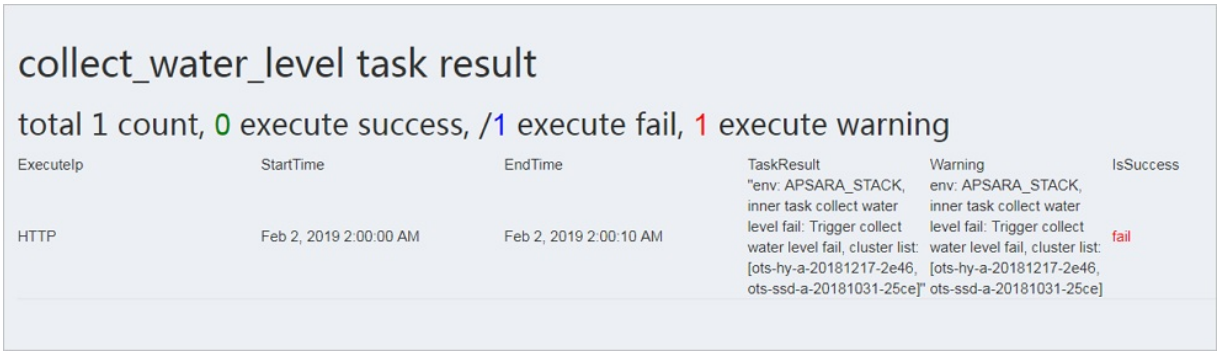

4.6.1.1.7. Platform audit

4.6.1.1.7.1. 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.

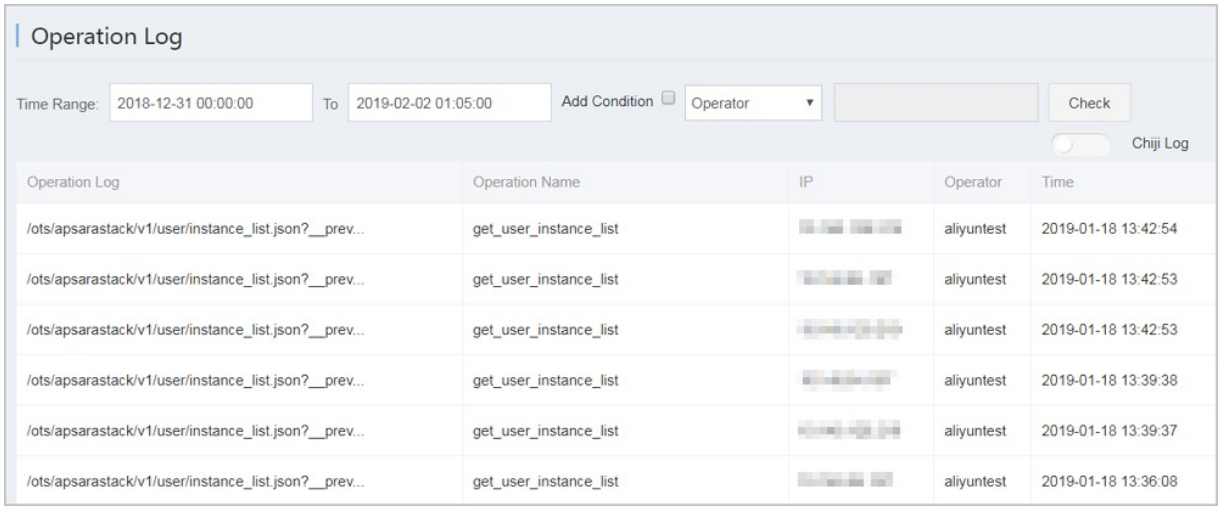

4.6.1.2. Cluster environments

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

4.6.1.3. 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

4.6.1.4. Pre-partition a table

4.6.1.4.1. 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
  ```

4.6.1.4.2. 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.

4.7. ApsaraDB RDS

4.7.1. Operations and Maintenance Guide

4.7.1.1. Architecture

4.7.1.1.1. O&M architecture

This topic describes the O&M architecture of ApsaraDB RDS.

O&M architecture

Category	Component	Description
Business Foundation System	ServiceTest	Automated regression testing for ApsaraDB RDS features.
	Abs and Service	The OpenAPI service for ApsaraDB RDS.
	Location Initialization	Initialization of the Location configurations.
	Database Initialization	DDL and DML configurations used to initialize ApsaraDB RDS Business Foundation System.
TianLong	dbinit	Initialization of TianLong.
	TianLongServer	The TianLong service.
	TunnelServer	An intermediate service that manages communications between ApsaraDB RDS and Elastic Compute Service (ECS).
ApsaraDB RDS for MySQL	InitCluster	Initialization of the cluster and host list.
	DbMySQL	Database node initialization, database installation, and RPM package management.
ApsaraDB RDS for MySQL Management Service	PengineMySQL	The task flow service related to ApsaraDB RDS for MySQL.
	RdsApiMySQL	API operations related to ApsaraDB RDS for MySQL.
	McNode	The messaging agent node.
	DbInit	DML configurations used to initialize ApsaraDB RDS for MySQL.
PolarDB	DbPolaro	The physical nodes of PolarDB.
	InitCluster	Initialization of the cluster and host list.
ApsaraDB RDS for PostgreSQL	DbPgSql	• Local SSDs: physical nodes of ApsaraDB RDS for PostgreSQL. • Standard SSDs: deployed on ECS instances.
	InitCluster	Initialization of the cluster and host list.
	DbMssql	Resources deployed on ECS instances.
	InitCluster	Initialization of the instance and host list.

Category	Component	Description
ApsaraDB RDS for SQL Server	EcsImport2012Ent	Initialization of ECS images of different SQL Server versions.
	EcsImport2016Ent	
	EcsImport2012Std	
	EcsImport2016Std	
	EcsImport2017Ag	
ApsaraDB RDS for SQL Server Management Service	DbInit	DML configurations used to initialize ApsaraDB RDS for SQL Server.
	PengineMssql	The task flow service related to ApsaraDB RDS for SQL Server.
	RdsApiMssql	API operations related to ApsaraDB RDS for SQL Server.
	McNode	The messaging agent node.
Core services	YumServer	Deployment of host nodes for database services.
	Robot	Automated O&M for database services.
	MC (MessageCenter)	The message center service.
	Aurora	The high-availability platform for database services.
	API	The backend core API service.
	Stat	The real executor of scheduled tasks for various monitoring and inspection activities.
	ApsaraDB Operations and Maintenance System	The O&M platform for database services.
	Performance Collection	Performance collection of earlier-version instances and hosts, and audit log collection of instances.
	Backup	Full backup and incremental backup of databases.
	TASK	Scheduling of the management task flow.
	Whitelists	IP address whitelists of instances.
	BackendLocation	The cluster and location management services.
	NameService	The name service.
	Database Initialization	Initialization of metadatabases.

Category	Component	Description
	Environment Configuration Initialization	Initialization of information of services such as Object Storage Service (OSS), Virtual Private Cloud (VPC), and Bakowner.

4.7.1.1.2. System architecture

This topic describes the system architecture of ApsaraDB RDS.

Backup system

ApsaraDB RDS can back up databases at any time and restore them to a specific point in time based on the backup policy, which makes the data more traceable.

- Automatic backup: ApsaraDB RDS provides various types of backup. You can flexibly configure the backup start time within off-peak hours.
- Manual backup: You can manually back up databases at any time.
- Cloned instance: A cloned instance is a new instance that has 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.
- Backup file download: During the retention period of backup files, you can log on to the ApsaraDB RDS console and download backup files to your computer.

Data migration system

ApsaraDB RDS provides Data Transmission Service (DTS) to help you migrate databases.

- ApsaraDB RDS allows you to migrate databases from one instance to another.
- ApsaraDB RDS provides professional tools and migration wizards to help you migrate data to or from ApsaraDB RDS instances.

Monitoring system

ApsaraDB RDS provides multi-dimensional monitoring services across the physical, network, and application layers to ensure business availability.

- Performance monitoring: ApsaraDB RDS monitors nearly 20 system performance metrics, such as disk capacity, IOPS, connections, CPU utilization, network traffic, transactions per second (TPS), queries per second (QPS), and cache hit ratio.
- SQL audit: The system records the SQL statements and related information sent to ApsaraDB RDS instances, such as the connection IP address, database name, access account, execution time, and number of records returned. You can use SQL audit to perform troubleshooting and check instance security.
- Threshold-based alerting: ApsaraDB RDS provides alert SMS notifications in the event of exceptions in instance status or performance.
- Web operation logging: The system records all modification operations in the ApsaraDB RDS console for administrators to check. These logs are retained for up to 30 days.

Control system

If a host or instance stops responding, it switches workloads over within 30 seconds after the high-availability (HA) component detects an exception. This ensures that applications run normally.

Task scheduling system

You can use the ApsaraDB 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.

4.7.1.2. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager Operations Console.

Prerequisites

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

  The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id*.ops.console.*intranet-domain-id*.
- A browser is available. We recommend that you use Google Chrome.

Procedure

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

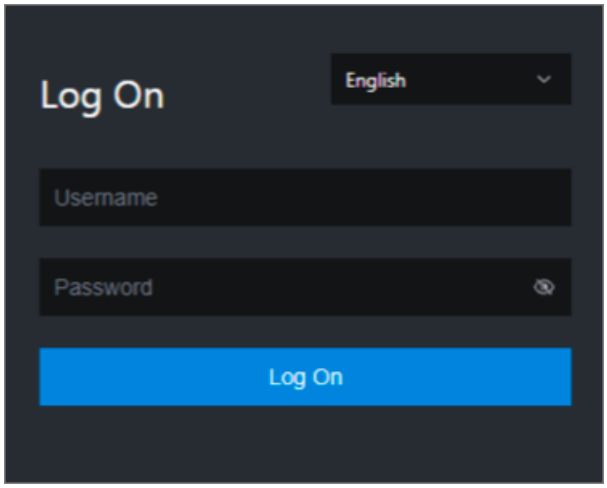

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains the following special characters: ! @ # $ %
- The password must be 10 to 20 characters in length.

4. Click **Log On**.

4.7.1.3. Manage instances

This topic describes how to manage ApsaraDB RDS instances. You can view instance details and user information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click O&M. In the left-side navigation pane, choose **Product Management > RDS**.
3. On the **Instance Management** tab of the **RDS** page, perform the following operations:
   - View instances

     View the instances that belong to your account on the **Instance Management** tab, as shown in Instance list.

     Instance list

     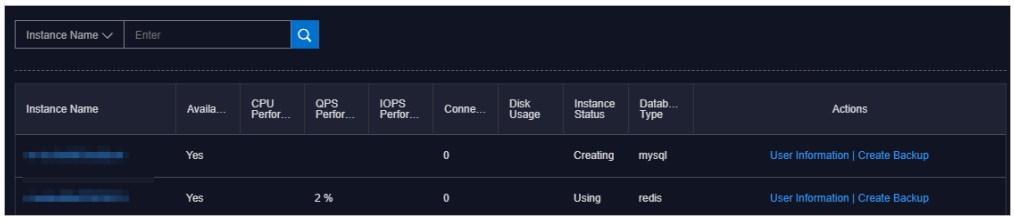

   - View instance details

     Click the ID of an instance to view its details, as shown in Instance details. You can switch your workloads between primary and secondary instances and query history operations on this page.

     > **Note** We recommend that you do not perform forced switchover, because it may cause data loss if data is not synchronized between the primary and secondary instances.

     Instance details

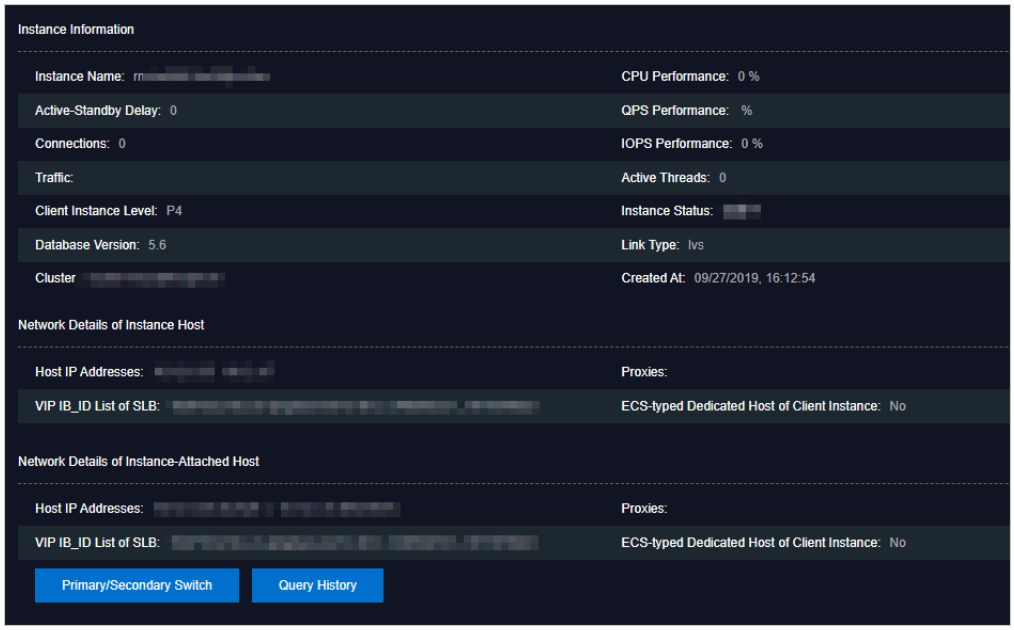

- **View user information**

  Click **User Information** in the **Actions** column corresponding to an instance, as shown in User Information.

  User Information

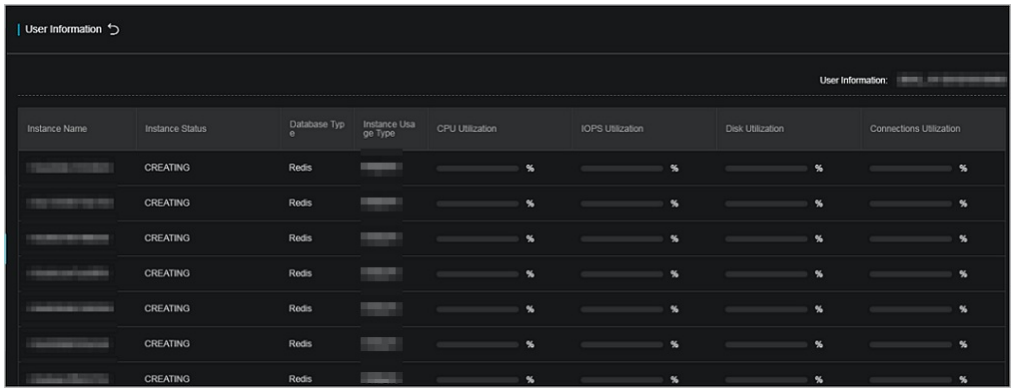

- **Create backups**

  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

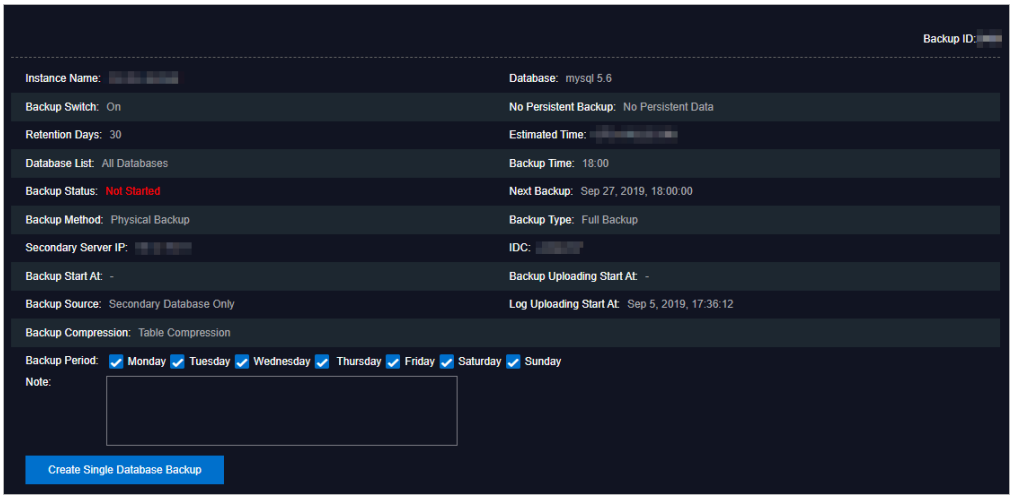

4.7.1.4. Manage hosts

This topic describes how to view and manage hosts.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the left-side navigation pane, choose **Product Management > RDS**.
3. On the **Host Management** tab of the **RDS** page, you can view information of all hosts.

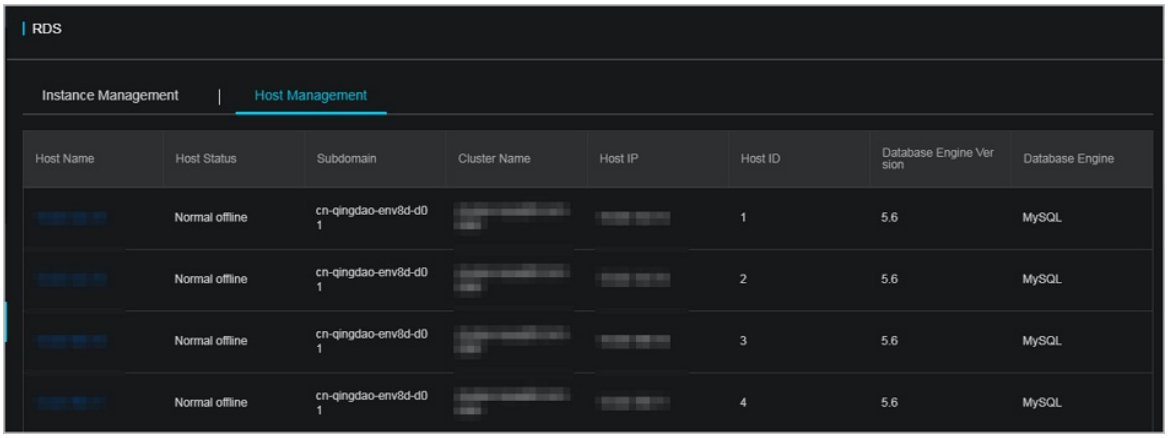

4. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.

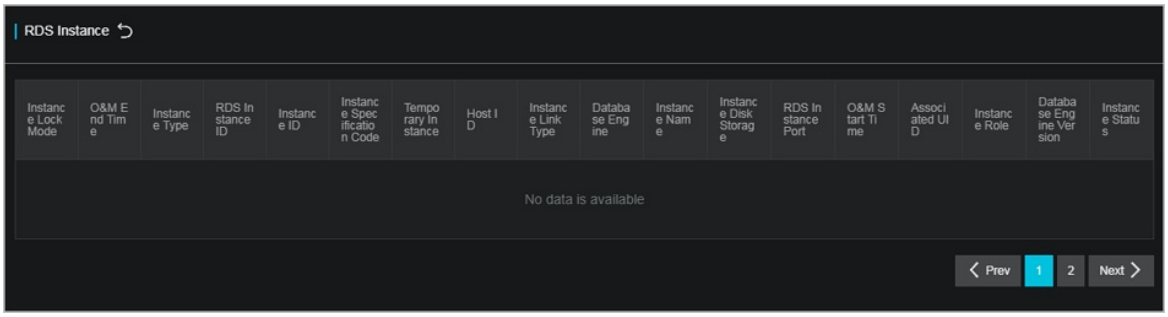

4.7.1.5. Security maintenance

4.7.1.5.1. 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.

4.7.1.5.2. 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.

4.7.1.6. Redline V4.3.3 O&M description

4.7.1.6.1. Services provided by Redline Enterprise

This topic describes the services provided by Redline Enterprise.

Redline Enterprise provides the following services:

- redline-perf.LogWorker#
- redline-perf.OpenApi#
- redline-perf.PerfMaster#
- redline-perf.PerfWorkerr#

> **Note**
>
> - redline-perf.DbInit is discontinued.
> - The collection agent of UE is deployed on physical machines in other modes.

4.7.1.6.2. Paths of files in the Docker container of redline-perf

This topic describes the paths of files in the Docker container of redline-perf.

PerfMaster, PerfWorker, and LogWorker

The following table describes the mapping relationships between the preceding data collection services and clusters.

Service	Cluster
PerfMaster	master
PerfWorker	perf
LogWorker	log

Paths:

- Root path: */home/admin/redline-perf/*
- Path of database data in the Docker container: */home/admin/redline-perf/ignite*
- Paths of database data on the host

Node	Path
PerfMaster node	*/cloud/data/redline-perf/PerfMaster#/perf-master/home/admin/redline-perf/ignite/*
PerfWorker node	*/cloud/data/redline-perf/PerfWorker#/perf-worker/home/admin/redline-perf/ignite/*
LogWorker node	*/cloud/data/redline-perf/LogWorker#/log-worker/home/admin/redline-perf/ignite/*

- Path of log data in the Docker container: */home/admin/redline-perf/log*

> **Note**
>
> - `app.log` is the main program log, which contains the log entries that record the extract, transform, load (ETL), storage, data retrieval API operation, and data aggregation to minute granularity.
> - `datax.log` is the synchronization module log, which contains the log entries that record the operations of synchronizing metadata from the metadatabase and pushing real-time tables to Cloud Monitor.

- Paths of log data on the host

Node	Path
PerfMaster node	*/cloud/data/redline-perf/PerfMaster#/perf-master/home/admin/redline-perf/log/*
PerfWorker node	*/cloud/data/redline-perf/PerfWorker#/perf-worker/home/admin/redline-perf/log/*
LogWorker node	*/cloud/data/redline-perf/LogWorker#/log-worker/home/admin/redline-perf/log/*

OpenAPI

This is an end-to-end gateway service.

Paths:

- Root path: */home/admin/dll-service-aliyun-com*
- Path of log data in the Docker container: */home/admin/dll-service-aliyun-com/logs*
- Path of log data on the host: */cloud/data/redline-perf/OpenApi#/open-api/home/admin/dll-service-aliyun-com/logs*

4.7.1.6.3. Perform environment checks

This topic describes how to perform environment checks.

Prepare the script for batch operations

1. Run the following commands on a jumper server to prepare the script for batch operations:

```
curl "http://localhost:7070/api/v3/column/m.ip?m.sr.id=redline-perf.PerfMaster%23" | grep ip | awk -F '"' '{print $(NF-1)}' > masterhosts
curl "http://localhost:7070/api/v3/column/m.ip?m.sr.id=redline-perf.PerfWorker%23" | grep ip | awk -F '"' '{print $(NF-1)}' > perfworkerhosts
curl "http://localhost:7070/api/v3/column/m.ip?m.sr.id=redline-perf.LogWorker%23" | grep ip | awk -F '"' '{print $(NF-1)}' > logworkerhosts
cat perfworkerhosts logworkerhosts masterhosts  | sort -u > allhosts
```

2. Run the following command to install parallel-ssh (pssh). Skip this step if pssh has been installed.

```
wget https://parallel-ssh.googlecode.com/files/pssh-2.3.1.tar.gz
# If the jump server cannot be connected to the Internet, you can download the installation file to another machine and upload the file to the jump server.
tar zxvf pssh-2.3.1.tar.gz
cd pssh-2.3.1
python setup.py install
```

Check node connectivity

After you install pssh, check whether all nodes are running normally.

1. (Required) Check the communication information between LogWorker nodes.
   i. Run the following command on the jumper server:

```
pssh -h logworkerhosts -p 1 -P 'docker ps |grep redline | grep LogWorker|awk '"'"'{ print $1}'"'"' | xargs -I ID docker exec ID curl 'http://127.0.0.1:8080/admin/makeBaseline?igniteCluster=log''
```

   ii. Check whether the number of nodes displayed in the ONLINE server nodes section is consistent with the number of deployed nodes. If the number is inconsistent or a message of NOT IN BASELINE is returned, fix this issue. For more information, see Fix connection failures.
2. (Required) Check the communication information between PerfWorker nodes.
   i. Run the following command on the jumper server:

```
pssh -h perfworkerhosts -p 1 -P 'docker ps |grep redline | grep PerfWorker|awk '"'"'{ print $1}'"'"' | xargs -I ID docker exec ID curl 'http://127.0.0.1:8080/admin/makeBaseline?igniteCluster=perf''
```

   ii. Check whether the number of nodes displayed in the ONLINE server nodes section is consistent with the number of deployed nodes. If the number is inconsistent or a message of NOT IN BASELINE is returned, fix this issue. For more information, see Fix connection failures.

3. (Required) Check the communication information between PerfMaster nodes.
   i. Run the following command on the jumper server:

   ```
   pssh -h masterhosts -p 1 -P 'docker ps |grep redline | grep PerfMaster|awk '"'"'{ print $1}'"'"' | xargs -I ID docker exec ID curl 'http://127.0.0.1:8080/admin/makeBaseline?igniteCluster=master''
   ```

   ii. Check whether the number of nodes displayed in the ONLINE server nodes section is consistent with the number of deployed nodes. If the number is inconsistent or a message of NOT IN BASELINE is returned, fix this issue. For more information, see Fix connection failures.
4. (Required) Check whether clusters are active.
   i. Run the following command on the jumper server:

   ```
   pssh -h allhosts -p 1 -P 'docker ps |grep redline | grep -v Open|awk '"'"'{ print $1}'"'"' | xargs -I ID sh -c "echo ID && docker exec ID tail -n 50 ./log/app.log |grep '"'deactived'"'"'
   ```

   ii. Check whether log entries that include deactived are returned. If yes, go to the corresponding Docker container and activate the clusters in the deactived state. For more information, see Activate clusters in the deactived state.

Check for program exceptions

(Required) Check application logs for exceptions.

1. Run the following command on the jumper server:

   ```
   pssh -h allhosts -p 1 -P 'docker ps |grep redline | grep -v Open|awk '"'"'{ print $1}'"'"' | xargs -I ID sh -c "echo ID && docker exec ID tail -n 100 ./log/app.log |grep -C 10 '"'Exception'"'"'
   ```

2. If exception logs are returned, fix this issue. For more information, see Troubleshoot program exceptions.
3. If this issue cannot be fixed by using the solutions described in Troubleshoot program exceptions, contact development personnel or reset node data. For more information, see Reset node data.

4.7.1.6.4. O&M operations

4.7.1.6.4.1. Scale in or out a cluster

This topic describes how to scale in or out a cluster.

Determine whether scale-out is required

- Check the diagnostic information of an instance in the Apsara Uni-manager Management Console. If the data records of the last 3 hours are displayed after more than 5 seconds, you may need to add nodes. More nodes can increase the response speed. You can balance the return on investment (ROI) of hardware based on your needs.
- Ten nodes can support 3,000 host nodes of ApsaraDB RDS for MySQL instances. A single ApsaraDB RDS for MySQL instance can contain more than one host node.

Add a node

1. Add a node and start it.

> **Note** You can add multiple nodes at a time.

2. Connect to a node that assumes a master role.
3. Run the `curl http://127.0.0.1:8080/admin/makeBaseline?apply=false` command to check whether the returned node information meets the expectation.
4. If yes, run the `curl http://127.0.0.1:8080/admin/makeBaseline?apply=true` command.

Remove a node

1. Remove a node.

> **Note** You can remove only a single node each time.

2. Connect to a node that assumes a master role.
3. Run the `curl http://127.0.0.1:8080/admin/makeBaseline?apply=false` command to check whether the returned node information meets the expectation.
4. If yes, run the `curl http://127.0.0.1:8080/admin/makeBaseline?apply=true` command.
5. Wait 30 minutes for data to be balanced among the remaining nodes and then continue to remove other nodes based on your needs.

4.7.1.6.4.2. Upgrade or restart a cluster

This topic describes how to upgrade or restart a cluster.

Procedure

1. Connect to a PerfMaster node and run the following commands to disable its database. If no `success` message is returned, repeat the following commands until a `success` message is returned.

```
curl "http://127.0.0.1:8080/admin/deactivateIgnite?igniteCluster=perf"
curl "http://127.0.0.1:8080/admin/deactivateIgnite?igniteCluster=log"
curl "http://127.0.0.1:8080/admin/deactivateIgnite?igniteCluster=master"
```

2. Wait 2 minutes and upgrade or restart the cluster after data is stored to disks.
3. After the upgrade or restart is complete, perform environment checks. For more information, see Perform environment checks.

4.7.1.6.4.3. Fix connection failures

This topic describes how to fix connection failures.

Fix the failure of missing nodes in the communication information

1. Connect to a missing node and run the following command:

```
ps aux|grep java|grep -v grep|awk '{ print $2}'|xargs kill -9
```

2. Wait 3 minutes and check node connectivity again.

Fix the NOT IN BASELINE failure in the communication information

Run the following command on a PerfMaster node:

```
curl "http://127.0.0.1:8080/admin/makeBaseline?igniteCluster=${cluster}&apply=true"
```

Note ${cluster} must be replaced with the cluster to which the node belongs, which can be master, perf, or log.

4.7.1.6.4.4. Activate clusters in the deactived state

This topic describes how to activate the clusters in the deactived state.

Run the following command on a PerfMaster node:

```
curl "http://127.0.0.1:8080/admin/activateIgnite?igniteCluster=${cluster}"
```

Note ${cluster} must be replaced with the cluster to which the node belongs, which can be master, perf, or log.

4.7.1.6.4.5. Troubleshoot program exceptions

This topic describes how to troubleshoot program exceptions.

Troubleshoot the exception of "partition has been lostPart"

LogWorker node

```
### Connect to a Docker container of LogWorker and run the following reset commands:
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=TABLE_LOG_AUDIT_SQLS"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=TABLE_LOG_ERROR_SQLS"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=TABLE_LOG_SLOW_SQLS"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=TABLE_LOG_AUDIT_SQL_TEMPLATES"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=TABLE_LOG_AUDIT_SQLS_ROUTE_INFO"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=CACHE_CODEC"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=CACHE_CODEC_LOCK"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=CACHE_CODEC_ATOMIC_LONG"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=ignite-sys-atomic-cache@default-ds-group"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=ignite-sys-atomic-cache@default-volatile-ds-group"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=log&cacheName=NOTIFY_LOG_IGNITE_SHUTTING_DOWN"
```

PerfWorker node

```
### Connect to a Docker container of PerfWorker and run the following reset commands:
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=CACHE_CODEC"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=CACHE_CODEC_LOCK"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=CACHE_CODEC_ATOMIC_LONG"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=NOTIFY_PERF_IGNITE_SHUTTING_DOWN"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=MONITOR_ALERT_HIS_CACHE"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=MONITOR_ALERT_STATUS_INFO_CACHE"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=CACHE_REDLINE_SNAPSHOT"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=seconds-level"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=perf&cacheName=minutes-level"
```

PerfMaster node

```
### Connect to a Docker container of PerfMaster and run the following reset commands:
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=CACHE_TS_CACHE_NAME"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=NOTIFY_MASTER_IGNITE_SHUTTING_DOWN"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=LOCK_MASTER_COMPETING"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=CLUSTER_COMMAND"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=CACHE_CODEC"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=CACHE_CODEC_LOCK"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=CACHE_CODEC_ATOMIC_LONG"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=NODE_EXECUTE_COMMAND_RECORD"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=IGNITE_PERF_LOG_METADATA"
curl "http://127.0.0.1:8080/admin/resetLostPartitions?igniteCluster=master&cacheName=CHECKPOINT_CACHE_NAME"
```

Troubleshoot the exception of "Failed to perform WAL operation"

If a message of Failed to perform WAL operation is returned, contact technical personnel or reset node data. For more information, see Reset node data.

4.7.1.6.4.6. Reset node data

This article describes how to reset node data.

Precautions

This operation deletes all data of a specific node. Monitoring data stored on the node is deleted and cannot be recovered. Proceed with caution.

Procedure

1. Connect to the node or Docker container whose data you want to reset and run the following command:

   ```
   touch log/HOLD_START && ps aux|grep java|grep -v grep|awk '{ print $2}'|xargs kill -9 && rm -rf ignite/* &
   & sleep 60 && rm -f log/HOLD_START &
   ```

2. Wait 3 minutes and run the following commands:

   ```
   curl "http://127.0.0.1:8080/admin/activateIgnite?igniteCluster=${cluster}"
   curl "http://127.0.0.1:8080/admin/makeBaseline?igniteCluster=${cluster}&apply=true"
   ```

   Note `${cluster}` must be replaced with the cluster to which the node belongs, which can be master, perf, or log.

4.7.1.6.4.7. Dump logs

This topic describes how to dump logs.

Context

Historical audit log data can be dumped for query.

Precautions

- This feature is available only for ApsaraDB RDS for PostgreSQL instances that use standard SSDs.
- Object Storage Service (OSS) buckets must be created to store audit logs. For more information about how to create a bucket, see Create a bucket in *Object Storage Service User Guide*.
- OSS buckets must be cleared on a regular basis if log dump is enabled.

Enable or disable log dump

1. Log on to the host on which you want to perform a log dump. Run the following command in the */root/redline* directory to obtain the log worker nodes:

   ```
   curl "http://localhost:7070/api/v3/column/m.ip?m.sr.id=redline-perf.LogWorker%23" | grep ip | awk -F '"
   ' '{print $(NF-1)}' > logworkerhosts
   ```

2. Enable log dump.
   - Enable log dump for all instances.

     ```
     pssh -h logworkerhosts 'curl "http://<The domain name of the logworker cluster in Apsara Infrastruc
     ture Management Framework>:8080/admin/logTransferStore?transType=oss&accessKey=AK&access
     Secret=SK&endpoint=http://oss-cn-neimeng-env30-d01-a.intra.env30.shuguang.com/&bucketName
     =lztest&region=cn-neimeng-env30-d01"'
     ```

   - Enable log dump for the instances that are included in IP address whitelists.

```
pssh -h logworkerhosts 'curl "http://<The domain name of the logworker cluster in Apsara Infrastructure Management Framework>:8080/admin/logTransferStore?/admin/logTransferStore?transType=ftp&accessKey=AK&accessSecret=SK&endpoint=oss-cn-beijing.aliyuncs.com&bucketName=rds-sqlaudit&filterType=WHITE&whiteList=["ppassql-HttpApiTest-INS-NAME-5"%2C"ppassql-HttpApiTest-INS-NAME-6"]"'
```

> **Note** In the preceding command, `ppassql-HttpApiTest-INS-NAME-5` and `ppassql-HttpApiTest-INS-NAME-6` are sample instances that are included in IP address whitelists for you to enable log dump. You can add actual instance names to the command based on your needs and connect multiple instance names with `%2C`.

3. Disable log dump.

```
pssh -h logworkerhosts 'curl "http://<The domain name of the logworker cluster in Apsara Infrastructure Management Framework>:8080/admin/logTransferStore?/admin/logTransferStore?transType=close&accessKey=&accessSecret=&endpoint=&bucketName=&region="'
```

Use log parsing tools

Log dump splits historical audit log data into multiple small files. To facilitate your query and analysis, you can use a log parsing tool. By default, each dump log contains data of up to 15 seconds in length or 100 MB in size.

Configuration command

```
java -jar perf-tools-0.0.1.jar --input <Log storage directory> --output <Directory of parsed logs> --begin <Start timestamp> --interval <Scan duration>
```

Parameters

Parameter	Description	Required
--input	The directory in which instance logs are stored.	Yes.
--output	The directory in which parsed logs are stored.	No. If this parameter is not specified, the dumped logs are stored in the standard output directory. **Note** The grep command can be used with a vertical bar (|) to query logs.
--begin	The start timestamp of the query. Unit: seconds.	Yes.
--interval	The time range of the query. Unit: seconds.	No. Default value: 3600. Maximum value: 10800.

Sample commands

- Parse logs

```
java -jar perf-tools-0.0.1.jar --input /logs/ppassql-HttpApiTest-INS-NAME-0 --output /logs/ppassql-HttpApiTest-INS-NAME-0/result.data --begin 1602592098 --interval 10800
```

- Use a grep keyword and a vertical bar (|) to query logs

```
java -jar ../perf-tools-0.0.1.jar --input /logs/ppassql-HttpApiTest-INS-NAME-0 --begin 1602592098 --interval 10800 | grep incomplete
```

4.8. AnalyticDB for PostgreSQL

4.8.1. Operations and Maintenance Guide

4.8.1.1. Overview

Purpose

. This document summarizes possible problems that you may encounter during O&M operations and provides solutions for you.

If you encounter a system problem that is not covered in this document, you can submit a ticket to Alibaba Cloud for technical support.

Requirements

You must possess basic IT skills and knowledge about topics such as computer networking, computer operation, problem analysis, and troubleshooting.

You also must pass the Alibaba Cloud system training to learn necessary knowledge about Alibaba Cloud systems, including but not limited to system principles, networking, features, and the use of maintenance tools.

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 from users.

Precautions

To ensure a stable system and avoid unexpected events, you must comply with 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. This prevents system faults that are caused by unauthorized operations.

- System security

  Before you perform system operations, you must be aware of their impacts.

  You must clearly record the details about the issues that you encounter during the operations. This helps you troubleshoot and handle the issues.

- 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.

Support

You can contact Alibaba Cloud technical support for help.

4.8.1.2. Architecture

The following figure shows the system architecture of AnalyticDB for PostgreSQL.

System architecture

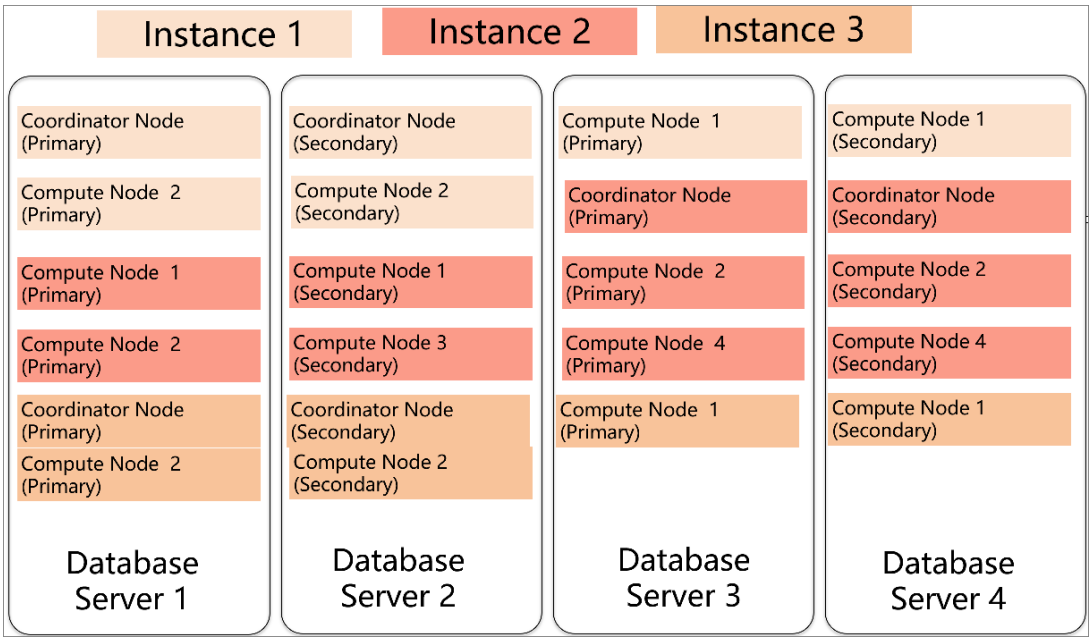

AnalyticDB for PostgreSQL consists of two major components: the coordinator node and compute nodes.

The coordinator node is used to access applications that are deployed on AnalyticDB for PostgreSQL. The coordinator node accepts connection and SQL query requests from clients and dispatches computing tasks to compute nodes. The system deploys a secondary node of the coordinator node on an independent physical server to replicate data from the primary node for failover. The secondary node cannot connect to compute nodes or accept external links.

Compute nodes are independent data nodes in AnalyticDB for PostgreSQL. Each compute node stores a part of data, and all compute nodes work together to execute computing tasks in parallel. Each compute node consists of a primary node and a secondary node for failover.

4.8.1.3. Routine maintenance

4.8.1.3.1. Check for data skew on a regular basis

You must check for data skew on a regular basis during maintenance to prevent the instances from becoming read-only due to excess data on some compute nodes.

You can use the following methods to check for and identify data skew. In these examples, SQL statements are used.

1. For individual tables or databases, check the space usage on each compute node to determine whether data is skewed.
   i. Execute the following statement to determine whether the data in a database is 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 on each compute node after the statement is executed. If the space occupied on one or more compute nodes is significantly greater than that on other compute nodes, the data in this database is skewed.

   ii. Execute the following statement to determine whether the data in a table is skewed:

   ```
   select pg_size_pretty(pg_relation_size('tblname')) from gp_dist_random('gp_id');
   ```

   The execution result of the preceding statement shows the space occupied by the tblname table on each compute node. If the space of the table on one or more compute nodes is significantly greater than that on other compute nodes, the data in this table is skewed. You must modify the distribution key to redistribute the data.
2. Use system views to determine whether data is skewed.
   i. Execute the following statement to check whether the storage space is skewed, which is similar to the preceding method that is used to check for space usage:

   ```
   select * from gp_toolkit.gp_skew_coefficients
   ```

   This view allows you to check the data volume of rows in a table. The larger the table, the more time it takes 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 determine whether the data is skewed:

   ```
   select * from gp_toolkit.gp_skew_idle_fractions
   ```

4.8.1.3.2. Execute VACUUM and ANALYZE statements

You can execute `VACUUM` and `ANALYZE` statements on a regular basis for frequently accessed tables and databases. You can also execute `VACUUM` and `ANALYZE` statements after you perform a large number of updates or write operations to prevent the operations from consuming excessive resources and storage space.

4.8.1.4. Security maintenance

4.8.1.4.1. 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.

4.8.1.4.2. 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, you must use complex passwords for your systems and devices and change these passwords on a regular basis.

4.9. KVStore for Redis

4.9.1. Operations and Maintenance Guide

4.9.1.1. Operations Guide

4.9.1.1.1. O&M tools

The Apsara Uni-manager Operations Console is a unified and intelligent O&M platform.

The platform provides the following features to manage ApsaraDB for Redis instances:

- Instance management: allows you to view instance details, instance logs, and user information.
- Host management: allows you to view and manage hosts.

4.9.1.1.2. System architecture

This topic describes the system components and their purposes.

Group	Component	Description
Redis resource services	ServiceTest	Automated regression of KVStore for Redis features.
	AbsCacheServiceAliyunCom and RedisaPhoenix	The OpenAPI service for KVStore for Redis.
	Location	Initialization of the Location configurations.
	DbInit	DDL and DML configurations used to initialize Redis resource services.
Redis services	InitCluster	Initialization of the cluster and host list.
	DbRedis	DB host node initialization, kernel installation, and RPM package management.

Group	Component	Description
Redis-Control services	PengineMySQL	The task flow service related to KVStore for Redis.
	RdsApiRedis	API operations related to KVStore for Redis.
	DbInit	DML configurations used to initialize KVStore for Redis.
Core services	DbBootstrap	Deployment of DB host nodes for database services.
	Robot	Automated O&M for database services.
	MessageCenter	The message center service.
	AuroraCluster	The high-availability platform for database services.
	API	The backend core API service.
	Stat-JobWorker	The real executor for scheduled task for various monitoring and inspections.
	Dukang	The O&M platform for database services.
	Brain	Performance collection of earlier-version instances and hosts, and audit log collection of instances.
	BakMaster and Bifrost	To improve the full backup and incremental backup of databases.
	TaskDispatcher	To manage the task flow scheduling service.
	SecGroup	The whitelist service.
	BackendLocation	The cluster and location management services.
	NameService	The name service.
	DBInit	Initialization of metadatabases.
	EnvInit	Initialization of information such as OSS, VPC, and Bakowner.

4.9.1.1.3. Server roles

This topic describes server roles and their respective projects, clusters, and services. This helps O&M personnel identify and locate faults.

Note You can view server roles by using Cluster Operations in Apsara Infrastructure Management Framework. For more information about Apsara Infrastructure Management Framework, see *Apsara Stack Operations and Maintenance Guide*.

Project	Cluster	Service	Server role	Description
rds	redis-control	rds-redis-control	Dbinit	DDL and DML configurations used to initialize KVStore for Redis.
			McNode	The messaging agent node.
			PengineRedis	The task flow service related to KVStore for Redis.
			RdsapiRedis	API operations related to KVStore for Redis.
	redis_7u	rds-redis	DbRedis	DB host node initialization, kernel installation, and RPM package management.
			initCluster	Initialization of the cluster and host list.

4.9.1.1.4. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager Operations Console.

Prerequisites

- The URL, username, and password that are required to log on to the Apsara Uni-manager Operations Console are obtained from the deployment personnel or administrators.

  The URL of the Apsara Uni-manager Operations Console is in the format of *region-id*.ops.console.*intranet-domain-id*.
- A browser is available. We recommend that you use Google Chrome.

Procedure

i. Open your browser.

ii. In the address bar, enter the URL. Then, press the Enter key.

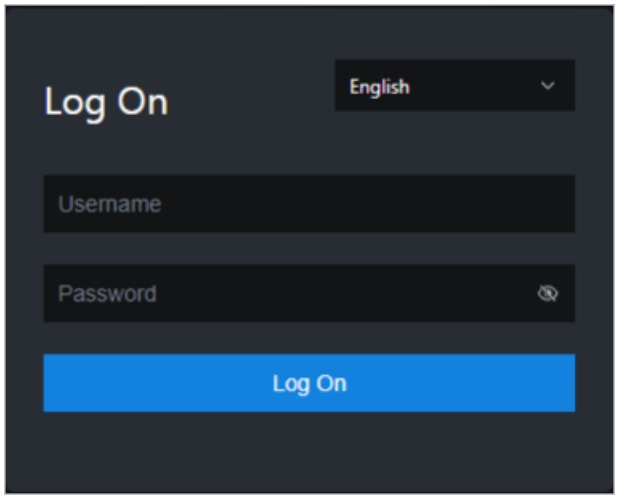

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

iii. Enter your username and password.

> **Note** To obtain the username and password that are used to log on to the Apsara Uni-manager Operations Console, contact the deployment personnel or administrators.

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

To enhance security, make sure that the password meets the following requirements:

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- The password is 10 to 20 characters in length.

iv. Click **Log On**.

4.9.1.1.5. Instance management

This topic describes how to manage KVStore for Redis instances. You can view instance details and user information.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > RDS**.
4. On the **Instance Management** tab of the RDS page, you can perform the following operations:

Operation	Description
View the details of an instance	A list of instances that belong to your account is displayed on the Instance Management tab. Click the ID of an instance to view the details of the instance. **Note** On the **Instance Details** page, you can click Primary/Secondary Switch and Query History.
View the user information of an instance	Find an instance and click **User Information** in the **Actions** column. On the User Information page, you can view the user information and key metrics of each instance.

4.9.1.1.6. Host management

This topic describes how to view resource usage and key performance metrics for KVStore for Redis hosts.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, click **O&M**.
3. In the left-side navigation pane, choose **Product Management > RDS**.
4. On the **Host Management** tab of the RDS page, you can view information of all hosts.

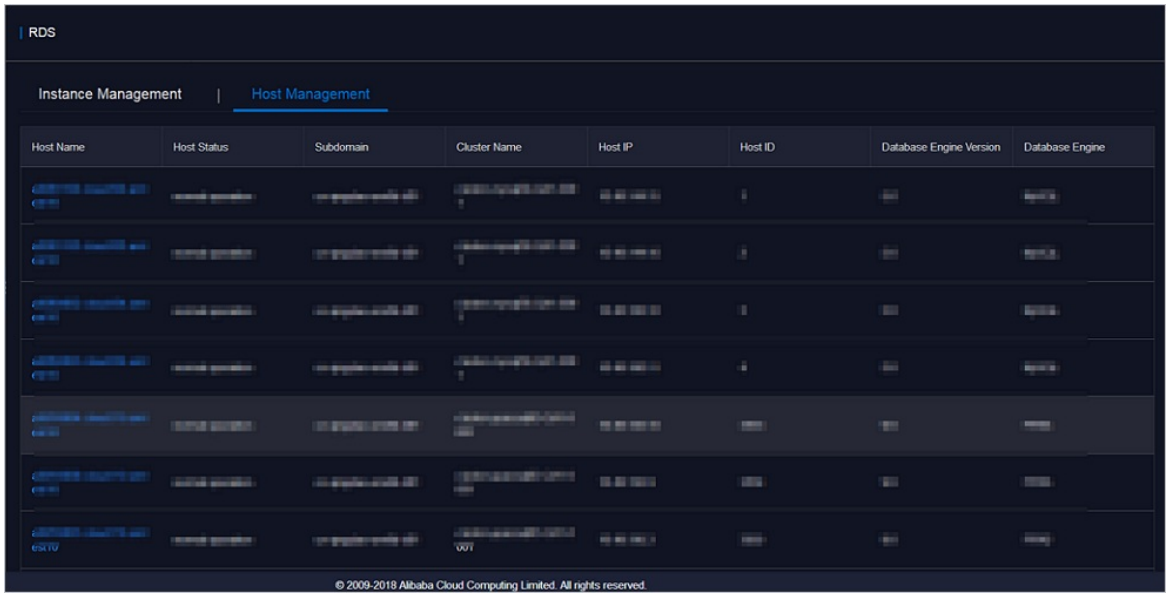

5. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.

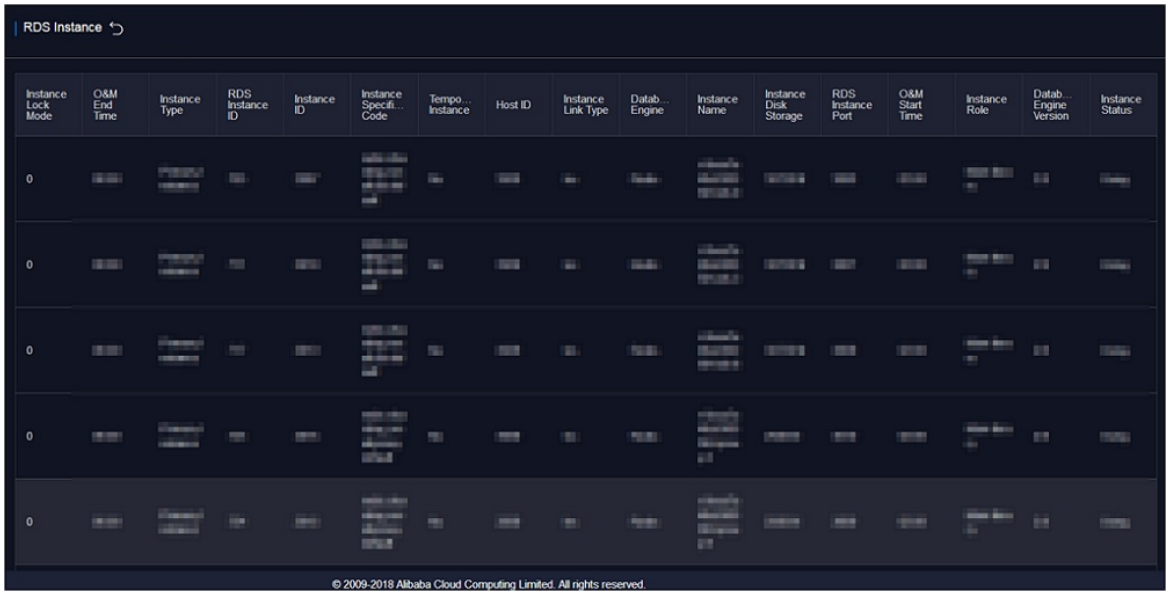

4.9.1.1.7. Security maintenance

4.9.1.1.7.1. 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.

4.9.1.1.7.2. 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.

4.10. ApsaraDB for MongoDB

4.10.1. Operations and Maintenance Guide

4.10.1.1. Operations Guide

4.10.1.1.1. Service architecture

4.10.1.1.1.1. System architecture

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.

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.

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.

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.

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.

4.10.1.1.2. ApsaraDB for MongoDB O&M overview

The Apsara Uni-manager 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.

4.10.1.1.3. Log on to the Apsara Uni-manager Operations Console

This topic describes how to log on to the Apsara Uni-manager Operations Console.

Prerequisites

- The endpoint of the Apsara Uni-manager Operations Console and the username and password used

to log on to the console are obtained from the deployment personnel or an administrator.

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

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

Procedure

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

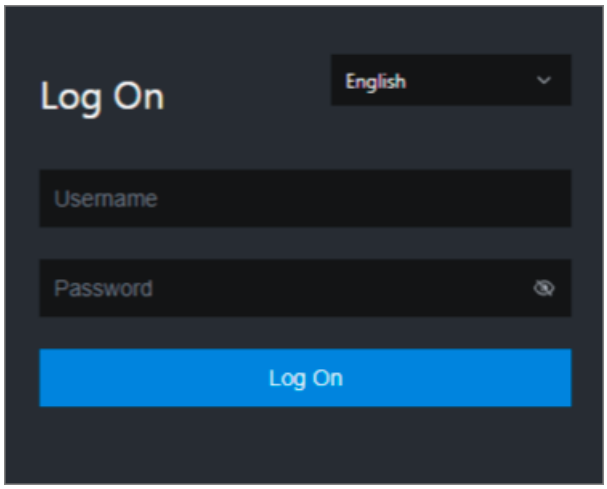

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains the following special characters: ! @ # $ %
- The password must be 10 to 20 characters in length.

4. Click **Log On**.

4.10.1.1.4. Security maintenance

4.10.1.1.4.1. 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.

4.10.1.1.4.2. 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.

4.11. Log Service

4.11.1. Operations and Maintenance Guide

4.11.1.1. 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 provides a user interface to manage Log Service. The Portal complies with the standard Java applications of Alibaba Cloud.

Terminal

1. Log on to the Apsara Uni-manager Operations Console.

   For more information, see the **ASAPI Reference > Log on to the API Tool console** topic in *Log Service Operation 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 destination cluster in the Clusters column, and then click **Operations**

in the Actions column.

7. On the **Services** tab, select the destination server role, for example, sls-backend-server.ServiceTest#, and then click **Terminal** in the Actions column.

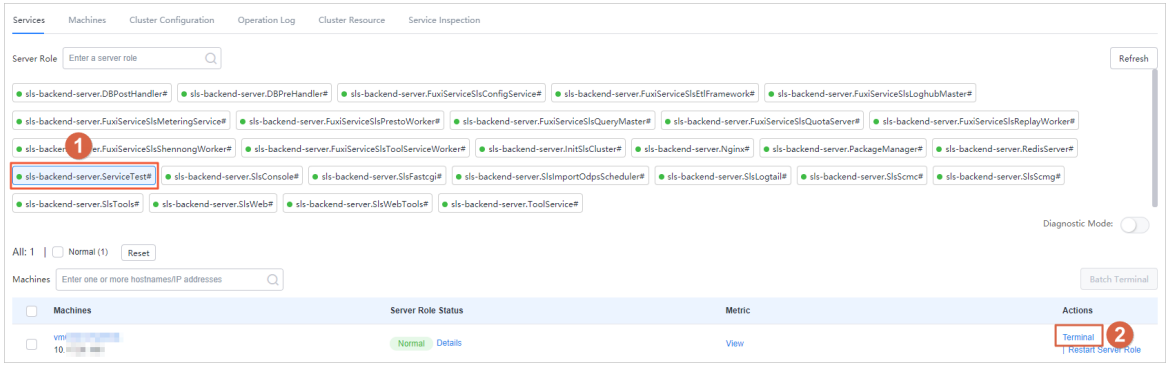

8. Log on to the server by using the terminal service and go to the related directory to view logs.

Portal

You can collect logs from your server and send the logs to the portal service. Then, you can query, 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. 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**.

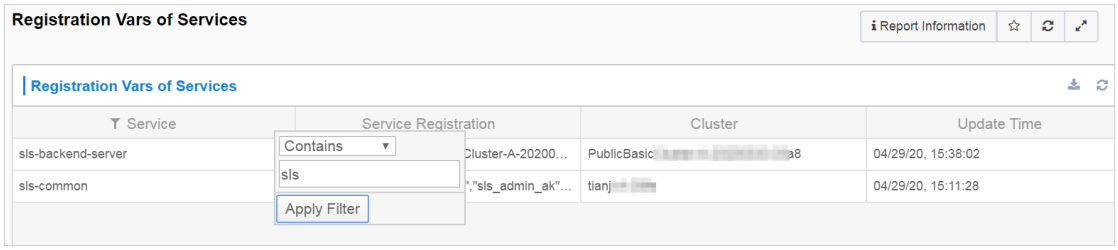

   v. Right-click the **Service Registration** column of the **sls-backend-server** service, and then select **Show More**.

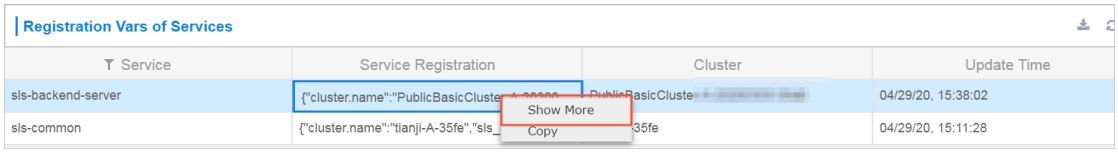

vi. On the **Details** page, find the endpoint of the Portal.

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 **Cluster** tab, find the destination cluster, and then choose > **Cluster Configuration File**.

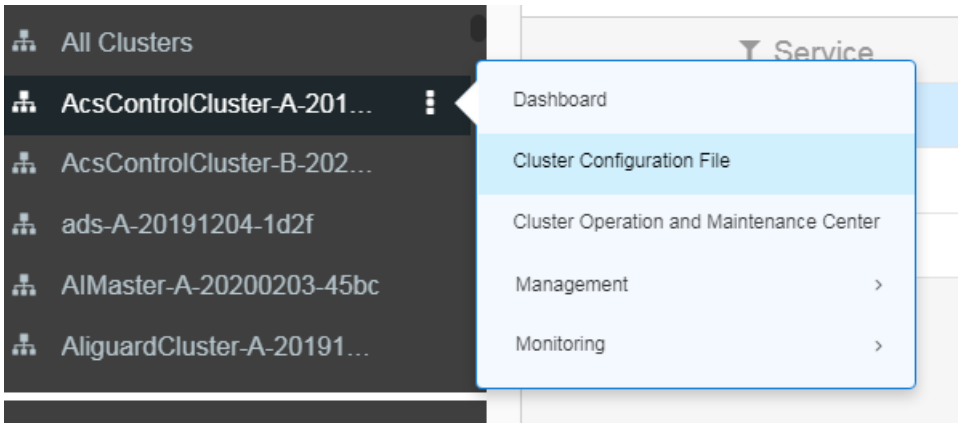

iv. Click **kv.conf** to obtain the AccessKey pair of the portal service.

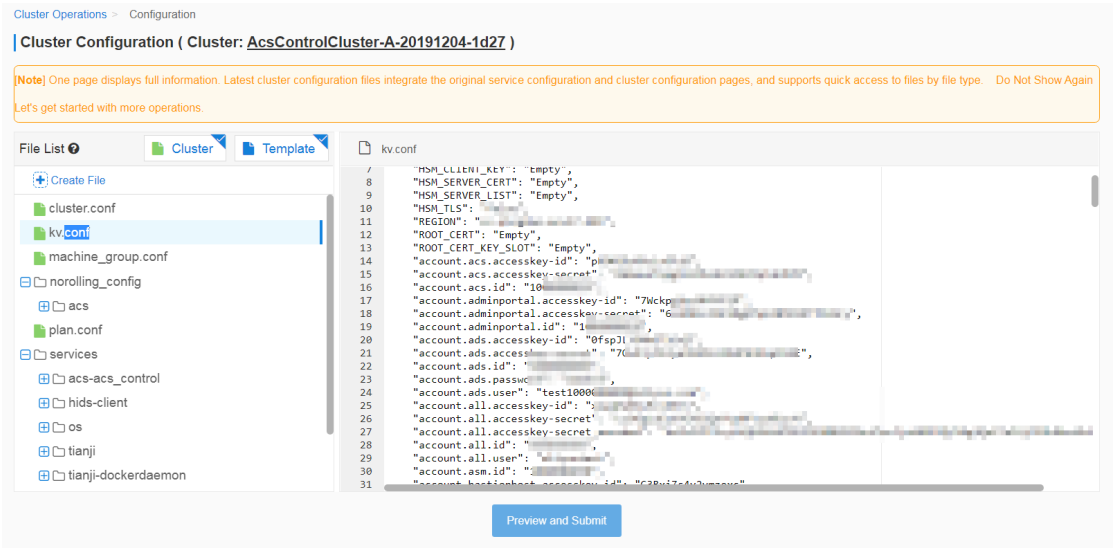

3. Log on to the Apsara Infrastructure Management Framework console.

   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 destination project and Logstore, and then query and analyze logs.

4.11.1.2. O&M

4.11.1.2.1. 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
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.

4.11.1.2.2. 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
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

Logstore	Log directory
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
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

Logstore	Log directory
sls_portal_api_audit	/apsara/sls/web/logsapidetail.log

4.12. Apsara Stack Security

4.12.1. Operations and Maintenance Guide

4.12.1.1. 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, username, and password that are required to log on to the Apsara Uni-manager Operations Console are obtained from the deployment personnel or administrators.

  The URL of the Apsara Uni-manager Operations Console is in the format of *region-id*.ops.console.*intranet-domain-id*.
- A browser is available. We recommend that you use Google Chrome.

Procedure

1. Open your browser.
2. In the address bar, enter the URL. Then, press the Enter key.

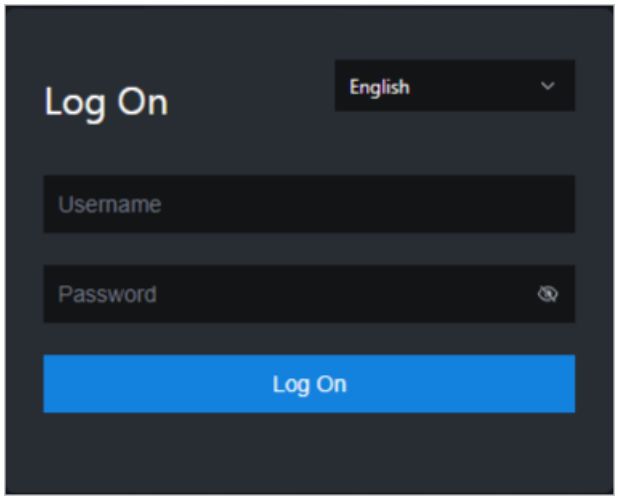

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

> **Note** To obtain the username and password that are used to log on to the Apsara Uni-manager Operations Console, contact the deployment personnel or administrators.

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

To enhance security, make sure that the password meets the following requirements:

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
- The password is 10 to 20 characters in length.

4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. In the **Apsara Stack O&M** section, click **Apsara Infrastructure Management Framework**.

4.12.1.2. Routine operations and maintenance of Server Guard

4.12.1.2.1. Check the service status

4.12.1.2.1.1. 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.

4.12.1.2.1.2. 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.file=/home/admin/aegiserverlite/.default/conf/logging.properties -Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager -server -Xms2g -Xmx2g -XX:PermSize=96m -XX:MaxPermSize=384m -Xmn1g -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=aegiserverlite -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 -Djava.endorsed.dirs=/opt/taobao/tomcat/endorsed -classpath /opt/taobao/tomcat/bin/bootstrap.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.tmpdir=/home/admin/aegiserverlite/.default/temp org.apache.catalina.startup.Bootstrap -Djboss.server.home.dir=/home/admin/aegiserverlite/.default -Djboss.server.home.url=file:/home/admin/aegiserverlite/.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*.

4.12.1.2.1.3. 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.

4.12.1.2.1.4. 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
   ```

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.

4.12.1.2.2. 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`

4.12.1.3. Routine operations and maintenance of Network Traffic Monitoring System

4.12.1.3.1. Check the service status

4.12.1.3.1.1. 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.

4.12.1.3.1.2. 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.

4.12.1.3.2. Common operations and maintenance

4.12.1.3.2.1. 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`

4.12.1.3.2.2. 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`

4.12.1.3.2.3. 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
   ```

4.12.1.3.2.4. 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
   > ```

4.12.1.4. Routine operations and maintenance of Anti-DDoS Service

4.12.1.4.1. Check the service status

4.12.1.4.1.1. 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.

4.12.1.4.1.2. 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

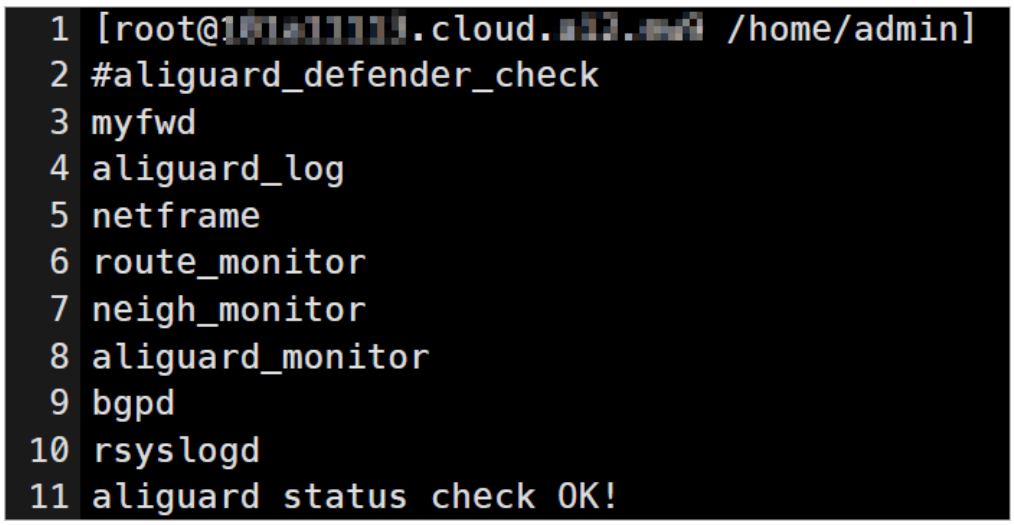

- If the error message shown in Reinjection route error message is returned, the reinjection route is incorrect.

  Reinjection route error message

  ```
  1 Error: route status error, we need two default routes to reinject the net flow!
  2 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.

   ```
   [root@          cloud.am54 /root]
   #lspci -v | grep Eth
   02:00.0 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
   02:00.1 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
   04:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
           Subsystem: Intel Corporation Ethernet Server Adapter X520-2
   04:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
           Subsystem: Intel Corporation Ethernet Server Adapter X520-2
   81:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
           Subsystem: Intel Corporation Ethernet Server Adapter X520-2
   81:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
           Subsystem: Intel Corporation Ethernet Server Adapter X520-2
   ```

4.12.1.4.2. Common operations and maintenance

4.12.1.4.2.1. 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.

4.12.1.4.2.2. 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.

4.12.1.5. Routine operations and maintenance of Threat Detection Service

4.12.1.5.1. Check the service status

4.12.1.5.1.1. 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.

4.12.1.5.1.2. 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.

4.12.1.5.2. 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
   ```

4.12.1.6. Routine operations and maintenance of WAF

4.12.1.6.1. Check the service status

4.12.1.6.1.1. 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 **SemaWafCluster**.
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.

4.12.1.6.1.2. 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 **SemaWafCluster**. 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.

4.12.1.7. Routine operations and maintenance of Security Audit

4.12.1.7.1. Check service status

4.12.1.7.1.1. Basic inspection

The basic inspection of Audit log 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 **Cluster** search box, enter **BasicThinCluster**.
4. Click the name of the cluster. The **Cluster Details** page appears.
5. On the **Services** tab, enter yundun-security-auditlog in the **Service** search box to search for the service. Then, check whether the service status is normal.

4.12.1.7.1.2. Advanced inspection: Check the status of the security-auditlog-app service

This topic describes how to check the status of the security-auditlog-app service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to two physical machines on which Auditlog is deployed.
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Cluster** search box, enter **BasicThinCluster**. Then, click the name of the cluster. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-security-auditlog** in the **Service** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Server Role** section, select **yundun-security-auditlog.SecirityAuditlogAPP#**.
   v. View the machine information. Click **Terminal** in the Actions column to log on to two physical machines on which Auditlog is deployed.
3. Log on to two Docker containers of auditlog-app.
   Run the `sudo docker exec -it $(sudo docker ps | grep auditlog-app | awk '{print $1}') bash` command.
4. Check the status of the security-auditlog process.
   Run the `ps aux | grep java | grep security-auditlog` command. If a process record is returned, the service is normal.
5. Check the health status.
   Run the `curl 127.0.0.1:3001/check.htm` command. If success is returned, the service is normal.
6. View logs.
   - View the Tomcat logs in */home/admin/security-auditlog/logs/jboss_stdout.log*.
   - View the audit logs in */home/admin/security-auditlog/logs/audit-exec.log*.
   - View the business error logs in */home/admin/security-auditlog/logs/biz-error.log*.
   - View the check error logs in */home/admin/security-auditlog/logs/check-error.log*.
   - View the scheduling task execution logs in */home/admin/security-auditlog/logs/job-exec.log*.
   - View the remote service call logs in */home/admin/security-auditlog/logs/remote-exec.log*.
   - View the service invocation logs in */home/admin/security-auditlog/logs/service-exec.log*.
   - View the system error logs in */home/admin/security-auditlog/logs/system-error.log*.
   - View the download task logs in */home/admin/security-auditlog/logs/task-exec.log*.
   - View other logs in */home/admin/security-auditlog/logs/main.log*.

4.12.1.7.1.3. Advanced inspection: Check the security-auditlog-syslog service status

This topic describes how to check the status of the security-auditlog-syslog service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to two physical machines on which Auditlog is deployed.
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Cluster** search box, enter **BasicThinCluster**. Then, click the name of the cluster. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-security-auditlog** in the **Service** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Server Role** section, select **yundun-security-auditlog.SecirityAuditlogAPP#**.
   v. View the machine information. Click **Terminal** in the Actions column to log on to two physical machines on which Auditlog is deployed.
3. Log on to two Docker containers of auditlog-syslog.
   Run the `sudo docker exec -it $(sudo docker ps | grep auditlog-syslog | awk '{print $1}') bash` command.
4. Check the syslog-ng process status.
   Run the `ps aux | grep syslog-ng | grep -v grep` command. If two process records are returned, the syslog-ng process is normal.
5. Check the status of port 2514 of the syslog-ng process.
   Run the `netstat -ano | grep 2514` command. If multiple records are returned, port 2514 of the syslog-ng process is normal.
6. Check the ilogtail process status.
   Run the `ps aux | grep ilogtail | grep -v grep` command. If two process records are returned, the ilogtail process is normal.
7. View logs.
   - View the ilogtail logs in */usr/local/ilogtail/ilogtail.LOG*.
   - View the syslog-ng logs in */var/log/messages*.

4.12.1.7.2. Restart Security Audit

Context

To restart Security Audit when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server of Auditlog.
2. Run the following command to find the image ID of Auditlog:
   `docker ps -a |grep service name`

3. Run the following command to go to the Docker container:

   ```
   docker exec -it [imageId] /bin/bash
   ```

4. Restart related services.
   - Restart the security-auditlog-app service.
     a. Run the following command to stop the current application process:

        ```
        kill -9 $(ps -ef | grep java | grep security-auditlog | grep -v grep | awk '{print $2}')
        ```

     b. Run the following command to restart the application process:

        ```
        /home/admin/security-auditlog/bin/jbossctl restart
        ```

     c. Run the following command to check whether the process has been successfully restarted:

        ```
        curl 127.0.0.1:7001/check.htm
        ```

        If the response is "success", the service is normal.
   - Restart the security-auditlog-syslog service.
     a. Run the following command to restart the syslog-ng process:

        ```
        service syslog-ng restart
        ```

     b. Run the following command to check whether the syslog-ng process is normal:

        ```
        ps aux | grep syslog-ng | grep -v grep
        ```

        If two records are returned, the syslog-ng process is normal.

     c. Run the following command to check whether port 2514 of the syslog-ng process is enabled:

        ```
        netstat -ano | grep 2514
        ```

        If multiple records are returned, port 2514 of the syslog-ng process is enabled.

     d. Run the following command to restart the ilogtaild process:

        ```
        /etc/init.d/ilogtaild stop
        /etc/init.d/ilogtaild start
        ```

     e. Run the following command to check whether the ilogtail process is normal:

        ```
        ps aux | grep ilogtail | grep -v grep
        ```

        If two records are returned, the ilogtail process is normal.

4.12.1.8. Routine operations and maintenance of Sensitive Data Discovery and Protection

4.12.1.8.1. Check the service status

4.12.1.8.1.1. Basic inspection

The basic inspection of Sensitive Data Discovery and Protection (SDDP) checks whether the service status is normal.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation page, choose **Operations > Cluster Operations**.
3. In the **Clusters** search box, enter **SddpVMCluster**.
4. Click the name of the target cluster. The **Cluster Details** page appears.
5. On the **Services** tab, enter yundun-sddp in the **Services** search box. Then, check whether the service status is normal.

4.12.1.8.1.2. Advanced inspection: Check the status of SddpService

This topic describes how to check the running status of SddpService.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical machines of Sensitive Data Discovery and Protection (SDDP).
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **SddpVMCluster**. Then, click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-sddp** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SddpService#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two security audit physical machines.
3. Log on to the two Docker containers of SddpService.

   Run the `sudo docker exec -it $(sudo docker ps | grep SddpService | awk '{print $1}') bash` command.
4. Check the process status of SddpService.

   Run the `ps aux | grep java | grep yundun-sddp-service` command. If a 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.acti
ve=cloud -server -Xms4g -Xmx4g -Xmn2g -XX:MetaspaceSize=256m -XX:MaxMetaspaceSize=512m -XX:MaxDirectMemorySize=1g
-XX:SurvivorRatio=10 -XX:+UseConcMarkSweepGC -XX:CMSMaxAbortablePrecleanTime=5000 -XX:+CMSClassUnloadingEnabled -X
X: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.defaultReadTime
out=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=70
01 -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/yu
ndun-sddp-service/target/yundun-sddp-service -Djava.endorsed.dirs= -Djava.io.tmpdir=/home/admin/yundun-sddp-servic
e/.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 **success** is returned, the service is normal.

```
#curl 127.0.0.1:7001/checkpreload.htm
"success"
```

6. View 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 frontend 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.

4.12.1.8.1.3. 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 machines of Sensitive Data Discovery and Protection (SDDP).
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **SddpVMCluster** and click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-sddp** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SddpData#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two security audit physical machines.
3. Log on to the two Docker containers of the SddpData service.
   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 a record is returned, the service is normal.
5. View logs.
   View logs in the */home/admin/yundun-sddp-data/logs/sddp.log* file.

4.12.1.8.1.4. 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 machines of Sensitive Data Discovery and Protection (SDDP).
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **SddpVMCluster**. Then, click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-sddp** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SddpPrivilege#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two security audit physical machines.
3. Log on to the two Docker containers of the SddpPrivilege service.

   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 a record is returned, the service is normal.
5. Check the health status.

   Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If **success** is returned, the service is normal.
6. View 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.

4.12.1.8.1.5. 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 machines of Sensitive Data Discovery and Protection (SDDP).
   i. In the left-side navigation pane, choose **Operations > Cluster Operations**.
   ii. In the **Clusters** search box, enter **SddpVMCluster**. Then, click the cluster name. The **Cluster Details** page appears.
   iii. On the **Services** tab, enter **yundun-sddp** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SddpLog#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two security audit physical machines.
3. Log on to the two Docker containers of the SddpLog service.

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` command. If a record is returned, the service is normal.

5. Check the health status.

   Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If **success** is returned, the service is normal.

6. View 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.

4.12.1.8.2. 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.

4.12.1.9. Routine operations and maintenance of Apsara Stack Security Center

4.12.1.9.1. Check service status

4.12.1.9.1.1. Basic inspection

The basic inspection of the secure-console service 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-secureconsole in the **Services** search box. Then, check whether the service status is normal.

4.12.1.9.1.2. Advanced inspection

This topic describes how to check the running status of the secure-console service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical machines.
   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-secureconsole** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SecureConsoleApp#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two security audit physical machines.
3. Log on to the two Docker containers of the secure-console service.
   Run the `sudo docker exec -it $(sudo docker ps | grep secureconsole | awk '{print $1}') bash` command.
4. Check the process status of the secure-console service.
   Run the `ps aux |grep console` command. If a record is returned, the service is normal.
5. Check the health status.
   Run the `curl 127.0.0.1:3014/check.htm` command. If `ok` is returned, the service is normal.
6. View logs.
   View Tomcat logs in the */home/admin/console/logs/jboss_stdout.log* file.

4.12.1.9.2. 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`
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
   ```

4.12.1.10. Routine operations and maintenance of secure-service

4.12.1.10.1. Check the service status

4.12.1.10.1.1. Basic inspection

The basic inspection of secure-service 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-secureservice in the **Services** search box. Then, check whether the service status is normal.

4.12.1.10.1.2. Advanced inspection: Check the status of secure-service

This topic describes how to check the running status of secure-service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Log on to the two physical machines.
   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-secureservice** in the **Services** search box. Then, click **Details**. The **Service Details** page appears.
   iv. In the **Service Role** search box, enter **SecureServiceApp#**.
   v. View the machine information, and click **Terminal** in the Actions column to log on to the two security audit physical machines.
3. Log on to the two Docker containers of secure-service.

   Run the `sudo docker exec -it $(sudo docker ps | grep secureservice | awk '{print $1}') bash` command.
4. Check the process status of secure-service.

   Run the `ps aux |grep secure-service` command. If a record is returned, secure-service is normal.

5. Check the health status.

   Run the `curl 127.0.0.1:3010` command. If `ok` is returned, the service is normal.
6. Run the following command to enter the Docker container:

   `sudo docker exec -it [imageId] /bin/bash`
7. View logs.
   - View Server Guard logs in the */home/admin/secure-service/logs/aegis-info.log* file.
   - View vulnerability analysis and scanning logs in the */home/admin/secure-service/logs/leakage-info.log* file.
   - View cloud intelligence logs in the */home/admin/secure-service/logs/threat-info.log* file.
   - View web attack logs in the */home/admin/secure-service/logs/web-info.log* file.

4.12.1.10.1.3. 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/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.

4.12.1.10.1.4. 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*.

4.12.1.10.2. 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
```

4.13. Key Management Service (KMS)

4.13.1. Operations and Maintenance Guide

4.13.1.1. O&M of KMS components

4.13.1.1.1. 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.

4.13.1.1.2. 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 Apsara Uni-manager Operations Console and the username and password used to log on to the console are obtained from a deployment engineer or an administrator.

  The URL of the Apsara Uni-manager Operations Console is in the format of *ops*.asconsole.*intranet-domain-id*.com.
- A browser is available. We recommend that you use Google Chrome.

Procedure

1. Open your browser.
2. In the address bar, enter the URL. Then, press the Enter key.

   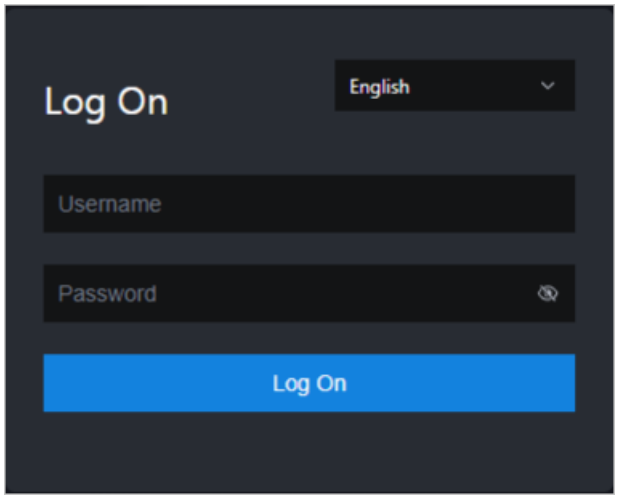

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

   To ensure security, your password must meet the complexity requirements. The password must be 10 to 20 characters in length. It must contain the following character types: uppercase letters, lowercase letters, digits, and special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).

4. Click **Log On** to log on to the Apsara Uni-manager Operations Console.

5. In the top navigation bar, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Apsara Stack O&M** section, click **Apsara Infrastructure Management Framework**.

4.13.1.1.3. KMS_HOST

This topic describes how to check the status of the KMS_HOST service.

Check whether the server role is normal

1. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the **Cluster Operations** page, search for the KMS cluster.

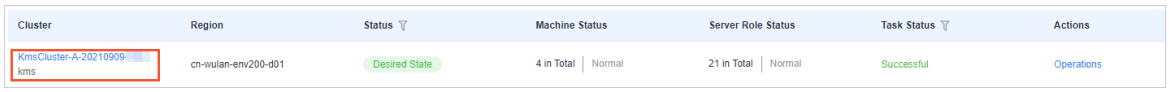

4. Click the cluster name in the Cluster column to go to the **Cluster Details** page.
5. On the **Services** tab, click **kms** in the Service column to go to the **Service Details** page.
6. On the **Service Details** page, check whether the KMS_HOST server role is at desired state.

   If the indicator for the kms.KmsHost# server role is green, the server role is at desired state.
7. On the **Service Details** page, click **kms.KmsHost#**. The information of the machines on which the KMS_HOST service is deployed is displayed in the lower part of the page. The IP addresses of the machines are required in subsequent steps.

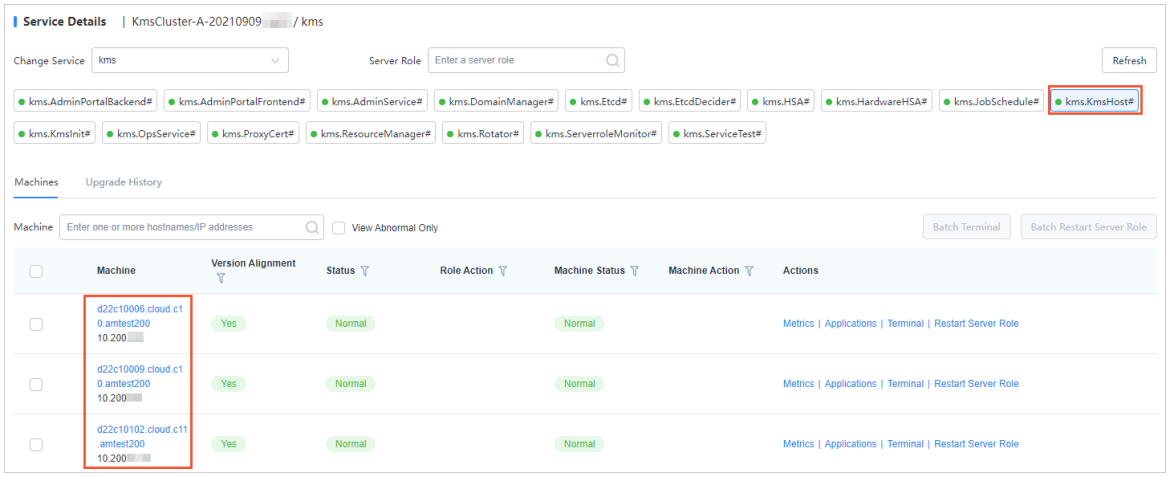

8. Click **Terminal** in the Actions column for a machine to log on to this machine.
9. Run the `curl http://ip:5555/status.html` command and check whether success is returned.

> **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 KMS_HOST service is deployed.

Troubleshooting

1. View logs in the */cloud/log/kms/KmsHost#/kms_host* directory.
2. Check whether the KMS_HOST service is normal.
   - If the KMS_HOST service abnormally exits after it starts, view debug logs in the *debug.log* file to troubleshoot the specific error.
   - If the KMS_HOST is running but does not function as expected, view status logs in the *status.log* file to troubleshoot the specific error.

Possible errors

Error	Troubleshooting
xxx selfCheck error **Note** *xxx* indicates a service on which the KMS_HOST service depends.	• Check whether the dependency configuration is valid. You can view debug logs in the *debug.log* file to troubleshoot the specific error. • Check whether the xxx service is normal.
exit code 1	View debug logs in the *debug.log* file to identify the cause of the abnormal exit.

4.13.1.1.4. HSA

This topic describes how to check the status of the HSA service.

Check whether the server role is normal

1. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the **Cluster Operations** page, search for the KMS cluster.

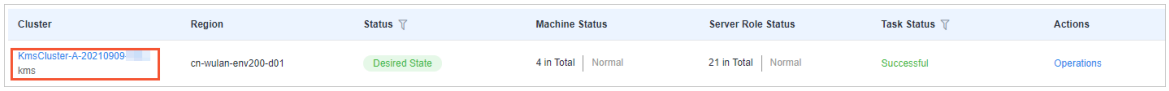

4. Click the cluster name in the Cluster column to go to the **Cluster Details** page.
5. On the **Services** tab, click **kms** in the Service column to go to the **Service Details** page.
6. On the **Service Details** page, check whether the HSA server role is at desired state.

   If the indicator for the kms.HSA# server role is green, the server role is at desired state.
7. On the **Service Details** page, click **kms.HSA#**. The information of the machines on which the HSA service is deployed is displayed in the lower part of the page. The IP addresses of the machines are required in subsequent steps.

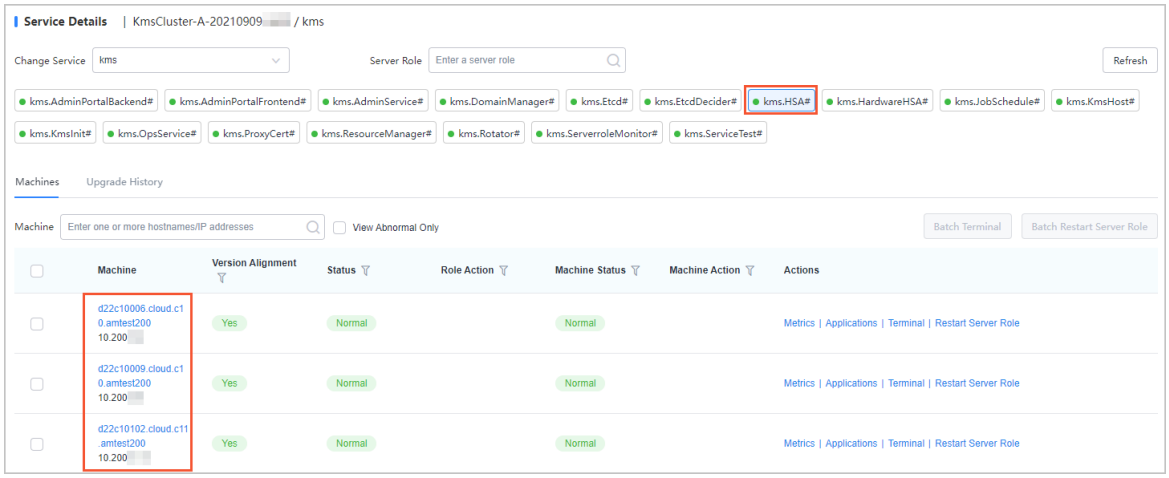

8. Click **Terminal** in the Actions column for a machine to log on to this machine.
9. Run the `curl http://ip:5555/status.html` command and check whether success is returned.

```
$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 HSA service is deployed.

Troubleshooting

1. View logs in the */cloud/log/kms/HSA#/hsa* directory.
2. Check whether the HSA service is normal.
   - If the HSA service abnormally exits after it starts, view debug logs in the *debug.log* file to troubleshoot the specific error.
   - If the HSA service is running but does not function as expected, view status logs in the *status.log* file to troubleshoot the specific error.

Possible errors

Error: `exit code 1`

Troubleshooting: View debug logs in the debug.log file to identify the cause of the abnormal exit. This error can be caused by one of the following reasons:

- The etcd service does not start as expected.
- The etcd service starts as expected, but its data is invalid.

> **Note** During disaster recovery, synchronization errors in the secondary cluster may cause this error.

4.13.1.1.5. etcd

This topic describes how to check the 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 are at desired state.

1. Log on to the Apsara Infrastructure Management Framework console. For more information, see Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. On the **Cluster Operations** page, search for the KMS cluster.

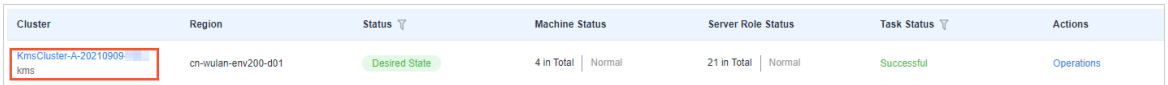

4. Click the cluster name in the Cluster column to go to the **Cluster Details** page.
5. On the **Services** tab, click **kms** in the Service column to go to the **Service Details** page.
6. On the **Service Details** page, check whether the Etcd and EtcdDecider server roles are at desired state.

   If the indicators for the kms.Etcd# and kms.EtcdDecider# server roles are green, the server roles are at desired state.

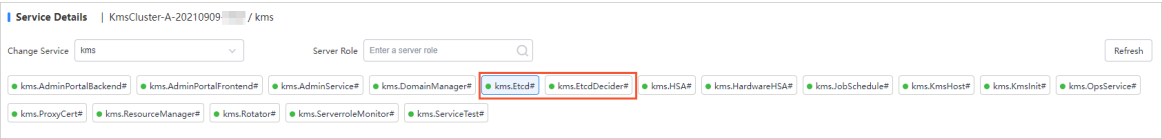

Troubleshooting

View logs in the */cloud/log/kms/Etcd#/etcd* and */cloud/log/kms/EtcdDecider#/decider* directories to troubleshoot errors of the etcd service.

Possible errors

Error	Troubleshooting
The startup parameters of the etcd service are invalid.	Find the correct settings of the startup parameters in the historical records of the *debug.log* file. Then, manually start the etcd service. **Note** Retain the error logs and request the technical support team to identify the cause.
Errors in the EtcdDecider server role during service updates cause errors in the etcd service.	In most scenarios, this issue occurs when a rolling task exists. You can analyze the issue and identify the cause based on the *debug.log* file of the EtcdDecider server role.

Error	Troubleshooting
The data directory of the etcd service is missing and the etcd service cannot start.	Use the Apsara Infrastructure Management Framework console to remove the abnormal etcd node from its server role group, and then add it back.

4.13.1.1.6. Rotator

4.13.1.1.6.1. Primary data center

This topic describes how to check the status of the rotator of the primary data center.

The rotator is a special component. Even if the server role of the rotator is at desired state in the Apsara Infrastructure Management Framework console, the rotator is not necessarily working normally.

Rotator exceptions do not have an impact on the API logic of Key Management Service (KMS).

In most scenarios, you must identify the cause of a rotator exception only when unexpected results are found, for example, when the data in ApsaraDB RDS does not meet expectations.

Check whether the rotator starts in primary data center mode

View the value of `current idc master` in the */cloud/log/kms/Rotator#/rotator/debug.log* file to check whether the rotator of the primary data center starts in primary data center mode, 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 work in distributed lock mode to ensure that only one node can work at a time. All the other nodes remain in the standby state.

View status logs in the */cloud/log/kms/Rotator#/rotator/status.log* file and check the status of each node.

Check the value of the RotatorState parameter in the status logs. Valid values: ExecuteWorker and TryLock.

- ExecuteWorker: The node is in the working state.
- TryLock: The node is in the standby state.

  Working state

  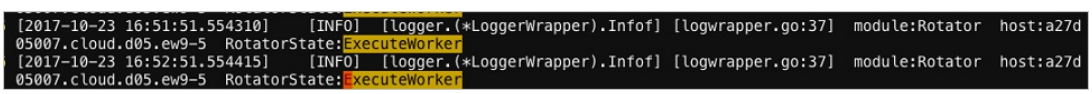

  Standby state

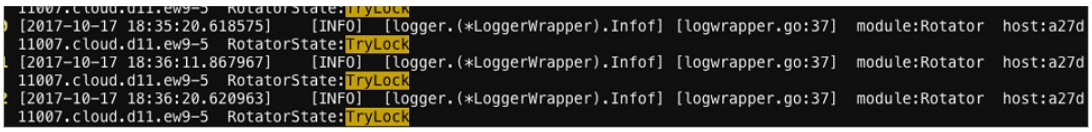

Possible errors

- Abnormal ApsaraDB RDS database access. Statistics collection and key deletion tasks cannot be executed.
- Abnormal HSA service. Key rotation tasks cannot be executed.
- Abnormal Log Service. Metering tasks cannot be executed.
- Abnormal etcd service. 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 be at desired state. However, when this occurs, the rotator may still appear to be at desired state in the Apsara Infrastructure Management Framework console.

4.13.1.1.6.2. 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.

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.

4.13.1.2. Log analysis

4.13.1.2.1. 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.

4.13.1.2.2. View logs by using request IDs

After you send a request to Key Management Service (KMS), KMS sends you a message that contains a request ID. This topic describes how to view logs by using request IDs.

Request IDs can be used in the following scenarios:

- You can view the KMS audit logs in the */cloud/log/kms/KmsHost#/kms_host/audit.log* file.

  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 the debug logs based on the value of request_id.

  Path to the on-premises logs: */cloud/log/kms/KmsHost#/kms_host/debug.log*

  > **Note** The */cloud/log/kms/KmsHost#/kms_host/debug.log* and *audit.log* files are stored on the same machine.

- If you need all details of a request, you can view detailed information in the trace logs.

  Path to the on-premises logs: */cloud/log/kms/KmsHost#/kms_host/debug.log*

  > **Note** The */cloud/log/kms/KmsHost#/kms_host/debug.log* and *audit.log* files are stored on the same machine.

- You can associate a cryptographic API operation with the trace logs of HSA by using the value of request_id.

  Path to the on-premises logs: */cloud/log/kms/HSA#/hsa/trace.log*

  > **Note** The */cloud/log/kms/KmsHost#/kms_host/trace.log* and *audit.log* files may be stored on different machines.

- You can retrieve log information based on other information.

  You can retrieve the information in the audit logs of KMS by using information other than request_id. If you need to associate the audit logs with other logs, you must use request_id.

4.13.1.2.3. Common KMS errors

4.13.1.2.3.1. 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.

4.13.1.2.3.2. 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.

4.13.1.2.3.3. 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.

4.13.1.2.3.4. Errors with HTTP status code 503

An error with HTTP status code 503 occurs when the user interrupts the connection or a dependency service of Key Management Service (KMS) is abnormal. We recommend that you handle the exception at the earliest opportunity.

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 causes:

  - The client of a user interrupts the connection in advance.
  - The client times out because KMS that functions as the server does not respond within the timeout period.

  You can check the trace logs 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 the error in the debug logs by using the value of request_id. We recommend that you contact the technical support of the dependency service for further assistance.

4.13.1.2.3.5. 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.

4.13.1.3. Log on to the KMS O&M platform

This topic describes how to log on to the KMS O&M platform.

Prerequisites

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

  The URL of the Apsara Uni-manager Operations Console is in the following format: ${region}.ops.console.${intranet-domain}.
- A browser is available. We recommend that you use Google Chrome.

Procedure

1. Open your browser.
2. In the address bar, enter the URL of the Apsara Uni-manager Operations Console in the ${region}.ops.console.${intranet-domain} format. Then, press the Enter key.

   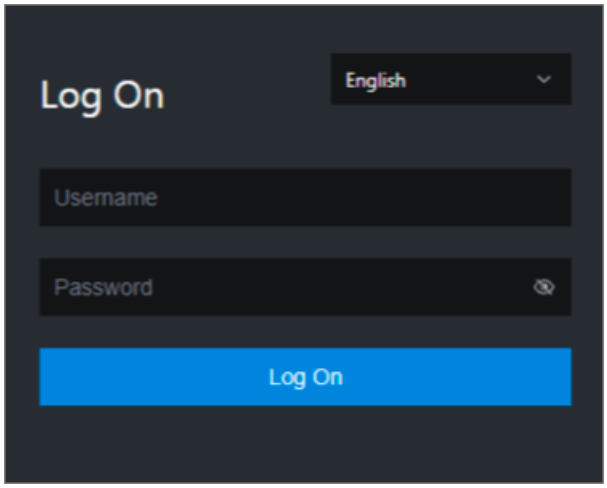

   Note You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

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

   To ensure security, your password must meet the complexity requirements. The password must be 10 to 20 characters in length. It must contain the following character types: uppercase letters, lowercase letters, digits, and special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
4. Click **Log On** to log on to the Apsara Uni-manager Operations Console.
5. In the top navigation bar, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Apsara Stack O&M** section, click **KMS**.

4.13.1.4. Deploy a managed HSM

This topic describes how to deploy a managed HSM.

Context

Managed HSM is an important feature of KMS to enable easy access to certified HSMs. An HSM is a highly secure hardware device that performs cryptographic operations and generates and stores keys. You can store the keys for your most sensitive Apsara Stack workloads and assets in HSMs that are managed on Apsara Stack.

> **Note** HSM refers to hardware security module.

Step 1: Register an HSM

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. Create an O&M event that is used to register an HSM.
   i. In the left-side navigation pane, choose **O&M Management > Resource**.
   ii. On the **Resource** tab, click **Register Resource**.
   iii. In the **Register Resource** dialog box, set the parameters.

Parameter	Description
Service Provider	Select Tass.
Service Provider ID	Enter Tass.
HSM IP Address	The IP address of the HSM. Enter the IP address that is planned by Deployment Planner, such as 10.4.XX.XX.
Additional Information	Enter `{"managerId":"","sk":"","managerSSLType":-1,"cryptoSSLType":-1,"firmwareVersion":"","rebootUniqueId":""}`.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

   iv. Click **OK**.
3. Query the details of the O&M event.
   i. In the left-side navigation pane, choose **O&M Management > Event**.

ii. On the **Event** tab, find the O&M event and check the status of the O&M event in the **Status** column.

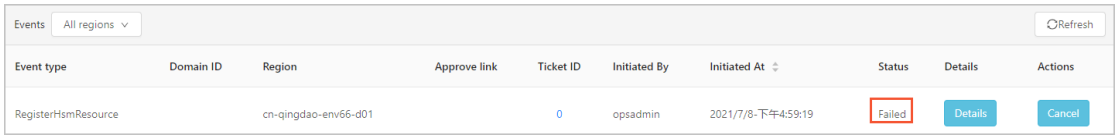

iii. Click **Details** to query the event details.

4. Query the task for the O&M event.
   i. In the left-side navigation pane, choose **O&M Management > Tasks**.
   ii. On the **Tasks** tab, check the status of the task for the O&M event and the number of times that the task was run.

Step 2: Upgrade the firmware of an HSM

> **Note** After you register at least three HSMs, you can upgrade the firmware of HSMs.

1. Query the firmware version of an HSM.
   i. In the left-side navigation pane, choose **O&M Management > Resource**.
   ii. On the **Resource** tab, find the HSM for which you want to upgrade the firmware and click **Other Details** in the **View Details** column.
   iii. In the dialog box that appears, obtain the firmware version of the HSM.
2. If the firmware version is not 1.00.06_GAWSGVH4.00.20.rv00, upgrade the firmware of the HSM.
   i. In the left-side navigation pane, choose **O&M Management > Resource**.
   ii. On the **Resource** tab, find the HSM for which you want to upgrade the firmware and click **Upgrade**.

iii. In the **Upgrade Firmware** dialog box, set the parameters.

Parameter	Description
Local File	Enter `/cloud/app/kms/HardwareHSA#/hardware_hsa/current/firmware/tapki_upgrade_AliKMS_GAWSGVH4.00.20.rv00./cloud/app/kms/HardwareHSA#/hardware_hsa/current/firmware/tapki_upgrade_AliKMS_GAWSGVH4.00.20.rv00.sig`.
OSS Bucket	Leave this parameter empty unless otherwise required.
OSS Object	Leave this parameter empty unless otherwise required.
Upgrade Type	Enter `NO_ZEROIZE`.
Firmware Version	Enter `1.00.06_GAWSGVH4.00.20.rv00`.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

iv. Click **OK**.

> **Note** After the firmware is updated, the system generates an O&M event. You can view the details about the event by choosing **O&M Management > Event**. You can also view the task status for the event by choosing **O&M Management > Tasks**.

Step 3: Create an HSM domain

1. In the left-side navigation pane, choose **O&M Management > Resource**.
2. On the **Resource** tab, click **Create Domain**.
3. In the **Create Domain** dialog box, set the parameters.

Parameter	Description
Service Provider	Select Tass.
Domain ID	Enter `kms_hsm_private`.
Maximum Number of Keys	Enter 100.
Maximum Number of Records in SSL Context	Enter 100.
Maximum Number of Cores	Enter 10.

Parameter	Description
Number of Topologies	Enter 3 in the single-data center scenario and 4 in the zone-disaster recovery scenario.
Number of Keys	Enter 2.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

4. Click **OK**.

> Note
> - After the HSM domain is created, the system generates an O&M event. You can view the details about the event by choosing **O&M Management > Event**. You can also view the task status for the event by choosing **O&M Management > Tasks**.
> - After the HSM domain is created, the HSM domain whose ID is `kms_hsm_private` appears in the domain list. The number of topologies is that specified by the **Number of Topologies** parameter.

Step 4: Scale out or in the HSM domain

> **Notice** Scale-outs and scale-ins are unconventional O&M operations that require risk assessment before proceeding.

1. In the left-side navigation pane, choose **O&M Management > Resource**.
2. On the **Resource** tab, find the HSM domain that you want to scale out or in and click **Change Topology**.
3. In the **Change Topology** dialog box, set the parameters.

Parameter	Description
Number of Topologies	Enter a positive number to scale out the HSM domain and a negative number to scale in the HSM domain. For example, 1 indicates that one HSM is added, and -1 indicates that one HSM is removed.
HSM IP Address	Leave this parameter empty unless otherwise required. If you enter an IP address, the HSM that uses this IP address is added or removed.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.

Parameter	Description
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

4. Click **OK**.

> **Note** After the HSM domain is scaled out or in, the system generates an O&M event. You can view the details about the event by choosing **O&M Management > Event**. You can also view the task status for the event by choosing **O&M Management > Tasks**.

What to do next

If you no longer need to use an HSM, you can bring the HSM offline.

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. In the left-side navigation pane, choose **O&M Management > Resource**.
3. On the **Resource** tab, find the HSM that you want to bring offline and click **Bring Offline**.
4. In the **Bring Offline** dialog box, set the **Ticket ID**, **Ticket Link**, **Force Offline**, **Restore Factory Settings**, and **Description** parameters.
5. Click **OK**.

4.13.1.5. Quota management

4.13.1.5.1. Manage throttling rules

This topic describes how to create and modify a throttling rule.

Create a throttling rule

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. In the left-side navigation pane, choose **O&M Management > Throttling**.
3. On the **Throttling** tab, click **Create Rule**.
4. In the **Create Rule** dialog box, set the parameters.

Parameter	Description
Resources	The name of the resource. Valid values: ○ s: applies to all API operations. ○ d: applies to the default API operation.

Parameter	Description
Source	The method to combine user and service IDs for your application. Valid values: ○ d: combines all users and service IDs. This is the default value. ○ d_d: combines the default user and service IDs.
Threshold	The threshold for throttling.
Throttling Policy	○ Standalone Throttling: Throttling is implemented on a single KMS API operation. ○ Standalone Token Bucket: Throttling is implemented on a single user.
Unit	The time unit for throttling.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

5. Click **OK**.

Modify a throttling rule

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. In the left-side navigation pane, choose **O&M Management > Throttling**.
3. On the **Throttling** tab, find the throttling rule that you want to modify and click **Modify** in the **Resources** column.
4. In the **Modify** dialog box, set the parameters.

Parameter	Description
Threshold	The threshold for throttling. **Note** If you want to change the threshold, you must change it for two rules. One rule has Resources set to s and Source to d. The other rule has Resources set to s and Source to d_d.

Parameter	Description
Throttling Policy	○ Standalone Throttling: Throttling is implemented on a single KMS API operation. ○ Standalone Token Bucket: Throttling is implemented on a single user.
Unit	The time unit for throttling.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

5. Click **OK**.

What to do next

After the throttling rule is created or modified, the system generates an O&M event. You can view the details about the event by choosing **O&M Management > Event**. You can also view the task status for the event by choosing **O&M Management > Tasks**.

4.13.1.5.2. Manage user quotas

This topic describes how to create and modify a user quota.

Create a user quota

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. In the left-side navigation pane, choose **O&M Management > User Quotas**.
3. On the **User Quotas** tab, click **Create Quota**.
4. In the **Create Quota** dialog box, set the parameters.

Parameter	Description
User ID	The ID of the user.
Maximum Number of Key Versions	The maximum number of key versions that can be created.
Maximum Number of Aliases	The maximum number of aliases that can be created.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.

Parameter	Description
Description	The description for your operation. Both Chinese and English are supported.

5. Click **OK**.

Modify a user quota

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. In the left-side navigation pane, choose **O&M Management > User Quotas**.
3. On the **User Quotas** tab, find the ID of the user whose quota you want to modify and click **Modify** in the **Actions** column.
4. In the **Modify** dialog box, set the parameters.

Parameter	Description
Maximum Number of Key Versions	The maximum number of key versions that can be created.
Maximum Number of Aliases	The maximum number of aliases that can be created.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

5. Click **OK**.

What to do next

After the user quota is created or modified, the system generates an O&M event. You can view the details about the event by choosing **O&M Management > Event**. You can also view the task status for the event by choosing **O&M Management > Tasks**.

4.13.1.5.3. Configure global quotas

This topic describes how to configure global quotas.

Procedure

1. Log on to the KMS O&M platform. For more information, see Log on to the KMS O&M platform.
2. In the left-side navigation pane, choose **O&M Management > Global Quotas**.
3. On the **Global Quotas** tab, find the HSM for which you want to configure global quotas and click **Modify** in the **Actions** column.
4. In the **Modify** dialog box, set the parameters.

Parameter	Description
Maximum Quantity	The maximum number of key versions that can be created.
Ticket ID	The ID of the change ticket in your change management system. If no tickets are created, enter 0.
Ticket Link	The link for the change ticket in your change management system. If no tickets are created, enter ignore.
Description	The description for your operation. Both Chinese and English are supported.

5. Click **OK**.

> **Note** After the global quota is modified, the system generates an O&M event. You can view the details about the event by choosing **O&M Management > Event**. You can also view the task status for the event by choosing **O&M Management > Tasks**.

4.14. Apsara Stack DNS

4.14.1. Operations and Maintenance Guide

4.14.1.1. 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.

4.14.1.2. Maintenance

4.14.1.2.1. 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.

4.14.1.2.2. 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.

4.14.1.2.3. 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.

4.14.1.3. DNS API

4.14.1.3.1. 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

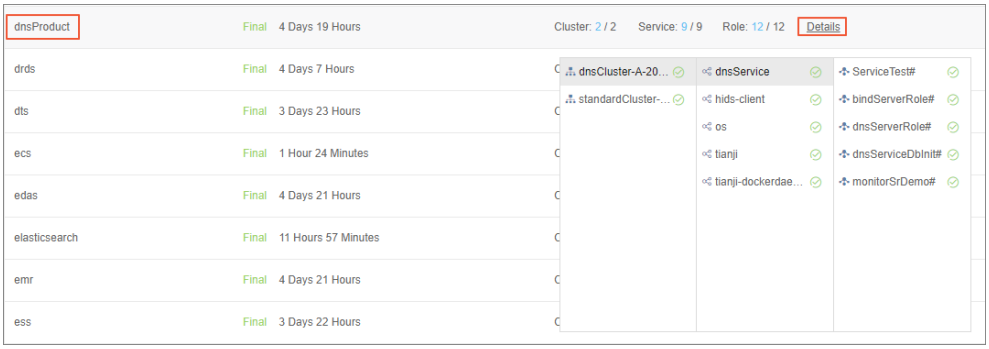

2. Obtain the IP addresses of servers where the API services are deployed.

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

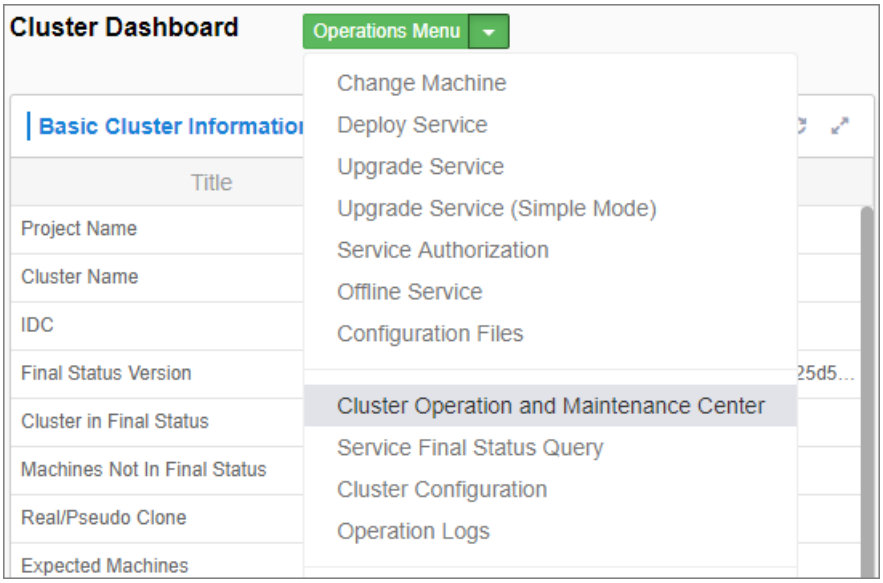

v. On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of servers that are deployed with the API service.

View the IP addresses of servers

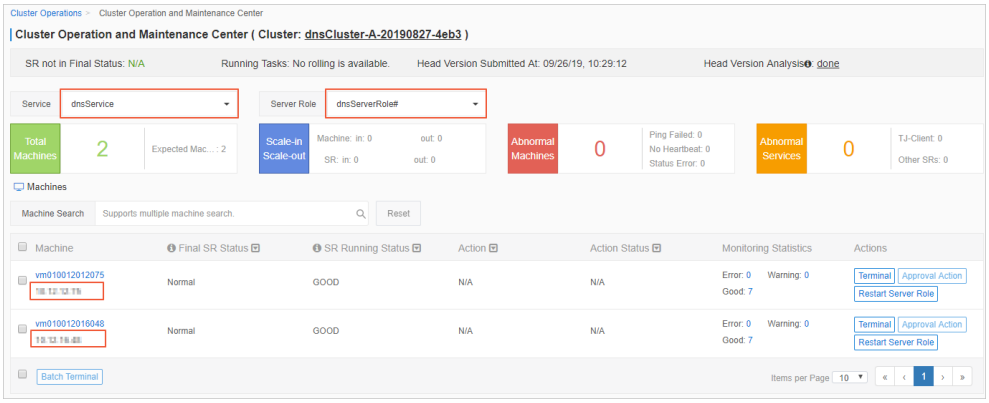

3. Log on to the DNS API server. Run the `curl http://localhost/checkpreload.htm` command, and check whether the command output is "success".
   i. Log on to the Apsara Infrastructure Management Framework console.
   ii. In the top navigation bar, choose **Operations > Machine Operations**.
   iii. Click **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

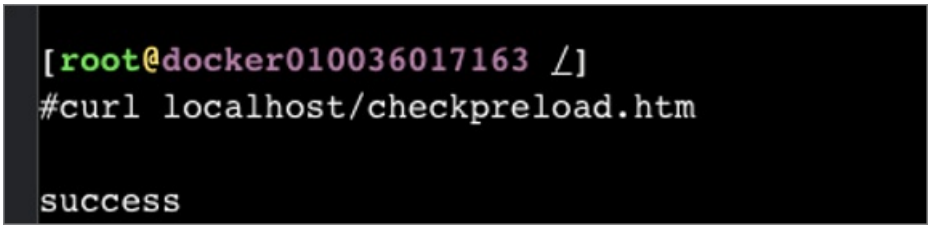

4.14.1.3.2. Troubleshooting

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.

4.14.1.4. DNS system

4.14.1.4.1. 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

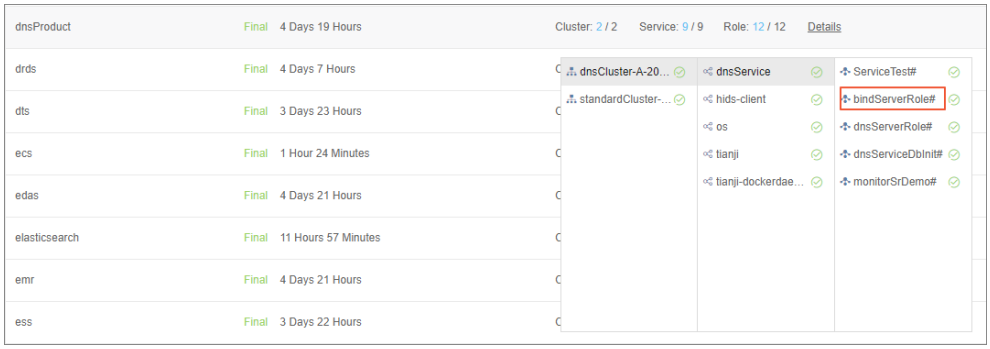

2. Obtain the IP addresses of the servers where DNS services are deployed.

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

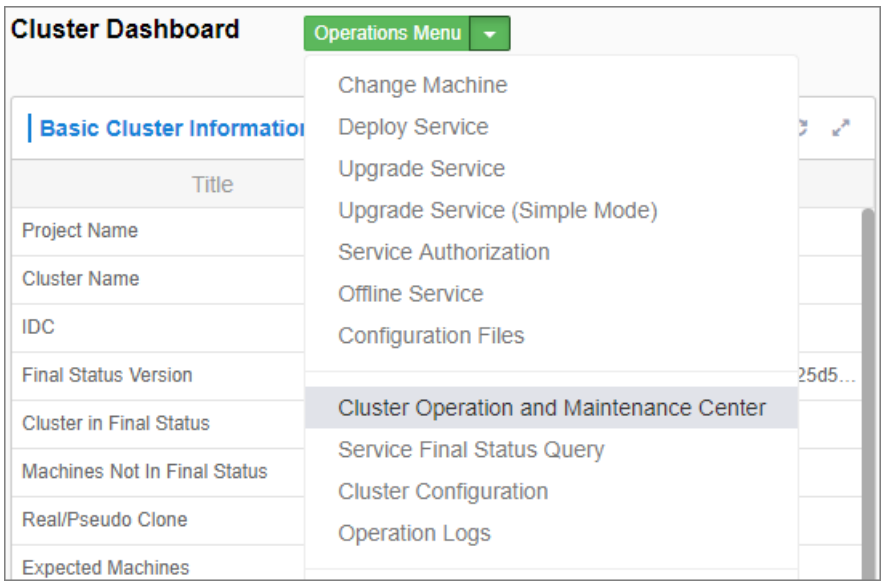

v. On the Cluster Operation and Maintenance Center page, view and obtain IP addresses of all the servers that are assigned with the bindServerRole# role.

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

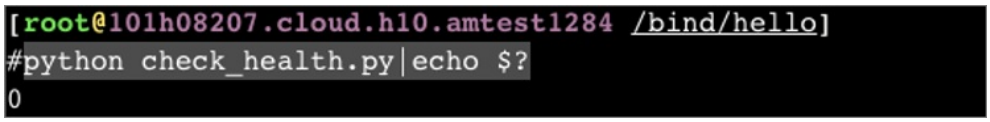

4.14.1.4.2. Troubleshooting

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.

4.14.1.4.3. 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.

4.14.1.5. 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.

4.14.1.6. 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.

4.15. API Gateway

4.15.1. Operations and Maintenance Guide

4.15.1.1. 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.

4.15.1.2. Routine maintenance

4.15.1.2.1. 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.

4.15.1.2.2. 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.

4.15.1.3. API Gateway O&M

4.15.1.3.1. System O&M

4.15.1.3.1.1. 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.

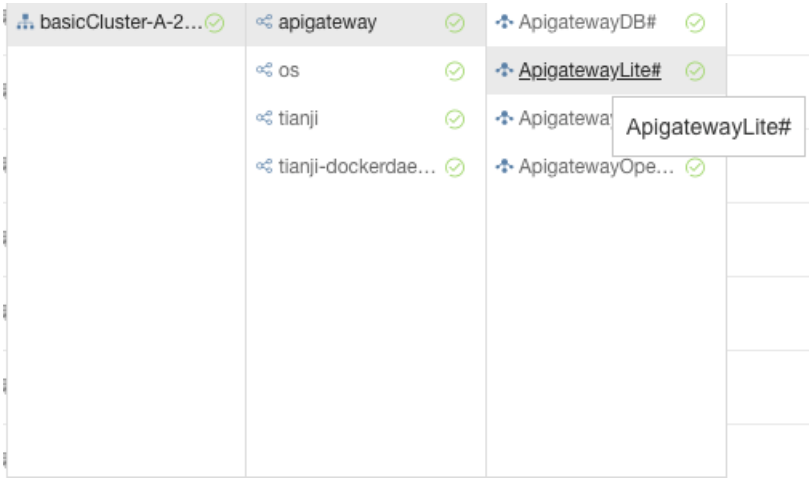

4.15.1.3.1.2. Check the service status of OpenAPI

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.

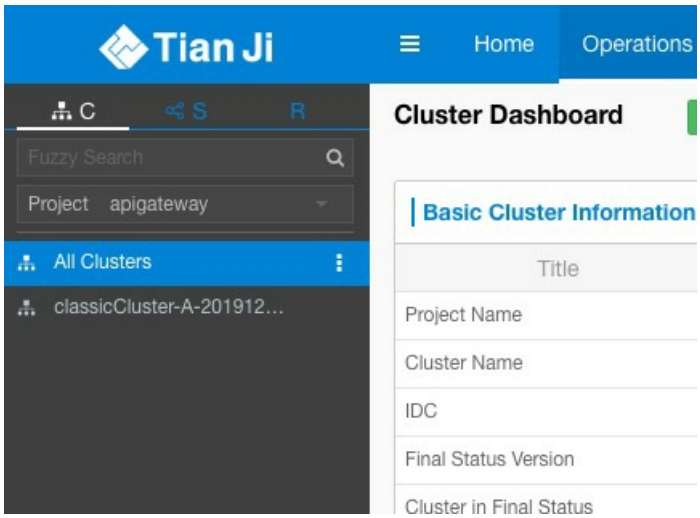

   iv. Place the pointer over the icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.
   v. In the **Service Instance List** section, click Details in the Actions column corresponding to the apigateway service instance.
   vi. In the **Server Role List** section, you can view the deployment status of each role.

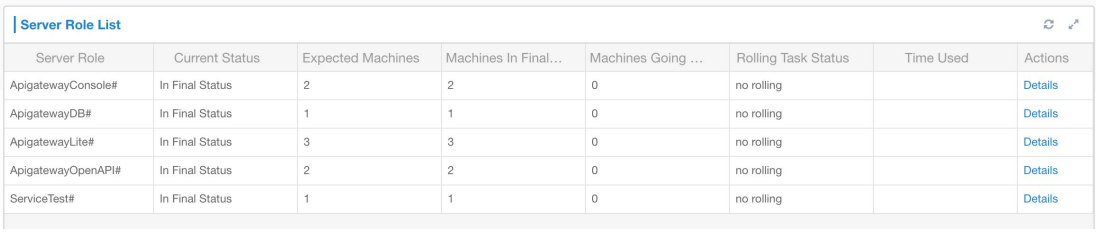

   vii. Click Details in the Actions column corresponding to the ApigatewayOpenAPI# role and view machine information of the role in the **Machine Information** section.

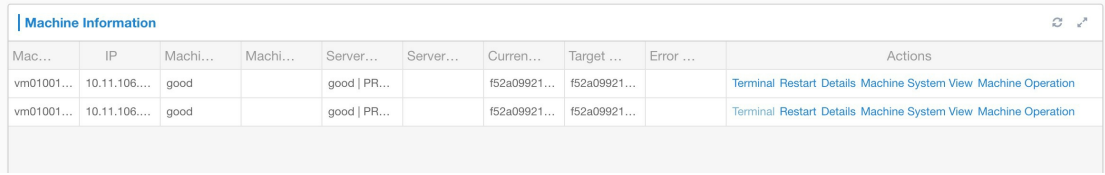

2. Click **Terminal** in the Actions column corresponding to a machine to log on to the machine.
3. Run the following command to find the container:

   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.

4.15.1.3.1.3. 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.

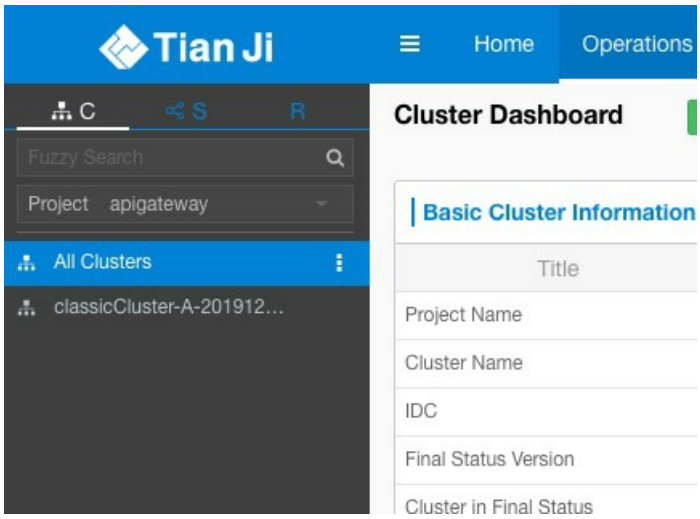

   iv. Place the pointer over the icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.
   v. In the **Service Instance List** section, click Details in the Actions column corresponding to the apigateway service instance.

vi. In the **Server Role List** section, you can view the deployment status of each role.

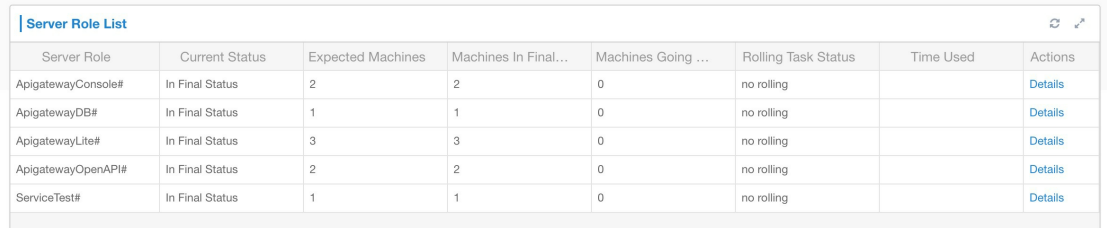

Server Role List

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.

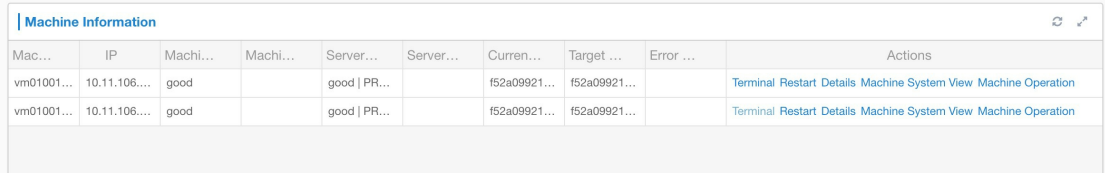

Machine Information

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.

4.15.1.3.1.4. 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.

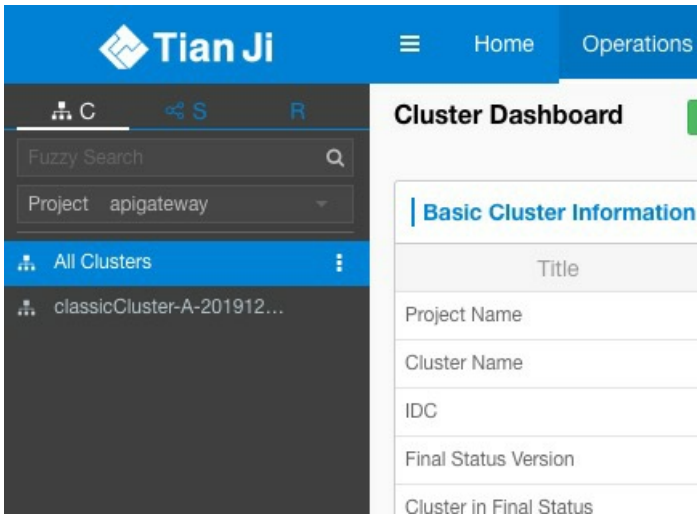

iv. Place the pointer over the icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.

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.

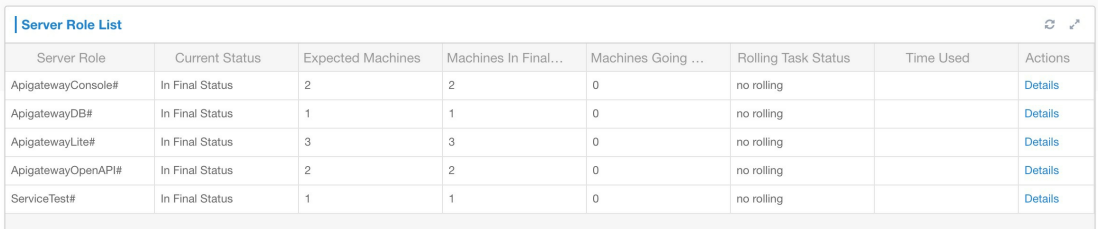

vii. Click Details in the Actions column corresponding to the ApigatewayLite# role and view machine information of the role in the **Machine Information** section.

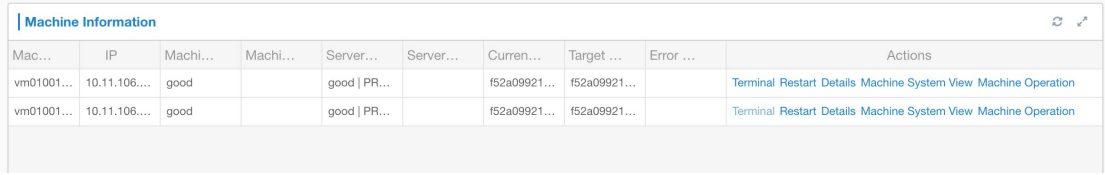

2. Click **Terminal** in the Actions column corresponding to a machine to log on to the machine.
3. Run the following command to check whether the `I'm fine, thank you, and you?` message is returned:

   curl -i http://localhost/status -H Host:status.taobao.com

4.15.1.3.1.5. 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.

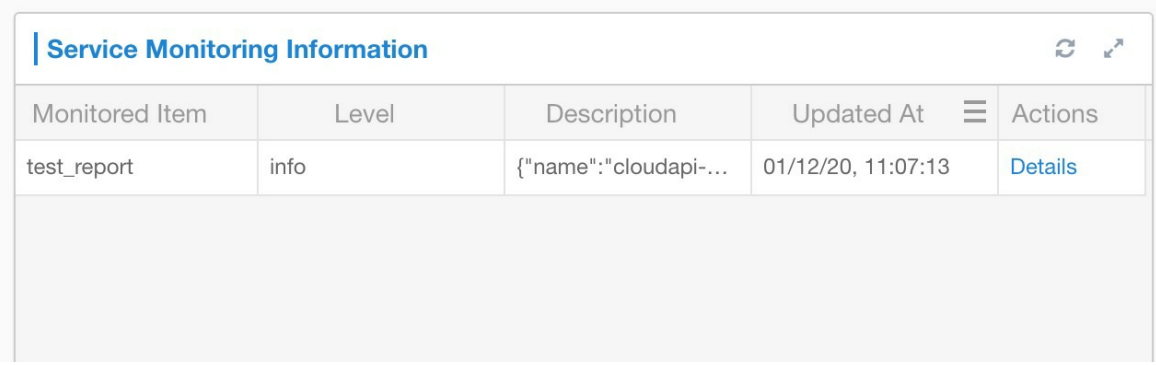

4.15.1.3.2. Troubleshooting

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.

4.15.1.4. 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.

5.Operations of middleware products

5.1. Message Queue for Apache RocketMQ

5.1.1. Operations and Maintenance Guide

5.1.1.1. O&M overview

This topic describes the product architecture and O&M architecture of Message Queue for Apache RocketMQ.

5.1.1.1.1. Product architecture

This topic is intended to help you understand the system architecture and components of Message Queue for Apache RocketMQ during O&M.

Message Queue for Apache RocketMQ provides components that are responsible for service discovery, data storage, authentication, configuration, log collection, and high availability.

System architecture

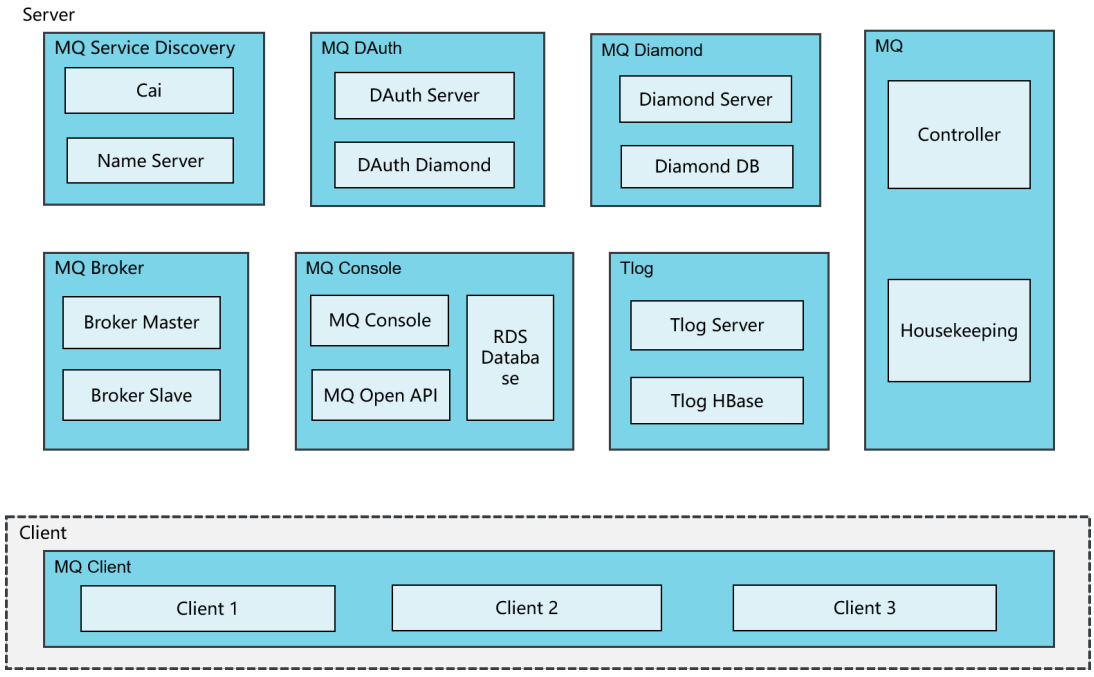

Message Queue for Apache RocketMQ consists of the following components:

Service discovery components

Message Queue for Apache RocketMQ provides Name Servers and Address Servers for service discovery.

An Address Server, also called Cai, is used to register and discover domain names for Name Servers and Diamond. A caller obtains domain names from the Address Server and locally caches the domains names.

A Name Server is a core component for flexible deployment and linear scaling of Message Queue for Apache RocketMQ. This component is used to register and search for message queues. A broker registers message queues with a Name Server and synchronizes the status of the broker to the Name Server. Producers and consumers can find specific brokers through Name Servers. After producers and consumers are connected to Name Servers, the producers and consumers can cache the information about the required brokers and message queues, and periodically update the cached information from the Name Servers.

Broker

A broker is a core component that is used to process message queues. Multiple brokers can be deployed in cluster, multi-replica, or primary/secondary mode to achieve high availability. Brokers provide message service capabilities that feature high availability, high stability, high efficiency, and linear scaling.

A broker registers topic and subscription information to Name Servers to provide services. Brokers are responsible for sending, receiving, and storing messages.

Console

The Message Queue for Apache RocketMQ console is a user interface (UI) that allows user access. All operations on message queues can be performed in the console. We recommend that you preferentially perform these operations in the console.

API

The Message Queue for Apache RocketMQ API provides a set of API operations that you can call over HTTP and HTTPS. This helps you use Message Queue for Apache RocketMQ and manage resources with ease. For example, you can call API operations to create topics, query group IDs, query messages, and query the status of consumers.

DAuth

DAuth provides unified logon services and access control for message queues. For example, it provides resource permission control, cross-account and RAM user access control, and resource authorization.

Diamond

Diamond is a configuration center that stores the configuration information of message queues, including VIP conversion rules and resource permission information.

Tlog

Tlog collects resource statistics and other key logs of Message Queue for Apache RocketMQ.

Controller

Controller is used to perform primary/secondary switchover for brokers to achieve high availability. It interacts with ZooKeeper to check the status of the brokers in a cluster. This way, it can ensure the consistency between the primary and secondary brokers.

Housekeeping

Housekeeping regularly checks the health status of the core link of Message Queue for Apache RocketMQ, covering core business, kernel parameters, application configurations, application metrics, and the console. If exceptions are detected, an alert is triggered at the earliest opportunity to ensure the high availability of the Message Queue for Apache RocketMQ cluster.

5.1.1.1.2. O&M architecture

This topic is intended to help you understand the O&M architecture of Message Queue for Apache RocketMQ.

The O&M operations of Message Queue for Apache RocketMQ rely on command lines.

O&M category	Description	O&M tool
Routine maintenance	Perform inspection and monitoring.	Command line: You can run commands to manually inspect the containers and components of Message Queue for Apache RocketMQ.
Shutdown maintenance	• Check and verify the statuses of containers and components. • Stop and start containers and components.	Command line
Troubleshooting	Handle component unavailability, service unavailability, and service discontinuity of Message Queue for Apache RocketMQ.	Command line

5.1.1.1.3. Updates

This topic describes the updates of Message Queue for Apache RocketMQ from V3.8.0 to V3.8.1 to help you get started with the updated version.

Optimization of resource isolation by instance

Message Queue for Apache RocketMQ provides instances for multi-tenancy isolation. Each user can purchase multiple instances and logically isolate them from each other.

To ensure the compatibility with the existing resources of existing users, Message Queue for Apache RocketMQ provides the following types of instances and namespaces:

- Default instances, which are compatible with the existing resources of existing users
  - This type of instance has no separate namespace. Resource names must be globally unique within and across all instances.
  - By default, an instance without a namespace is automatically generated for the existing resources of each existing user. If no existing resources are available, you can create at most one instance without a namespace.

- You can configure the endpoint, which can be obtained from the **Instances** page in the Message Queue for Apache RocketMQ console.

```
// Recommended configuration:
properties.put(PropertyKeyConst.NAMESRV_ADDR, "xxxx");
// Compatible configuration, which is not recommended. We recommend that you update this configur
ation to the recommended configuration:
properties.put(PropertyKeyConst.ONSAddr, "xxxx");
```

- New instances
  - A new instance has a separate namespace. Resource names must be unique within an instance but can be the same across different instances.
  - You can configure the endpoint, which can be obtained from the **Instances** page in the Message Queue for Apache RocketMQ console.

```
// Recommended configuration:
properties.put(PropertyKeyConst.NAMESRV_ADDR, "xxx");
```

  - A RocketMQ client must be updated to the following latest versions for different programming languages:
    - Java: V1.8.7.1.Final
    - C and C++: V2.0.0
    - .NET: V1.1.3

Optimization of resource application

Previously, Message Queue for Apache RocketMQ resources consisted of topics, producer IDs, and consumer IDs. Each two of the resources have a many-to-many relationship, which was difficult to comprehend. Each time you created a topic, you must associate the topic with a producer ID and a consumer ID. This process was too complex for medium- and large-sized enterprise customers.

To optimize user experience and help new users get started, the resource application process has been simplified.

The resource application process has been optimized in the following aspects:

- Topic management, which is unchanged
  - You need to apply for a topic. A topic is used to classify messages. It is the primary classifier.
- Group management
  - You do not need to apply for a producer ID. Producer IDs and consumer IDs are integrated into group IDs. In the Message Queue for Apache RocketMQ console, the Producers module has been removed. The Producers and Consumers modules have been integrated into the Groups module.
  - You do not need to associate a producer ID or consumer ID with a topic. Instead, you need only to apply for a group ID and associate it with a topic in the code.
  - Compatibility:
    - The list of producer IDs is no longer displayed. This does not affect the current services.
    - The consumer IDs that start with CID- or CID_ and that you have applied for can still be used and can be set in the PropertyKeyConst.ConsumerId or PropertyKeyConst.GROUP_ID parameter of the code.
- Sample code

> Note
>
> - We recommend that you update a RocketMQ client to the following latest versions for different programming languages:
>   - Java: V1.8.7.1.Final
>   - C and C++: V2.0.0
>   - .NET: V1.1.3
> - Existing producer IDs or consumer IDs can still be used and do not affect the current services. However, we recommend that you update your instance configuration to the recommended configuration.

- Recommended configuration: Integrate producer IDs and consumer IDs into group IDs.

```
// Set the PropertyKeyConst.GROUP_ID parameter. The original PropertyKeyConst.ProducerId and Pro
pertyKeyConst.ConsumerId parameters are deprecated.
properties.put(PropertyKeyConst.GROUP_ID, "The original CID-XXX or the GID-XXX");
```

- Compatible configuration: Use a producer ID to identify a producer and a consumer ID to identify a consumer.

```
// When you create a producer, you must set the PropertyKeyConst.ProducerId parameter.
properties.put(PropertyKeyConst.ProducerId, "The original PID-XXX or the GID-XXX");
// When you create a consumer, you must set the PropertyKeyConst.ConsumerId parameter.
properties.put(PropertyKeyConst.ConsumerId, "The original CID-XXX or the GID-XXX");
```

5.1.1.2. High-risk operations

5.1.1.2.1. Levels of O&M operations

O&M operations are classified into the following three levels based on the business scenarios and risk degree: G1, G2, and G3.

When you perform O&M operations as Level-1 and Level-2 support engineers, take note of the following rules:

G1: You can follow the instructions in related documents, without the need to apply for changes. Such operations will not affect the business.

G2: Before you can follow the instructions in related documents, you must apply for changes and obtain approval from the product line. Such operations will not affect the business.

G3: Before you can follow the instructions in related documents, you must submit a ticket to apply for changes and obtain approval from the product line and customer. Such operations may affect the business.

5.1.1.2.2. High-risk operations and files

This topic describes the high-risk operations that you must pay special attention to and the high-risk files that cannot be randomly deleted during O&M.

High-risk operations

High-risk operations are the G3-level O&M operations defined in the "Levels of O&M operations" topic. Follow the corresponding instructions when you perform such operations.

High-risk files

High-risk files are the files that cannot be randomly deleted. When you process these files, take note of the following rules:

- In principle, the files that have not dynamically grown since they are generated cannot be deleted. If you need to delete these files such as the legacy packages that are no longer used after upgrade, you must confirm this operation with the product line and customer.
- Do not delete configuration files. All content in the *conf* directory are protected.
- The files in the application package directory cannot be deleted.
- The log files that are in use cannot be deleted. For example, the *mw.log* file is in use and the *mw-2018-07-22.log* file is used to archive historical data. In this case, do not delete the *mw.log* file.
- A broker manage its high-risk files. You cannot manually delete these files.

High-risk files

File	Purpose	Description
/home/admin/store/commitlog/*	Stores the messages that are submitted by producers. These files are data files.	• The mapedFileSizeCommitLog parameter in the */home/admin/rmq/conf/broker.conf* file specifies the size of each file. By default, each file is 1 GB in size. • The deleteWhen parameter in the */home/admin/rmq/conf/broker.conf* file specifies the time to delete outdated data. By default, outdated data is deleted at 04:00 every morning. • The diskMaxUsedSpaceRatio parameter in the *home/admin/rmq/conf/broker.conf* file specifies the space usage threshold for deleting outdated data. By default, outdated data is deleted when the storage usage reaches 75%. • The leanFileForciblyEnable parameter in the *home/admin/rmq/conf/broker.conf* file specifies whether to forcibly delete outdated files when the disk is full.
/home/admin/store/config/consumerOffset.json	Stores the consumer offsets.	This file is managed by brokers and cannot be manually deleted.

File	Purpose	Description
/home/admin/store/config/delayOffset.json	Stores the consumer offsets of retry queues.	This file is managed by brokers and cannot be manually deleted.
/home/admin/store/config/subscriptionGroup.json	Stores the consumer group configurations.	This file is managed by brokers and cannot be manually deleted.
/home/admin/store/config/topics.json	Stores the topic configurations.	This file is managed by brokers and cannot be manually deleted.

5.1.1.3. O&M preparation

This topic describes the logon portal, account, permissions, and tools required for O&M.

O&M preparation

Item	Purpose	Description
Secure Shell (SSH) clients such as *MobaXterm* and *PuTTY*	Used to log on to the servers where related components reside.	The account used to log on to the servers must be granted permissions to perform related operations. However, we recommend that you do not use the root or admin account.
Account	Used to log on to the Message Queue for Apache RocketMQ console or the servers where related components reside.	• You can obtain the username and password for logging on to the console from the technical support engineers of Message Queue for Apache RocketMQ. • The account used to log on to the servers must be granted permissions to perform related operations. However, we recommend that you do not use the root or admin account.

5.1.1.4. Routine maintenance

Message Queue for Apache RocketMQ inspection and monitoring cover two aspects: components and services.

Component inspection and monitoring

Monitor the statuses of the following Message Queue for Apache RocketMQ components: Address Servers, Name Servers, brokers, Controller, Housekeeping, and the console.

Service inspection and monitoring

Inspect and monitor the core link of Message Queue for Apache RocketMQ during operation, covering core business, kernel parameters, application configurations, application metrics, and the console.

5.1.1.4.1. Component inspection and monitoring

Component inspection and monitoring for Message Queue for Apache RocketMQ cover the following aspects:

- Inspection: Periodically perform a dial test on URLs or ports to determine whether Message Queue for Apache RocketMQ components are normal. Inspections in HTTP, TCP, ping, and JDBC modes are supported.
- Monitoring: Collect logs from clients by using Tlog. Aggregate container monitoring data and custom metrics to measure the status of system operation.

5.1.1.4.1.1. Manual inspection

Linux commands are used for manual inspection. You can set specific parameters based on the actual environment.

Name server inspection

1. Log on to the container of a name server.
2. Run the `ps aux|grep -i NamesrvStartup` command to check whether the state of the Java process of the name server is normal.
3. Run the `ss -s` command to check the number of connections. If the number of connections is more than half of the maximum number of connections that are allowed for the name server, submit a ticket to contact O&M engineers.

Console inspection

The Message Queue for Apache RocketMQ console is a web system deployed in Tomcat. Perform the following steps to perform routine inspection:

1. Log on to the container of the Message Queue for Apache RocketMQ console.
2. Run the `netstat -an |grep 7001` command to check whether the Tomcat process exists and whether Port 7001 is enabled.
3. Check whether the NGINX process exists and whether Port 80 properly forwards packets to Port 7001.
4. Check the */home/admin/logs/ons-api/*.log* file to determine whether a major exception exists.
5. Check whether an unwanted log file needs to be cleared from the */home/admin/ons-api/logs* directory.
6. Check whether the disk size is sufficient in */home/admin*.

Controller inspection

1. Log on to the container of a controller.
2. Run the `ps aux|grep -i FailoverStartUp` command to check whether the state of the Java process of the controller is normal.
3. Run the `grep 'take the leadership of role' /home/admin/logs/rocketmqlogs/failover_controller.log` command to determine which of the two controllers is managing the failover.
4. Check the files in the `/home/admin/logs/rocketmqlogs/*.log` path to determine whether an obvious exception error exists.

Broker inspection

1. Log on to the container of a broker.
2. Run the `ps aux|grep brokerStartup` command to view the process information.
3. Run the `ss -s` command to check the number of connections. If more than half of the connections are in use, contact after-sales engineers.
4. View the logs of the broker in the */home/admin/logs/rocketmqlogs/* directory.
5. Run the `df -h` command. Make sure that the disk usage is less than 80%.
6. Check whether the following configurations exist in the */home/admin/store* directory:
   - */home/admin/store/config/consumerOffset.json*: consumer offsets.
   - */home/admin/store/config/delayOffset.json*: consumer offsets of scheduled messages.
   - */home/admin/store/config/subscriptionGroup.json*: consumer group configurations.
   - */home/admin/store/config/topics.json*: topic configurations.
   - */home/admin/store/consumequeue*: consumption index file. It is stored by topic.
   - */home/admin/store/index*: index file for key-based message query.
7. Check the commit log after a message is sent.
   i. Check whether a file is generated in the */home/admin/store/commitlog* directory.
   ii. Run the *du -sh /home/admin/store/commitlog/** command to check whether new data is written to the file.
   iii. Run the `sh /home/admin/rmq/bin/mqadmin clusterlist` command to view the message sending and receiving status of all brokers in the cluster. Check the message sending and receiving status of the current broker.

Address Server inspection

1. Log on to the container of an Address Server.

   An Address Server is an NGINX process module. By default, HTTP port 8080 is enabled.
2. Run the `curl -v 'http://localhost:8080/rocketmq/nsaddr4broker-internal'` command to check the configurations of Name Servers on the Address Server. If the configurations of Name Servers are normal, the status code 200 and the content in */home/admin/cai/htdocs/rocketmq/nsaddr4broker-internal* are returned.

Diamond Server inspection

1. Log on to the container of a Diamond Server.
2. Run the `df -lh` command to check the disk space. This prevents exceptions caused by full disk storage.
3. Run the `top` command to verify that the memory usage and CPU utilization are normal.
4. Run the `ping jmenv.tbsite.net -c 3` command to verify that Address Servers can be reached with a low latency. For example, Address Servers can be reached within 10 ms over internal networks.
5. Run a script to cyclically ping all Diamond Servers. Make sure that the Diamond Servers can connect to each other with a low latency.
6. Run the `tsar --nginx --load --cpu -l -i3` command to verify that the load on Diamond Servers is in a controllable range. The single-server queries per second (QPS) limit is approximately 4,000. We recommend that the single-server QPS be less than 500, the number of processes that are being executed be less than the number of CPU cores, and the CPU utilization be less than 20%.

DAuth inspection

DAuth is a web system deployed in Tomcat. Perform the following steps to perform routine inspection:

1. Log on to the container of DAuth.
2. Check whether the Tomcat process exists and whether Port 8080 is enabled.
3. Check whether the NGINX process exists and whether Port 80 properly forwards packets to Port 8080.
4. Check the */home/admin/dauth/logs/*.log* file to determine whether a major exception exists.
5. Check whether an unwanted log file needs to be cleared from the */home/admin/dauth/logs* directory.
6. Run the *curl 'localhost:8080'|grep Unsupported* command. If the command output contains the Unsupported keyword, DAuth is normal.
7. Check whether the disk size is sufficient in */home/admin*.
8. DAuth also serves as a client connected to Diamond Servers. Pay attention to the following aspects during inspection:
   - Check the */home/admin/logs/diamond-client/diamond-client.log* file to determine whether a major exception exists.
   - Run the `curl 'http://jmenv.tbsite.net:8080/diamond-server/diamond-unit-spas'` command to obtain the server list of the Diamond Spas unit. Ping each server. Make sure that the latency is less than 10 ms.

Tlog console inspection

- Log on to the container of the Tlog console.
- Run the `ps aux|grep tomcat` command to check the Tomcat process.
- Run the `df -h` command to check the disk space.
- Run the `curl -v 'http://localhost'` command to check whether the access is normal.

5.1.1.4.2. Service inspection and monitoring

The Housekeeping component of Message Queue for Apache RocketMQ is responsible for service inspection.

Housekeeping periodically inspects various metrics of Message Queue for Apache RocketMQ, and records alert information in the */home/admin/logs/aliware-mq-hkagent/warning.log* file in the container of Housekeeping. The metrics of Message Queue for Apache RocketMQ are fixed and cannot be changed.

5.1.1.4.3. Logs

Logs are one of the most important O&M methods for Message Queue for Apache RocketMQ. They are used for monitoring and troubleshooting. Therefore, routine log maintenance is required.

Message Queue for Apache RocketMQ logs include basic monitoring logs, service monitoring logs, container monitoring logs, and Java virtual machine (JVM) monitoring logs. For more information, see Log reference.

5.1.1.5. O&M commands

Backend O&M commands are provided for experienced O&M engineers. We recommend that you perform related operations in the Message Queue for Apache RocketMQ console.

5.1.1.5.1. Overview

Message Queue for Apache RocketMQ provides a backend O&M tool to facilitate routine O&M. The O&M tool is named sh mqadmin in the /home/admin/rmq/bin/ directory.

To run O&M commands, you must log on to the container of a Name Server or broker.

1. Run the docker exec -it <Name Server or broker container ID> /bin/bash command to log on to the container of the Name Server or broker.
2. Run O&M commands in the /home/admin/rmq/bin directory.

Command examples

Run the cd command in the mq-broker container to go to the bin directory, and then use the O&M tool to run relevant commands.

- Display all supported commands: `./mqadmin`
- Display the help information (usage and parameters) of the clusterList command: `./mqadmin help clusterList`
- Run the clusterList command to show the cluster list and details: `./mqadmin clusterList -m`

The clusterList command uses the following syntax: `mqadmin clusterList [-h] [-i <arg>] [-m] [-n <arg>]`. When you use the clusterList command, take note of the following rules:

- `-h`, `-i`, `-m`, and `-n` are short IDs of parameters. If a short ID is provided, the corresponding option and parameter are specified.
- The short IDs enclosed in brackets indicate optional parameters.
- arg enclosed in angle brackets indicates that the parameter must be set to a specific value, for example, `./mqadmin clusterList -i 3 -m`.

Resource naming

The naming rules for the resources contained in instances with or without namespaces have been changed in Message Queue for Apache RocketMQ V3.8.1 and later. Take note of the following rules:

- Topic naming
  - Instances with namespaces: instanceid%topic
  - Instances without namespaces: topic
- Group ID naming
  - Instances with namespaces: instanceid%gid
  - Instances without namespaces: gid

For information about instance namespaces, see Updates.

5.1.1.5.2. View all O&M commands

mqadmin

Command	Description
updateTopic	Create a topic.
deleteTopic	Delete a topic.
topicList	View the topic list information.
topicCLusterList	View the cluster of a topic.
topicRoute	View the topic route information.
updateTopicPerm	Update the read and write permissions for a topic.
topicStatus	View topic statistics.
brokerStatus	View the status of a broker.
wipeWritePerm	Revoke the specified permissions from a broker.
updateBrokerConfig	Update the configurations of a broker.
producerConnection	View the connection status of a producer.
consumerConnection	View the connection status of a consumer.
brokerConsumerStats	View the consumption status of all group IDs of a broker.
updateSubGroup	Create or modify a consumer group.
deleteSubGroup	Delete a consumer group.
consumerProgress	View the consumption status of a consumer group.
consumerStatus	Obtain the consumption progress of a consumer.
queryMsgById	Query messages by message ID.
queryMsgByKey	Query messages by message key.
queryMsgByOffset	Query messages by message offset.
printMsg	View the messages in the specified period under a topic.
clusterList	View the cluster information.
resetOffsetByTime	Reset the offset.

5.1.1.5.3. updateTopic

updateTopic

Parameter	Required	Description
-b	Yes if -c is empty	The address of the broker on which the topic resides.
-c	Yes if -b is empty	The name of the cluster on which the topic resides. You can run the clusterList command obtain the cluster list.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.
-p	No	The permission of the topic to be created. Valid values: W, R, and WR.
-r	No	The number of readable queues. Default value: 8.
-w	No	The number of writable queues. Default value: 8.
-t	Yes	The name of the topic.

5.1.1.5.4. deleteTopic

This command is high-risk. Proceed with caution.

deleteTopic

Parameter	Required	Description
-c	Yes	The name of the cluster on which the topic resides. You can run the clusterList command to obtain the cluster list.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.
-p	Yes	The name of the topic.

5.1.1.5.5. updateSubGroup

updateSubGroup

Parameter	Required	Description
-b	Yes if -c is empty	The address of the broker on which the topic resides.
-c	Yes if -b is empty	The name of the cluster on which the topic resides. You can run the clusterList command to obtain the cluster list.
-g	Yes	The name of the consumer group.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.
-d	No	Specifies whether to allow consumption in broadcasting mode.

Parameter	Required	Description
-m	No	Specifies whether to allow consumption starting from the minimum position of the queue. Default value: false.
-q	No	The number of retry queues for each consumer group. Messages to be consumed are placed in a retry queue.
-r	No	The maximum number of consumption retries. If the maximum number of consumption retries is exceeded for a message, the message is delivered to the dead-letter queue and no more delivery attempts are made.
-s	No	Specifies whether to enable consumption.

5.1.1.5.6. deleteSubGroup

This command is high-risk. Proceed with caution.

deleteSubGroup

Parameter	Required	Description
-b	Yes if -c is empty	The address of the broker on which the topic resides.
-c	Yes if -b is empty	The name of the cluster on which the topic resides. You can run the clusterList command to obtain the cluster list.
-g	Yes	The name of the consumer group.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.7. updateBrokerConfig

This command is high-risk. Proceed with caution.

updateBrokerConfig

Parameter	Required	Description
-b	Yes if -c is empty	The address of the broker on which the topic resides.
-c	Yes if -b is empty	The name of the cluster on which the topic resides. You can run the clusterList command to obtain the cluster list.
-k	Yes	The name of the broker configuration item.
-v	Yes	The value of the broker configuration item.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.8. topicList

topicList

Parameter	Required	Description
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.9. topicRoute

topicRoute

Parameter	Required	Description
-t	Yes	The name of the topic.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.10. queryMsgById

queryMsgById

Parameter	Required	Description
-i	Yes	The ID of the message.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.11. queryMsgByOffset

queryMsgByOffset

Parameter	Required	Description
-o	Yes	The offset of the message.
-i	Yes	The ID of the queue.
-b	Yes	The address of the broker on which the topic resides.
-t	Yes	The name of the topic.

5.1.1.5.12. queryMsgByKey

queryMsgByKey

Parameter	Required	Description
-k	Yes	The key content set in the message.
-t	Yes	The name of the topic.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.13. consumerConnection

consumerConnection

Parameter	Required	Description
-g	Yes	The group ID of the consumer group.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.14. consumerProgress

consumerProgress

Parameter	Required	Description
-g	Yes	The group ID of the consumer group. The format of the value varies depending on whether the instance has a namespace. • If the instance has a namespace, the value is in the format of instanceid%gid. • If the instance does not have a namespace, the value is in the format of gid.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.15. clusterList

clusterList

Parameter	Required	Description
-m	Yes	Specifies whether to display cluster statistics.
-i	No	The interval between executions, in seconds.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.16. consumerStatus

consumerStatus

Parameter	Required	Description
-g	Yes	The name of the consumer group.
-i	No	The IP address of the consumer.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.
-s	No	Specifies whether to display the stack information.

5.1.1.5.17. updateTopicPerm

This command is high-risk. Proceed with caution.

updateTopicPerm

Parameter	Required	Description
-b	Yes if -c is empty	The address of the broker.
-c	Yes if -b is empty	The name of the cluster.
-p	Yes	The type of permissions. A value of 2 indicates write-only permissions. A value of 4 indicates read-only permissions. A value of 6 indicates read and write permissions.
-t	Yes	The name of the topic.
-n	Yes	The addresses of Name Servers. The value is in the format of ip:port;ip:port.

5.1.1.5.18. topicClusterList

topicClusterList

Parameter	Required	Description
-n	Yes	The address of the Name Server. Example: 192.168.0.1:9876.
-t	Yes	The name of the topic.

5.1.1.5.19. brokerStatus

brokerStatus

Parameter	Required	Description
-b	Yes if -c is empty	The address of the broker. Example: 11.164.190.48:10911.
-c	Yes if -b is empty	The name of the cluster.
-n	Yes	The address of the Name Server. Example: 192.168.0.1:9876.

5.1.1.5.20. printMsg

printMsg

Parameter	Required	Description
-b	Yes	The beginning of the time range to query. Specify the time in the format of yyyy-MM-dd#HH:mm:ss:SSS.
-e	Yes	The end of the time range to query. Specify the time in the format of yyyy-MM-dd#HH:mm:ss:SSS.
-d	No	Specifies whether to display the message body.
-s	No	The tag of the message. The value is in the format of TagA || TagB. Default value: *.
-c	No	The encoding format, such as UTF-8 or GBK.
-t	Yes	The name of the topic.

5.1.1.5.21. brokerConsumeStats

brokerConsumeStats

Parameter	Required	Description
-b	Yes	The address of the broker. Example: 11.164.190.48:10911.
-t	No	The timeout period, in milliseconds. Default value: 50000.

5.1.1.5.22. producerConnection

producerConnection

Parameter	Required	Description
-g	Yes	The group ID of the producer. The format of the value varies depending on whether the instance has a namespace. • If the instance has a namespace, the value is in the format of instanceid%gid. • If the instance does not have a namespace, the value is in the format of gid.
-t	Yes	The name of the topic.

5.1.1.5.23. resetOffsetByTime

This command may cause messages to be repeatedly consumed or skipped without consumption.

resetOffsetByTime

Parameter	Required	Description
-g	Yes	The group ID of the consumer. The format of the value varies depending on whether the instance has a namespace. • If the instance has a namespace, the value is in the format of instanceid%gid. • If the instance does not have a namespace, the value is in the format of gid.
-s	Yes	The time. Specify the time in the format of yyyy-MM-dd#HH:mm:ss:SS.
-t	Yes	The name of the topic.

5.1.1.5.24. topicStatus

topicStatus

Parameter	Required	Description
-t	Yes	The name of the topic.

5.1.1.5.25. wipeWritePerm

This command is high-risk. Proceed with caution.

wipeWritePerm

Required	Description
Yes	The name of the broker.

5.1.1.6. Capacity assessment

5.1.1.6.1. Name Server connections

Metrics

- **Current resource usage**:

  Log on to the mq-namesrv container and run the ss -s command.

- **Benchmark and resource usage threshold**:

  If more than half of the connections are in use, contact O&M engineers.

5.1.1.6.2. Broker connections

Metrics

- **Current resource usage**:

  Log on to the mq-broker container and run the ss -s command.

- **Benchmark and resource usage threshold**:

  If more than half of the connections are in use, contact O&M engineers.

5.1.1.6.3. TPS

Metrics

- **Current resource usage**:

  Log on to the mq-namesrv or mq-broker container, run the sh /home/admin/rmq/bin/mqadmin clusterlist command, and then summarize the input and output TPS.

- **Benchmark and resource usage threshold**:

  For two primary brokers and two secondary brokers, contact O&M engineers to obtain the threshold for the total sending and receiving TPS.

  Brokers can be linearly resized. For example, if two primary brokers and two secondary brokers are used, you must consider a resizing plan when the total sending and receiving TPS exceeds half of the threshold.

5.1.1.6.4. Disk usage

Metrics

- **Current resource usage**:

  Log on to the mq-broker container and run the df -h command to check the disk usage in /home/admin/store or /metadata.

- **Benchmark and resource usage threshold**:

  If the disk usage exceeds 70%, you must consider scaling out or scaling up the disk.

5.1.1.7. Log reference

You can check logs to view the status of each Message Queue for Apache RocketMQ component or locate faults during O&M. This topic lists the paths of all the log files of Message Queue for Apache RocketMQ components and provides relevant descriptions.

Component logs

Component	File	Description
	/home/admin/logs/rocketmqlogs/broker.log	Keyword: Register broker to name server Check whether a broker is registered: The information "Register broker to name server {ip} OK" indicates that the broker has been registered on a Name Server. That is, the broker can be found in the cluster list.
	/home/admin/logs/rocketmqlogs/remoting.log	Find the IP address of the client to determine when the network is disconnected. Many Message Queue for Apache RocketMQ projects involve remote calls. This log file is rarely used.
	/home/admin/logs/rocketmqlogs/failover_controller.log	Provide the information about primary/secondary switchovers, especially the switchover time of the state machine.
	/home/admin/logs/rocketmqlogs/filter.log	Provide the filtered control results. A tag-like expression is used in Message Queue for Apache RocketMQ to configure filter conditions.
	/home/admin/logs/rocketmqlogs/lock.log	Keyword: tryLockBatch Provide the binding relationship between consumers and queues. One queue can be allocated only to one consumer instance. The unlockBatch log is generated when a consumer instance is discontinued.

Component	File	Description
	/home/admin/logs/rocketmqlogs/time.log	Provide the information about scheduled messages. Scheduled messages are first stored in the database on the server. Then, the database polls the messages to detect the messages whose sending time is approaching, and sends the messages by using the broker. Keywords: Scan add infly If a scheduled message is not delivered or not punctual, you can search for "Scan add infly: msgID" and view the point in time at which the message is pulled from the database and put into the broker for delivery. You can view the point in time at which the message is pulled in the time.log file based on the message ID. Then, you can refer to the trace to check when and where the process is blocked and slowed down.
	/home/admin/logs/rocketmqlogs/transaction.log	Provide the information about transactional messages and in-doubt transactional messages. These messages are first stored in the database on the server. Some messages are transactions in the doubted state. In-doubt transactions that are not committed will be scanned by the database and the producer will be regularly asked for confirmation on these transactions.

Component	File	Description
Broker	/home/admin/logs/rocketmqlogs/stats.log	Record the transactions per second (TPS) of the current broker to facilitate query for historical TPS. In addition, the resource report is generated by aggregating these logs. Keywords: BROKER_PUT_NUMS: the TPS of writing messages per minute. BROKER_GET_NUMS: the TPS of reading messages per minute. This log file can be used for troubleshooting of a message reading or writing failure at a specific point in time. It can be used with Tlog to generate reports.
	/home/admin/logs/rocketmqlogs/store.log	Provide macro statistics on TPS. Keywords: PAGECACHERT totalput: the TPS of writing messages. putmessagedistributetime: the number of messages that can be stored on the disk within the required duration. The messages are classified by time consumption. Messages are considered normal if they can be stored on the disk within 10 ms. This log file can be used to troubleshoot low TPS. It is usually used to analyze latency during stress testing. Serial writing of commitLog can be detected with ease.
		Provide the information about primary/secondary switchovers. The broker plays the role of SYNC_MASTER. The state machine has the following states: 1. orignal-state: NULL_STATE

Component	File	Description
	/home/admin/logs/rocketmqlogs/dup.log	2. singal-master-state: ONLY_MASTER 3. async-sync-state: ASYNC_STATE 4. semi-sync-state: HALF_SYNC 5. sync-state: SYNC_STATE The following content describes the state transition of the state machine after the service is started: The state machine is in the original state (NULL_STATE) when the service is started. After the primary broker is started, the state machine transits to the ONLY_MASTER state. Start the secondary broker and asynchronously synchronize messages to the secondary broker. The primary broker continues to write requests. When the message synchronization offset is less than 100 messages, hold the write requests to the primary broker, and enable the secondary broker in the HALF_SYNC state to continue asynchronous replication. The state machine transits to the synchronous replication state (SYNC_STATE) when all messages are written to the secondary broker. Controller determines which broker is primary or secondary. In Apsara Stack, approximately 99.99% of brokers are deployed in primary/secondary synchronous mode. Keyword: DuplicationStatus voteDupoffset identifies an offset. If it continues to grow, the consumer offset is normal.

Component	File	Description
	/home/admin/logs/rocketmqlogs/dup-stats.log	Provide performance monitoring data such as the number of tasks, the writing latency of the secondary broker, and the writing latency of the on-premises broker.
	/home/admin/logs/rocketmqlogs/dup-study.log	Provide the asynchronous replication status of the primary and secondary brokers.
	/home/admin/logs/rocketmqlogs/broker_default.log	Provide watermark data. This log file is of little significance for troubleshooting.
	/home/admin/logs/rocketmqlogs/rocketmq.log	This log file is of little significance for troubleshooting.
	/home/admin/logs/spas/spas_sdk.log	Provide authentication, permission, validation, and authorization information. Obtain the UID of the Apsara Stack tenant account based on the AccessKey pair to determine the matched topic in the database. Check whether authorization is obtained. Permission rules: The AccessKey ID and AccessKey secret of the account used to create the topic in the console are used by the client port.
Name Server	/home/admin/logs/spas/spas_sdk.log	Provide authentication, permission, validation, and authorization information. Obtain the UID of the Apsara Stack tenant account based on the AccessKey pair to determine the matched topic in the database. Check whether authorization is obtained. Permission rules: The AccessKey ID and AccessKey secret of the account used to create the topic in the console are used by the client port.

Component	File	Description
	/home/admin/logs/rocketmqlogs/namesrv.log	N/A
Console	/home/admin/logs/spas/spas_sdk.log	Provide authentication, permission, validation, and authorization information. Obtain the UID of the Apsara Stack tenant account based on the AccessKey pair to determine the matched topic in the database. Check whether authorization is obtained. Permission rules: The AccessKey ID and AccessKey secret of the account used to create the topic in the console are used by the client port.
	/home/admin/rmq/ons-api/register.log	Allow users to query the information about the following operations: Add, delete, modify, and query topics. Add, delete, modify, and query group IDs.
	/home/admin/rmq/ons-api/rest.log	Record exceptions of console operations.
Address Server	/home/admin/cai/logs/error.log	Provide execution error logs of Address Servers.
	/home/admin/cai/logs/access.log	Provide access logs of Address Servers.
Controller	/home/admin/logs/rocketmqlogs/failover_controller.log	Provide the primary/secondary switchover logs of brokers. After the logs are generated, you need to check and analyze whether the operations are normal.
Housekeeping	/home/admin/logs/aliware-mq-hkagent/warning.log	Record the inspection results of the inspection tool provided by Message Queue for Apache RocketMQ. After the logs are generated, you need to check and analyze whether the operations are normal.

6.Operations of big data products

6.1. Apsara Big Data Manager (ABM) platform

6.1.1. User Guide

6.1.1.1. 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
- RealtimeCompute
- 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. ABM also allows you to update big data services, customize alert configurations, and view the O&M history.

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

6.1.1.2. Common operations

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

Search for a project quickly

You can quickly search for a project based on the project name.

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

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

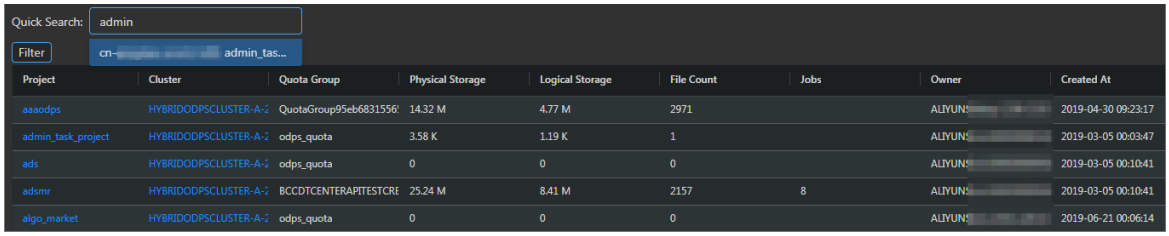

Example:

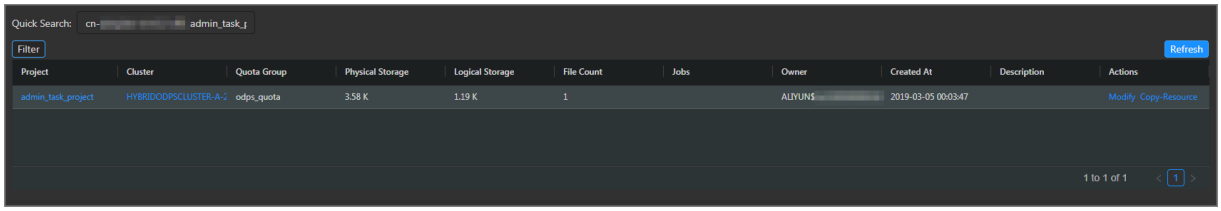

Filter projects

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

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

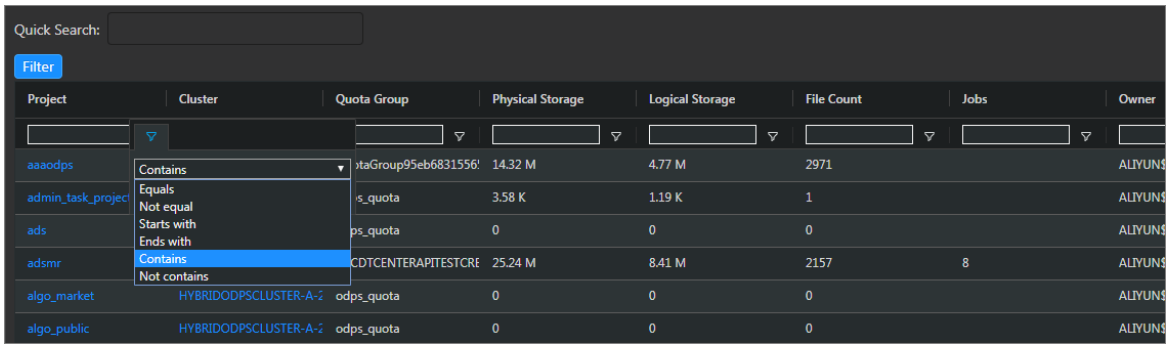

Optional filtering methods include:

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

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

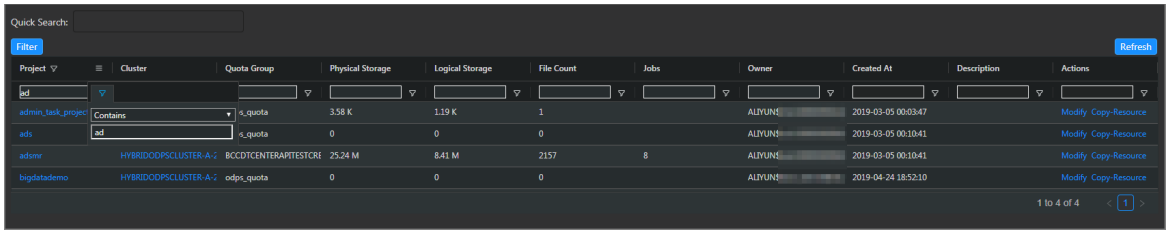

5. If the filtering result is not accurate, you can continue performing this operation on other columns.

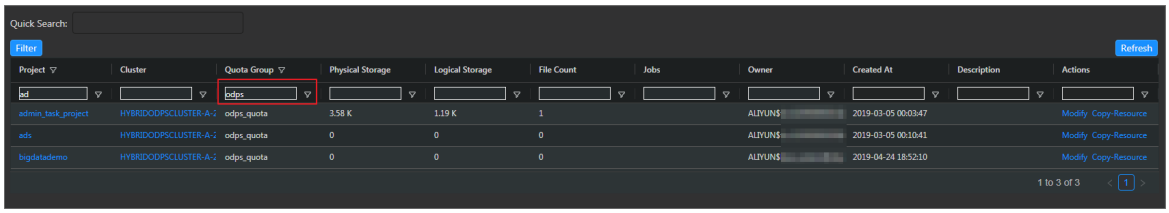

After you set the filter conditions for the projects, the **Filter** button is highlighted. If you need to cancel filtering, click the highlighted **Filter** button.

Search for items

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

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

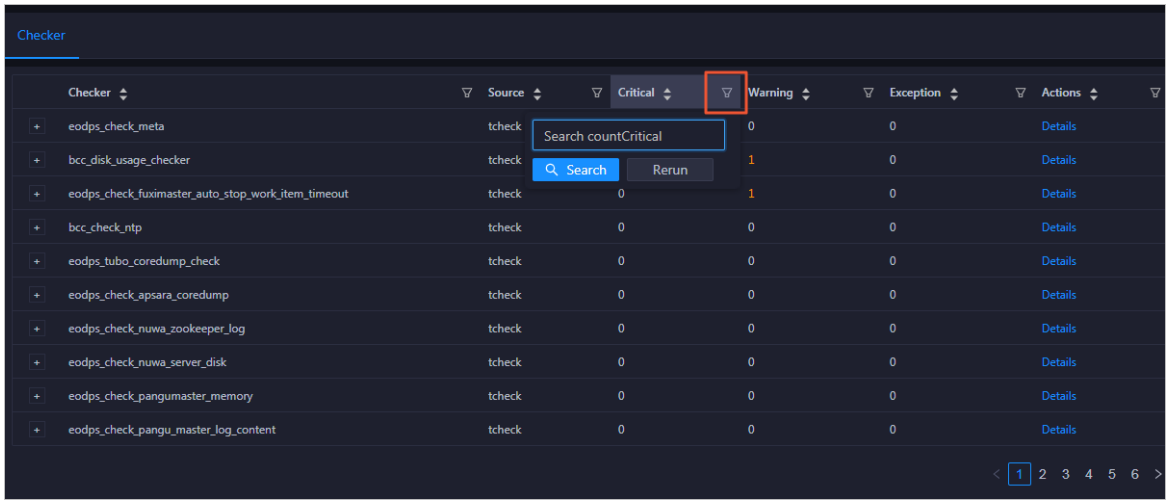

3. Click **Search**. The checkers that meet the requirements appear.
4. If the search result is not accurate, you can continue performing this operation on other columns.

Customize a column

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

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

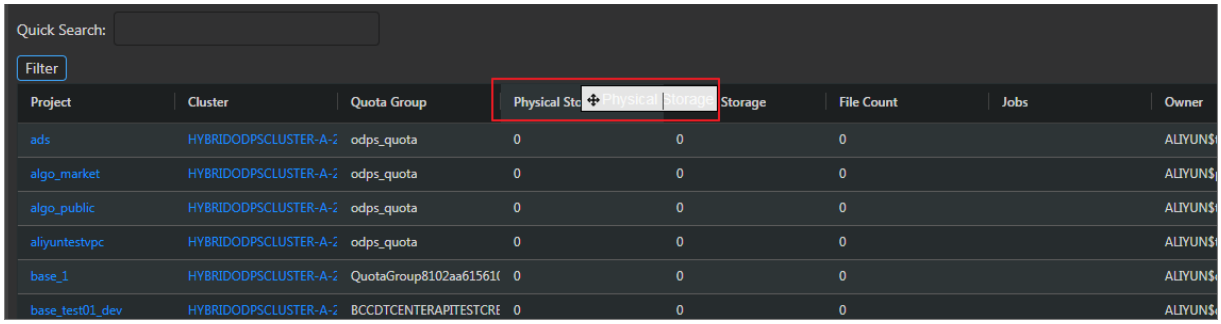

You can click in a column heading to customize the column.

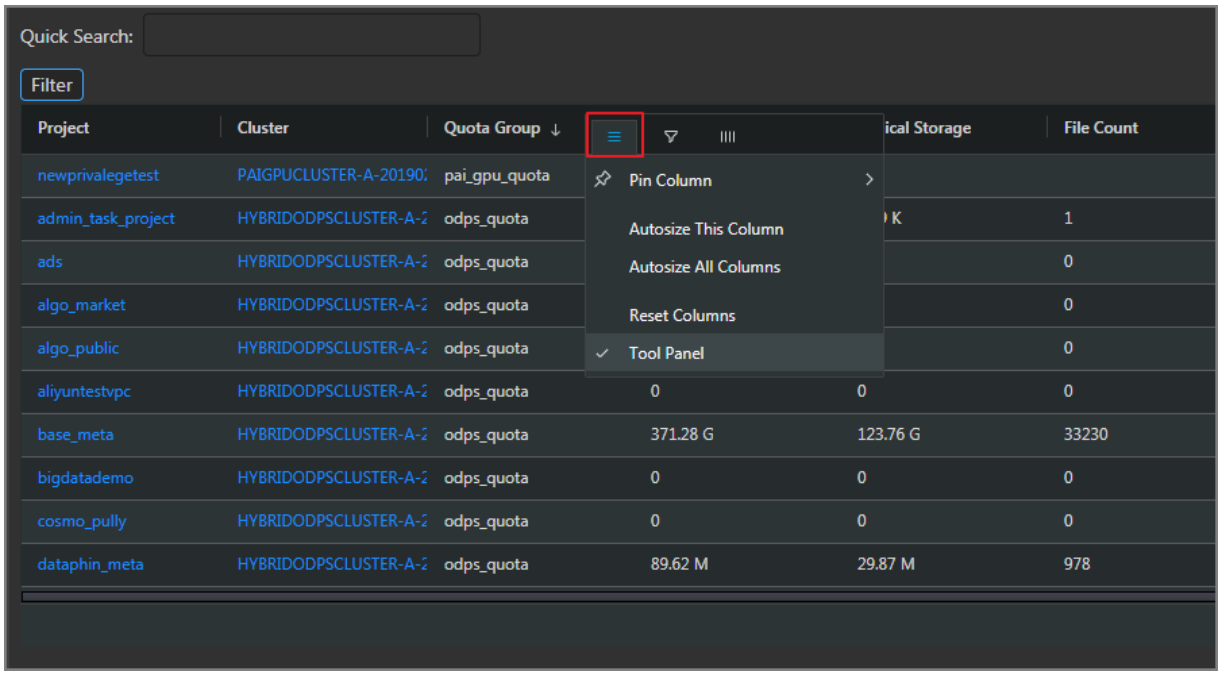

- **Pin Column**: allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.
- **Autosize This Column**: allows you to adjust the width of a column automatically.
- **Autosize All Columns**: allows you to adjust the width of all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

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

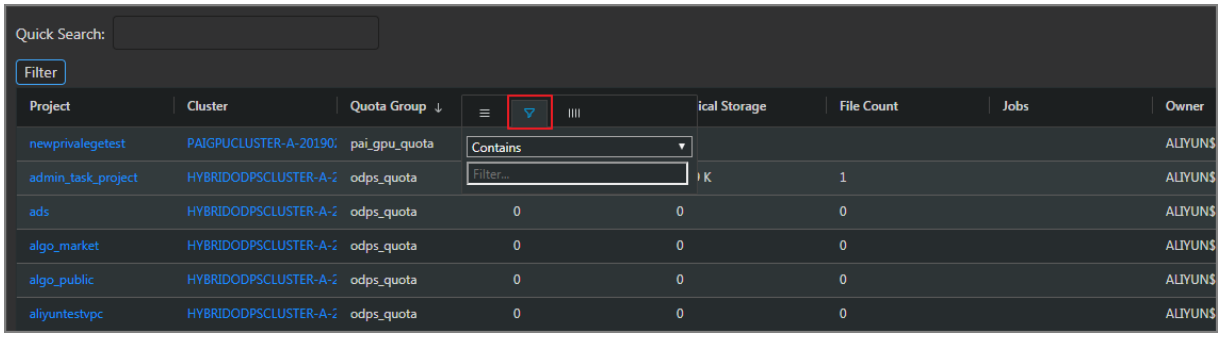

Click in a column heading and select the columns to be displayed.

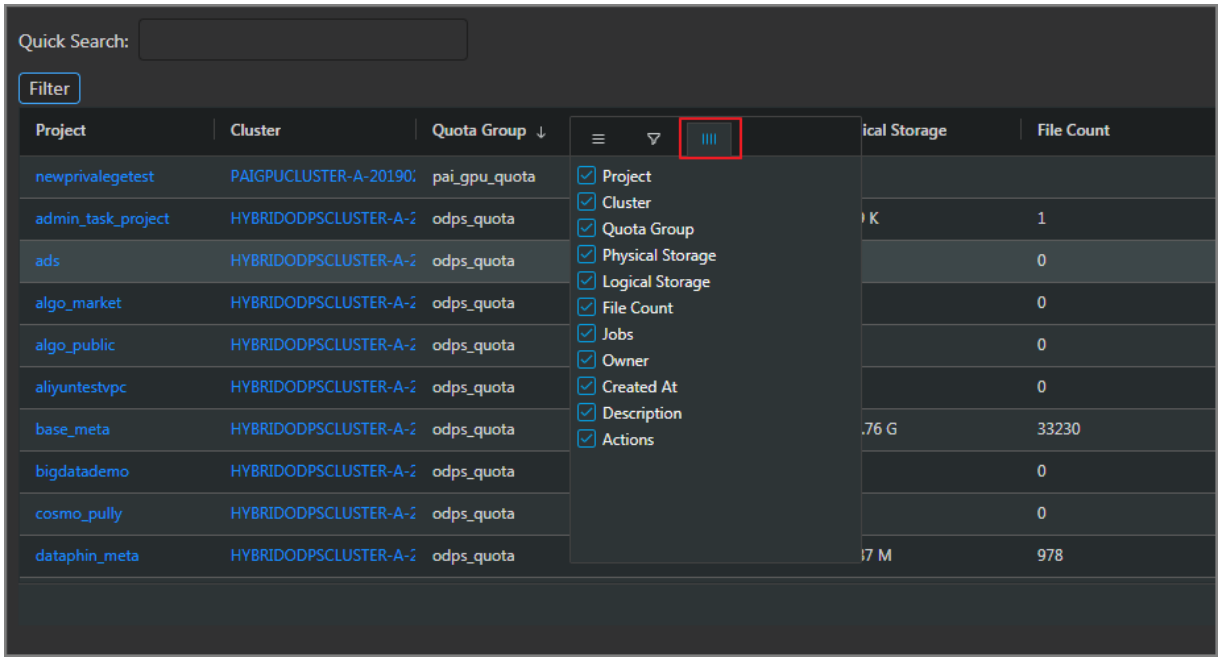

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

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

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

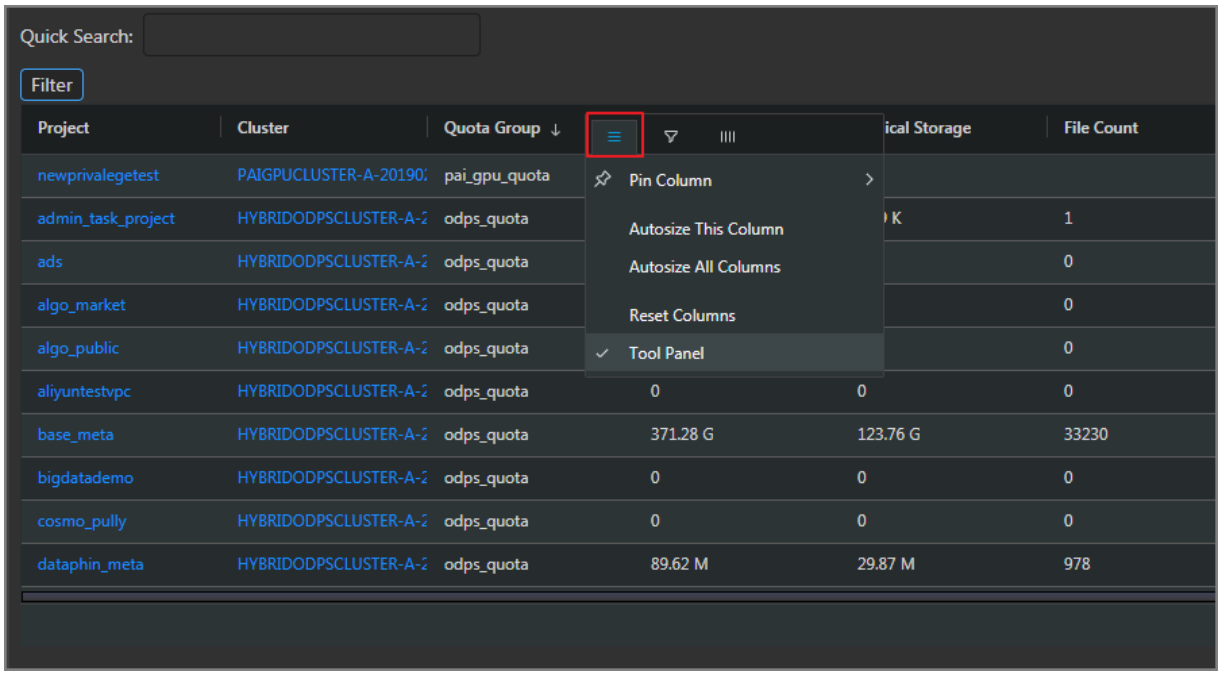

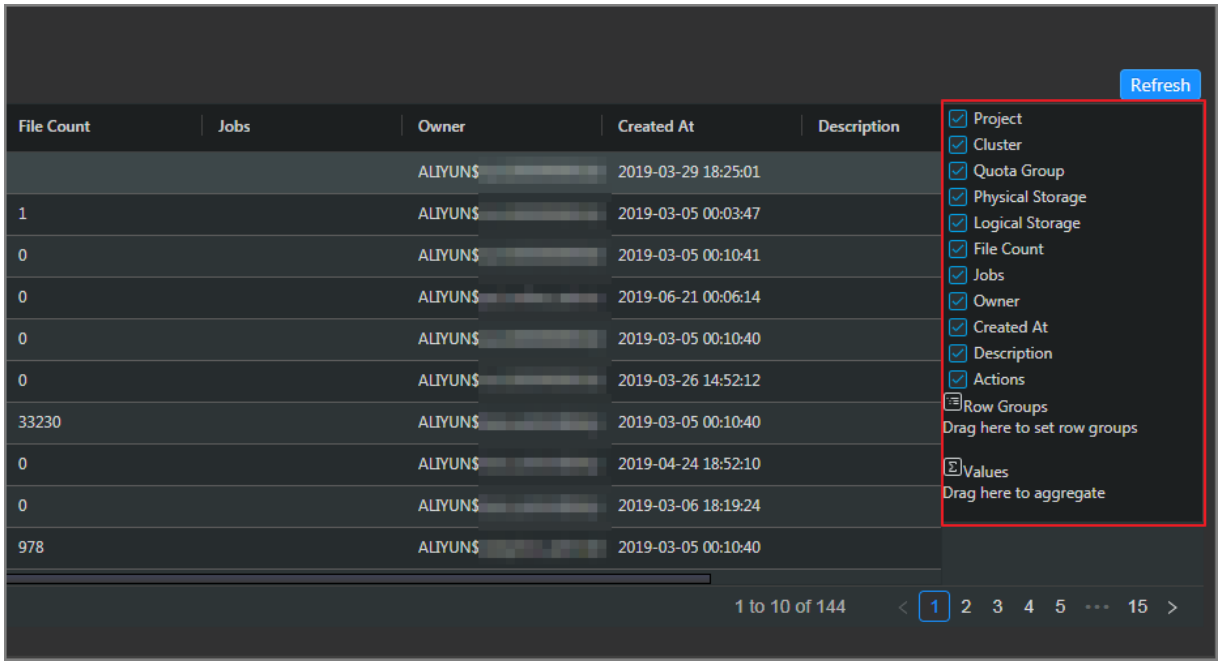

Sort projects based on a column

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

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

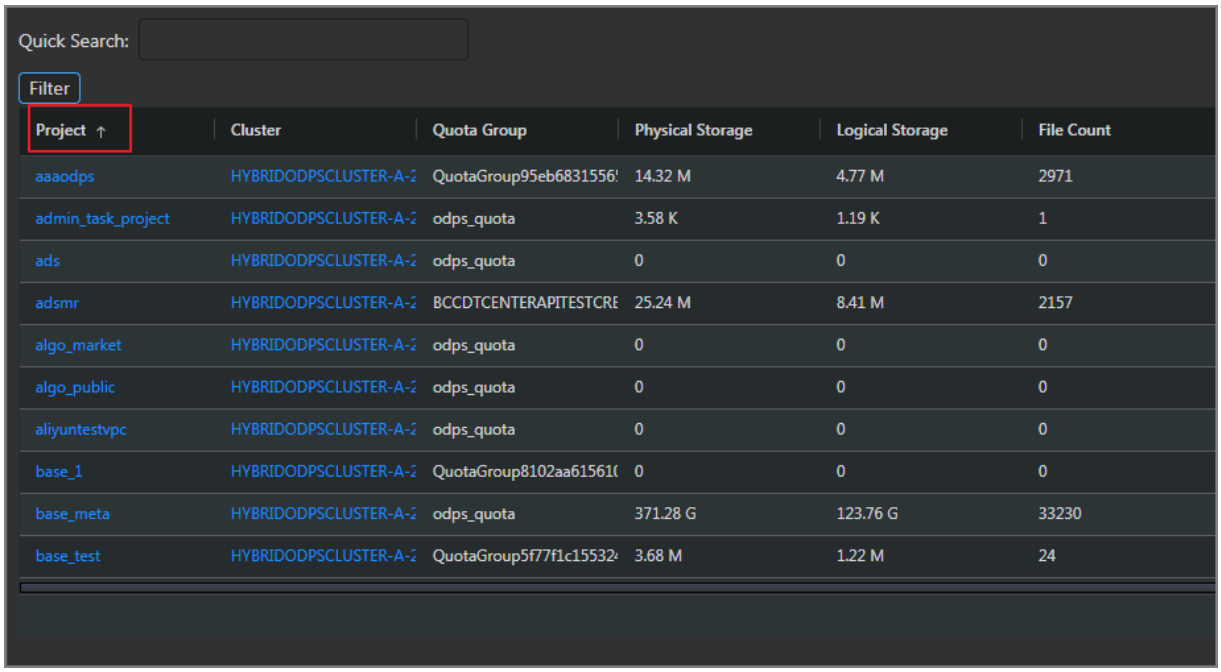

Sort items based on a column

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

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

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

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

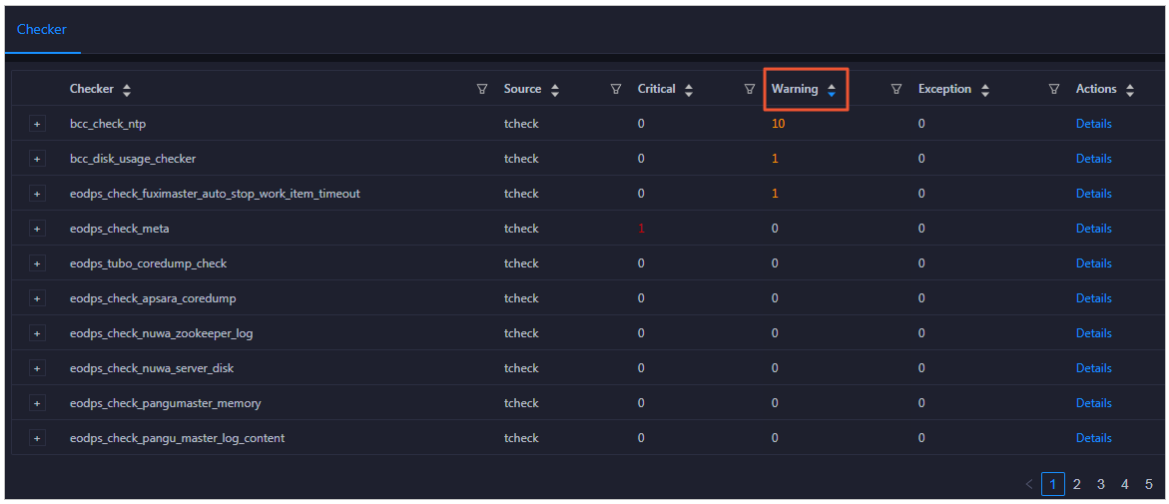

The highlighted up arrow indicates that the checkers are sorted in ascending order. The highlighted down arrow indicates that the checkers are sorted in descending order.

Trend chart 1

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

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

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

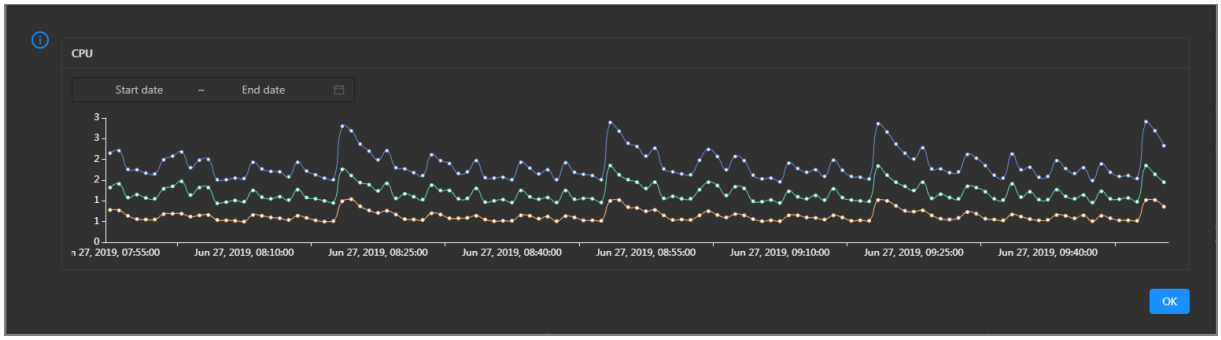

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

6.1.1.3. Quick start

6.1.1.3.1. Log on to the ABM console

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

Prerequisites

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

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

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

Procedure

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

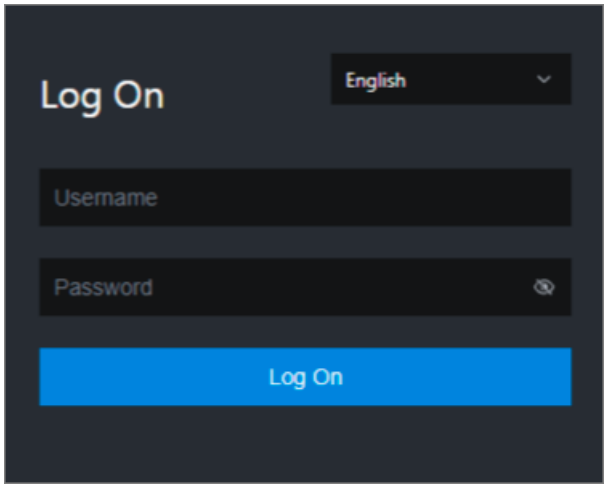

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.
4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Big Data Services** section, choose **General-Purpose O&M > Apsara Big Data Manager**.

6.1.1.3.2. 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.

6.1.1.3.3. View the trace dashboards

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

Prerequisites

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

Background information

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

Procedure

1. Log on to the ABM console.

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

2. View and clear service alerts.

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

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

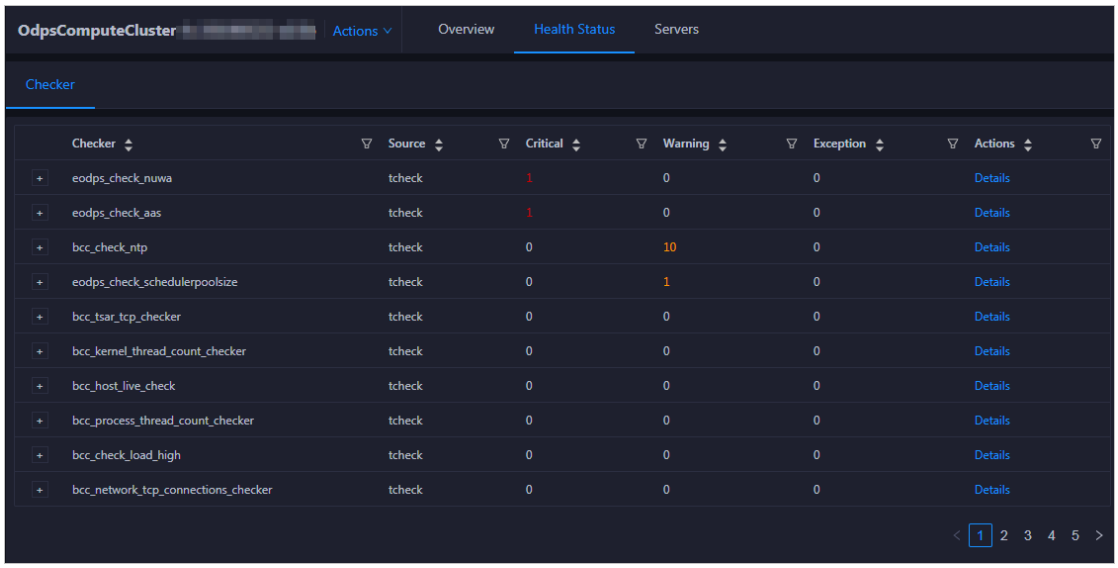

   In the **Health** On the page that appears, you can view all check items of the product.

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

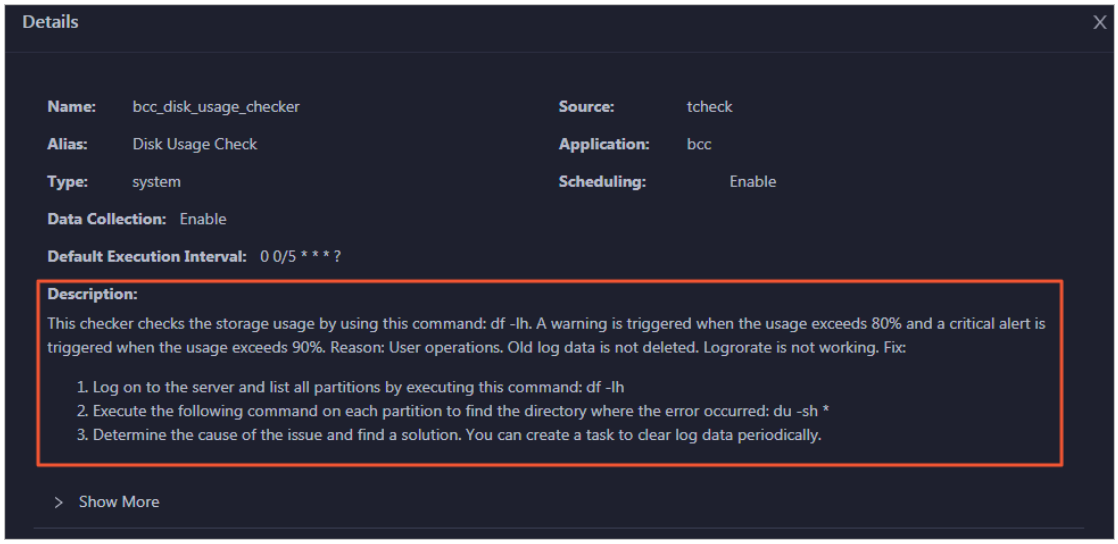

iii. Log on to the hosts on which the alerts are detected to handle the alerts.

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

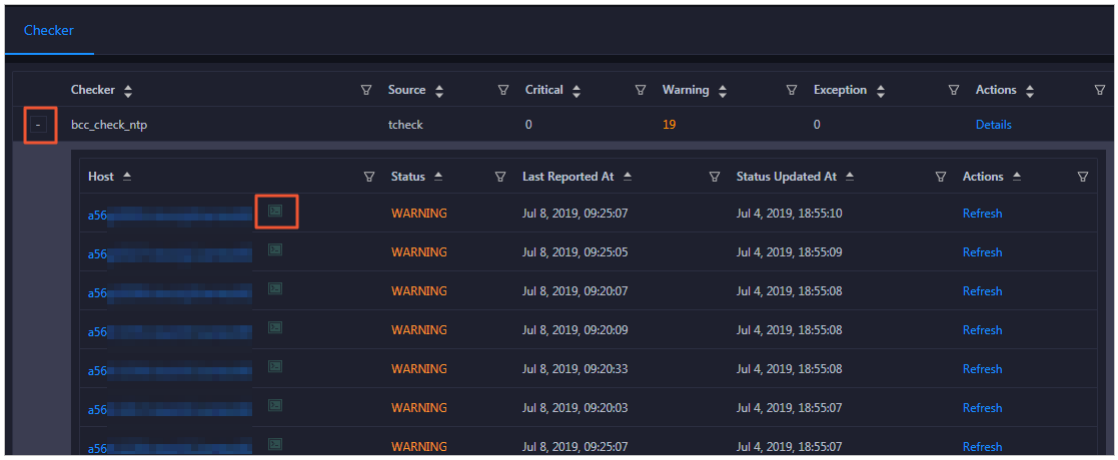

iv. In the newly opened **TerminalService** On the page that appears, select a host on the left to log on.

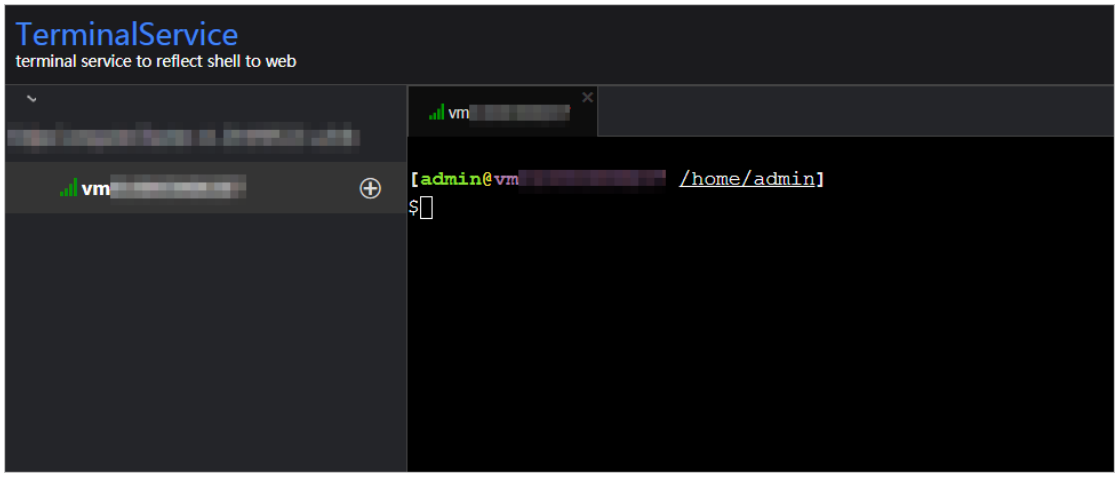

3. In the **Dashboard** On the page that appears, click the **MaxCompute** , view **MaxCompute** .

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

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

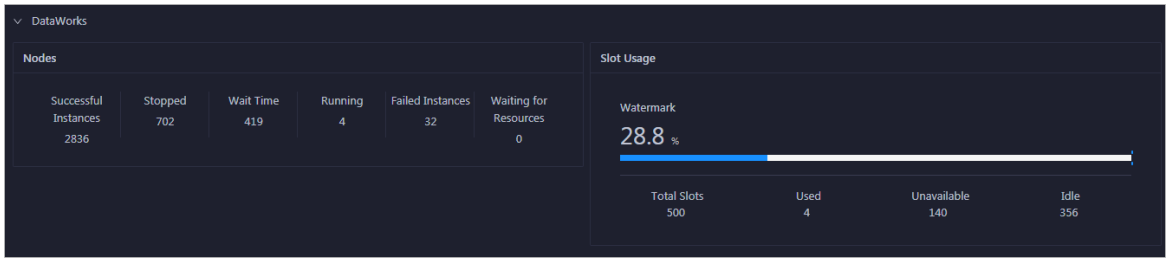

DataWorks The section displays the node scheduling overview, slot resource overview, and the cumulative trend chart of task completion in the DataWorks cluster.

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

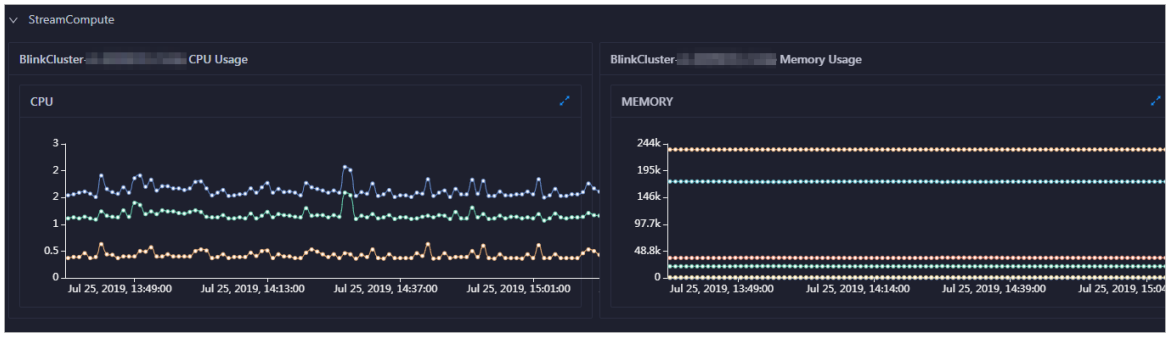

RealtimeCompute The section displays the trend charts of TPS and FAILOVER for RealtimeCompute cluster jobs, and the trend charts of CPU and memory usage.

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

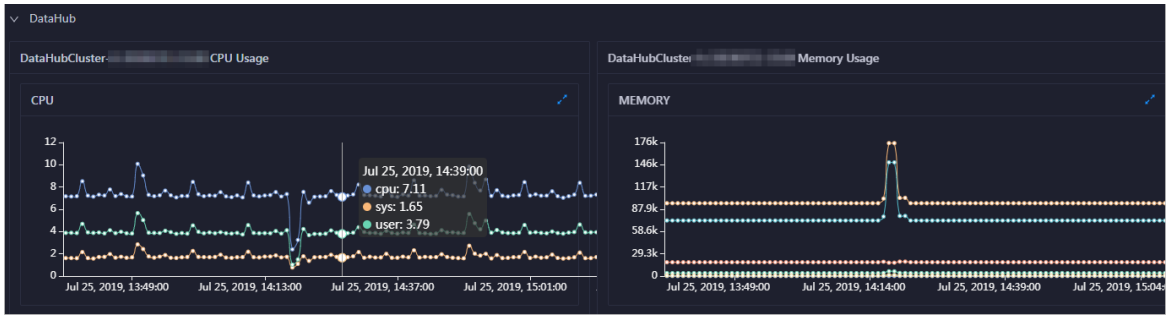

DataHub section displays the trend charts of read /write latency, number of read /write records, read /write QPS, read /write byte traffic, CPU level, and memory level of the DataHub cluster.

6.1.1.3.4. 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.

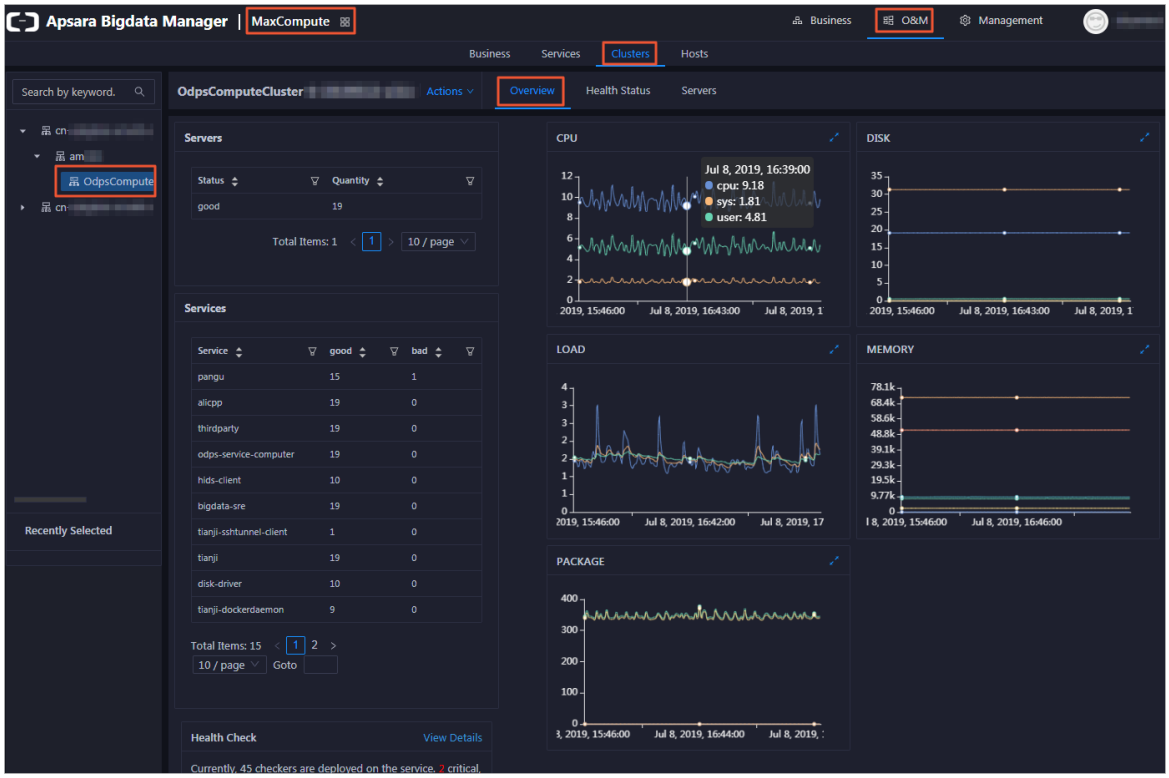

On the **Overview** page, you can view the host status, service status, health check result, and health check history of the selected cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster.

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.

6.1.1.3.5. 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.

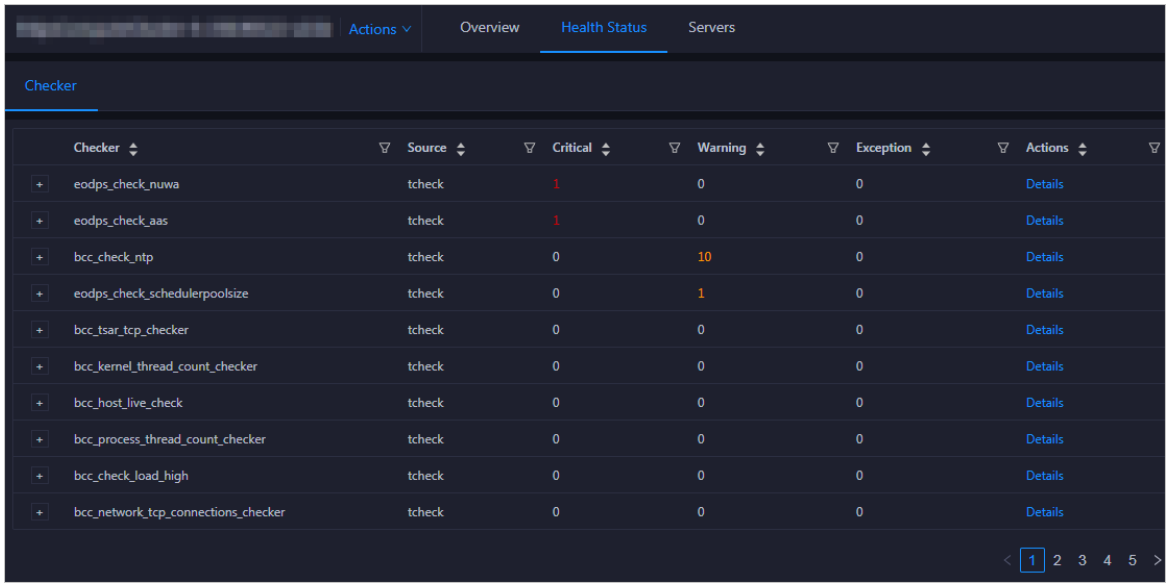

4. On the **Health Status** page, click **+** to expand a checker with alerts. You can view all hosts where the checker is run.

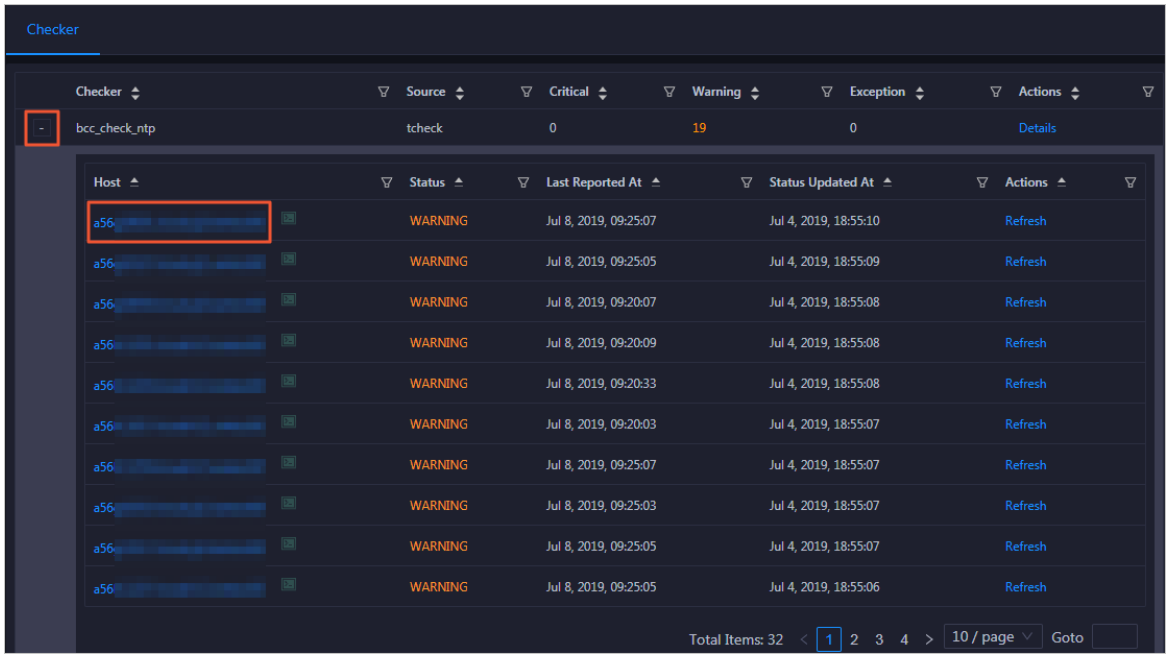

5. Click a hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

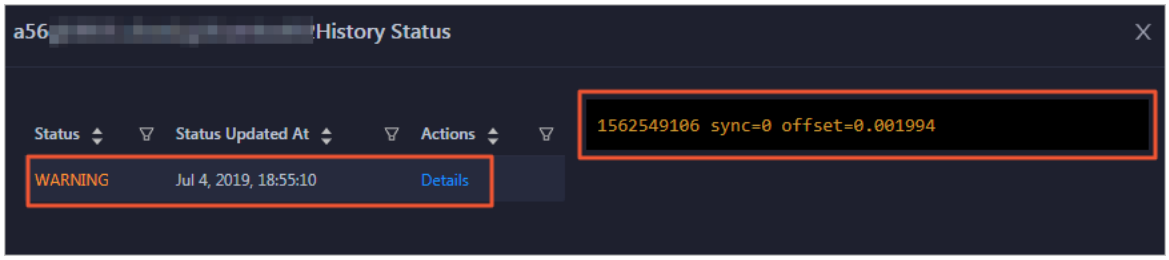

6. On the **Health Status** page, click **Details** in the Actions column of the checker to view the schemes to clear the alerts.

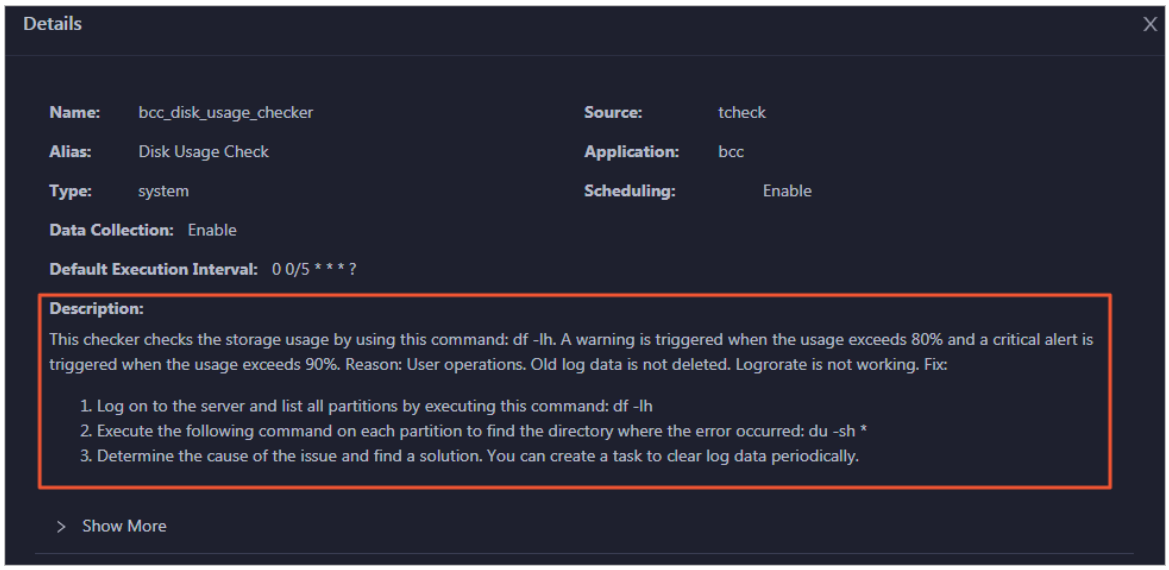

7. Clear the alerts according to the schemes.

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.

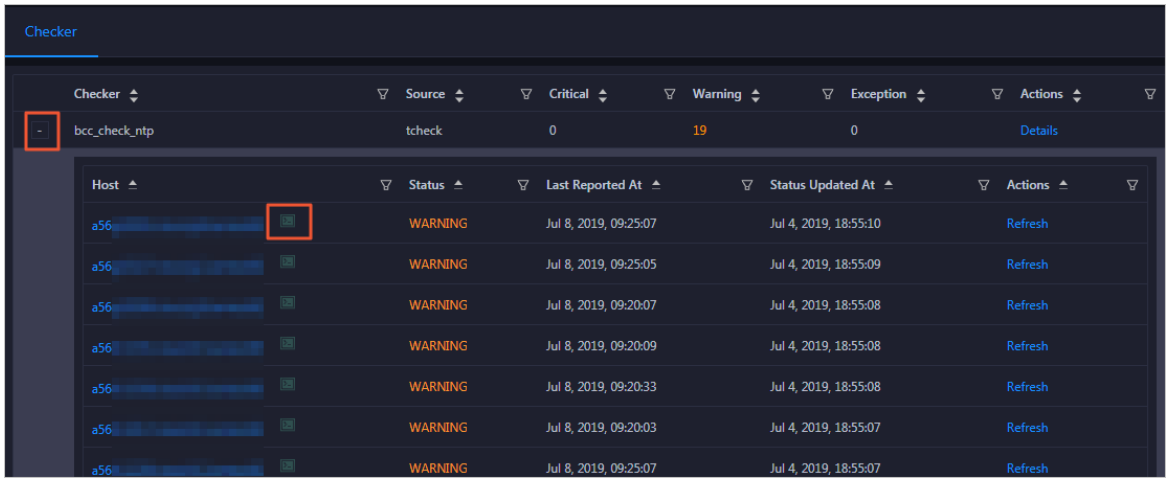

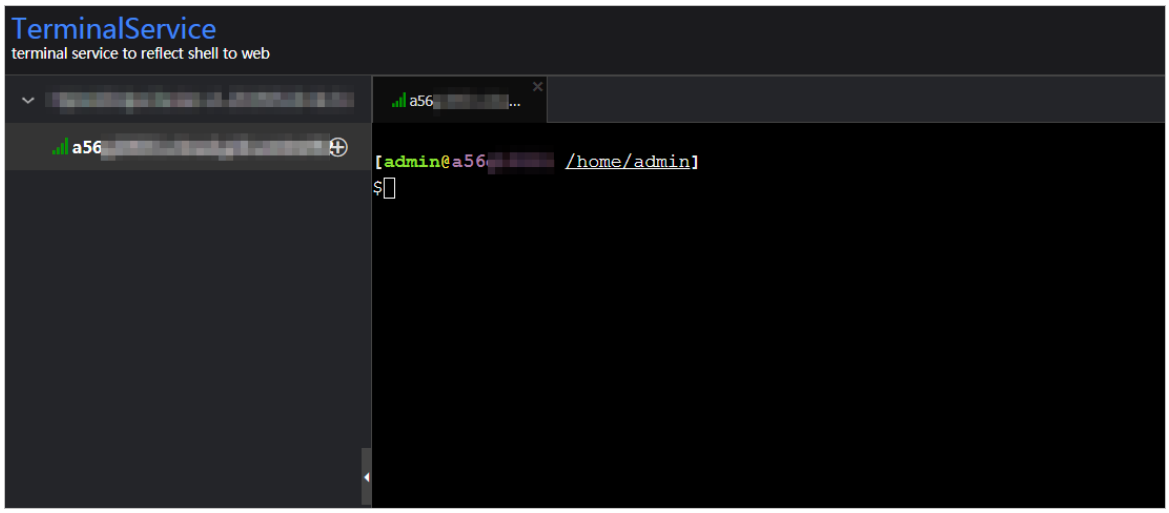

8. After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

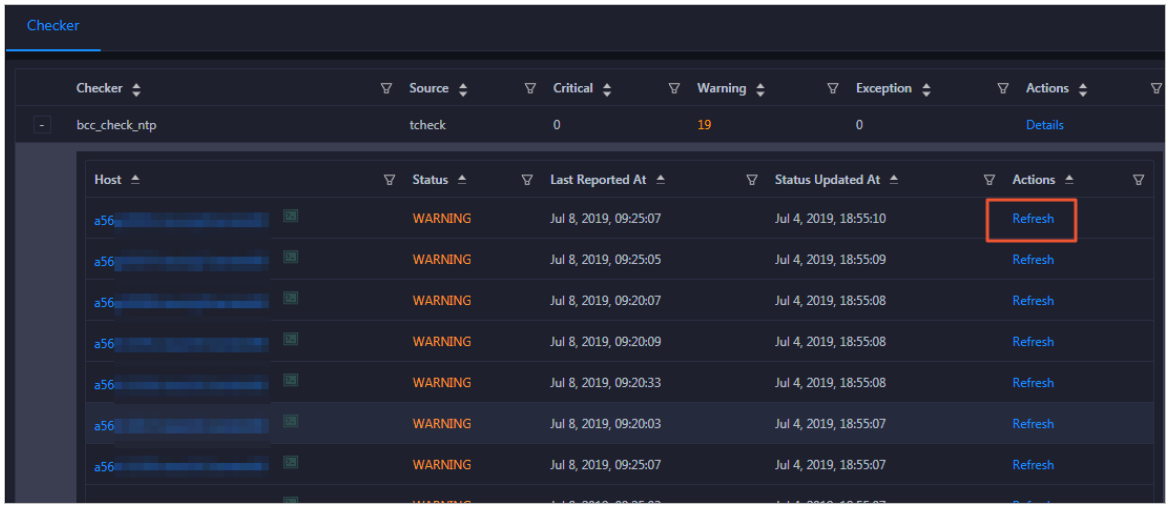

6.1.1.4. ABM

6.1.1.4.1. ABM dashboard

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

Go to the Dashboard tab

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

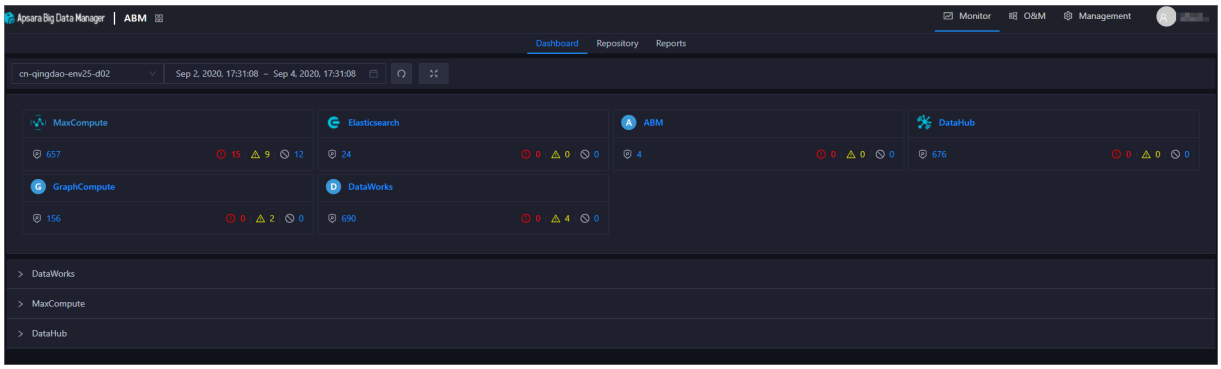

In the upper-left corner of the **Dashboard** tab, you can select a region from the drop-down list to view the cluster status of each big data service in the region.

View and handle the alerts of various services

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

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

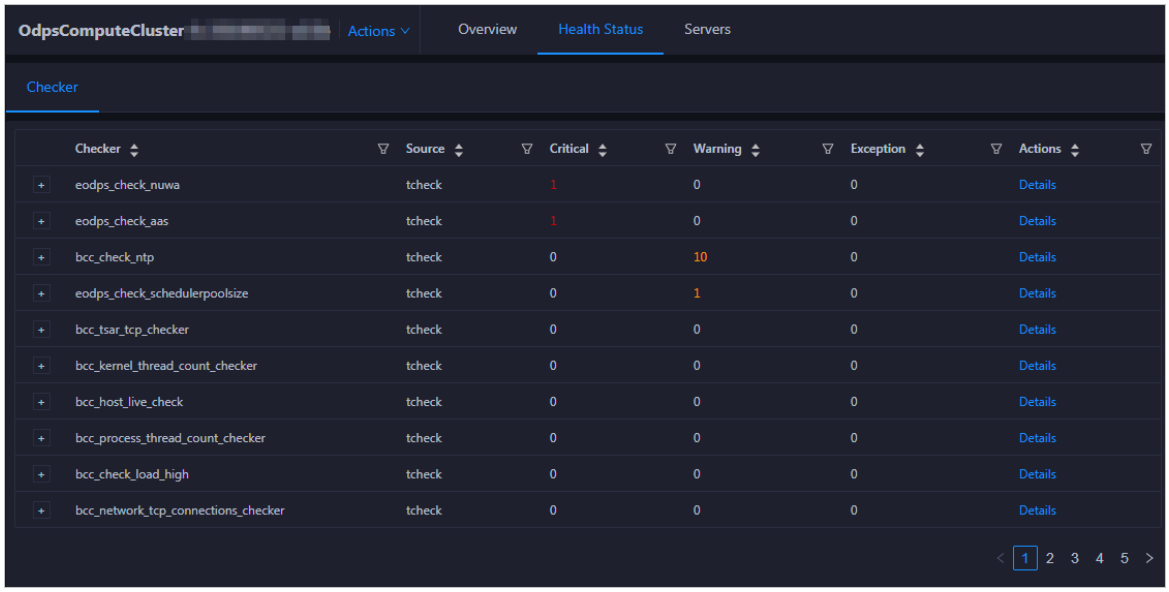

On the **Health Status** page, you can view all the check items of the service.

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

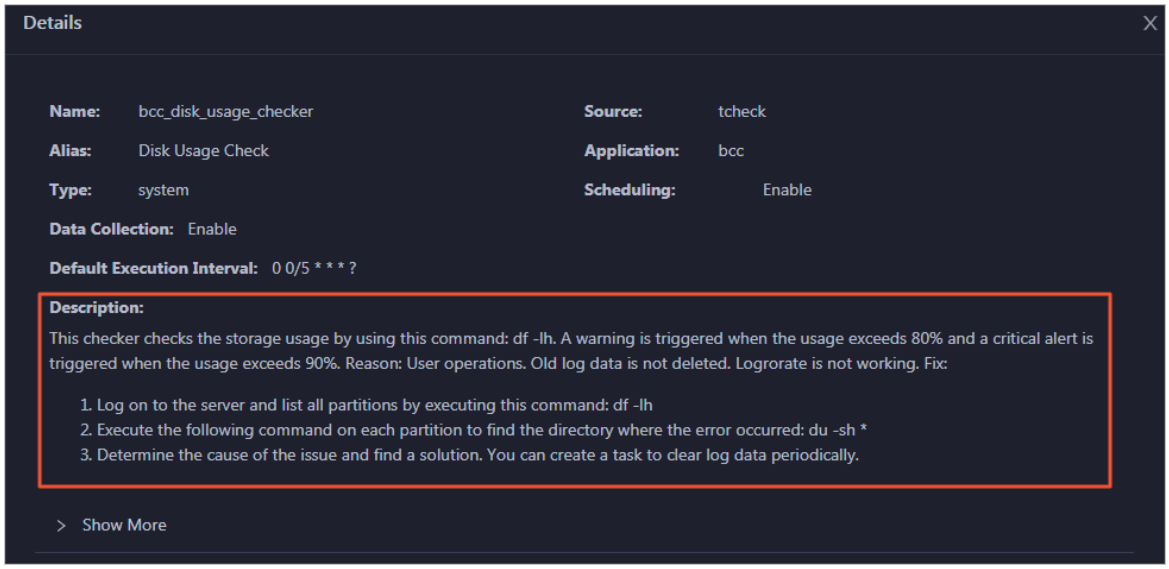

3. Log on to the hosts on which the alerts are detected to handle the alerts.

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

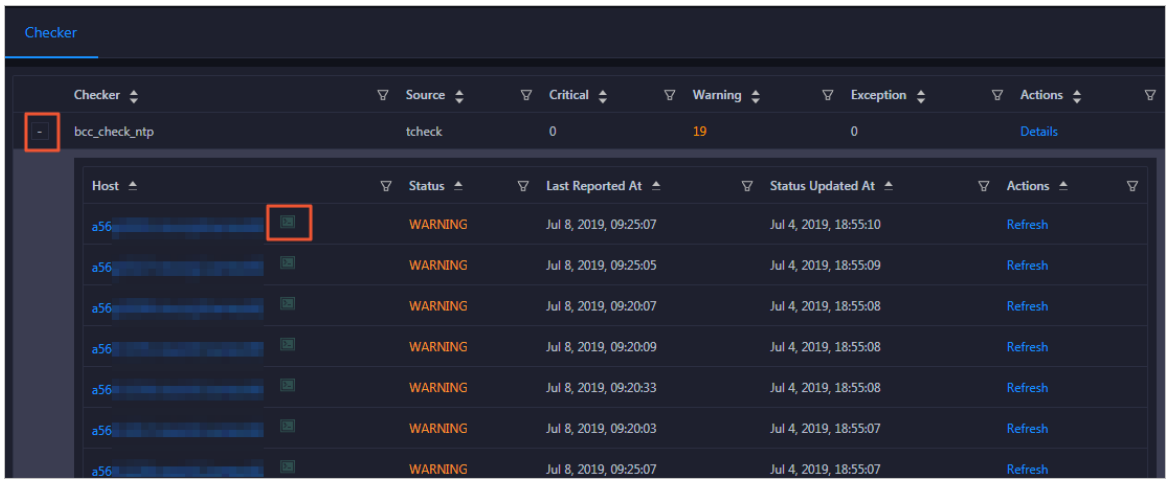

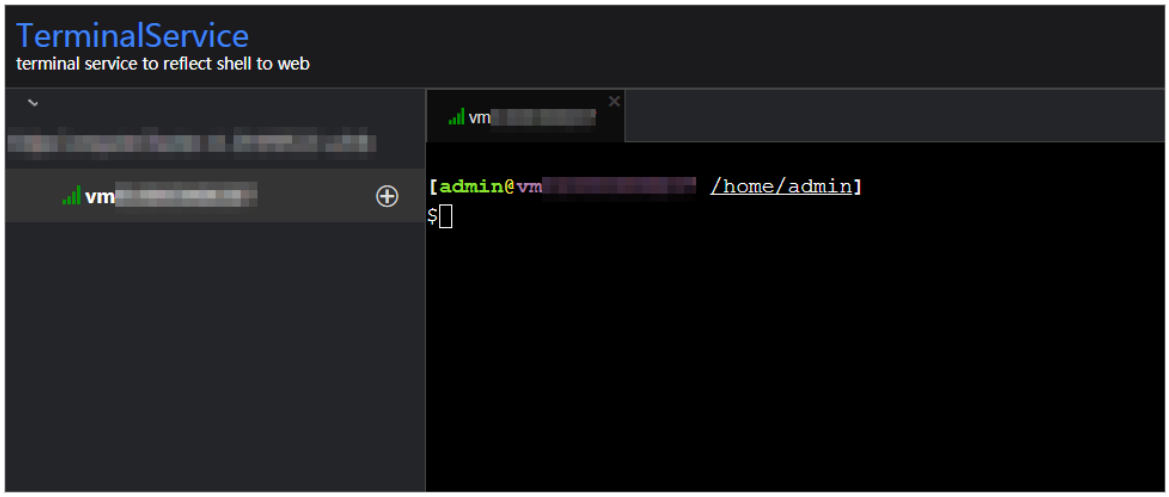

View key indicators of MaxCompute

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

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

View key indicators of DataWorks

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

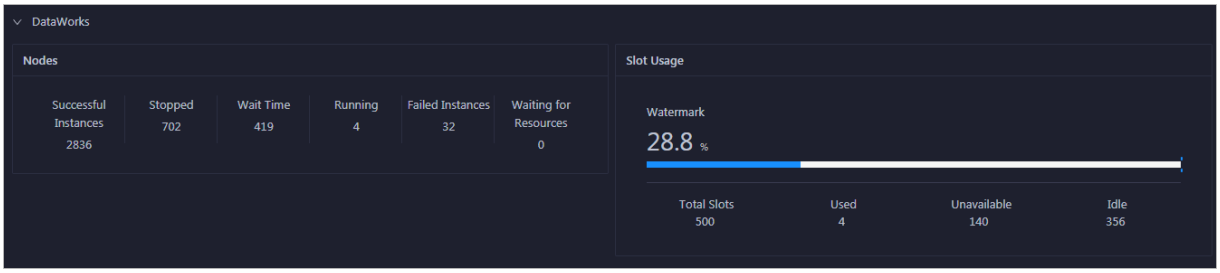

In the **DataWorks** section, you can view the node scheduling and slot usage of a DataWorks cluster. You can also view the trend chart of the total number of daily finished tasks.

View key indicators of Realtime Compute for Apache Flink

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

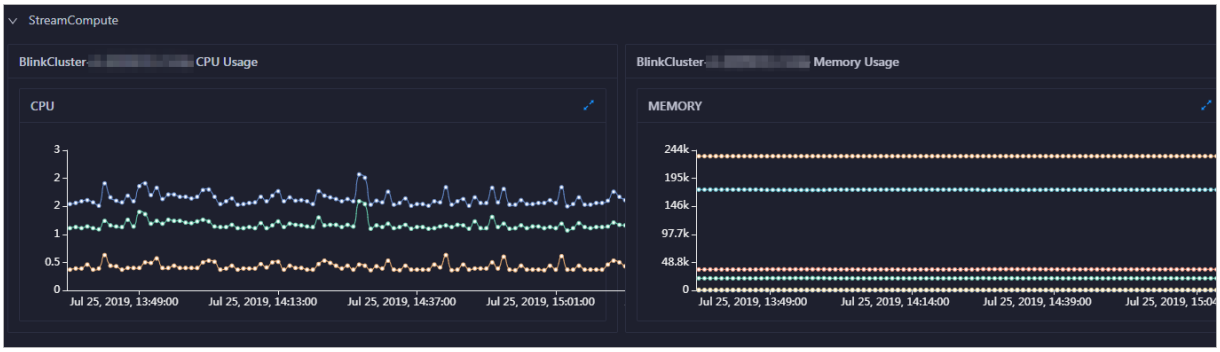

In the **Realtime Compute** section, you can view the trend charts of the transactions per second (TPS), failover rate, CPU utilization, and memory usage for a Realtime Compute for Apache Flink cluster.

View key indicators of DataHub

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

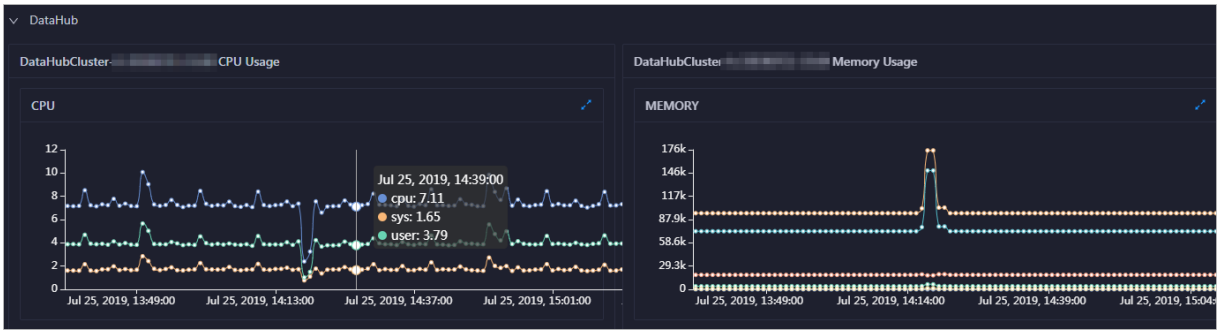

In the **DataHub** section, you can view the trend charts of the read/write latency, read/write records, read/write queries per second (QPS), and read/write throughput. You can also view the trend charts of CPU utilization and memory usage of a DataHub cluster.

Enable and disable auto-refresh

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

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

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

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

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

Display the dashboard in full-screen

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

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

6.1.1.4.2. 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.

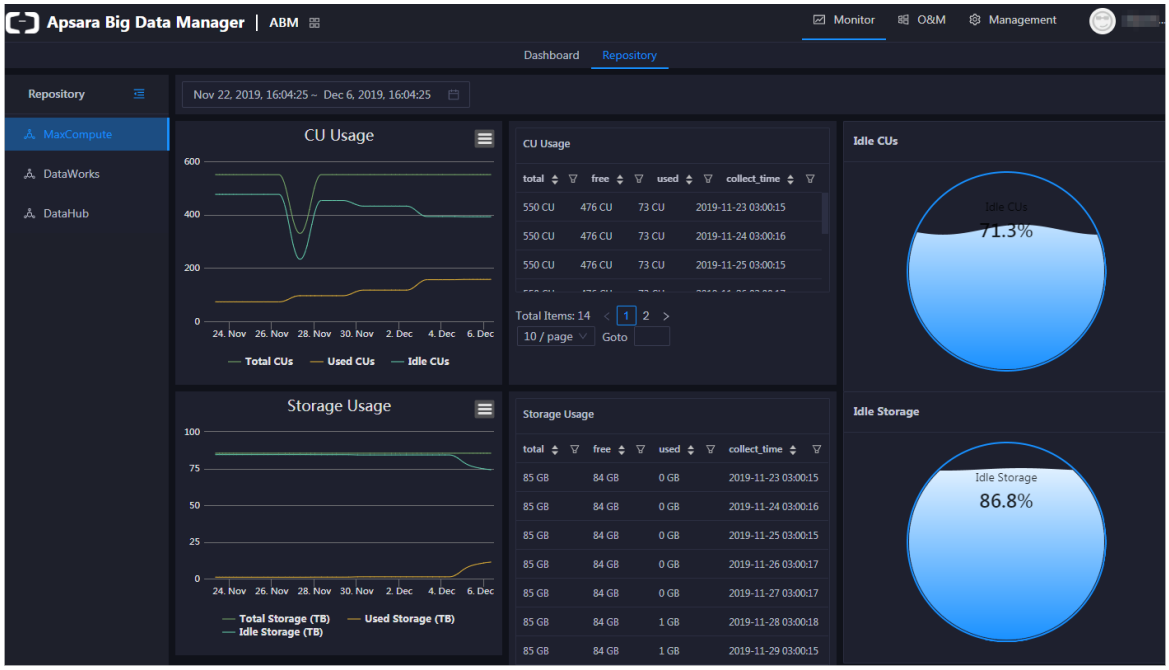

View the resource usage in MaxCompute

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.

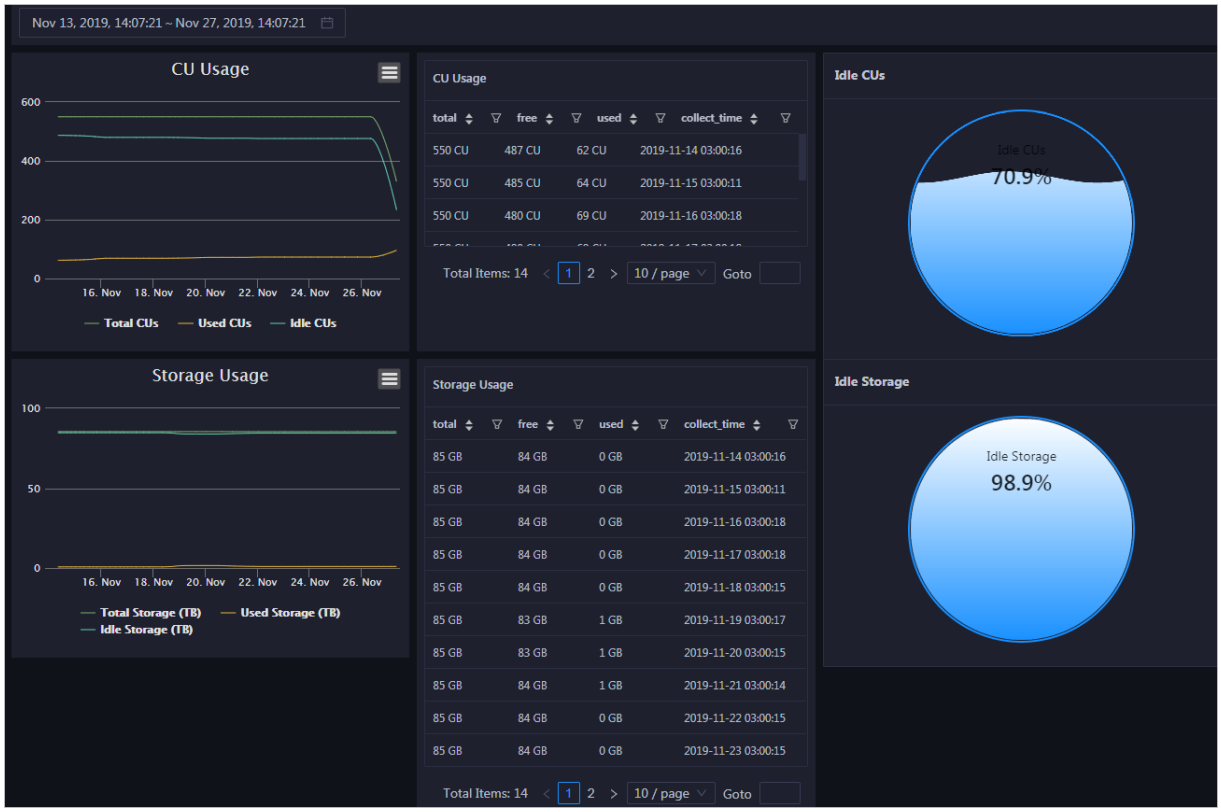

For MaxCompute, the Repository page displays the trend charts of CU and storage usage, records of CU and storage usage, and proportions of idle CUs and storage.

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.

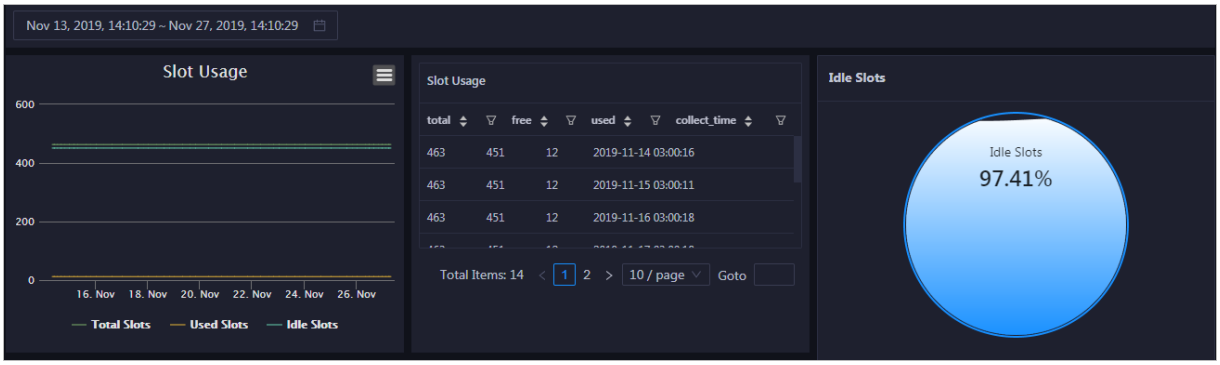

For DataWorks, the Repository page displays the trend chart of slot usage, records of slot usage, and proportion of idle slots.

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.

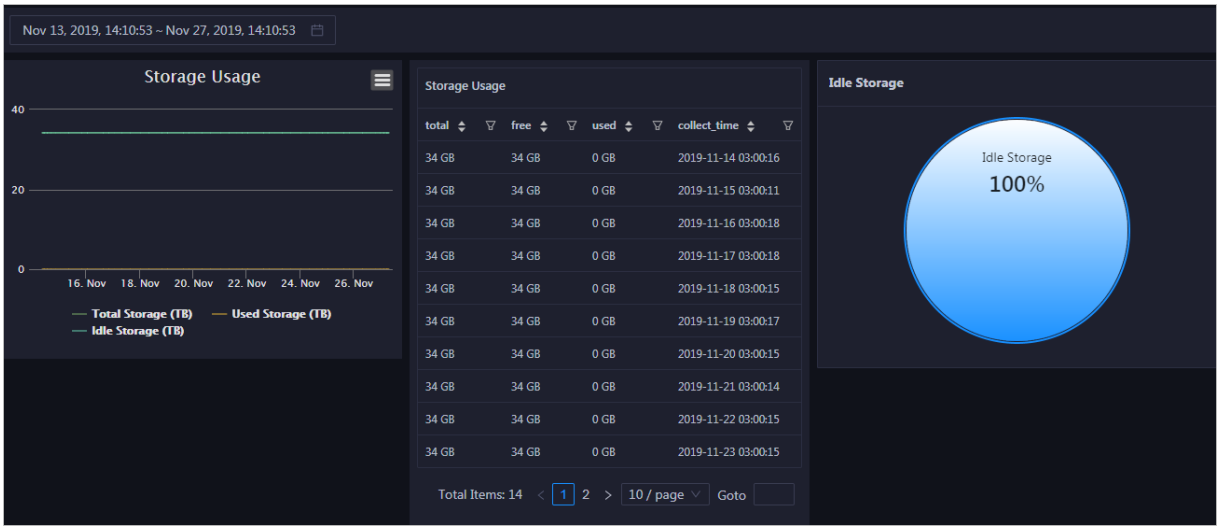

For DataHub, the Repository page displays the trend chart of storage usage, records of storage usage, and proportion of idle storage.

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.

6.1.1.4.3. ABM O&M overview

This topic describes the O&M modules of 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 exist). You can also log on to a host and perform manual checks on the host.
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 exist). You can also log on to a host and perform manual checks on the host.

Go to the Apsara Big Data Manager O&M page

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and 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**.

6.1.1.4.4. Service O&M

6.1.1.4.4.1. Service overview

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

Entry

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

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

CPU

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

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

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

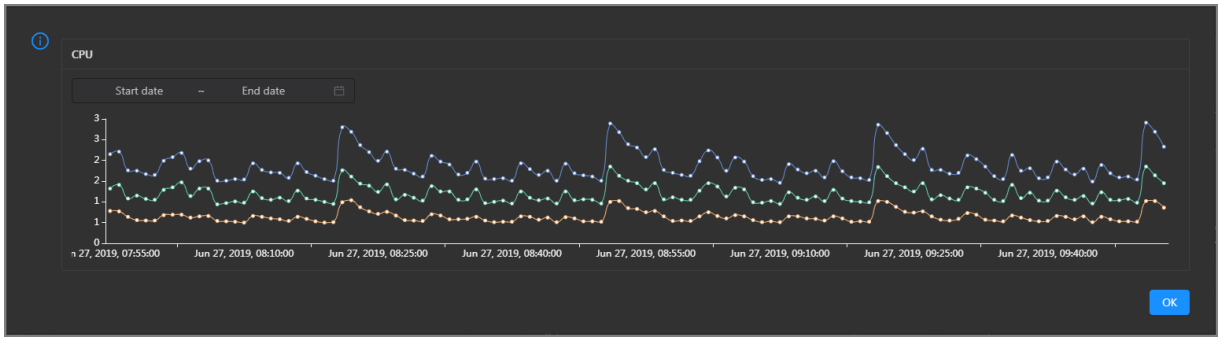

DISK

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

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

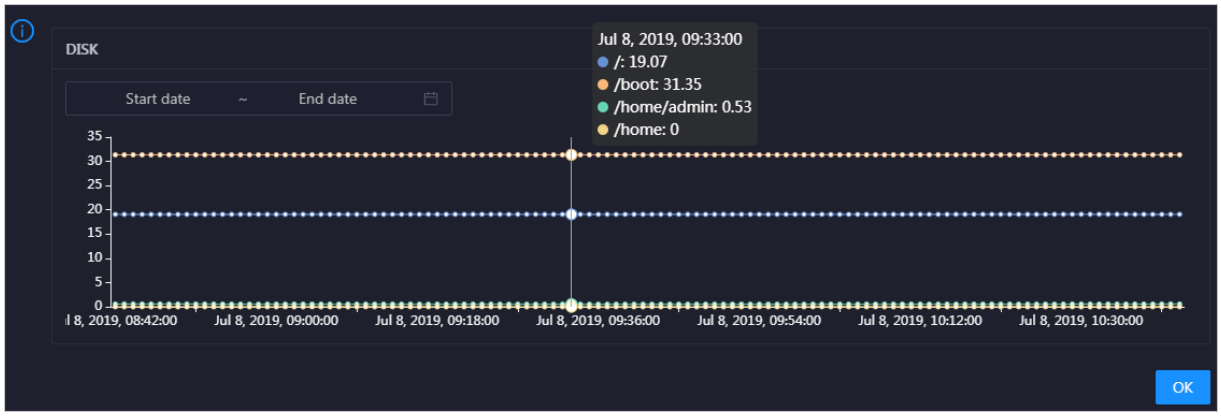

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

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

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

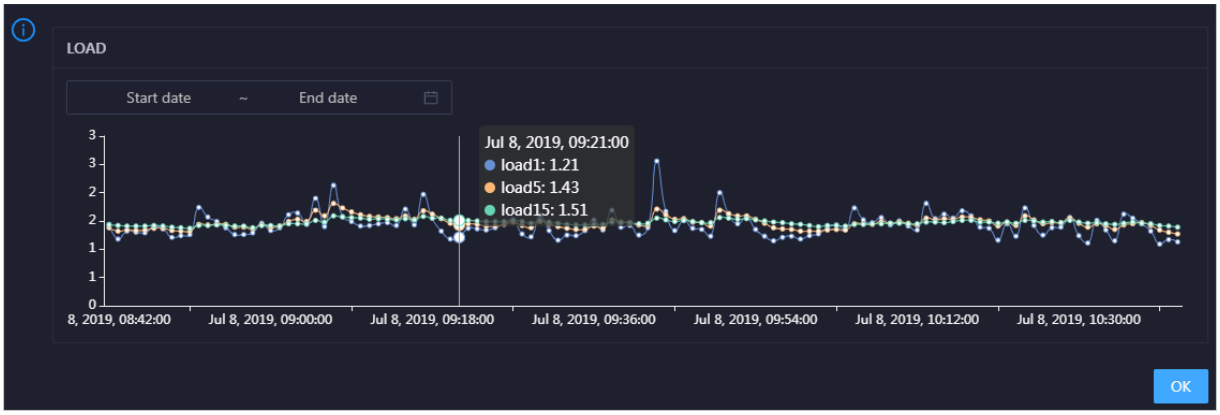

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

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

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

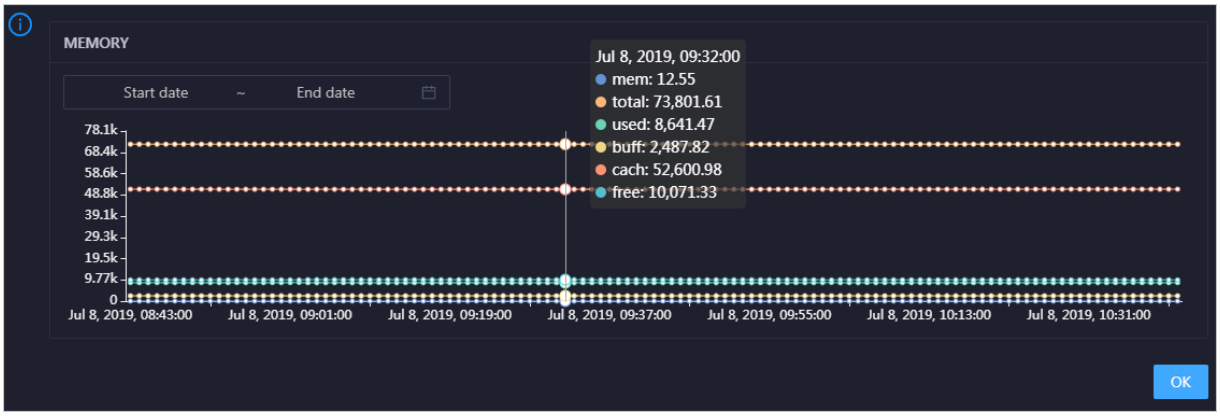

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

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

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

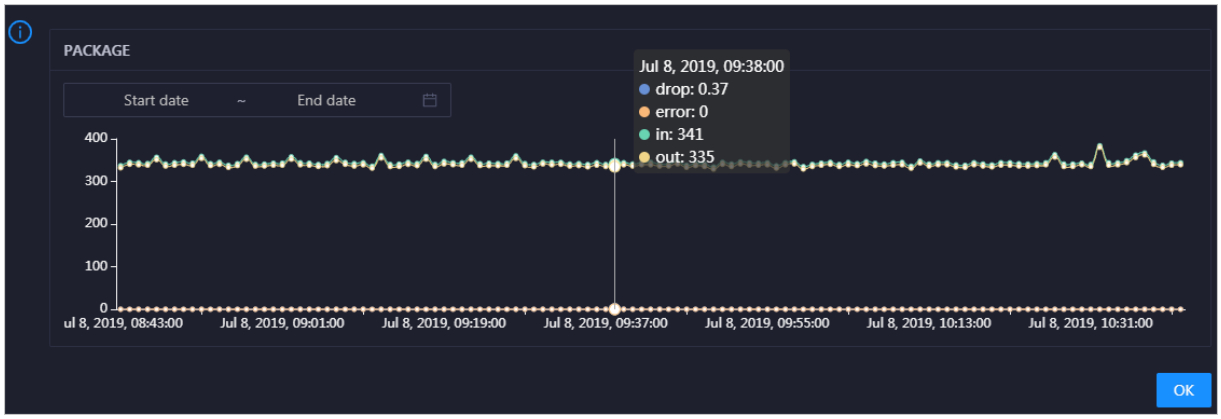

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

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

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

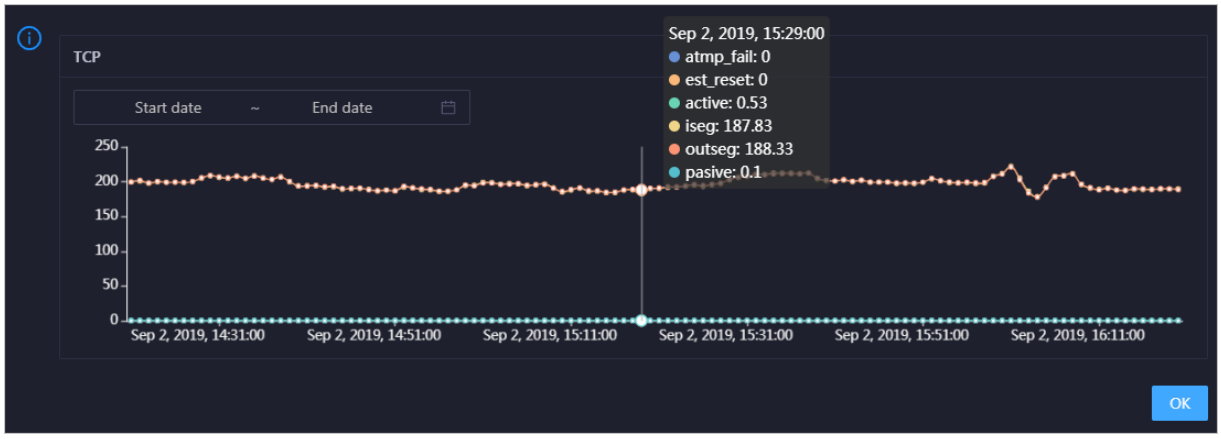

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

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

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

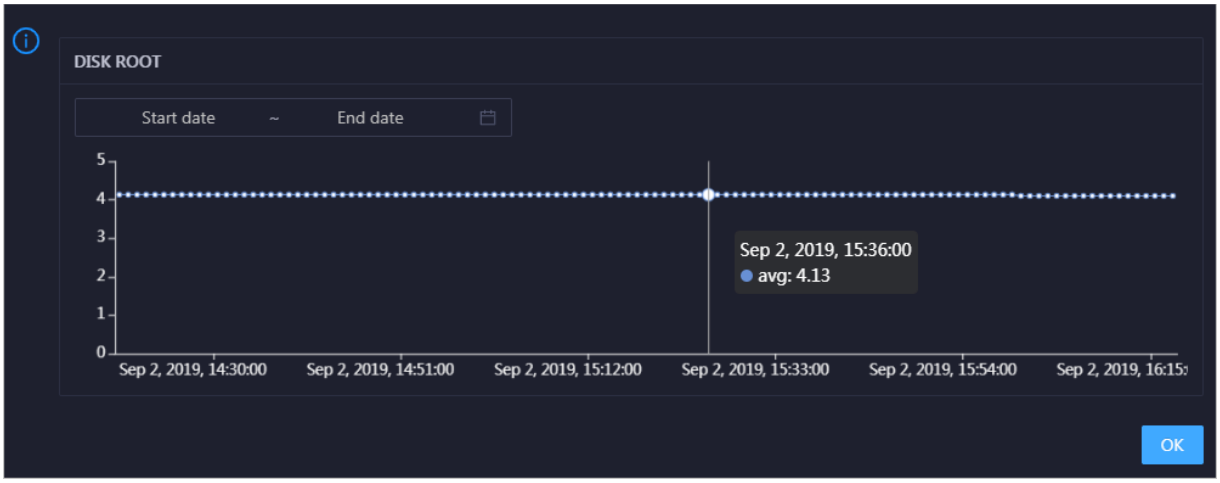

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

6.1.1.4.4.2. 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 for the service appears.

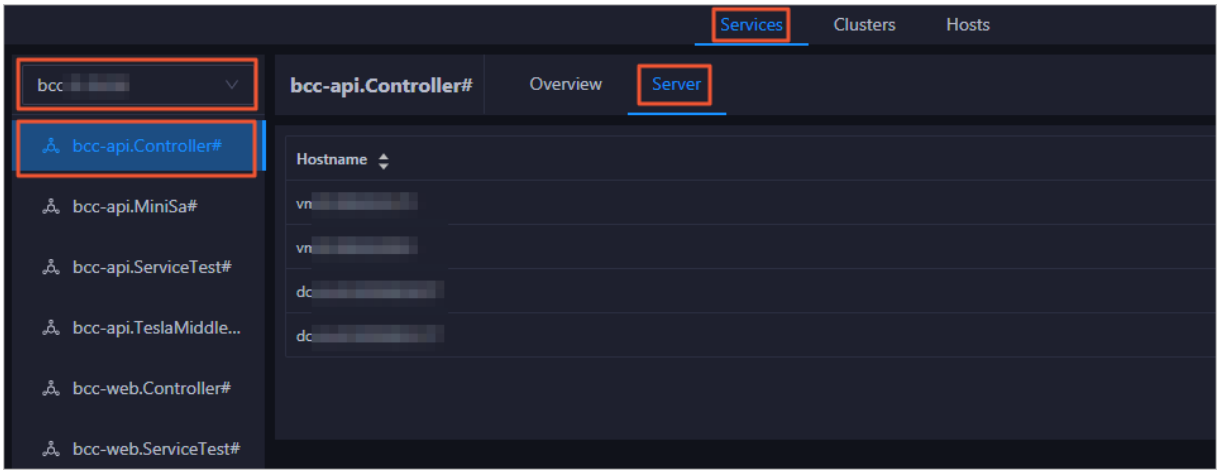

On the **Server** page, you can view the hosts where the selected service is run.

6.1.1.4.5. Cluster O&M

6.1.1.4.5.1. 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.

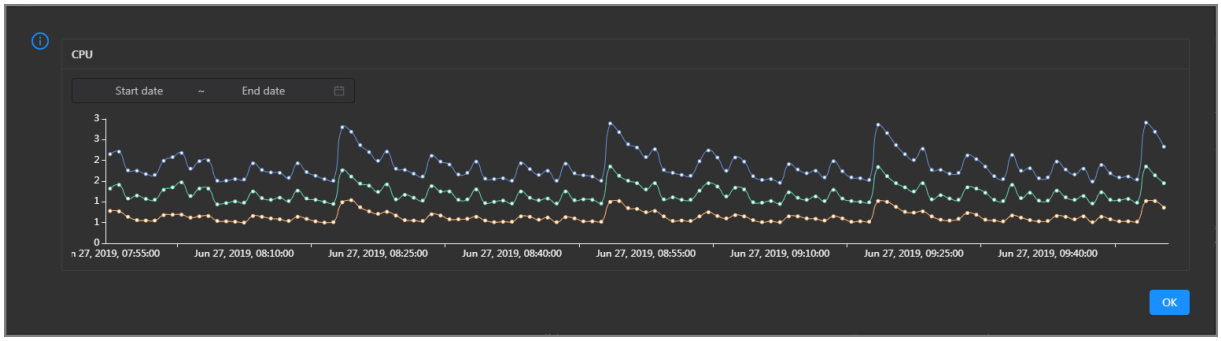

DISK

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

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

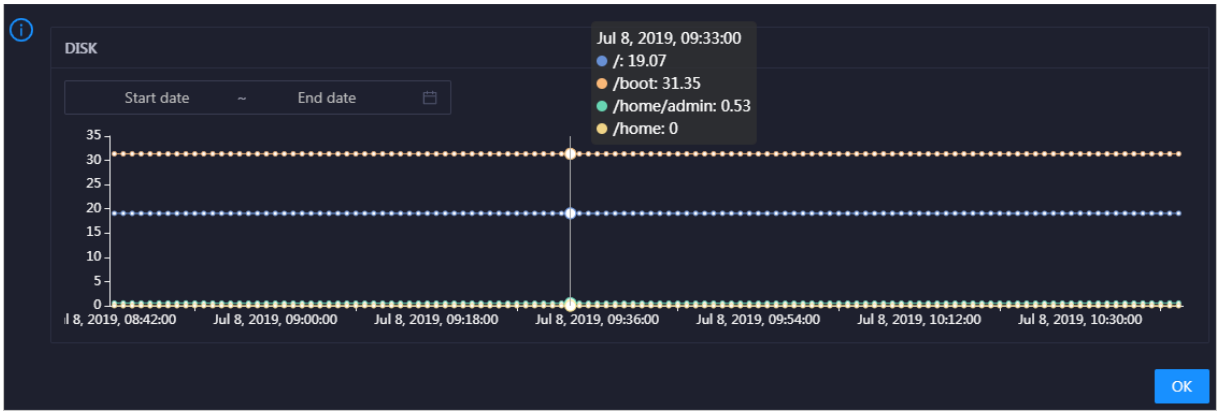

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

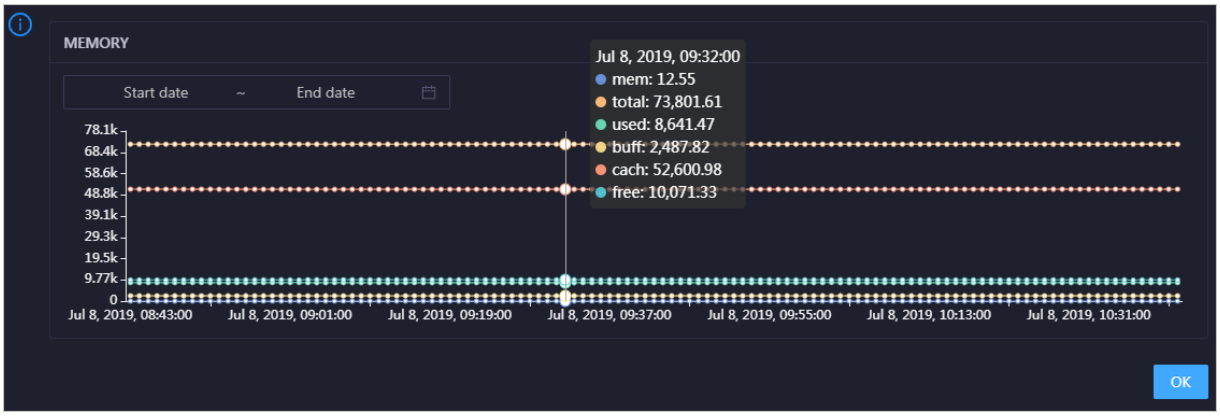

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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.

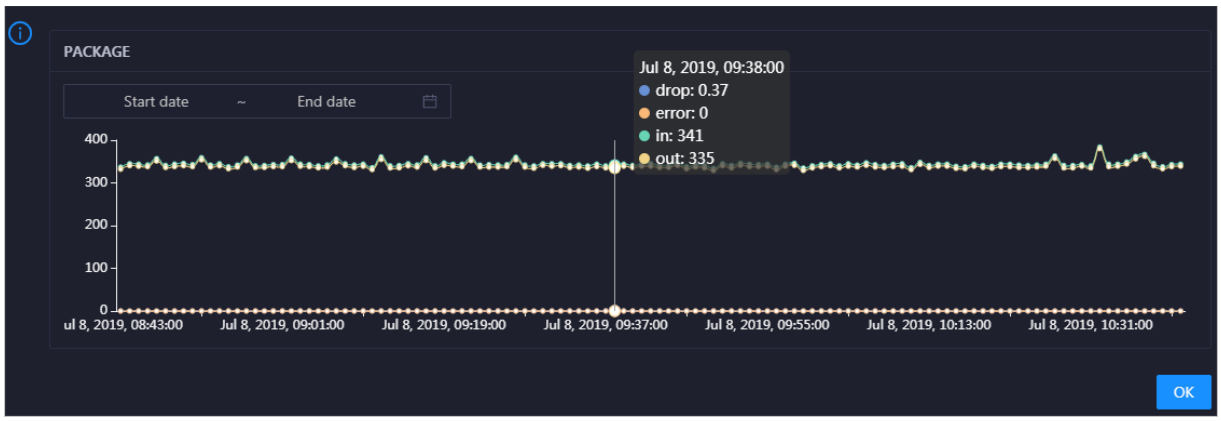

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

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.

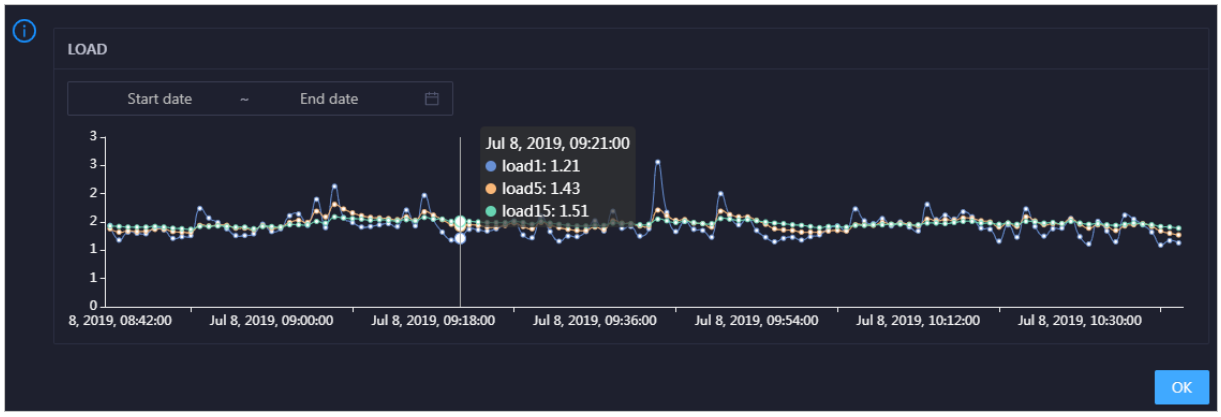

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

6.1.1.4.5.2. Cluster health

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

Entry

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

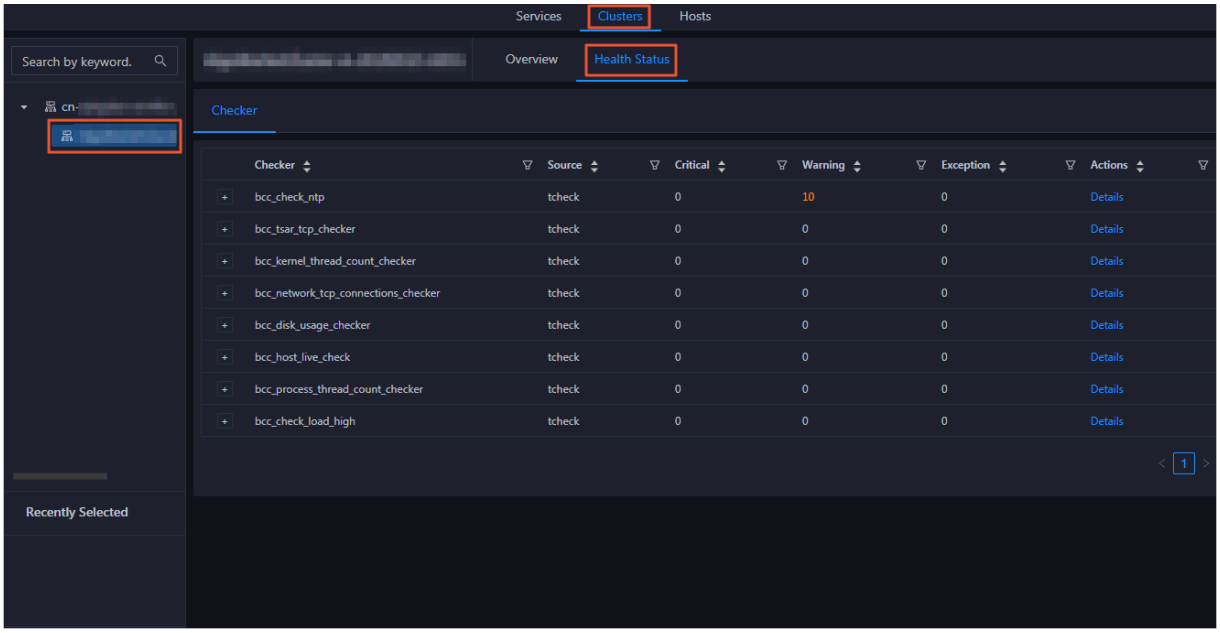

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are represented in different colors. You must handle the alerts in a timely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

View checker details

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

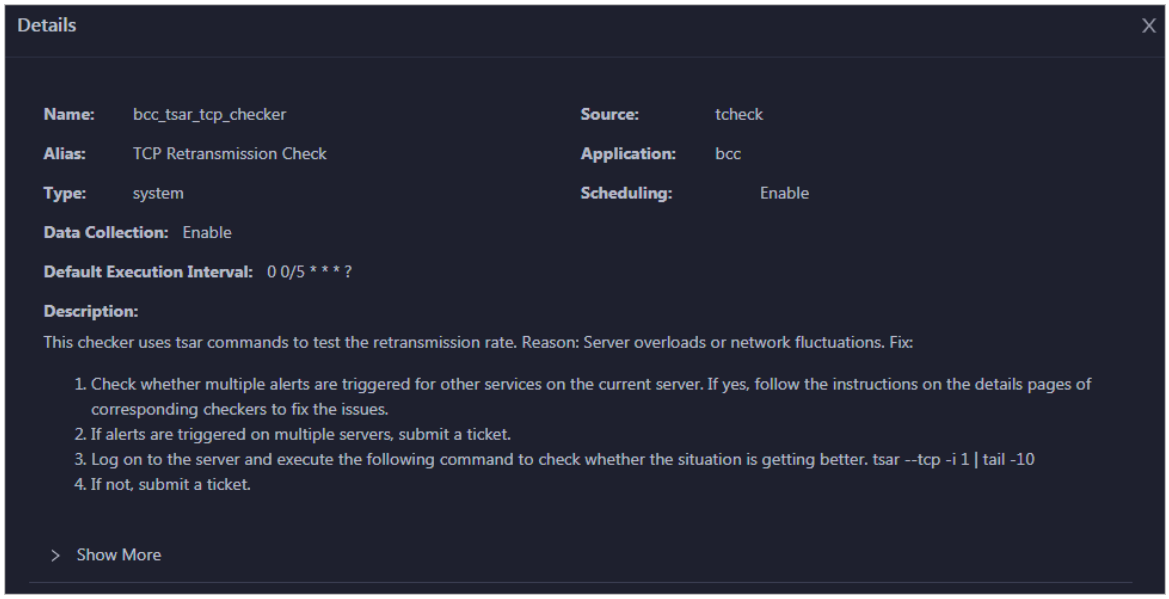

   The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

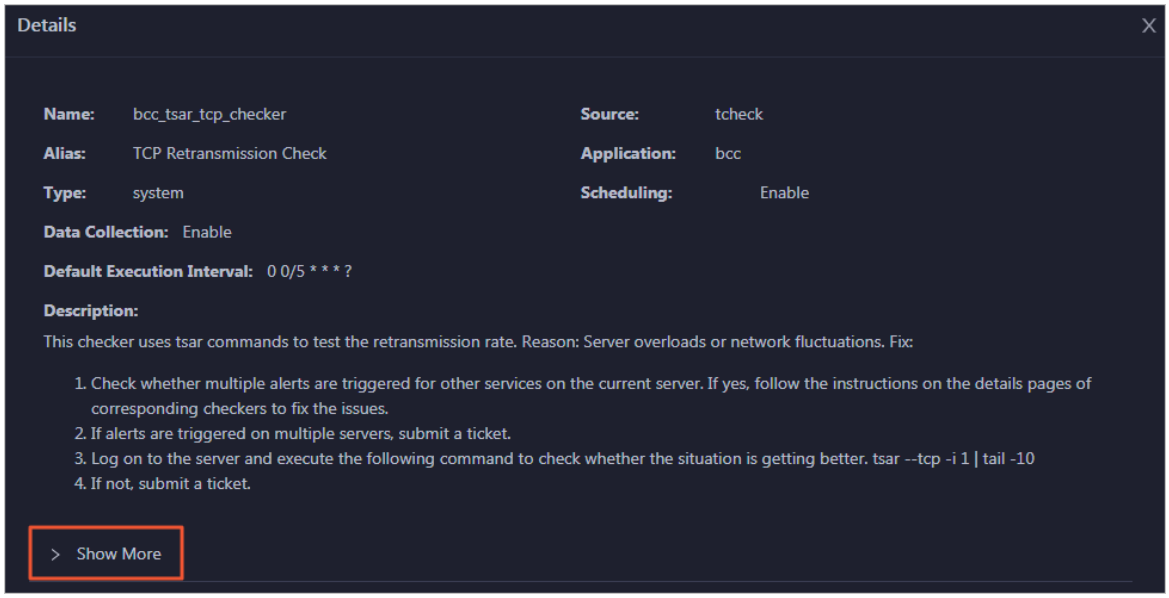

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

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

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

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

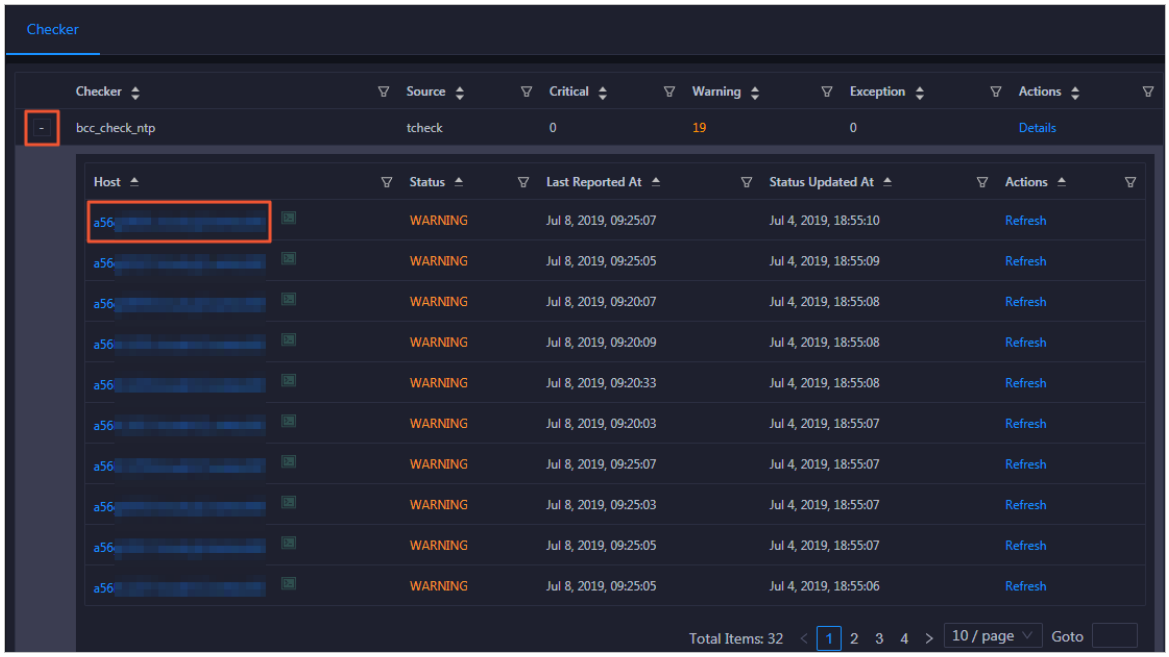

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

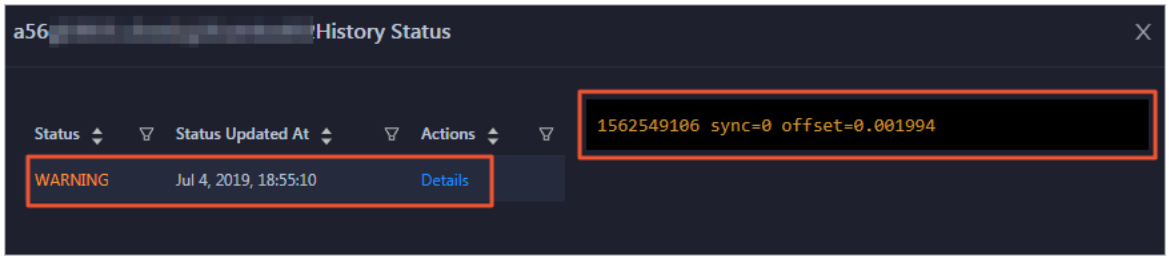

Clear alerts

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

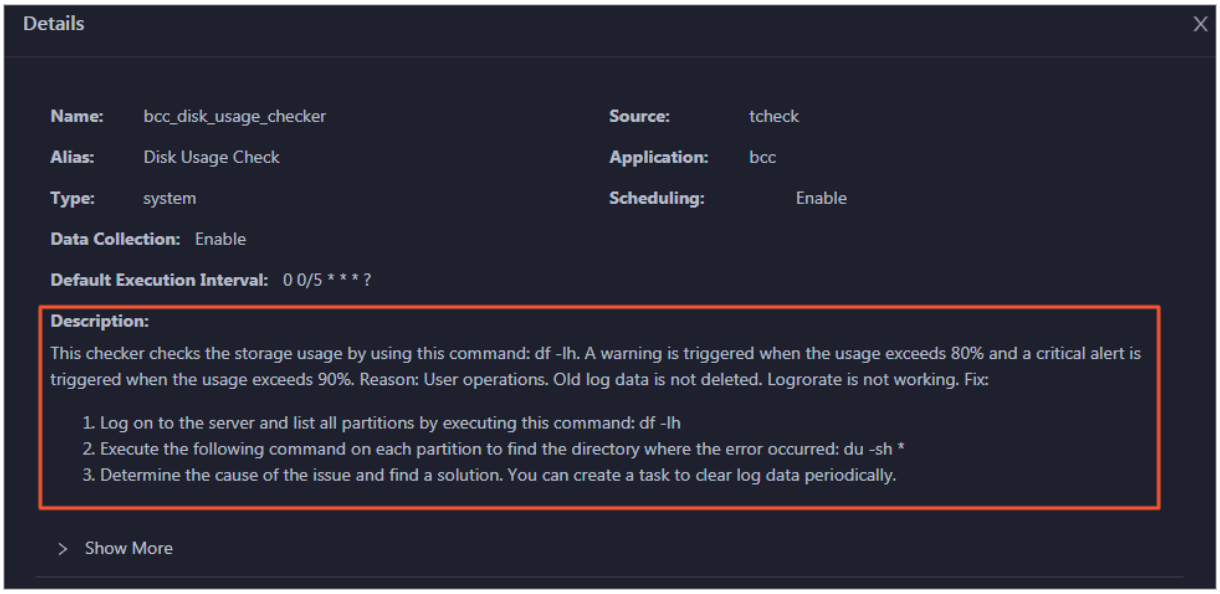

Log on to a host

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

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

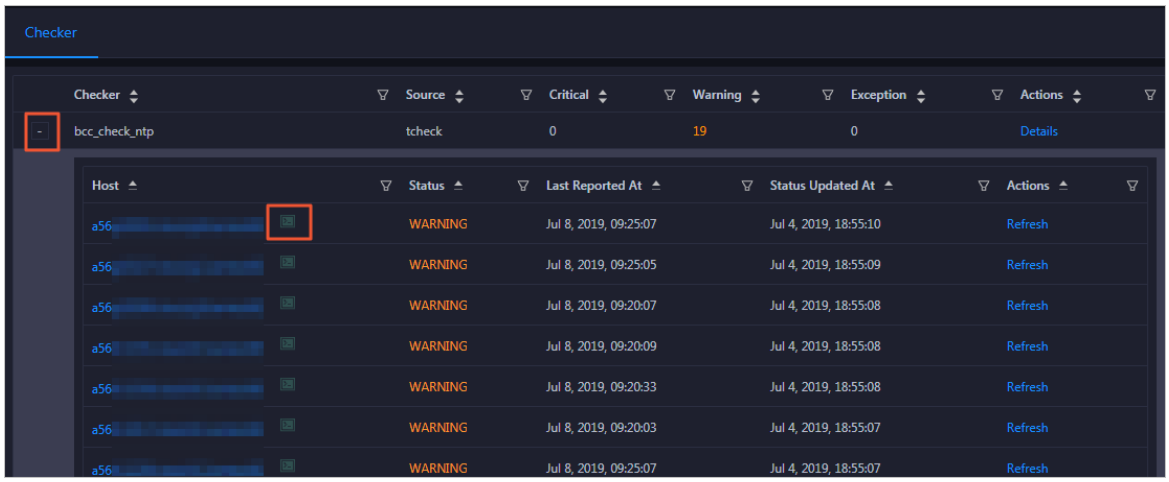

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

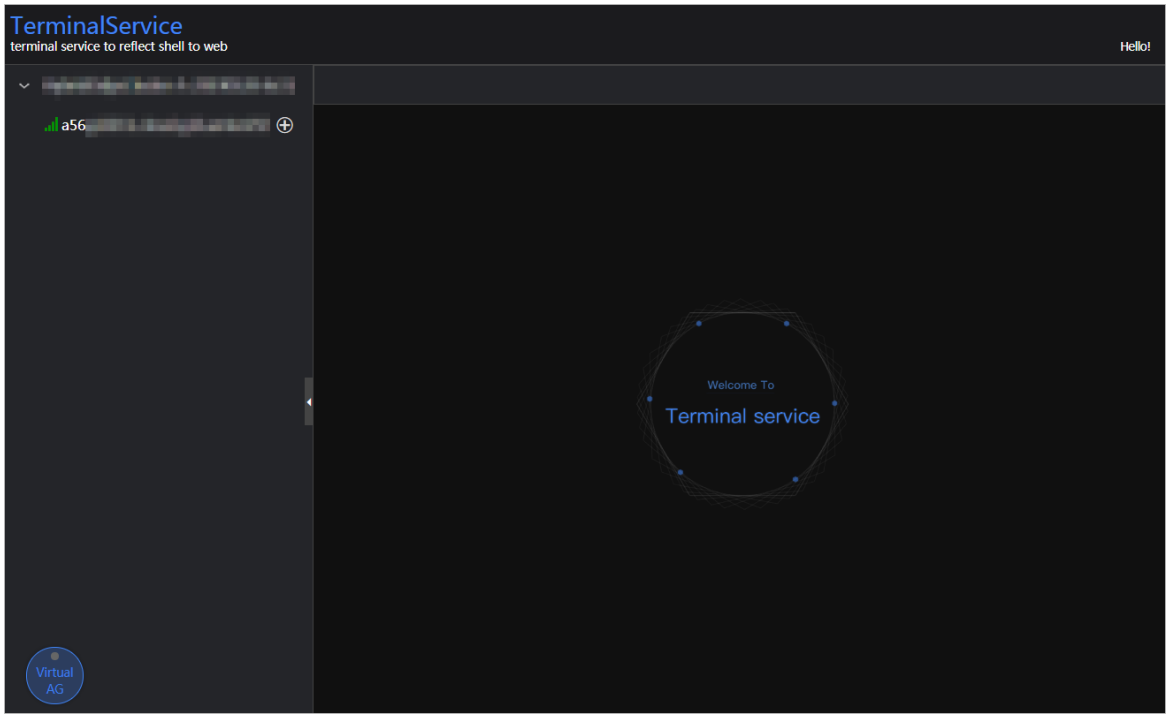

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

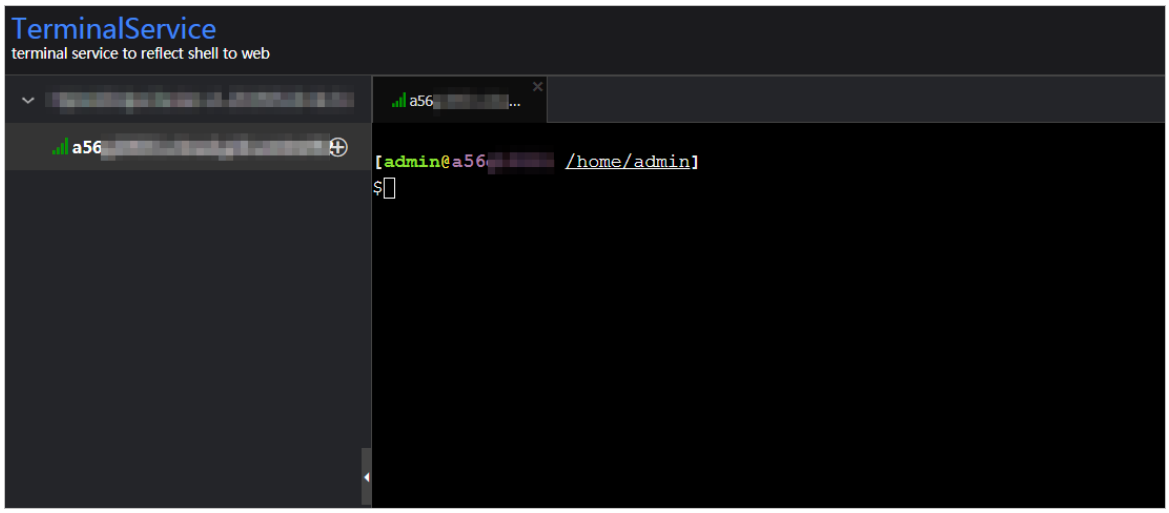

Run a checker again

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

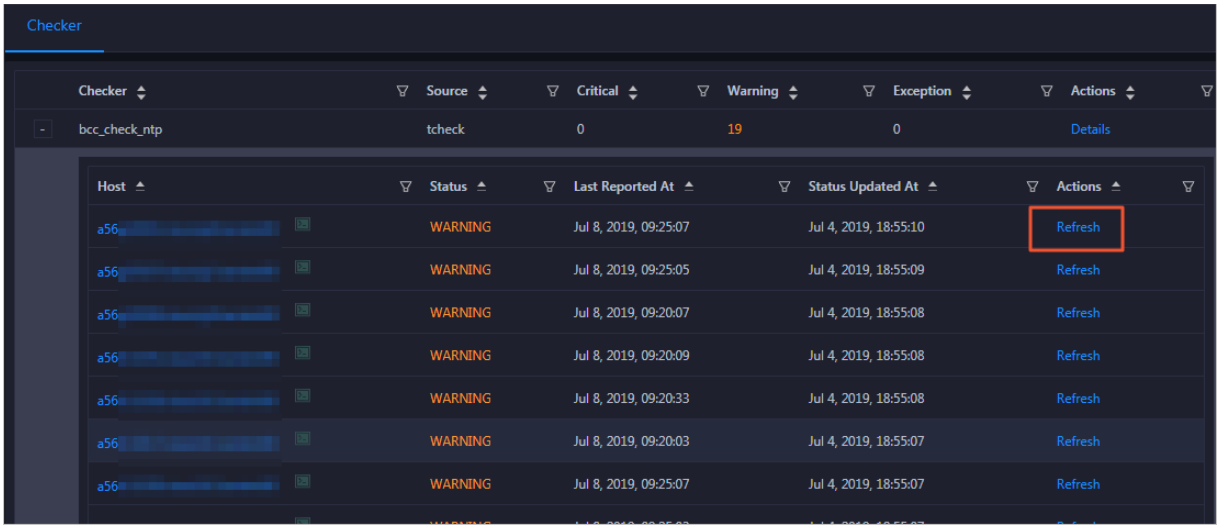

6.1.1.4.5.3. Restore environment settings

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 (,).

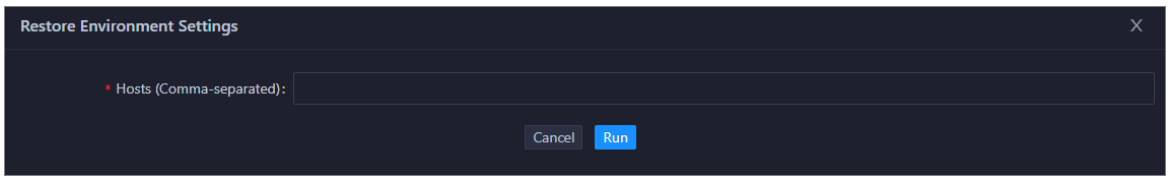

6. Click **Run**.
7. Check the execution status.

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

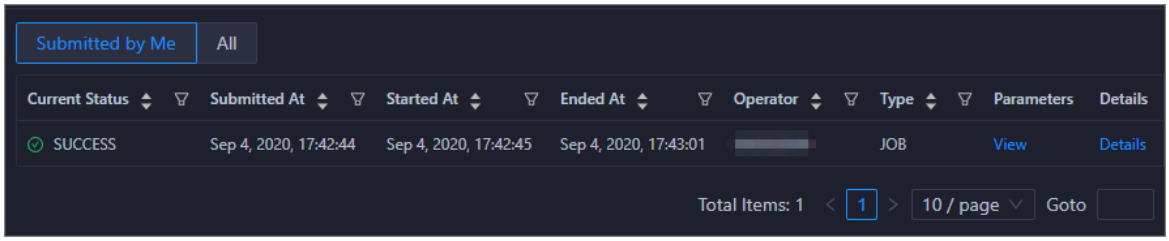

It requires a long time to restore environment settings. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeded. **FAILED** indicates that the execution failed.

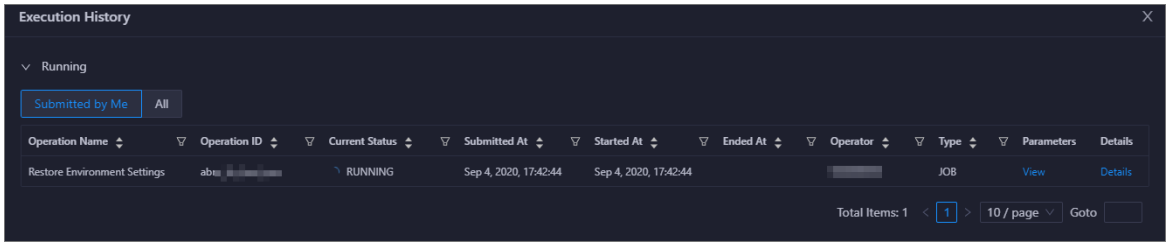

8. If the status is RUNNING, click **Details** in the Details column to view the steps and progress of restoration.

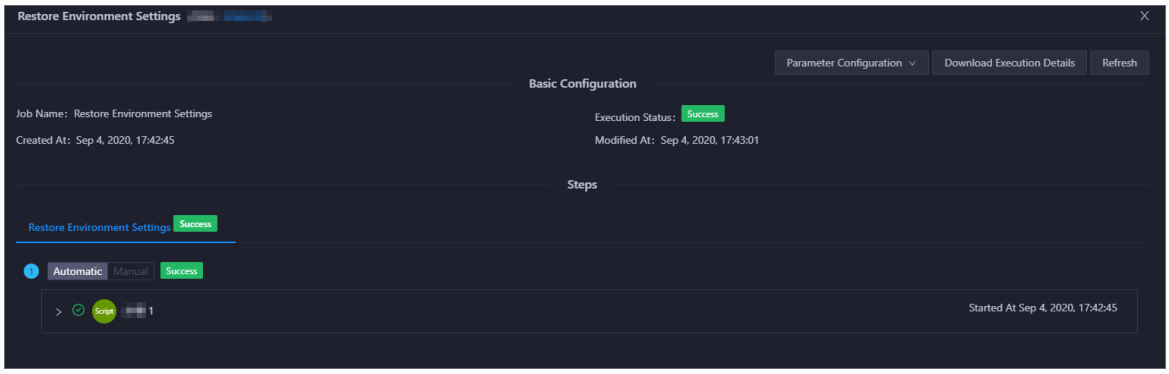

9. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of the failure to restore environment settings.

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.

6.1.1.4.6. Host O&M

6.1.1.4.6.1. Host overview

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

Entry

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

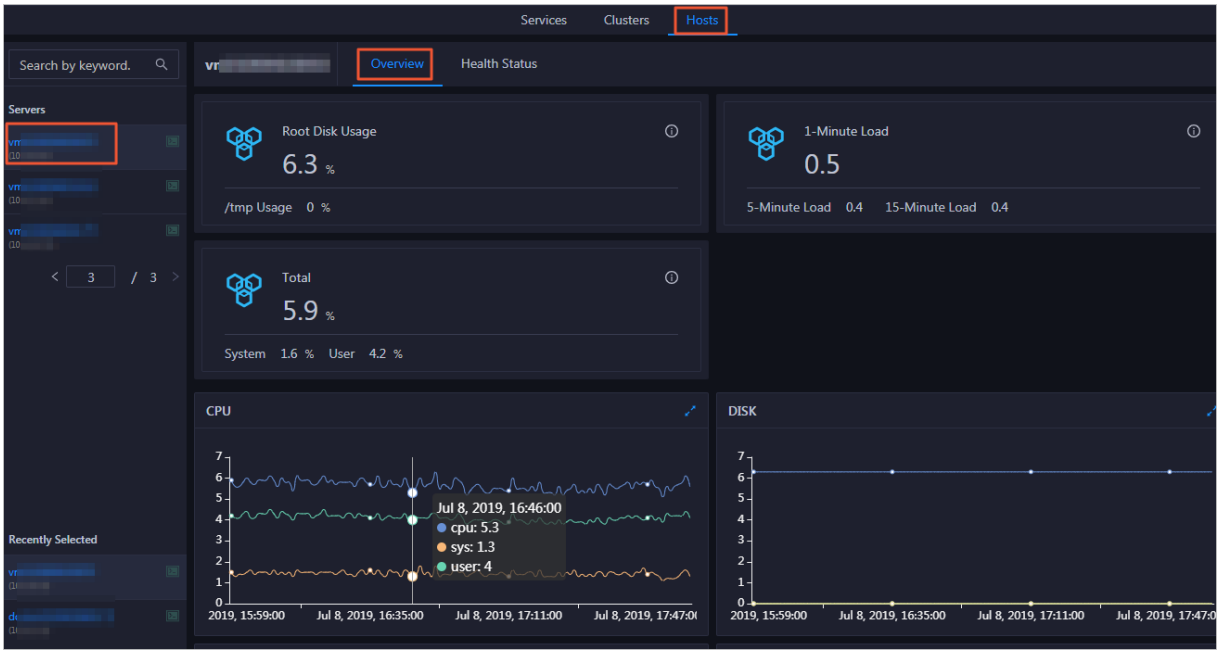

Root Disk Usage, Total, and 1-Minute Load

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

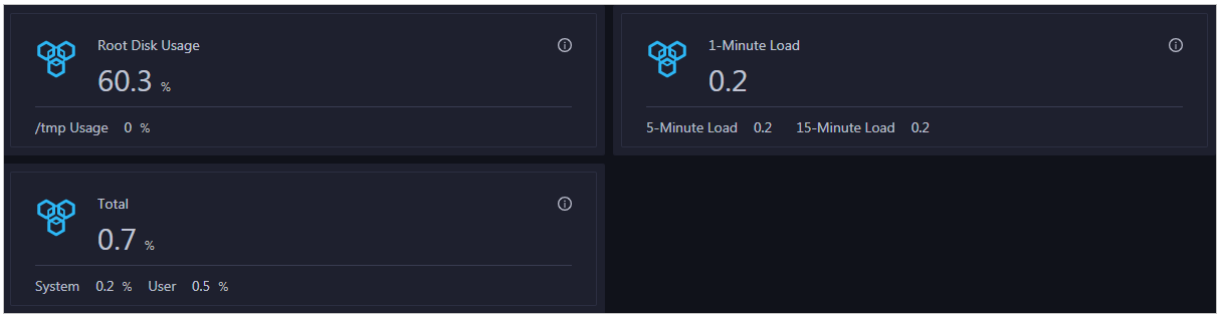

CPU

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

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

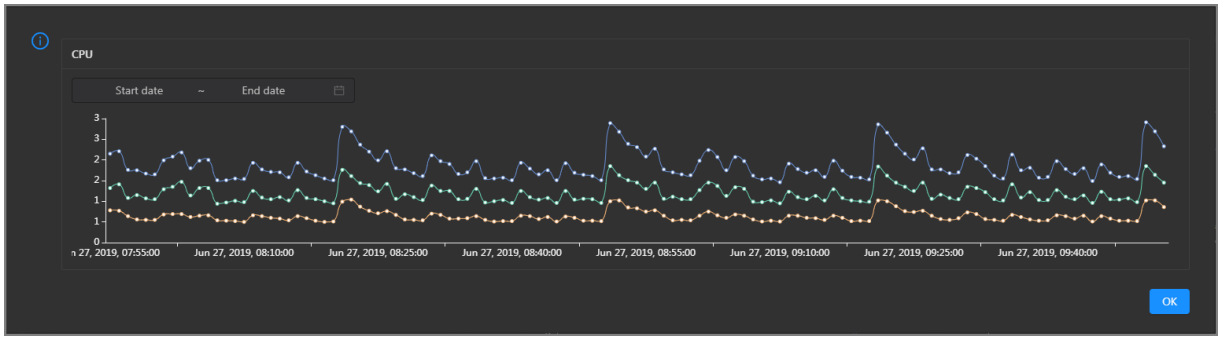

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

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

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

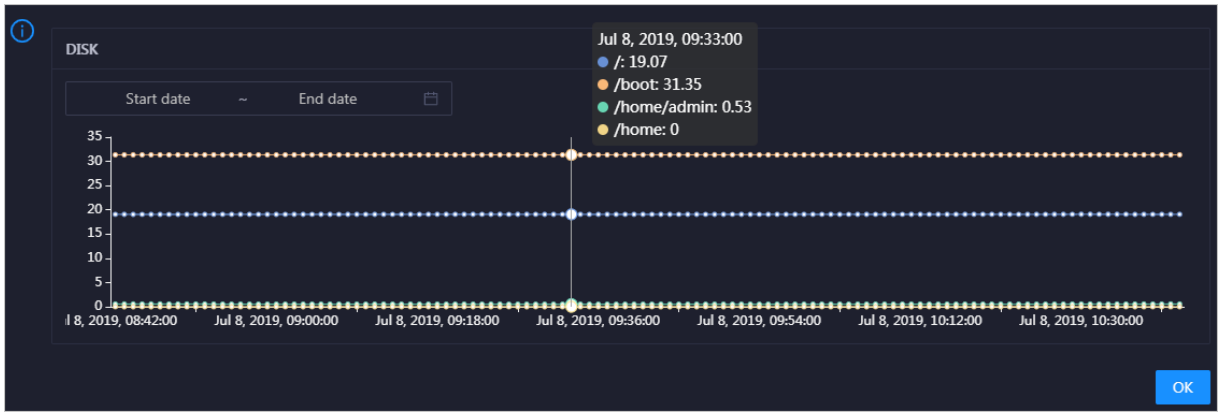

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

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

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

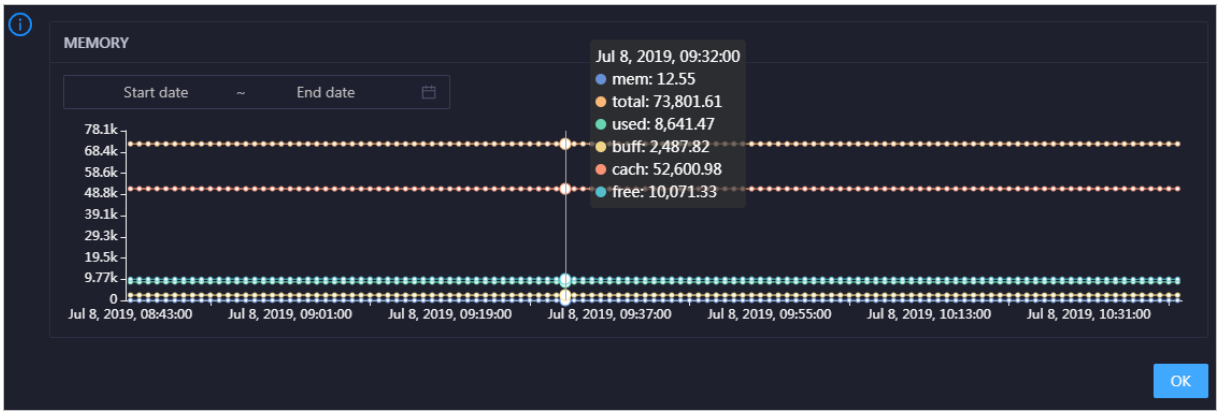

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

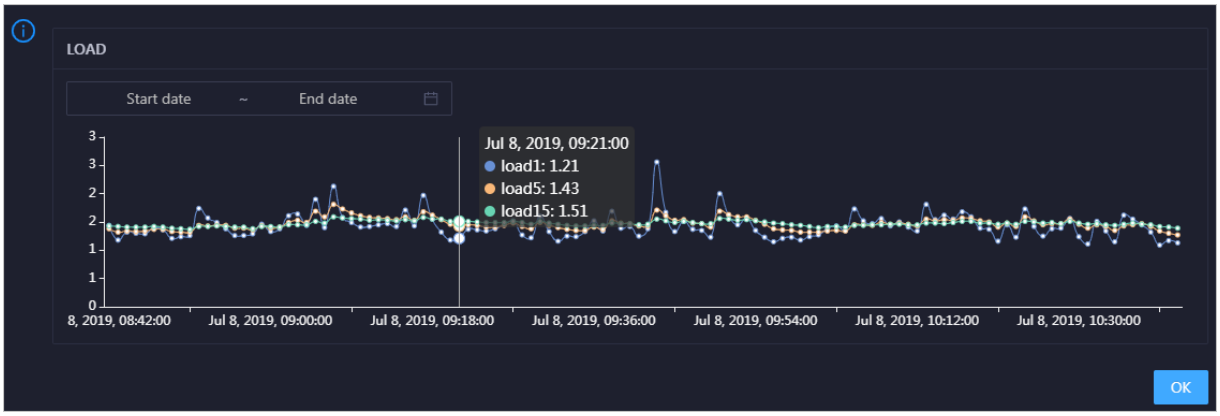

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

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

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

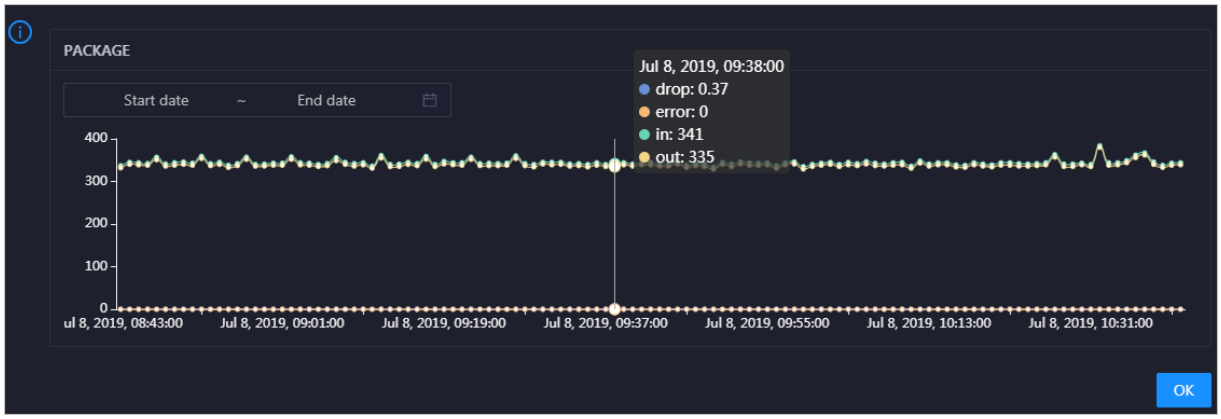

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

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

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

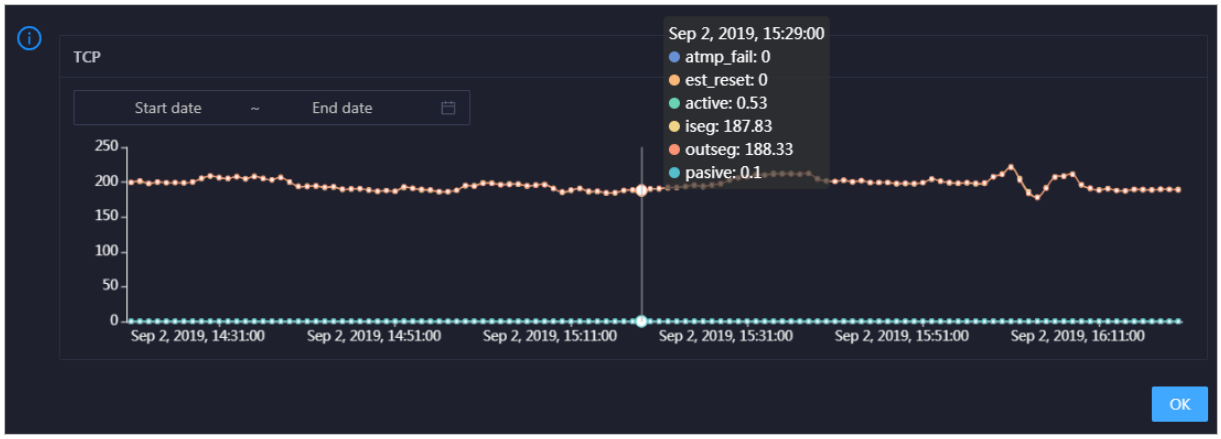

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

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

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

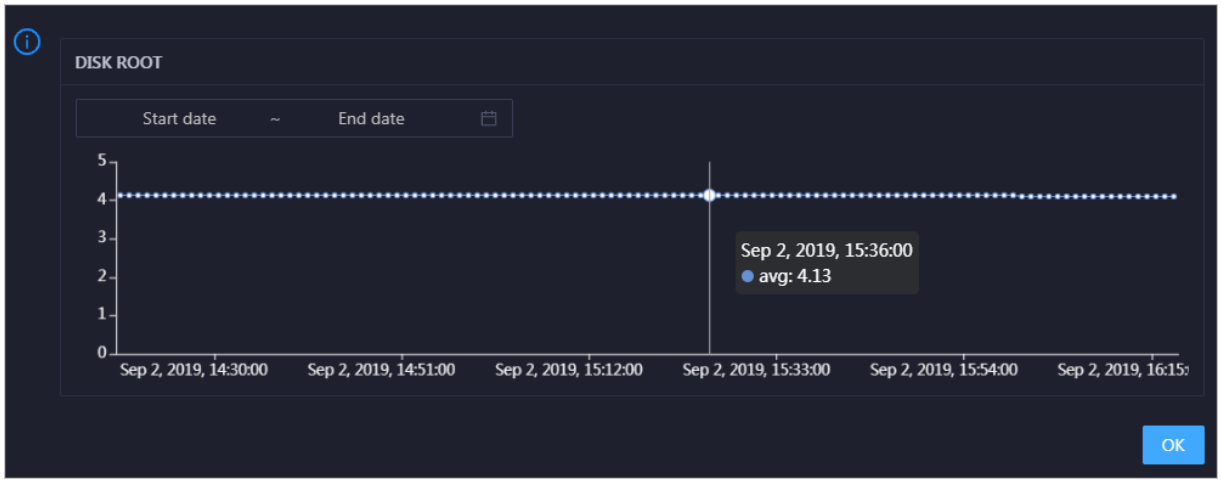

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

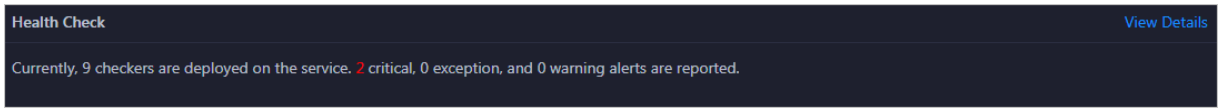

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

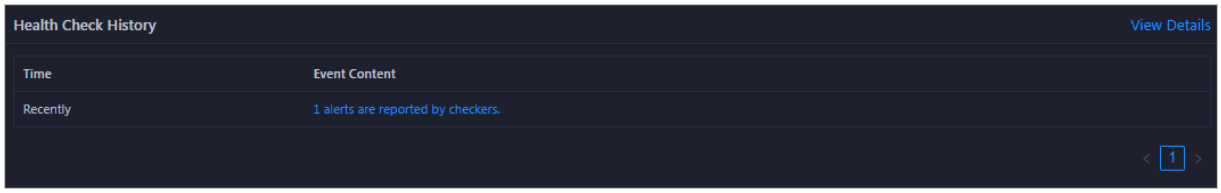

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

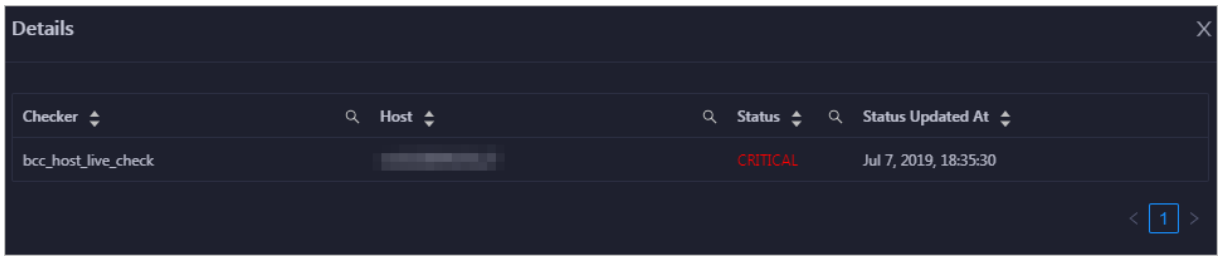

6.1.1.4.6.2. Host health

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

Entry

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

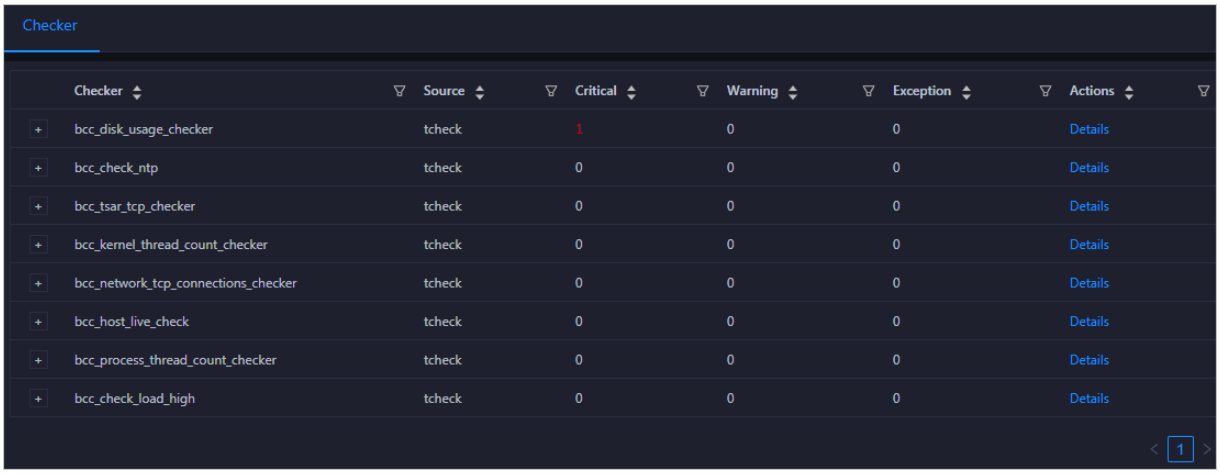

On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

View checker details

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

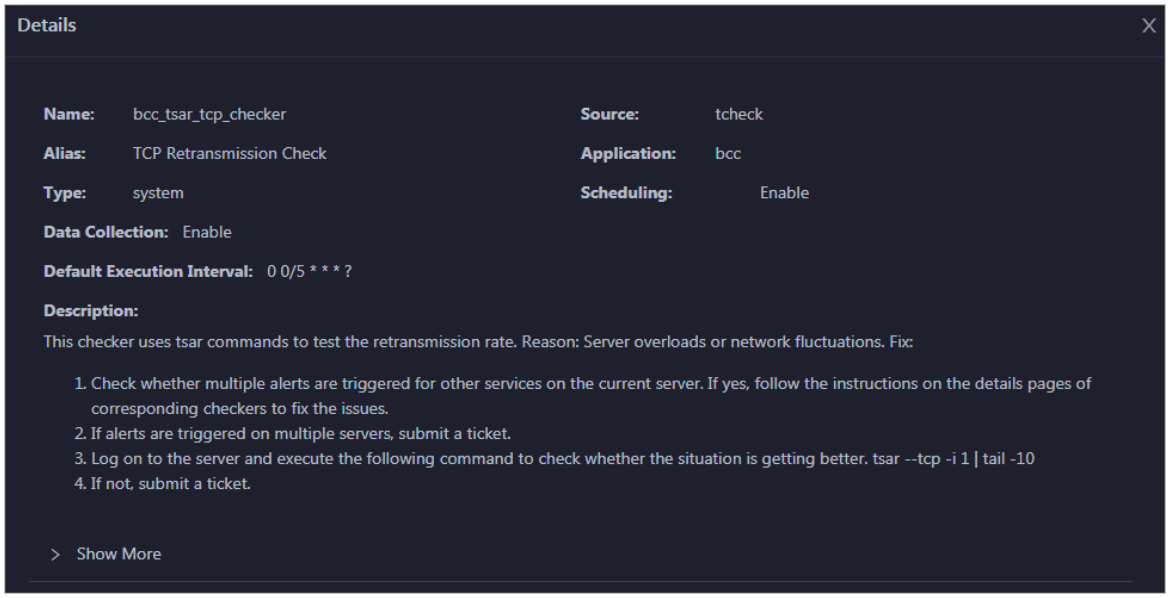

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

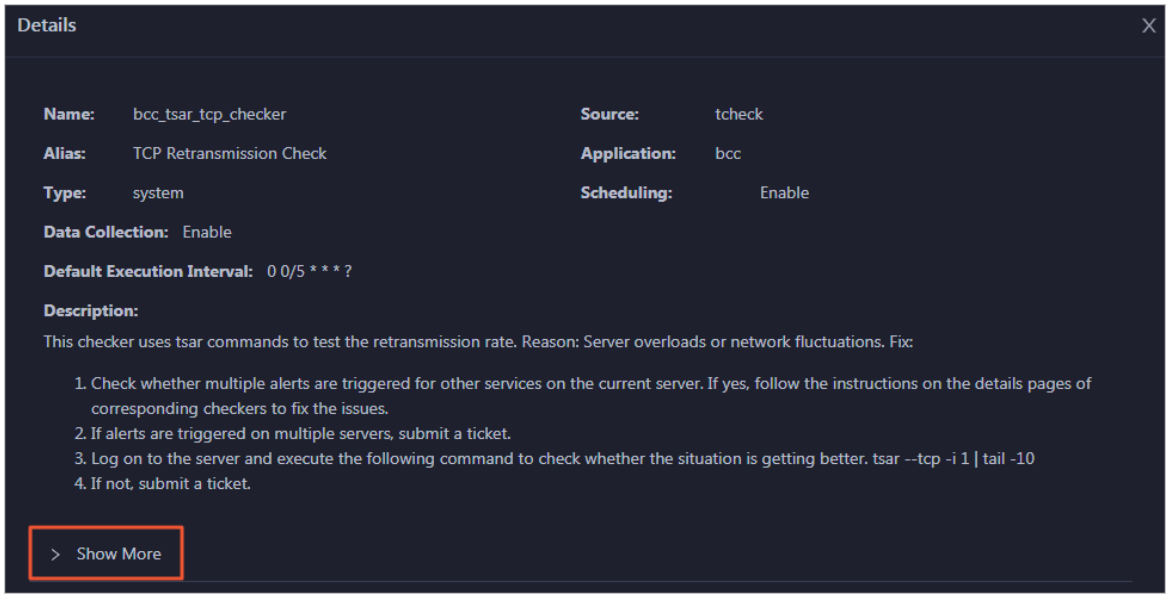

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

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

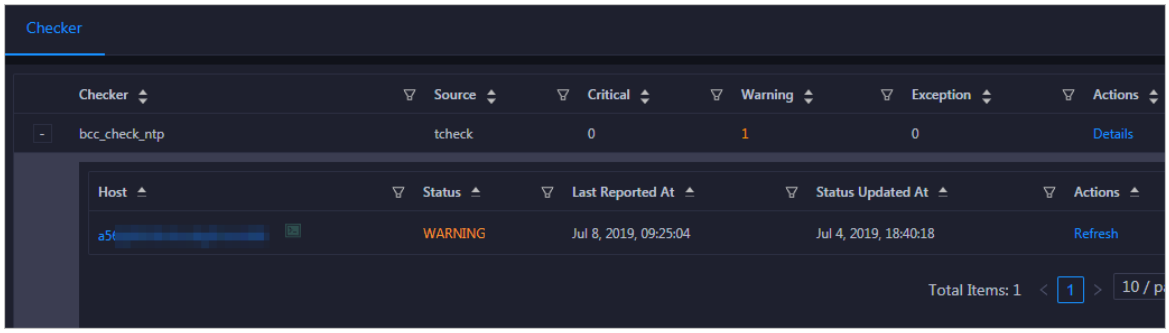

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

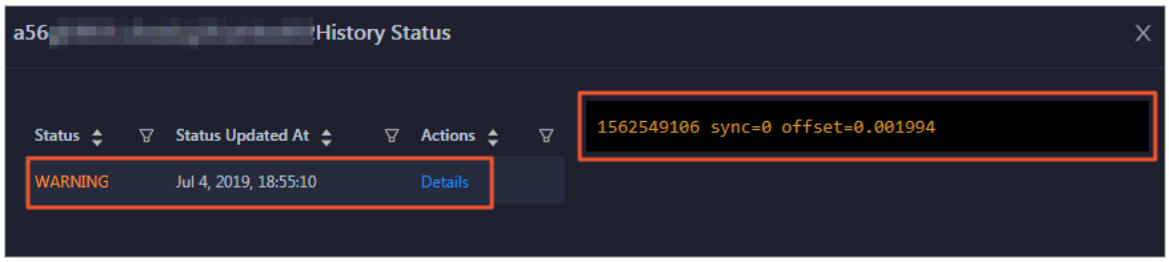

Clear alerts

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

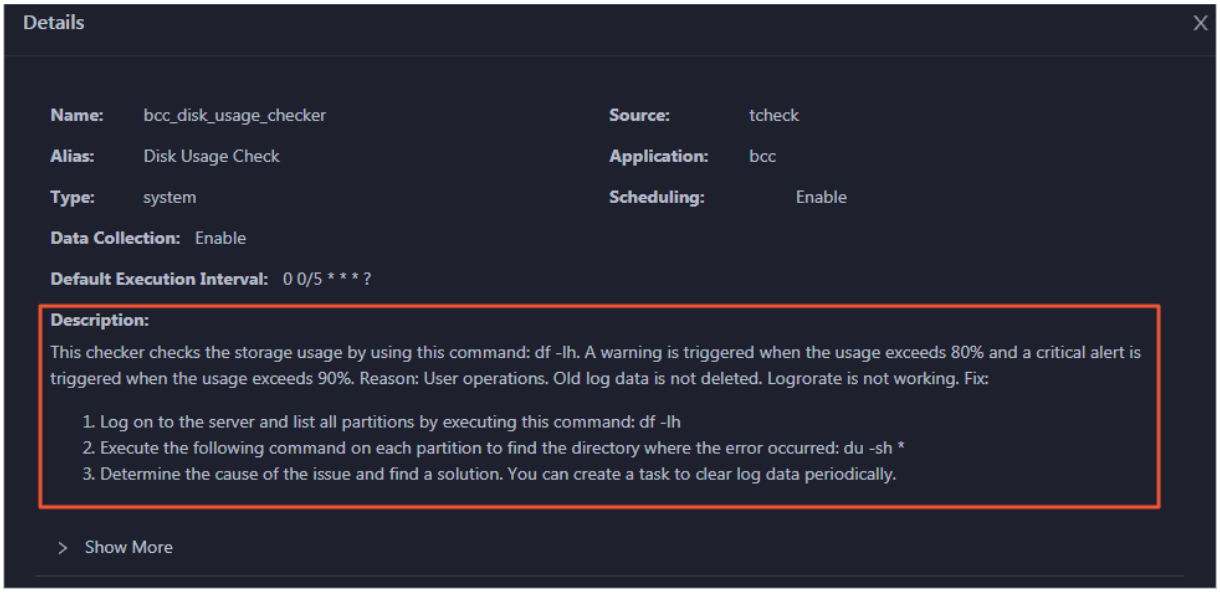

Log on to a host

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

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

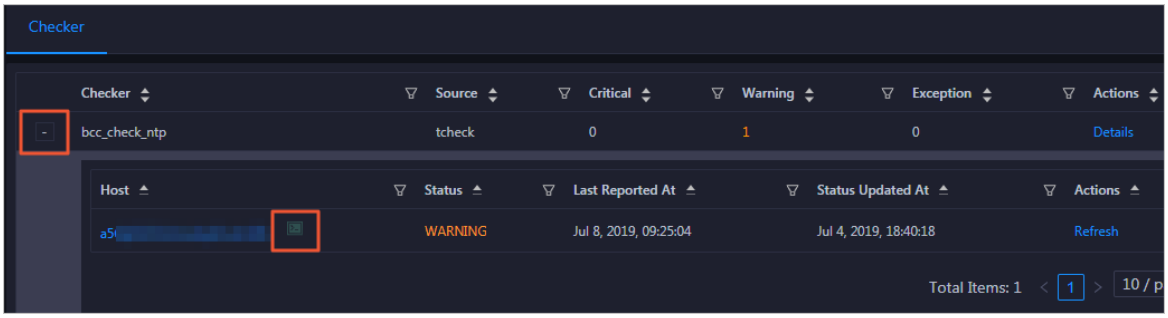

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

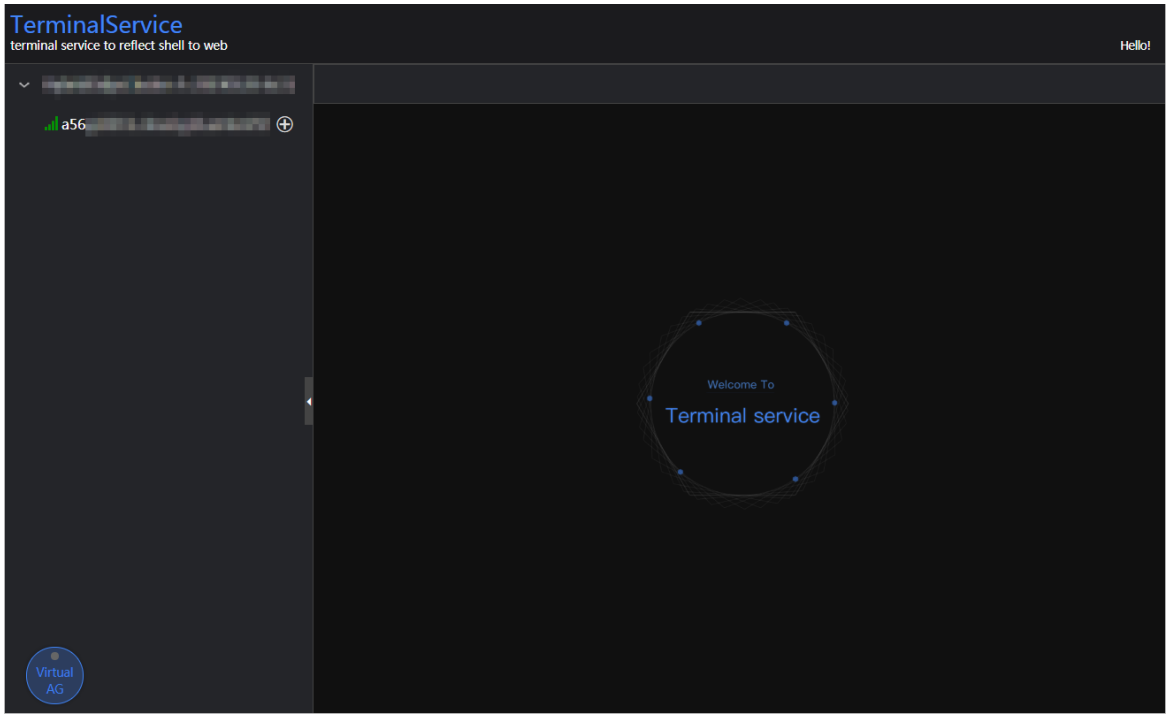

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

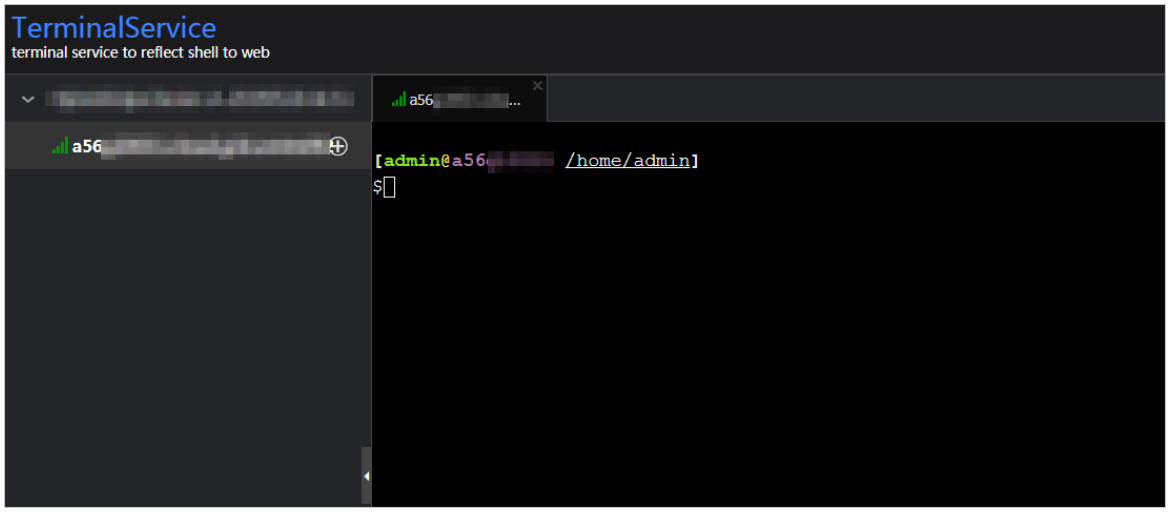

Run a checker again

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

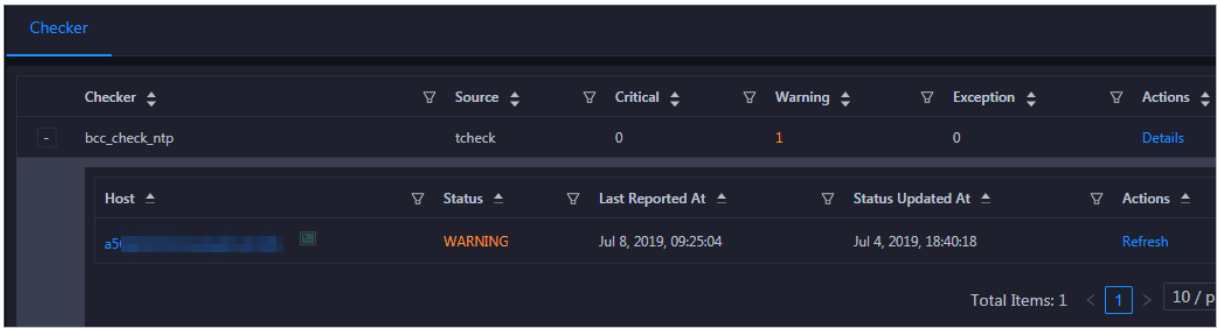

6.1.1.5. MaxCompute

6.1.1.5.1. Project details

The Apsara Big Data Manager (ABM) console shows your MaxCompute projects and project details. You can view the project overview, jobs, storage, configuration, quota group, and tunnel, as well as information about resource analysis, storage encryption, and cross-cluster replication.

Go to the project details page

1. Log on to the Apsara Big Data Manager 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. Choose **Projects** > **Project List**. Click the name of a project to view its details.

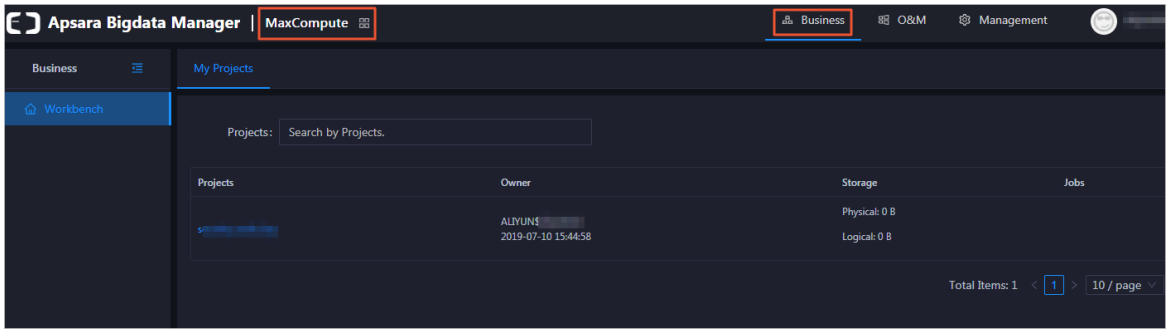

Overview

On the **Overview** tab, you can view the following information about the 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, status, start time, priority, and type. You can also view the operational logs of a job to locate errors during job running.

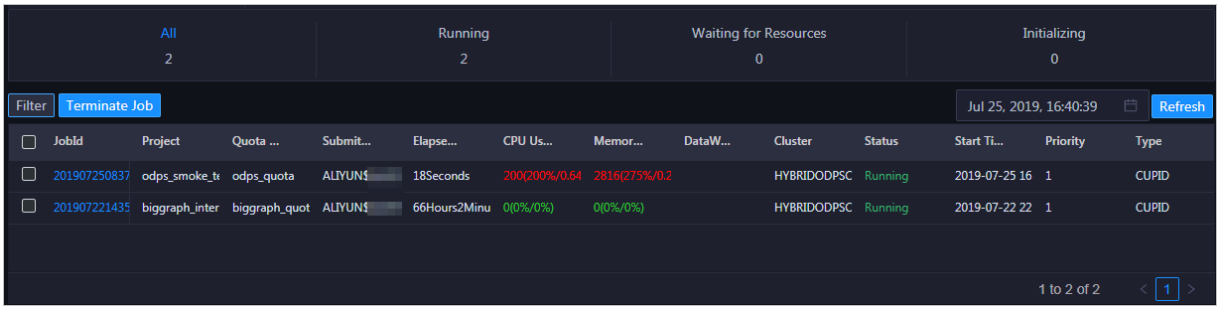

You can perform the following operations on the Jobs tab:

- 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**.

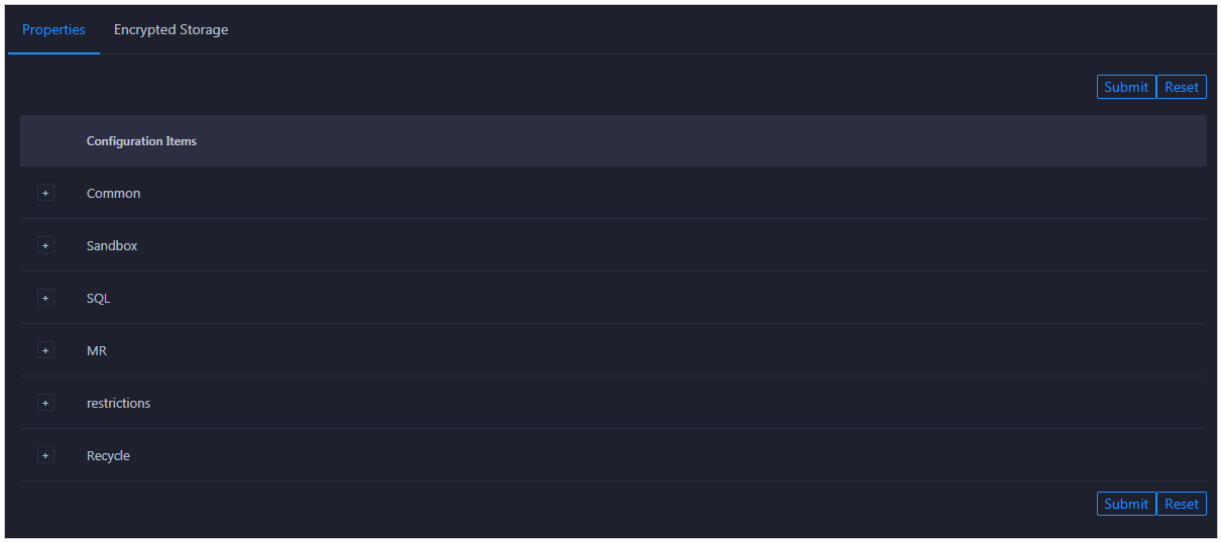

On the **Authorize Package for Metadata Repository** tab, you can install the package and perform package-based authorization.

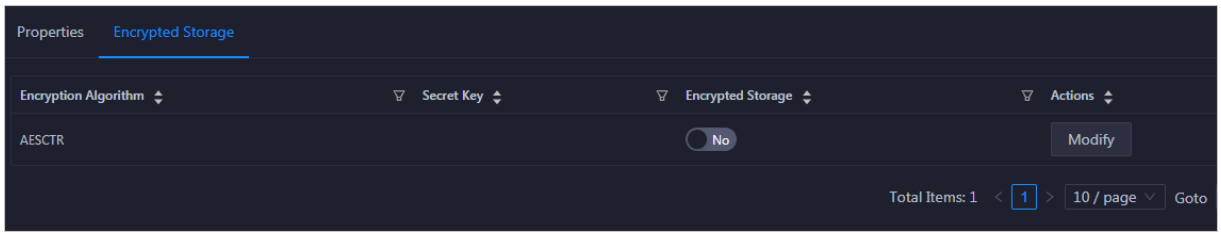

Quota Groups

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

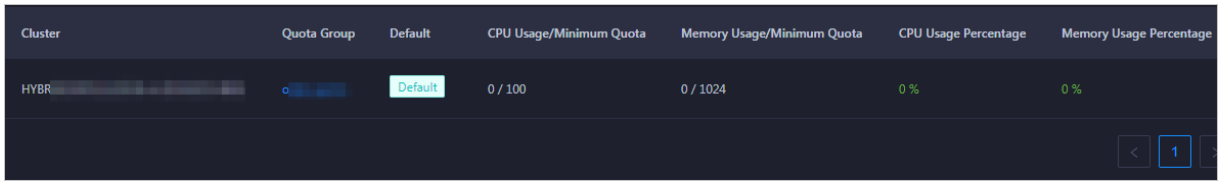

To view details about a quota group, click the quota group name in the **Quota** column.

> **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.

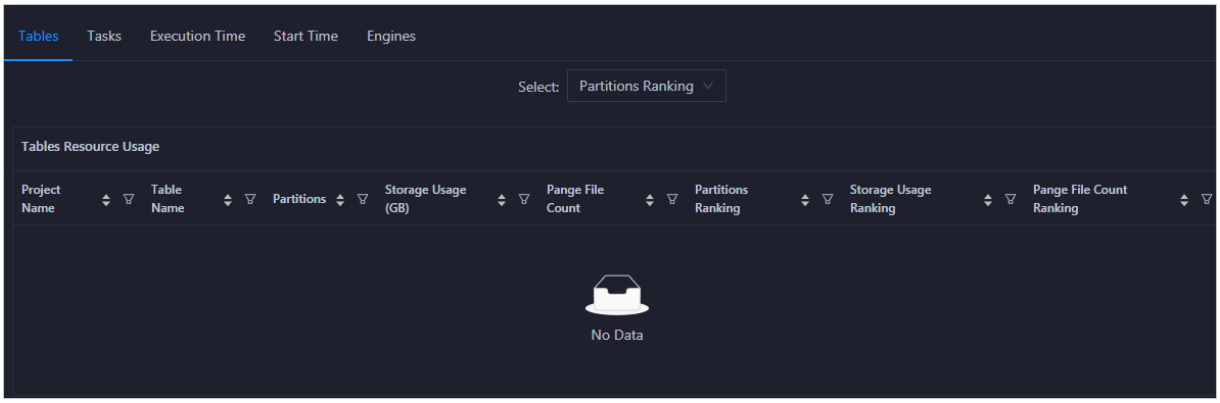

Encryption at Rest

On the **Encryption at Rest** tab, you can encrypt data by using the following encryption algorithms: AES-CTR, AES256, RC4, and SM4.

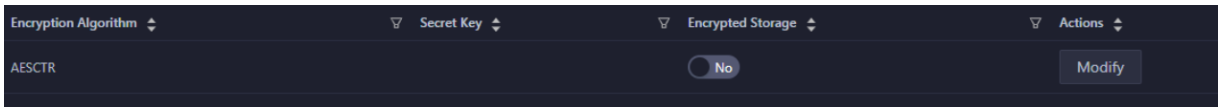

Cross-cluster Replication

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.

For example, 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** sub-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.

6.1.1.5.2. Business O&M

6.1.1.5.2.1. 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 search for and filter 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.
- Quota Groups: shows the quota groups of all projects in a MaxCompute cluster. It allows you to create and modify quota groups. You can also view details about quota groups and enable period management for quota groups.
- Jobs: shows information about jobs in a MaxCompute cluster. You can search for and filter jobs. You can also view the operational logs, terminate running jobs, 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 Apsara Big Data 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 **Business** tab. In the left-side navigation pane, choose **Projects > Project List**.

6.1.1.5.2.2. Project management

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 set 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.

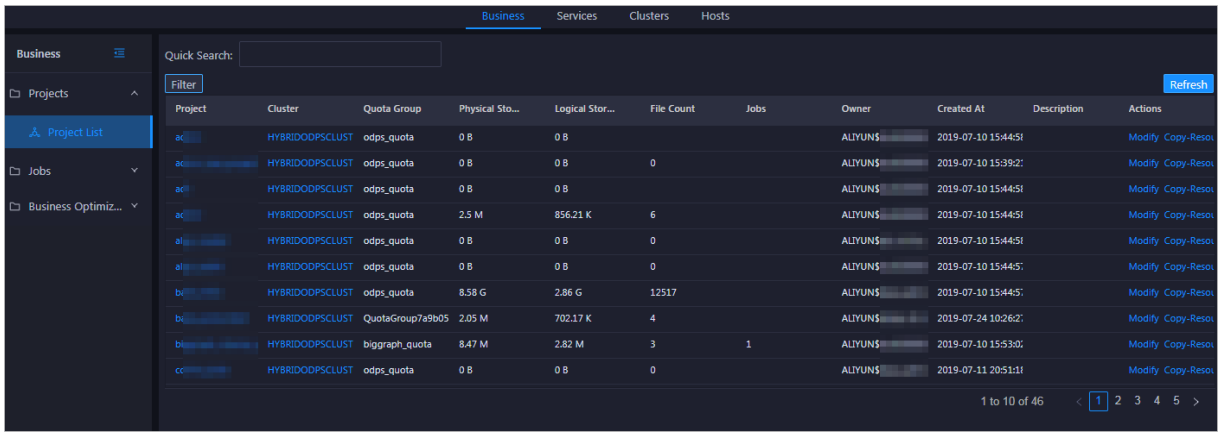

The **Project List** page shows the detailed information about all projects in a cluster. You can view the name, cluster, used storage, storage quota, storage usage, number of files, owner, and creation time of a project.

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 project for which you want to change the quota group, click **Actions** in the Actions column, and select **Change Default Quota Group**. In the **Change Default Quota Group** pane, configure parameters.

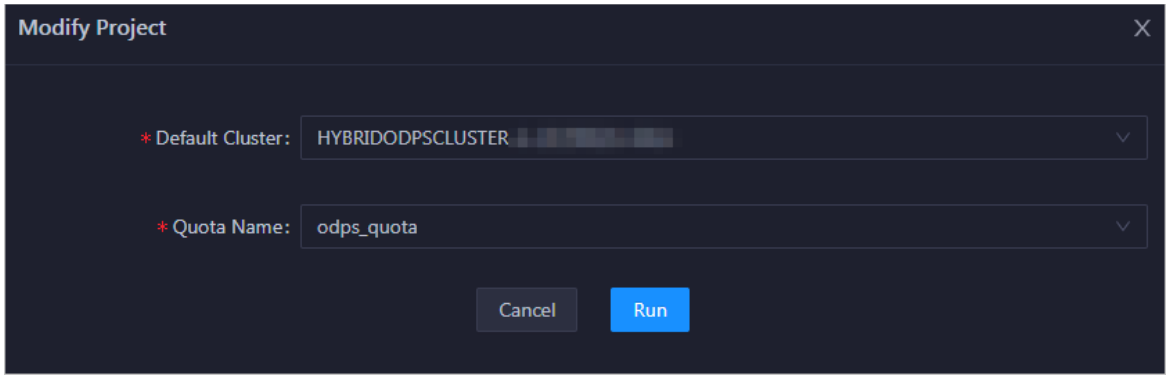

Parameters:

- **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 configure 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 project for which you want to modify the storage quota, click **Actions** in the Actions column, and select **Modify Storage Quota**. In the **Change Storage Quota** pane, configure 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 reason for the modification

2. After you configure 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 project for which you want to configure resource replication, click **Actions** in the Actions column, and select **Resource Replication**. In the **Copy Resource** pane, configure parameters.

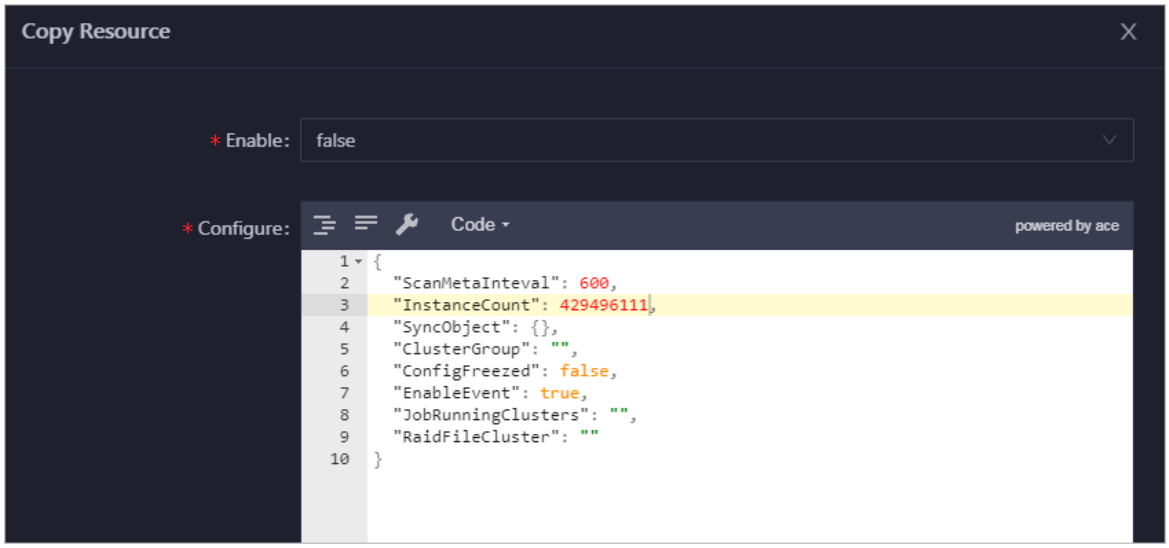

Parameters:

- **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 for a project. In most cases, the default settings are used. If you want to modify the settings, consult O&M engineers.

2. After you modify code in the **Configure** field, click **Compare Versions** to view the differences, which are highlighted.

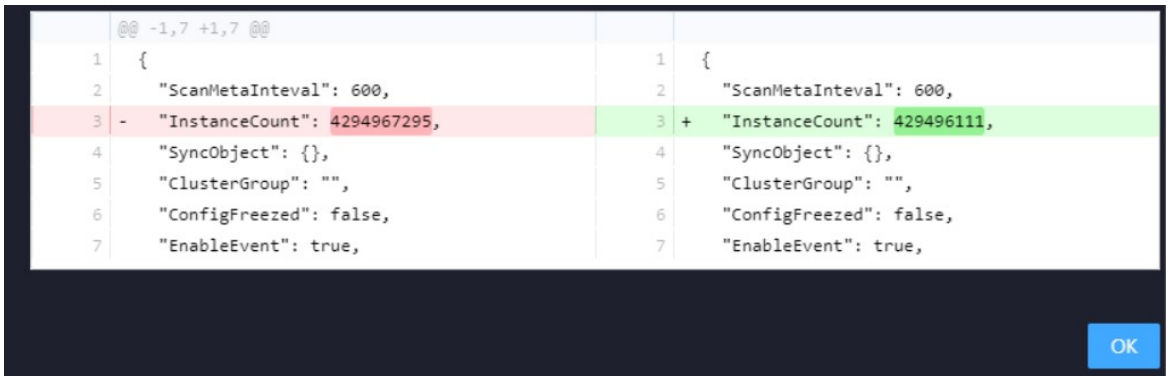

3. Click **Run**.

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 Big Data 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 required 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** panel, specify **Encryption Algorithm**, **region**, and **project**.

   > **Note** AES-CTR, AES256, RC4, and SM4 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 Create Rule message, 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** panel, 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**.
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	{
	"odps.merge.cross.paths": "true",
	"odps.idata.useragent": "odps encrypt key rotate via force mergeTask",
	"odps.merge.max.filenumber.per.job": "10000000",
	"odps.merge.max.filenumber.per.instance": "10000",
	"odps.merge.failure.handling": "any",
	"odps.merge.maintain.order.flag": "true",
	"odps.merge.smallfile.filesize.threshold": "4096",
	"odps.merge.quickmerge.flag": "true",
	"odps.merge.maxmerged.filesize.threshold": "4096",
	"odps.merge.force.rewrite": "true",
	"odps.merge.restructure.action": "hardlink"
	}

8. (Optional)View the history of 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.

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 Big Data 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. For more information about how to create a workspace, see *Create a workspace* in *DataWorks User Guide*.

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.

Procedure

1. Log on to the Apsara Big Data Manager 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. Choose **Projects > Project List**.
4. Click the name of a 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** message, click **Run**. A message appears, indicating that the permissions are granted.
7. Click **Install** in the Actions column. In the **Install Package** message, 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.

Disaster recovery

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

Entry

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

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

The **Disaster Recovery** page consists of the following elements:

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

  > **Note** The total number of projects in the primary cluster is the same as that in the standby cluster.

- **Resource Synchronization Status**: specifies whether to run the scheduled task to synchronize resources between the primary and standby MaxCompute clusters. When you turn on this switch, resources are scheduled between the primary and standby clusters every 30 minutes.
- **View Resource Synchronization History**: allows you to view the execution history of the scheduled resource synchronization task.

Prerequisites for switchover

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

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

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

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

   If the **Resource Synchronization Status** switch is turned on, the scheduled resource synchronization task is enabled. If the switch is turned off, the task is disabled.

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

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

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

1. Log on to the ABM console.

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

   Set the **NowBccApiOneIp** parameter to the IP address of any Docker container in a path of the current ABM cluster, for example, **bcc > bcc-api > controller# > #Docker#xx.xx.xx.xx**. Set the **NewBccApiOneIp** parameter to the IP address of any Docker container in the same path of the standby ABM cluster.
6. Click **Run** in the upper-right corner and confirm the risks of running the job.
7. Click **Confirm**. The job running page appears.
8. Click **Start** at the top of the page.

Start switchover

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

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

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

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

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

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

   i. Click **Restart Standby Cluster**. It takes about 20 seconds to restart the standby cluster. When the standby cluster is restarted, the value of the **MaxCompute Cluster Status** parameter changes from **Abnormal** to **Normal**.

   ii. Click **Restart Frontend Server**. It takes about 20 seconds to restart the front-end server. When the front-end server is restarted, a success message appears.

iii. Click **Test adminTask** to check whether the MaxCompute service is normal. If the test is passed, the clusters are switched. The **Next Step** button becomes operable, and the **Switching...** message disappears.

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

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

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

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

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

   You must select the projects with unsynchronized data, and click **Download Unsynchronized Data of Selected Projects** to download the data to a local device so that you can manually fill in the missing data later based on the statistics. Only after the unsynchronized data is downloaded does the **Next Step** button become operable.

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

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

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

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

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

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

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

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

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

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

12. After enabling resource replication, click **Next Step**. The Complete Wizard page appears.
13. Click **Back** in the upper-left corner. The primary and standby clusters have been switched.

6.1.1.5.2.3. Job management

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 Big Data Manager (ABM). You can also view Logview information about jobs, terminate jobs, and collect job logs.

View job snapshots

You can view job snapshots by day in the last week. The information about a job snapshot includes the job ID, project, quota group, submitter, running duration, minimum CPU utilization, and maximum CPU utilization. It also includes the 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, choose **Jobs > Job Snapshots**. The **Job Snapshots** page appears.

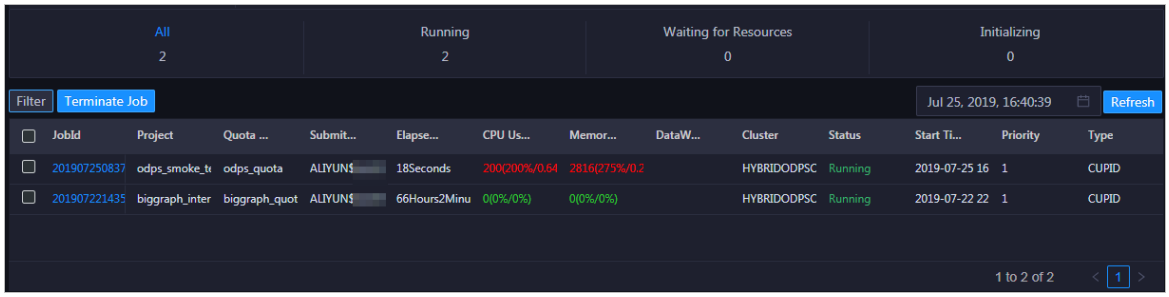

2. In the upper-right corner, select the date and time to view job snapshots by day.

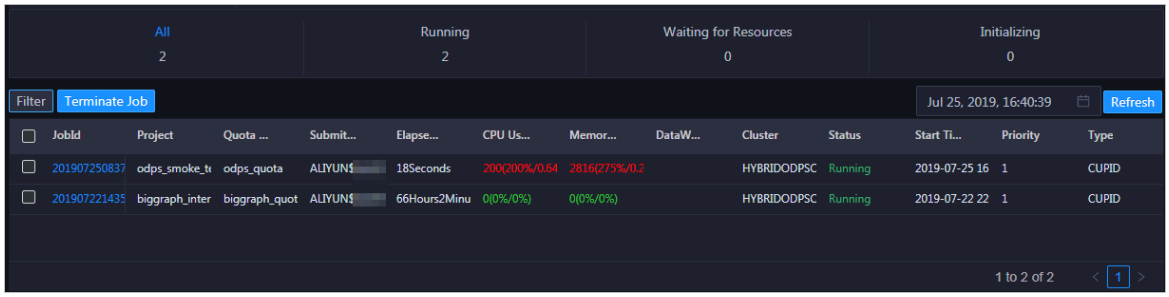

3. Click **All**, **Running**, **Waiting for Resources**, or **Initializing** to view job snapshots on the specified date.
4. Find the required snapshot and click **Logview** in the Actions column. In the dialog box that appears, click **Run** to view Logview information about the job.

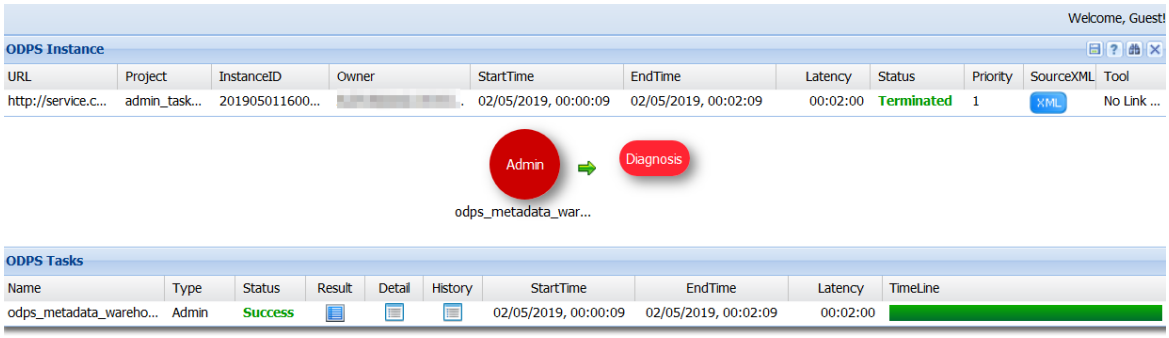

Terminate jobs

1. In the left-side navigation pane of the **Business** tab, choose **Jobs > Job Snapshots**. The **Job Snapshots** page appears.

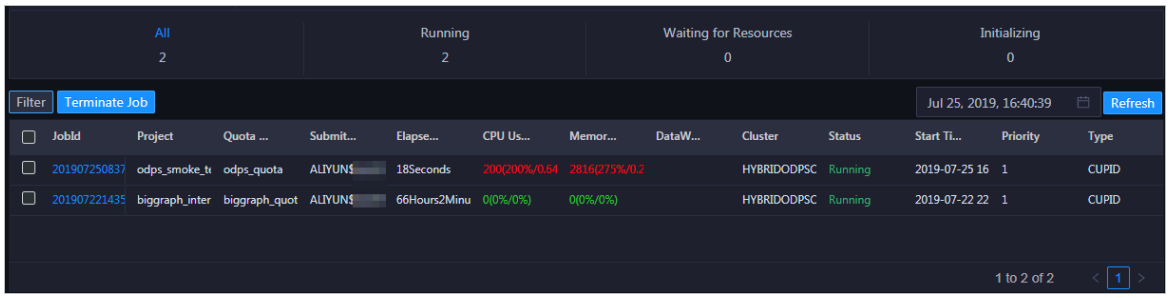

2. Select one or more jobs and click **Terminate Job** above the snapshot list. In the panel that appears, view information about the job or jobs that you want to terminate.

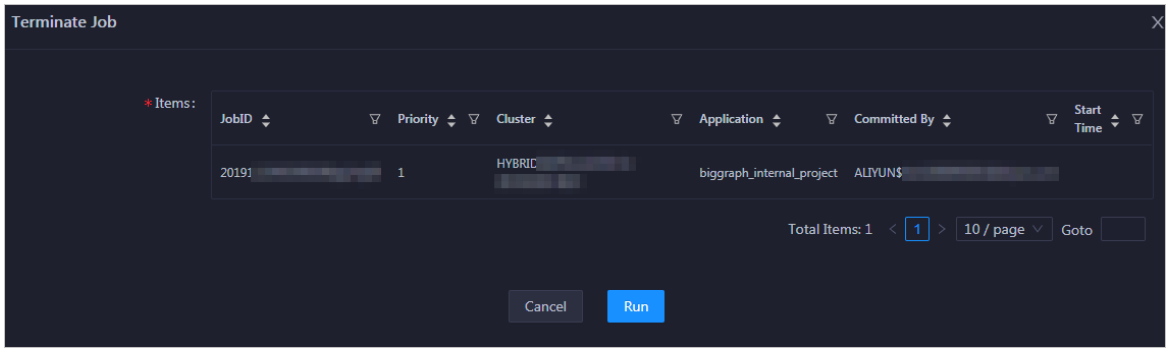

3. Click **Run**. A message appears, indicating the running result.

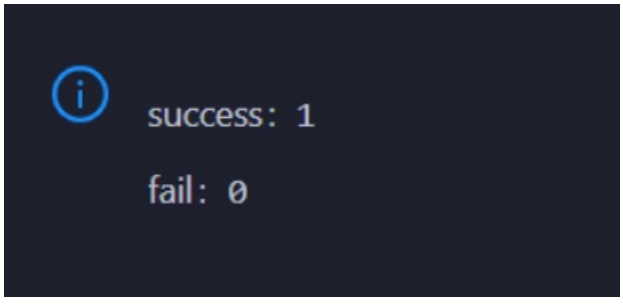

Collect job logs

If 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, choose **Jobs > Job Snapshots**. The **Job**

Snapshots page appears.

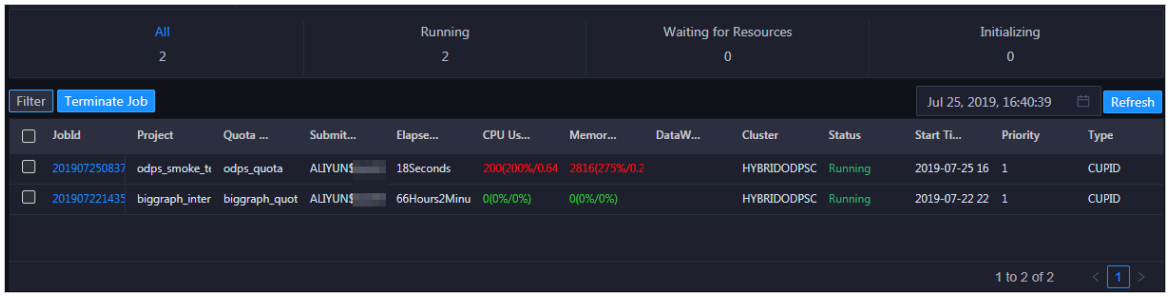

2. In the upper-right corner of the **Job Snapshots** page, choose **Actions > Collect Job Logs**.
3. In the **Collect Job Logs** panel, configure the parameters.

   The following table describes the parameters.

Parameter	Description
Target Service	The service from which you want to collect job logs.
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 **Execution History** panel, 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** panel, 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.

6.1.1.5.2.4. Business optimization

Merge small files

Excessive small files in a MaxCompute cluster occupy a lot of memory resources. Apsara Big Data 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, choose **Business Optimization > File Merging**. The **Merge Tasks** tab appears.

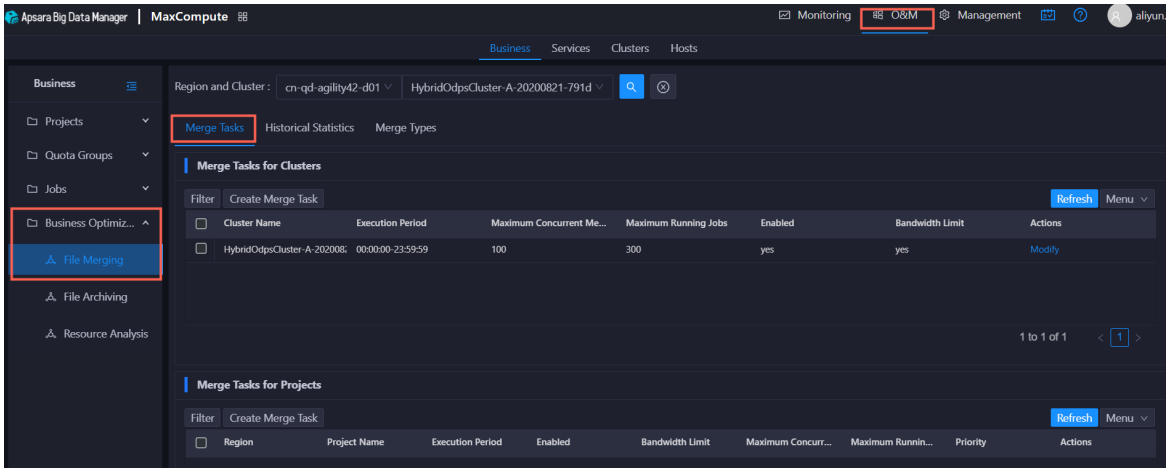

2. In the **Merge Tasks for Clusters** section, click **Create Merge Task**. In the Modify Merge Task for Cluster panel, specify the required parameters.

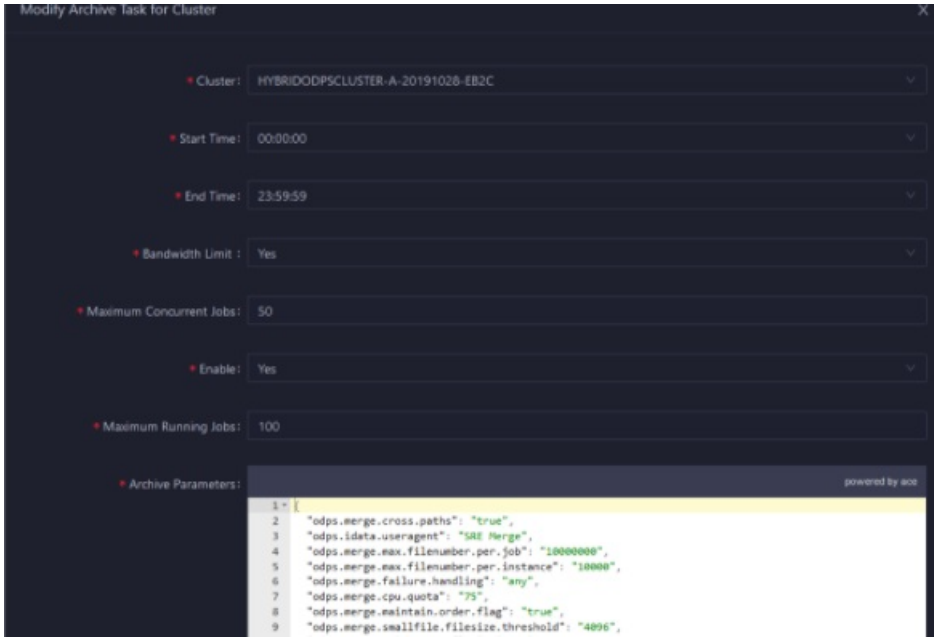

The following table describes the parameters.

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, choose **Business Optimization > File Merging**. The **Merge Tasks** tab appears.

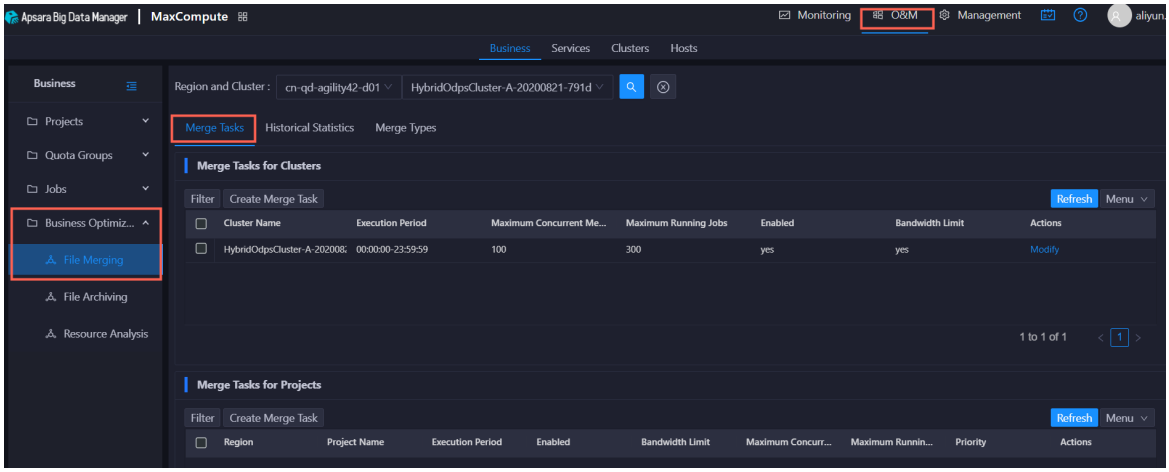

2. In the **Merge Tasks for Projects** section, click **Create Merge Task**. In the Modify Merge Task for Project panel, specify the required parameters.

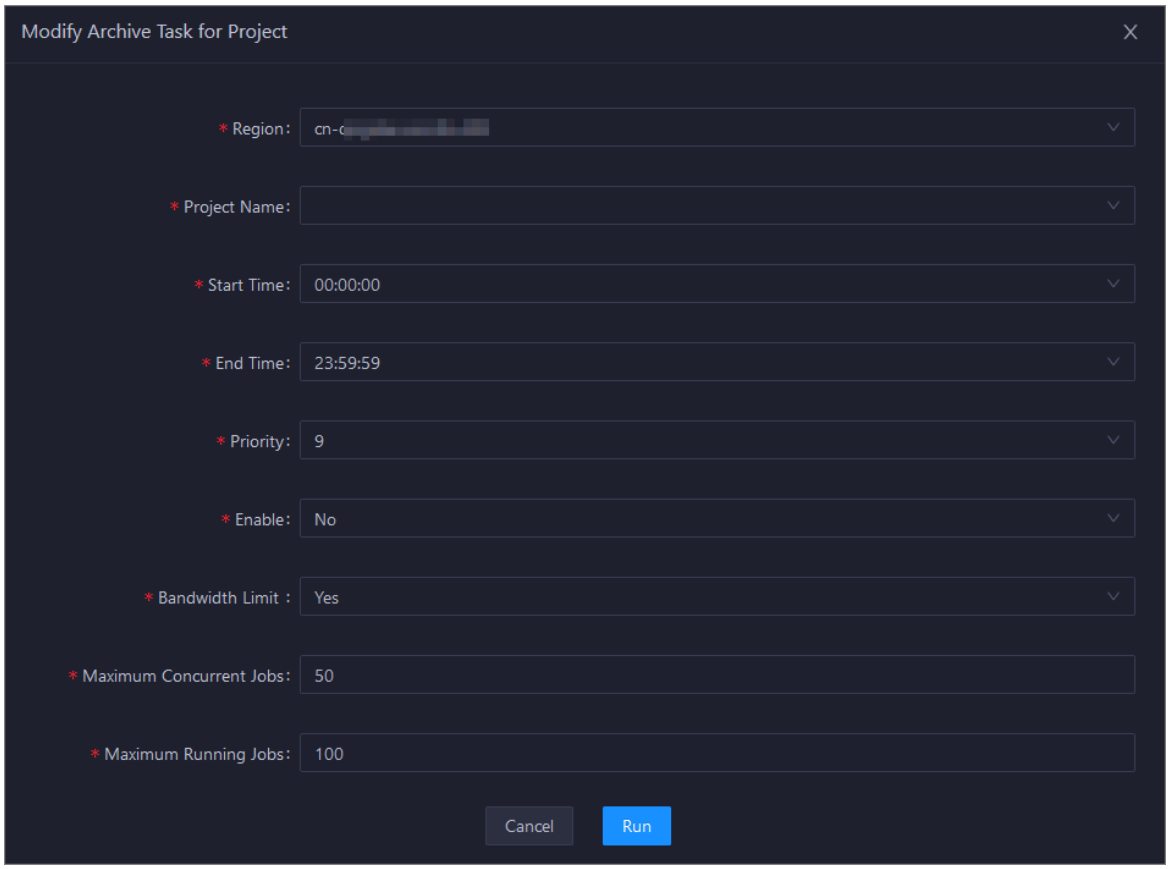

The following table describes the parameters.

Parameter	Description

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 where the selected project resides 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 where the selected project resides at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the cluster where the selected project resides, 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, choose **Business Optimization > 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.

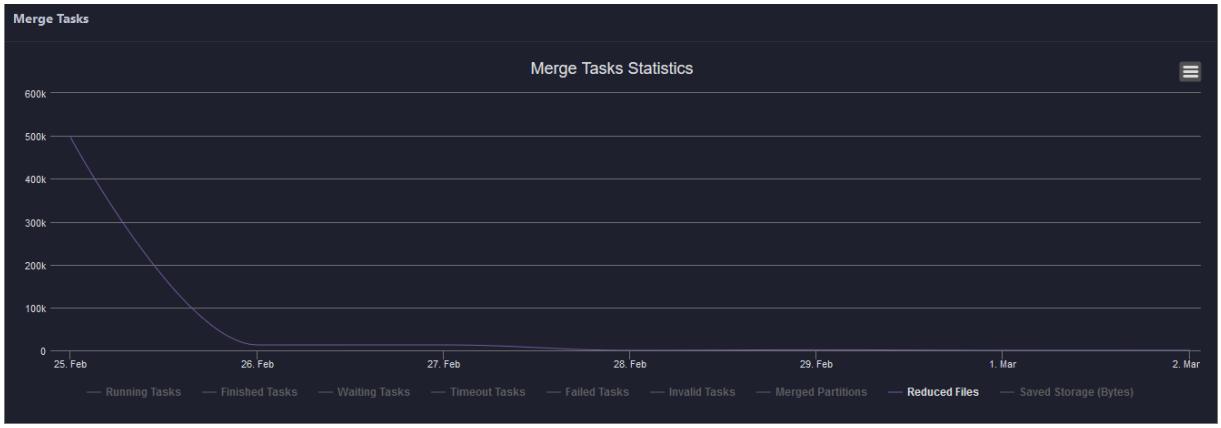

Merge Tasks for Clusters and Merge Tasks for Projects

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.

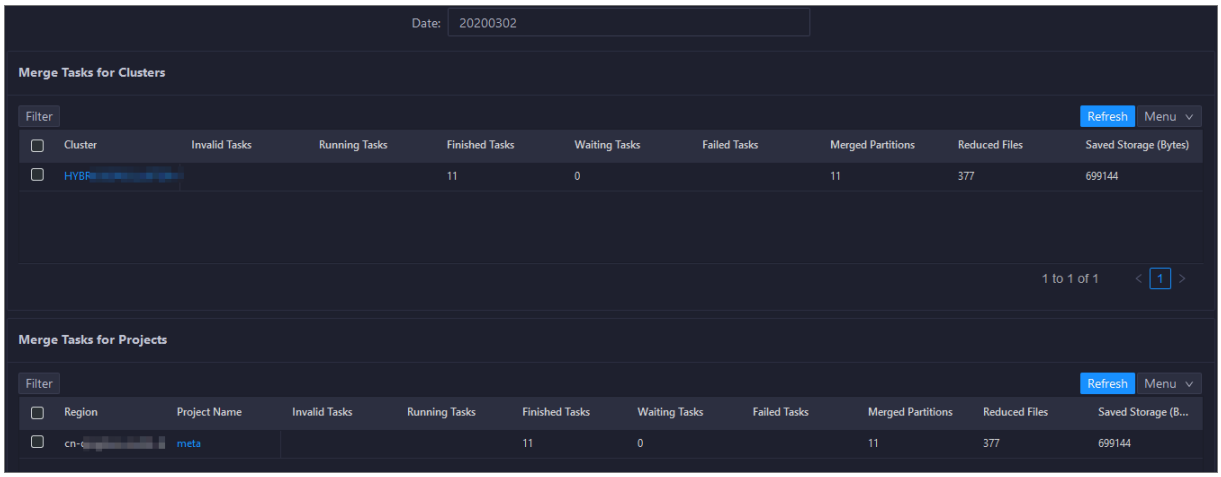

Manage merge types

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > 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 Modify Merge Type panel, specify the required parameters.

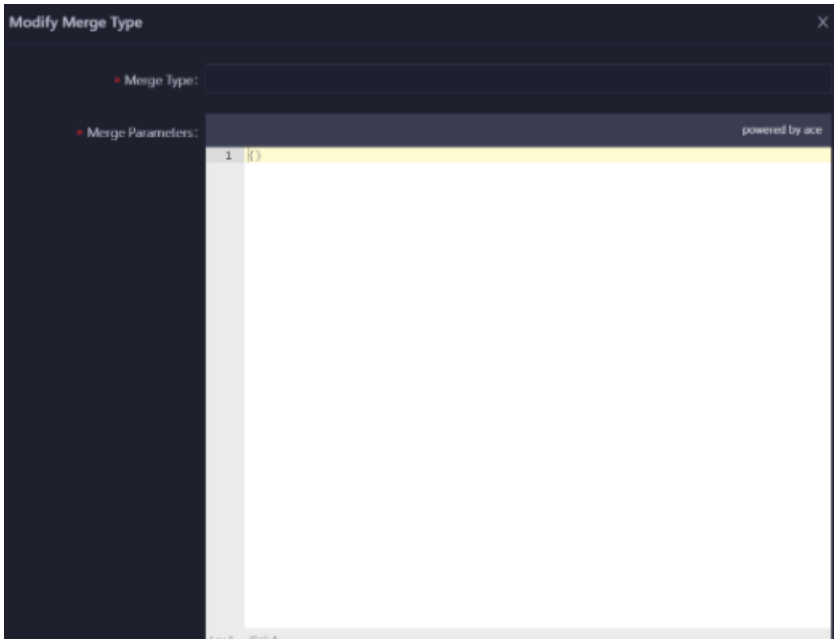

The following table describes the parameters.

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.

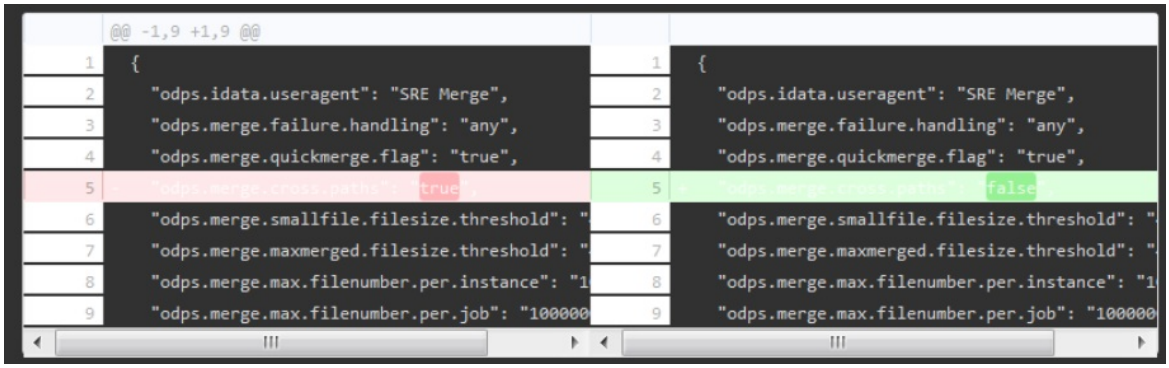

3. Click **Run**.

   The newly created merge type appears in the list of merge types.

Compress idle files

Apsara Big Data 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, choose **Business optimization > File Archiving**. The **Archive Tasks** tab appears.

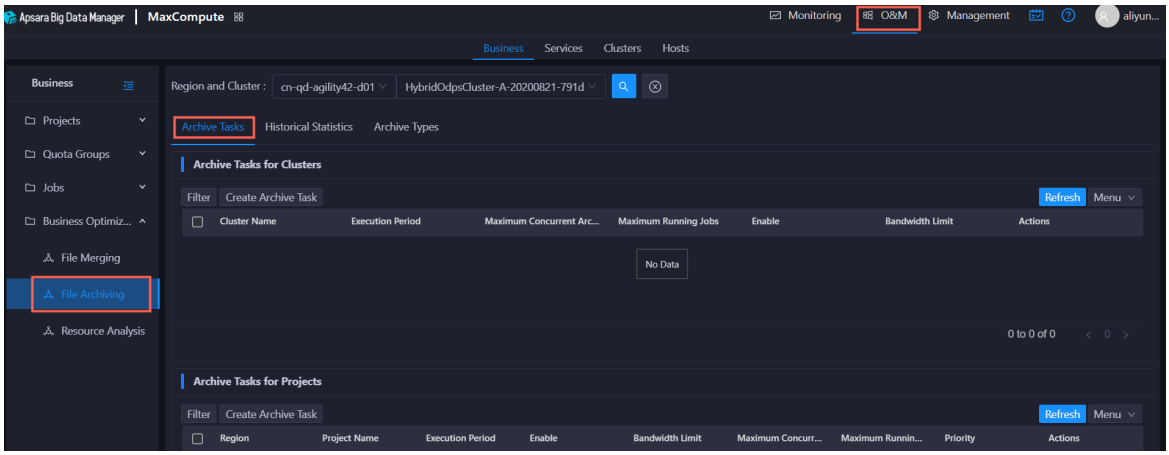

2. In the **Archive Tasks for Clusters** section, click **Create Archive Task**. In the Modify Archive Task for Cluster panel, specify the required parameters.

   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**.
Enable	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, choose **Business Optimization > File Archiving**. The **Archive Tasks** tab appears.

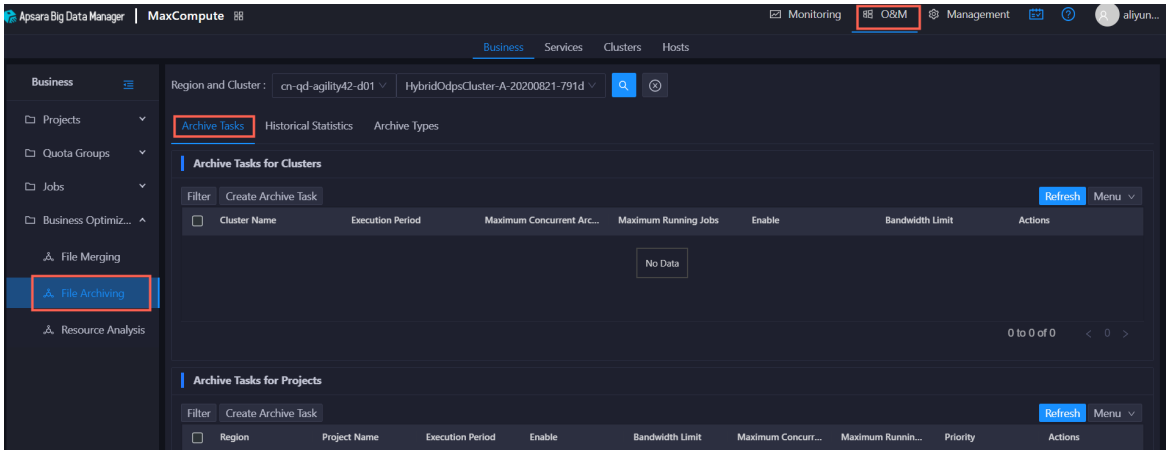

2. In the **Archive Tasks for Projects** section, click **Create Archive Task**. In the Modify Archive Task for Project panel, specify the required parameters.

   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 where the selected project resides at the same time. This parameter is valid only when **Bandwidth Limit** is set to **No**.
Enable	Specifies whether the task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run for the cluster where the selected project resides at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the cluster where the selected project resides, 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, choose **Business Optimization > 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, choose **Business Optimization > 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 Modify Archive Type panel, specify the required parameters.

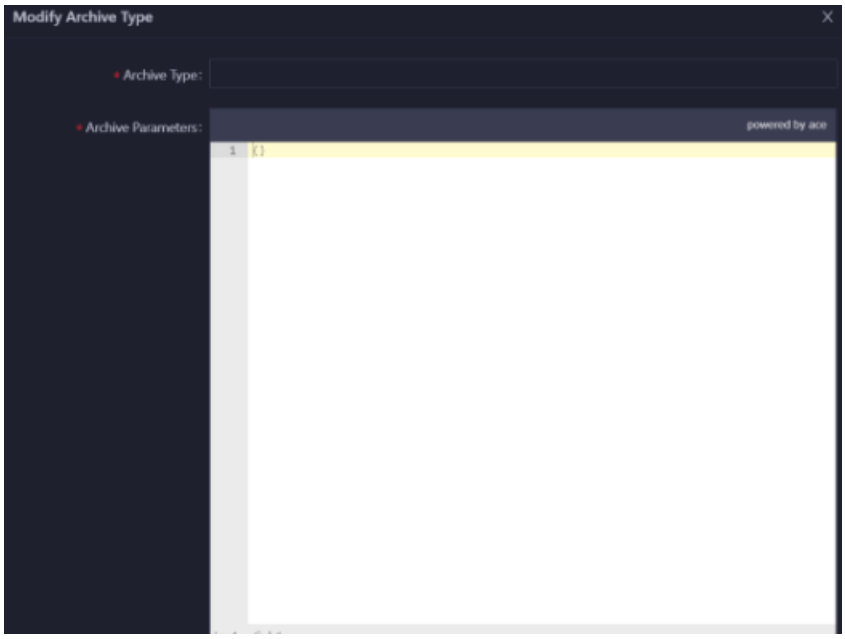

The following table describes the parameters.

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.

Analyze resources

Apsara Big Data Manager (ABM) allows you to analyze the resources for MaxCompute clusters on different tabs in the ABM console. This way, you can better understand the data storage in MaxCompute. The tabs include Tables, Projects, Tasks, Execution Time, Start Time, and Engines.

Tables

On the Tables tab, you can view the detailed information about all tables in each project, including Partitions, Storage Usage (GB), Pangu File Count, Partitions Ranking, Storage Usage Ranking, and Pangu File Count Ranking. You can sort tables by partition quantity, physical storage usage, and file quantity of Apsara Distributed File System.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. The **Tables** tab appears.

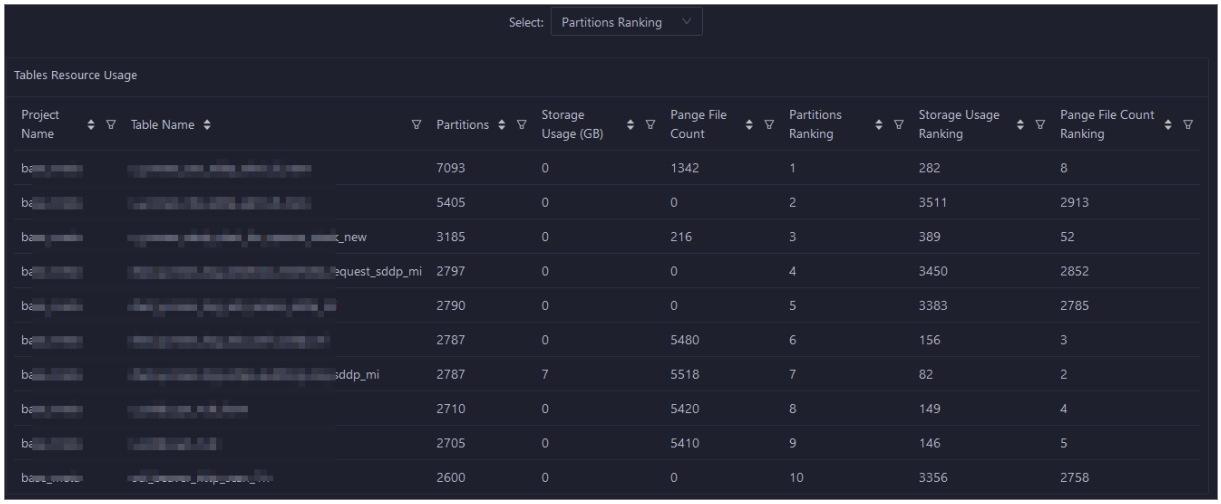

Project Name	Table Name	Partitions	Storage Usage (GB)	Pange File Count	Partitions Ranking	Storage Usage Ranking	Pange File Count Ranking
ba		7093	0	1342	1	282	8
ba		5405	0	0	2	3511	2913
ba	_new	3185	0	216	3	389	52
ba	equest_sddp_mi	2797	0	0	4	3450	2852
ba		2790	0	0	5	3383	2785
ba		2787	0	5480	6	156	3
ba	sddp_mi	2787	7	5518	7	82	2
ba		2710	0	5420	8	149	4
ba		2705	0	5410	9	146	5
base_		2600	0	0	10	3356	2758

Projects

On the Projects tab, you can view the detailed information about storage for each project, including Pangu File Count, Storage Usage (GB), CU Usage, Total Memory Usage, Tasks, Tables, Idle Storage, and daily and weekly increases in percentage of these items.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Projects** tab.

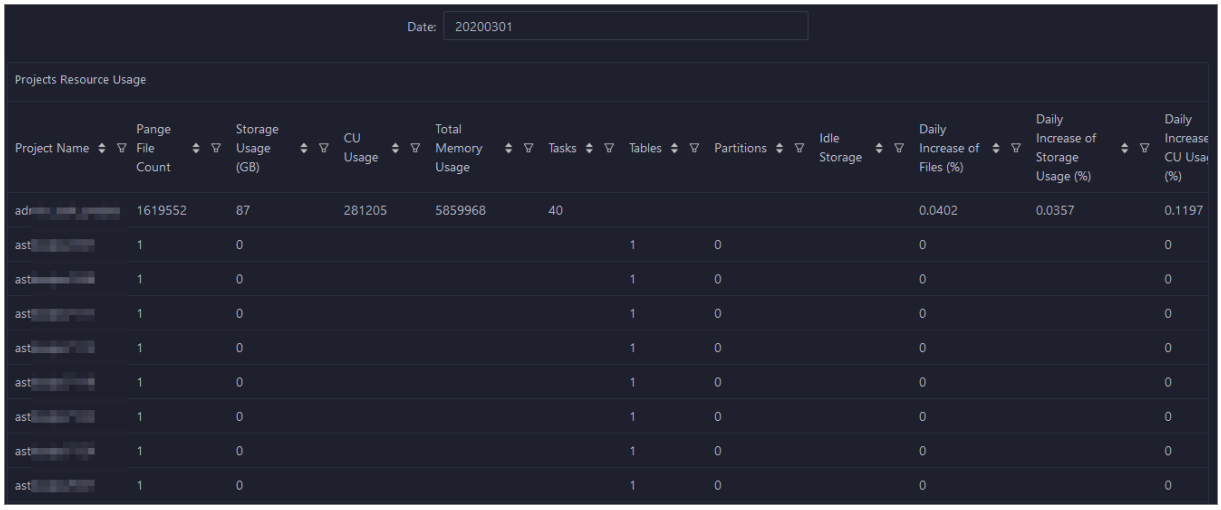

Tasks

On the Tasks tab, you can view the detailed information about all tasks in each project, including instanceid, Status, CU Usage, Start Time, End Time, Execution Time (s), CU Usage Ranking, and SQL Statements.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Tasks** tab.

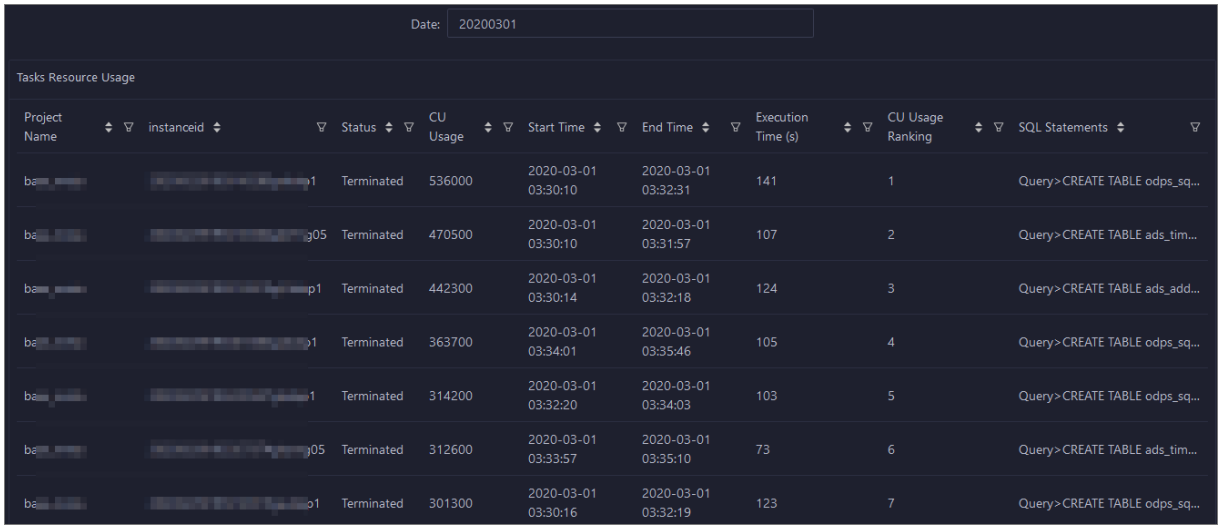

Execution Time

On the Execution Time tab, you can view the numbers of tasks whose execution time is within different time ranges in each project. The metrics include Less than 5 Minutes, Less than 15 Minutes, Less than 30 Minutes, Less than 60 Minutes, and More than 60 Minutes. The Execution Time chart displays the trend lines of task quantity in different colors by day.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Execution Time** tab.

Start Time

On the Start Time tab, you can view the numbers of tasks started in different time periods for each project. The time interval is 30 minutes. The Tasks chart displays the trend line of the number of tasks started in a specified time period by day.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Start Time** tab.

Engines

On the Engines tab, you can view the trend lines of performance statistics of tasks in each project in the Task Performance Analysis chart. The performance metrics include cost_cpu, cost_mem, cost_time, input_bytes, input_bytes_per_cu, input_records, input_records_per_cu, output_bytes, output_bytes_per_cu, output_records, and output_records_per_cu.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Engines** tab.

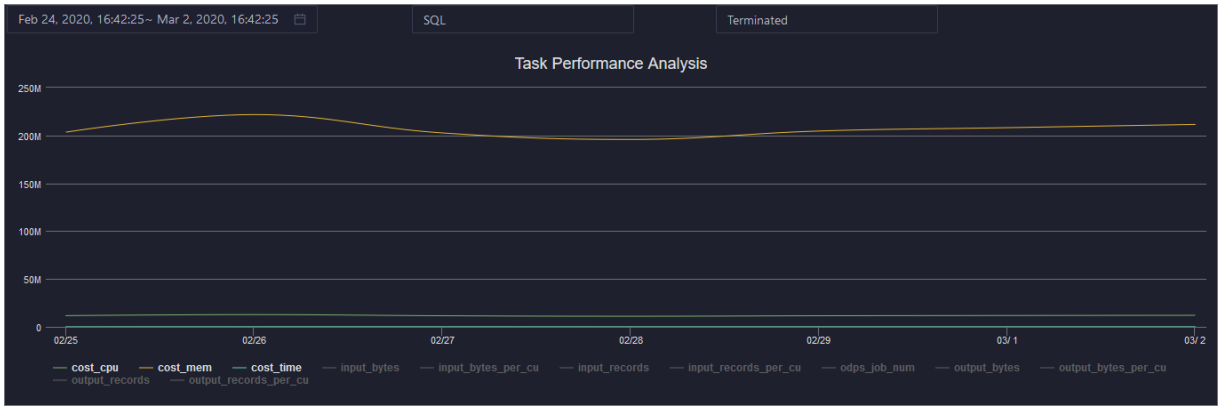

6.1.1.5.3. Service O&M

6.1.1.5.3.1. Control service O&M

O&M features 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 cause of 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 Big Data 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.

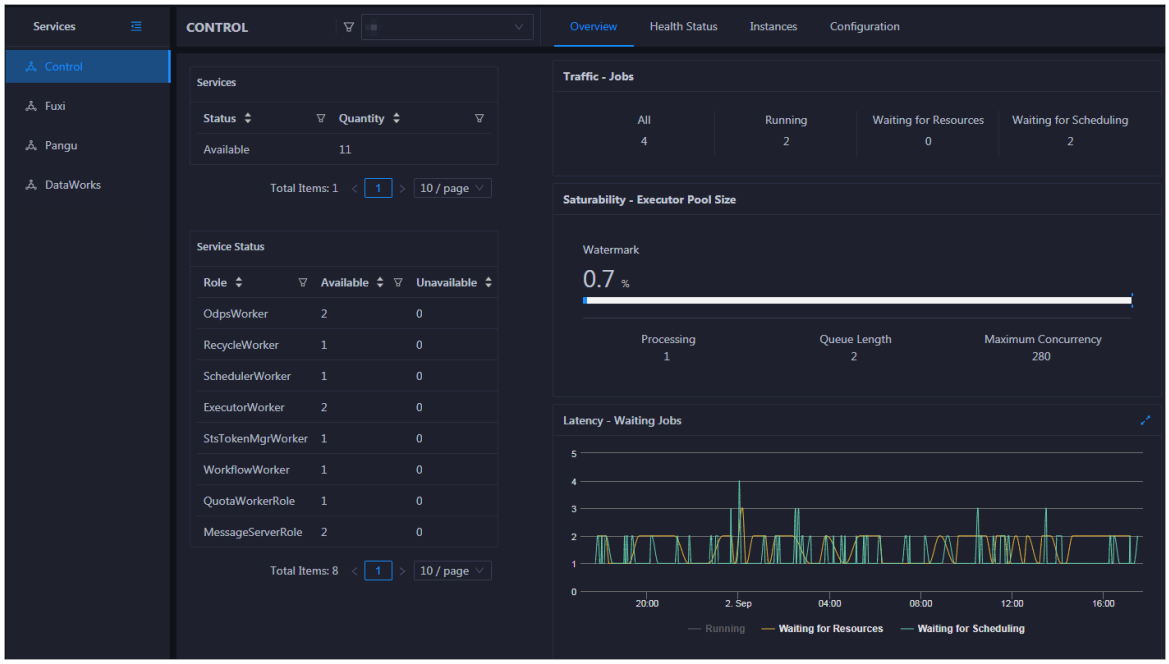

Control service overview

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.

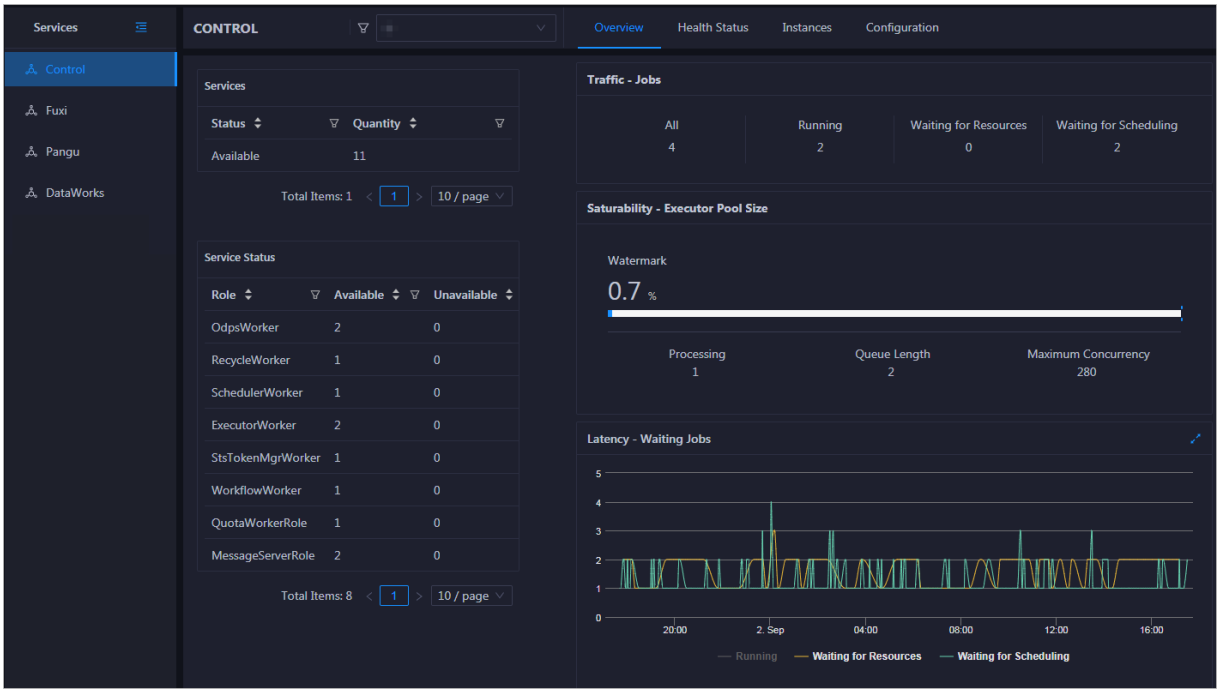

On the **Overview** page, you can view the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

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.

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.

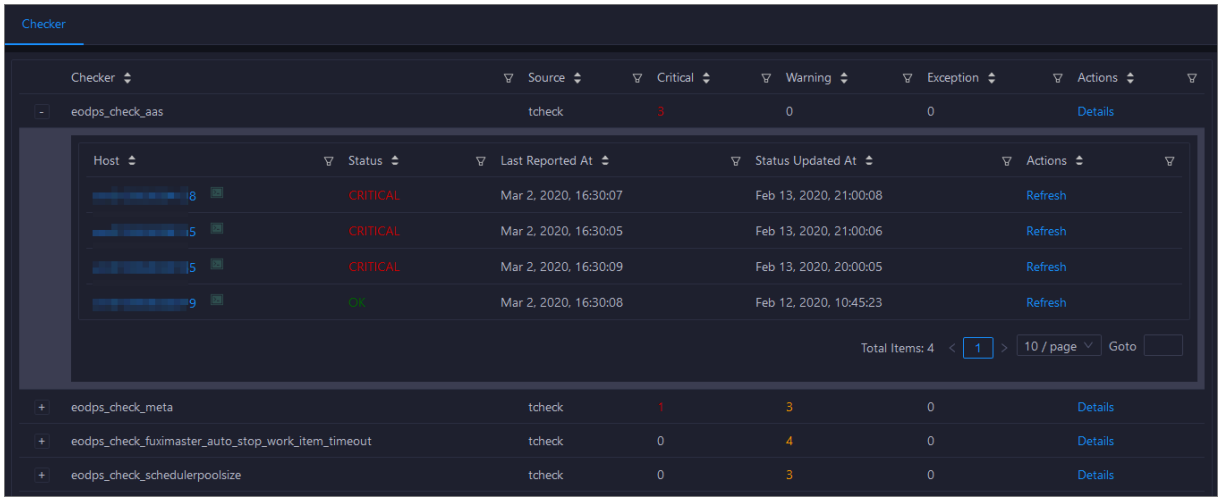

On the **Health Status** page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical**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

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.

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.

Metadata warehouse for the control service

This topic describes how 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.

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

In the left-side navigation pane of the **Services** tab, click **Control**. On the page that appears, click the **Metadata Repository** tab.

The **Metadata Repository** tab displays the completion time of the output tasks of the metadata warehouse and the trend chart of the consumed time for running tasks. The time displayed in the **Completed At** column indicates the time when an output task is complete. The time displayed in the **Collected At** column indicates the last time at which ABM obtains the status of output tasks.

Stop or start a server role

Apsara Big Data 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, choose **Actions** > **Stop Service Role**.
5. In the Stop Service Role panel, select a server role that 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 panel 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, choose **Actions** > **Start Service Role**.
5. In the Start Service Role panel, select a server role that 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 panel 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 panel, click **Details** in the Details column of the task to view the details.
2. In the Start Service Role panel, 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.

Start AdminConsole

AdminConsole is a management platform of MaxCompute. It is disabled by default. Apsara Big Data 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 page that appears, choose **Actions** > **Start Admin Console**.
6. In the **Start Admin Console** panel, 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** 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 **Start Admin Console** in the upper-right corner to view the execution history.
2. In the Execution History panel, 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.

Change the number of managed service roles

Apsara Big Data Manager (ABM) allows you to change the number of managed service roles in MaxCompute. If nodes of a managed service role are heavily loaded or a managed node is added, you can use this feature to change the number of managed service roles.

Procedure

1. In the left-side navigation pane of the Services tab, click Control.
2. In the upper-right corner of the tab that appears, choose Actions > Change Managed Service Role Amount.

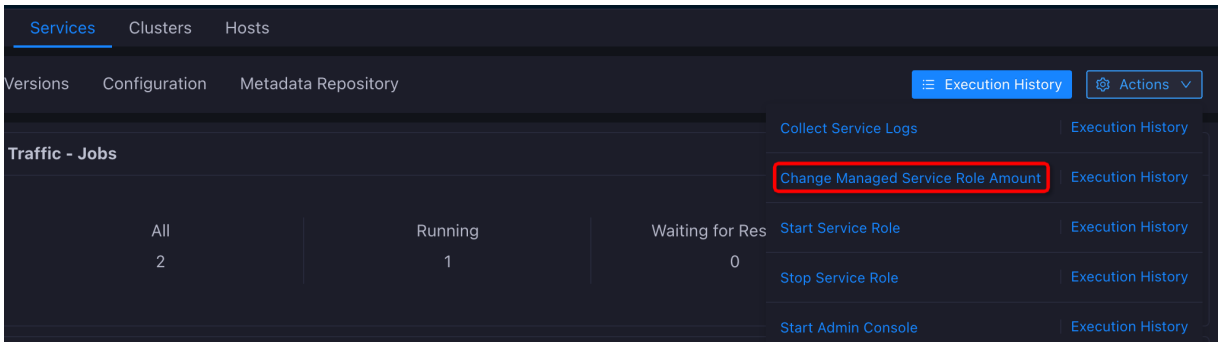

3.In the Change Managed Service Role Amount dialog box, configure the number of roles for each service.

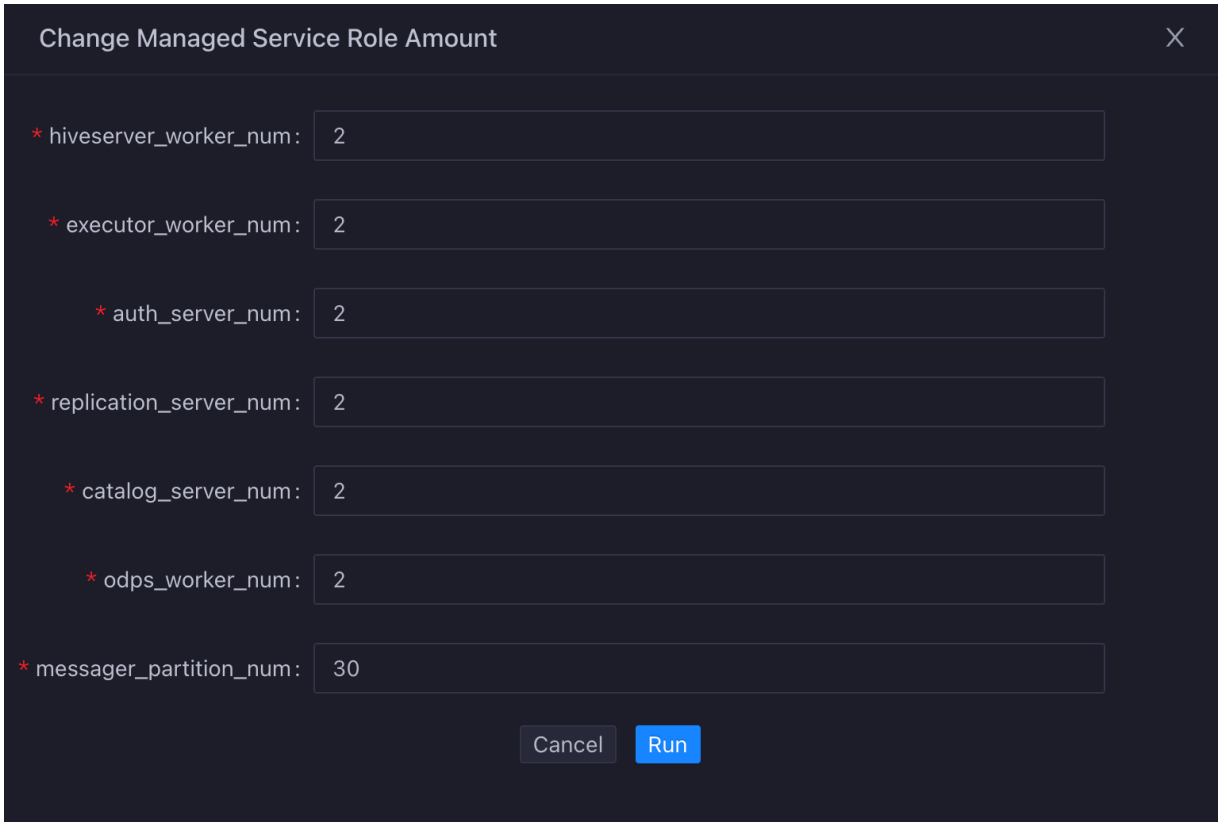

4. Click Run. A message appears, which indicates that the configurations have been submitted.

After you change the number of managed service roles, you must restart the services in MaxCompute for the change to take effect. For more information, see Stop or start a server role.

Collect service logs

Apsara Big Data 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 page that appears, choose **Actions > Collect Service Logs**.
6. In the **Collect Service Logs** panel, specify the required parameters.

   The following table describes the parameters.

Parameter	Description

Parameter	Description
Target Service	The service from which you want to collect service logs. Select a 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 panel, 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.

6.1.1.5.3.2. Job Scheduler O&M

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: displays 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: displays 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: displays information about the master nodes and server roles of Job Scheduler and allows you to restart the master nodes.

- Compute Nodes: displays 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 Big Data 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.

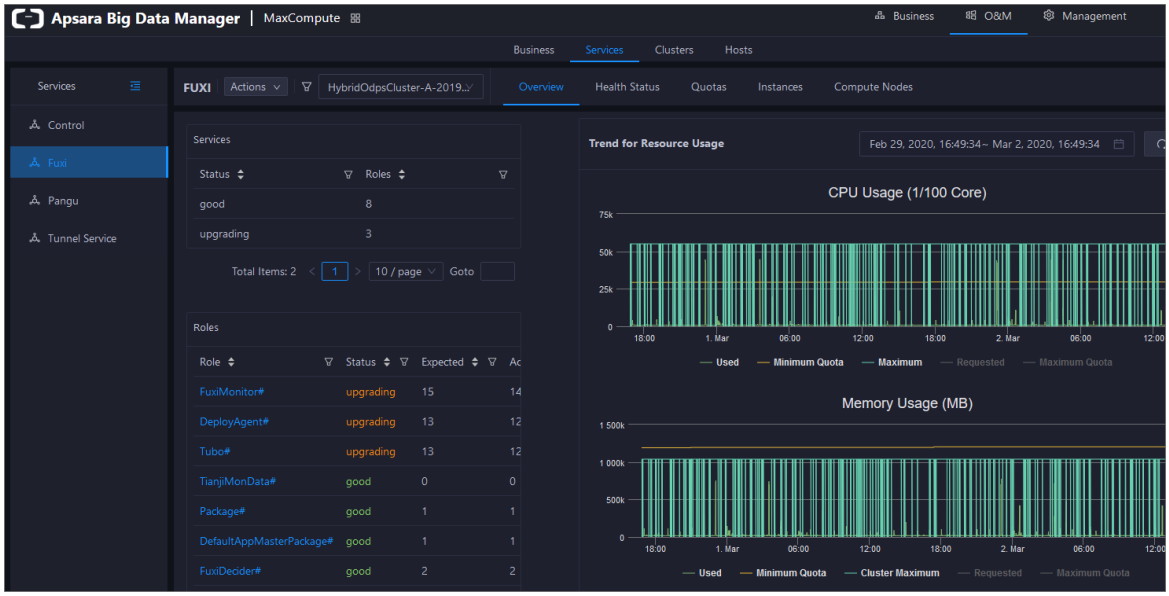

Overview

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.

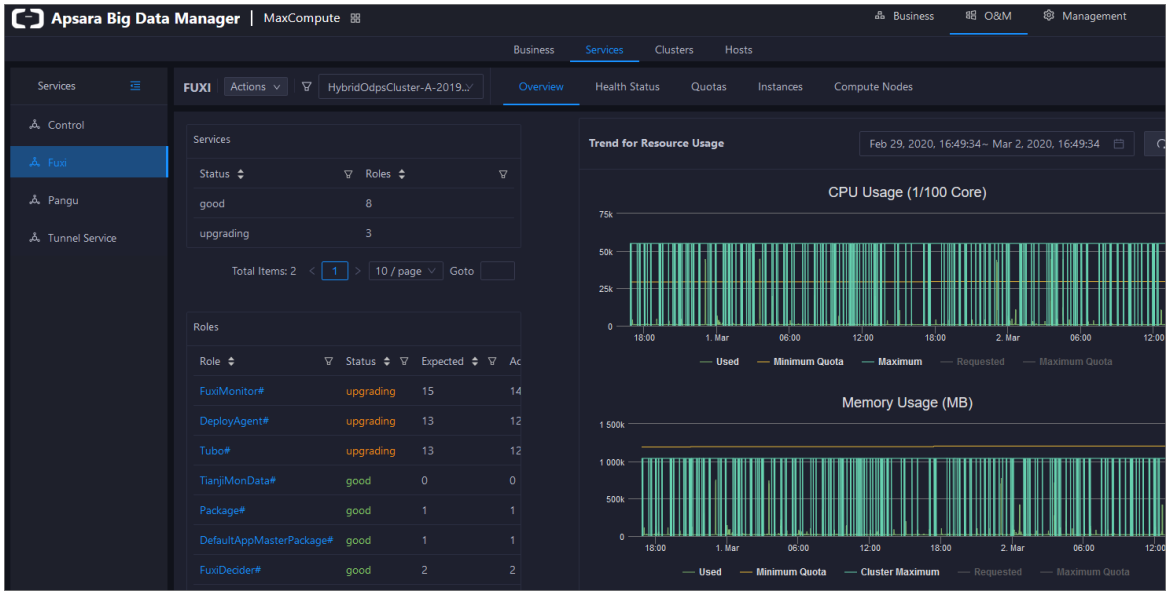

The **Overview** tab shows the key operating information of Job Scheduler. The information includes the service overview, service status, resource usage, compute node overview, and the trend charts of CPU utilization and memory usage.

Services

This section shows the numbers of available services, unavailable services, and services that are being updated.

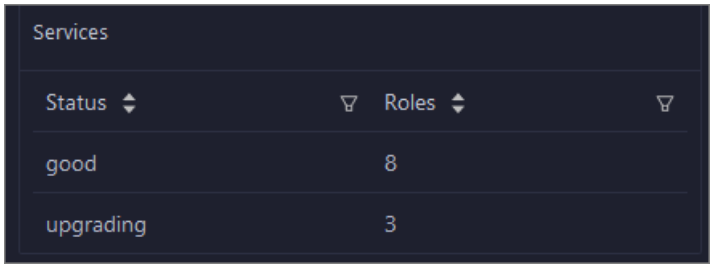

Roles

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.

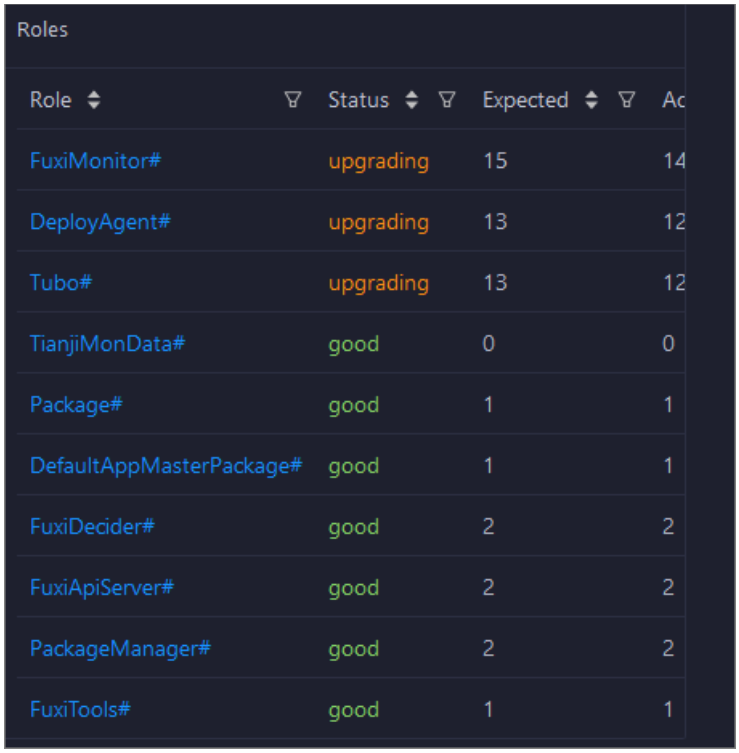
Roles

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

Click the name of a 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.

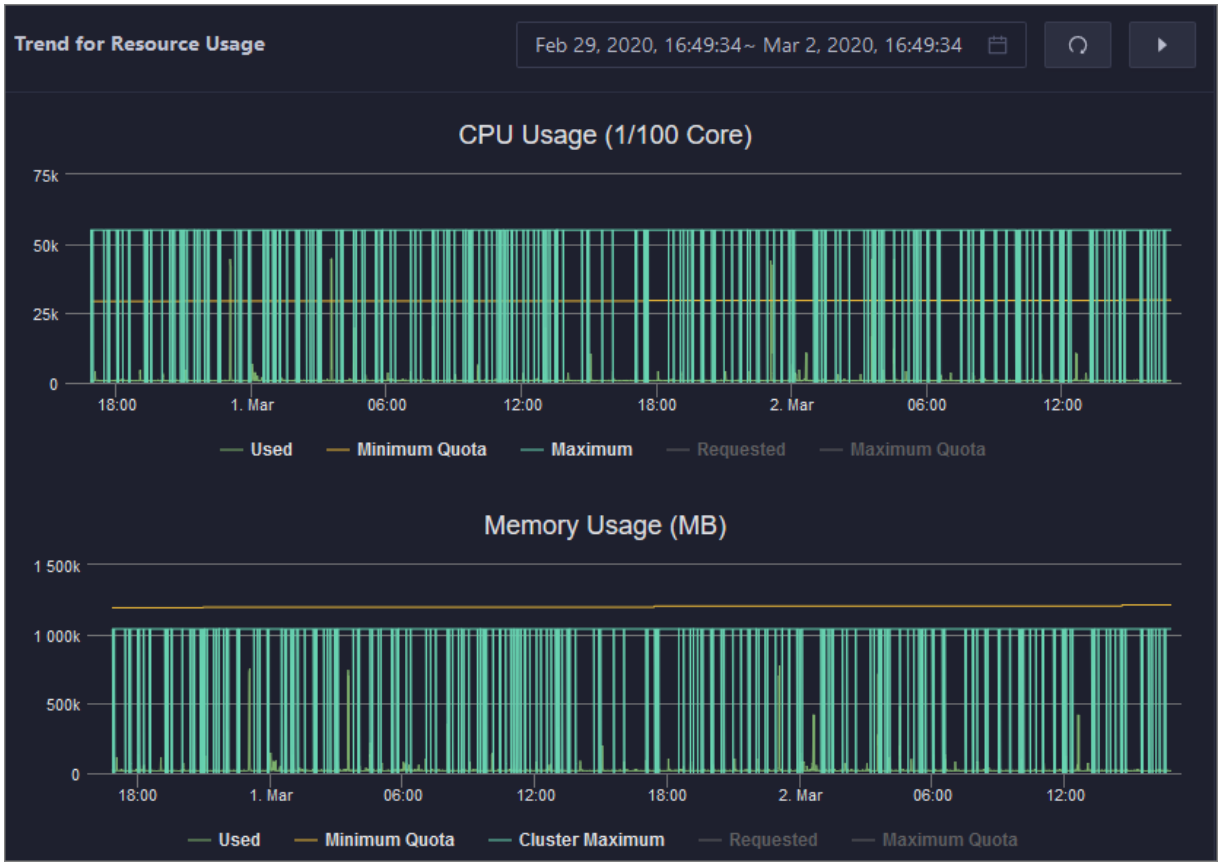

Saturability - Resource Usage

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.

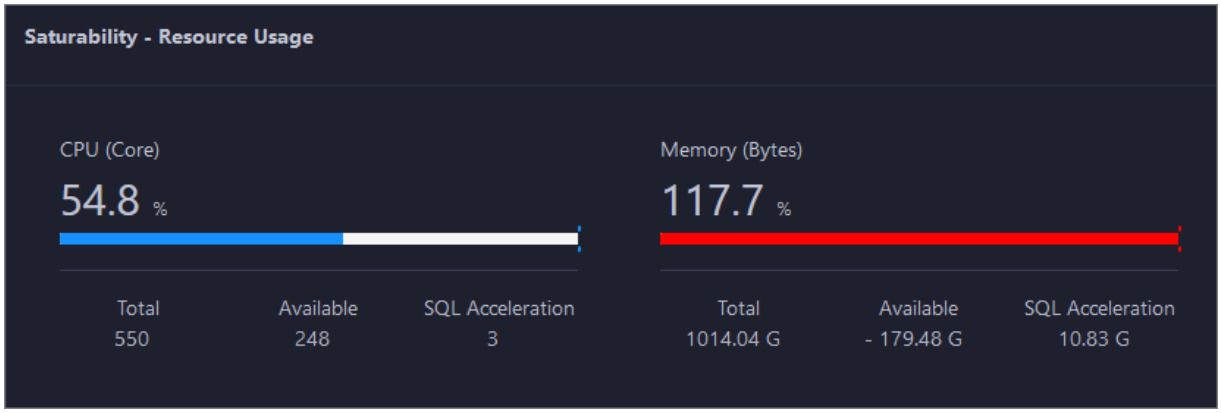

Compute Nodes

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.

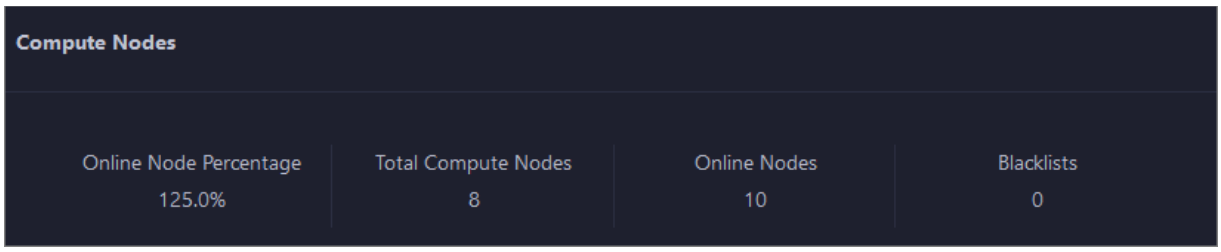

Job Scheduler health

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.

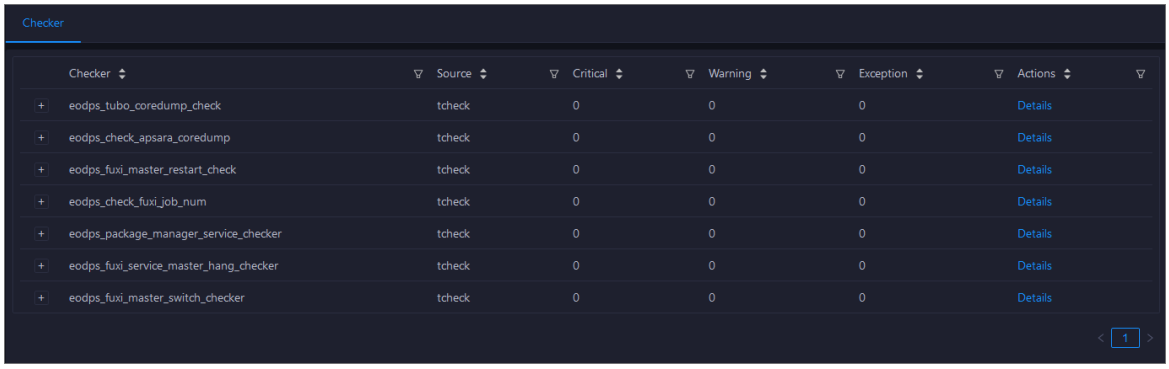

On the **Health Status** page, you can view all checkers of the Job Scheduler service and the check results for all hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

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.

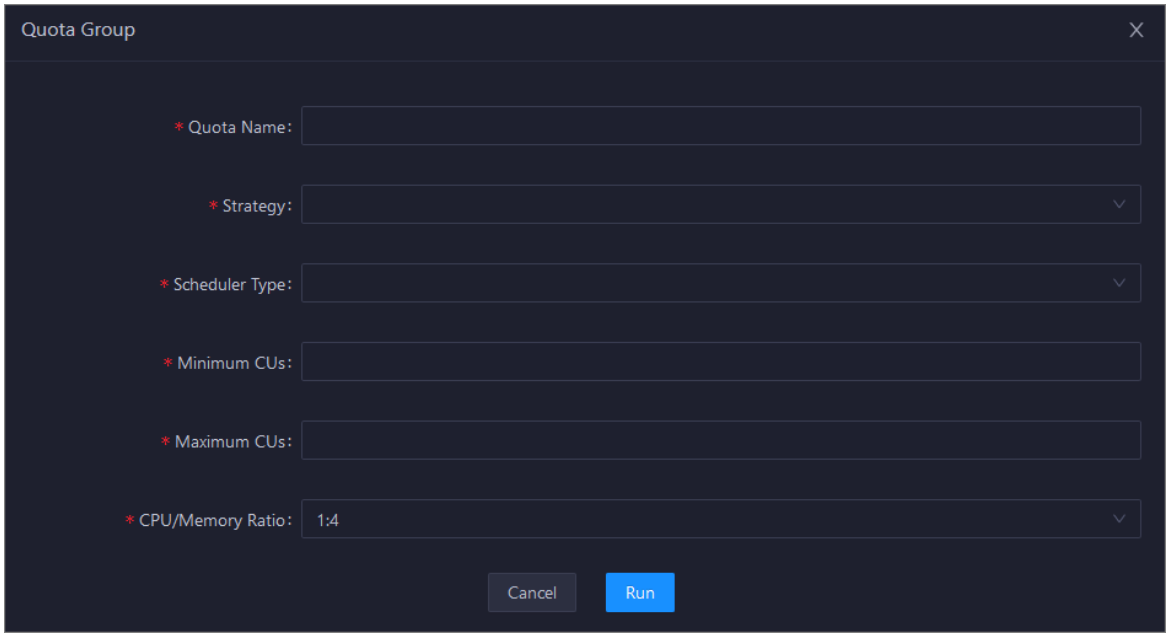

3. Click **Run**.

   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

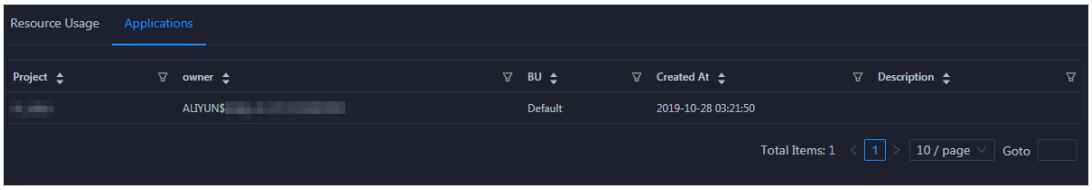

Modify a quota group

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.

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.

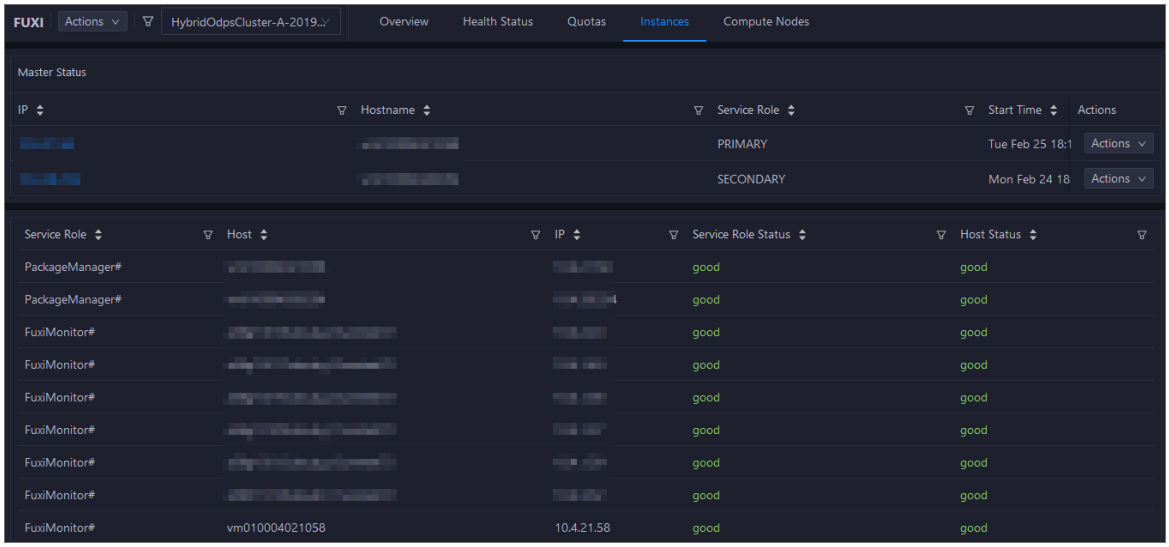

The **Instances** tab shows information about the master nodes and server roles of Job Scheduler. The information about the master nodes includes the IP address, hostname, server role, and start time. The information about a server role includes the role name, hostname, role status, and host status.

Supported operations

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

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.

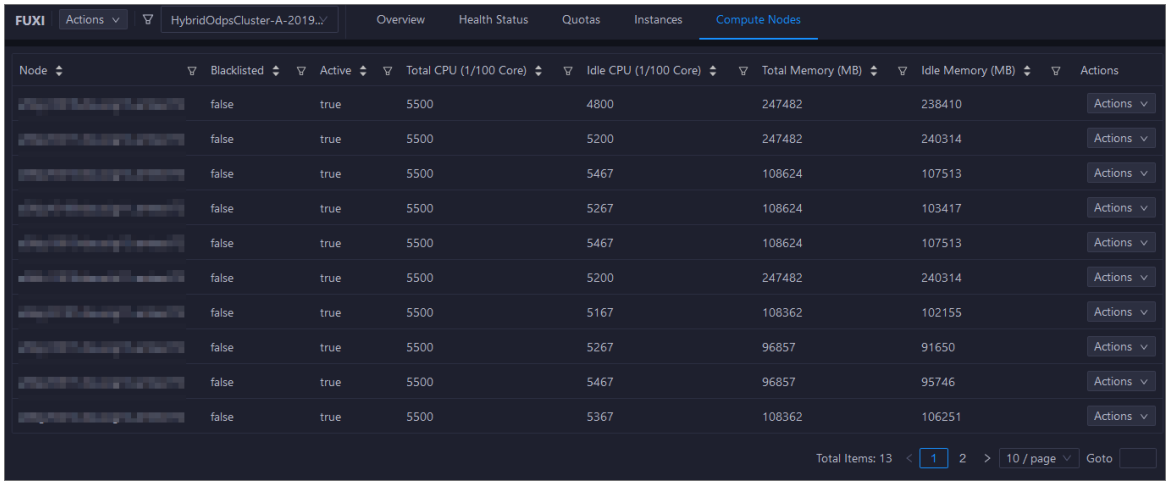

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including the total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active.

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.

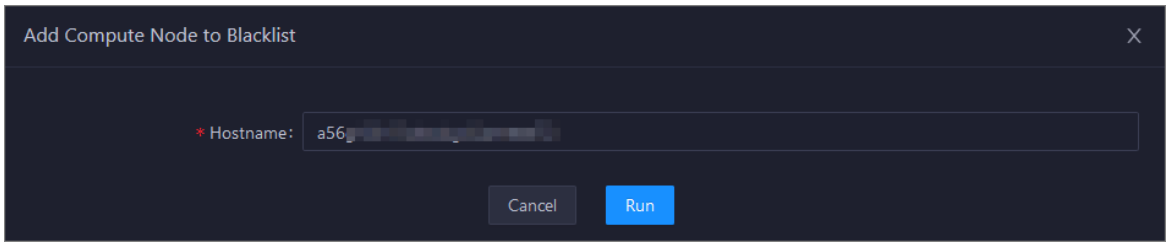

The value of the **Hostname** parameter is automatically filled. You do not need to specify a value for this parameter.

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

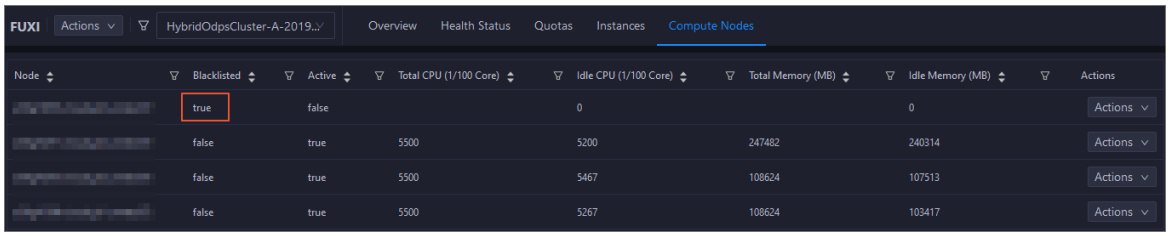

Enable and disable SQL acceleration

You can enable or disable SQL acceleration for Job Scheduler in the Apsara Big Data 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, choose **Actions** > **Enable SQL Acceleration**.
3. In the Enable SQL Acceleration panel, set the **WorkerSpans** parameter.

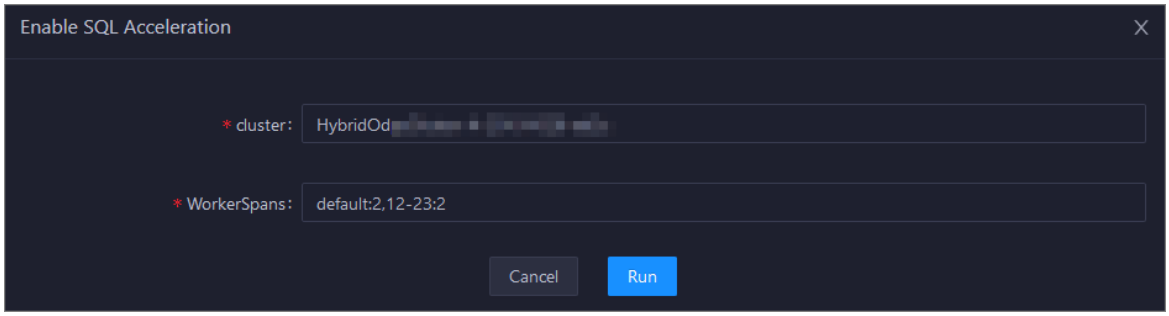

WorkerSpans: the default resource quota of the cluster and the resource quota for a specific period. Default value: **default:2,12-23:2**.

> **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, choose **Actions** > **Disable SQL Acceleration**.
3. In the Disable SQL Acceleration panel, 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 Execution History panel, view the execution history of enabling SQL acceleration.

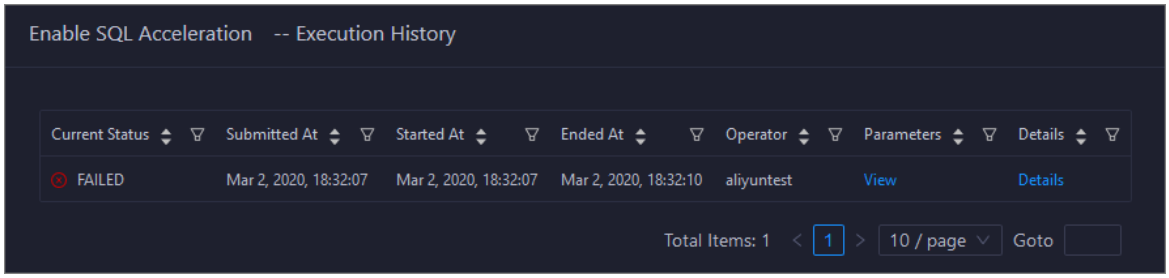

The execution history shows the current status, submission time, start time, end time, and operator of each execution.

4. If the execution fails, click **Details** in the Details column to identify the cause of the failure.

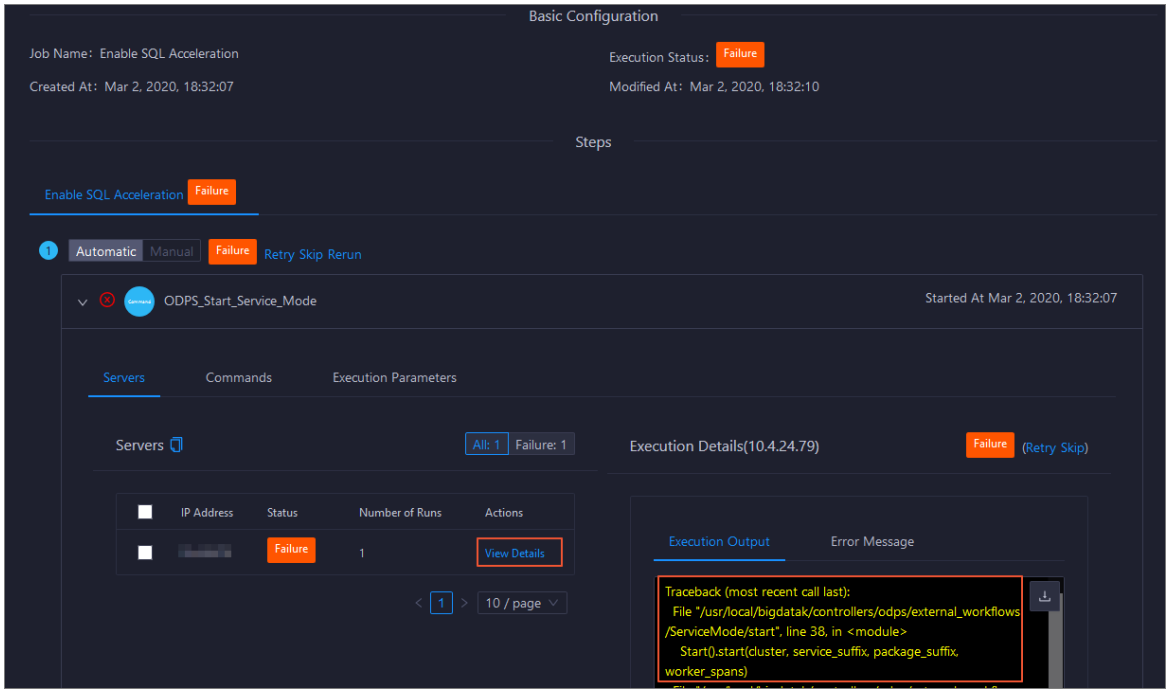

Restart a master node of Job Scheduler

Job Scheduler is the resource management and task scheduling system of the Apsara distributed operating system. Apsara Big Data 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-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**. Then, click the **Instances** tab.
5. On the **Instances** tab, choose **Actions** > **Restart Fuxi Master Node** in the Actions column of a primary or secondary master node.
6. In the **Restart Fuxi Master Node** panel, click **Run**. The **Restart Fuxi Master Node** panel appears.

Step 2: View the execution status or progress

1. In the **Restart Fuxi Master Node** panel, check the execution history of restarting master nodes.

   The **Restart Fuxi Master Node** panel displays 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) 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. In the **Restart Fuxi Master Node** panel, 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.

6.1.1.5.3.3. Apsara Distribute File System O&M

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 Big Data 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.

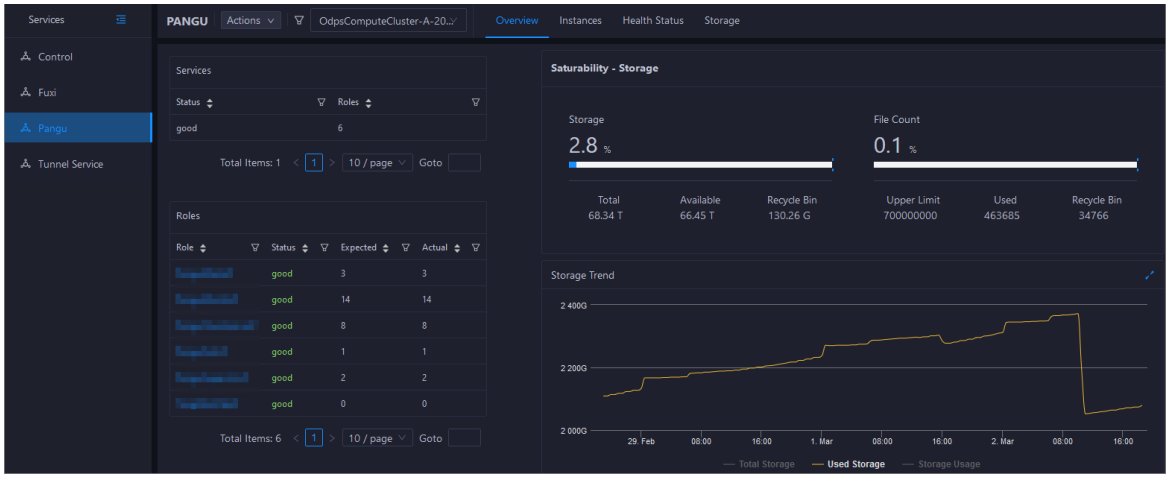

Overview

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.

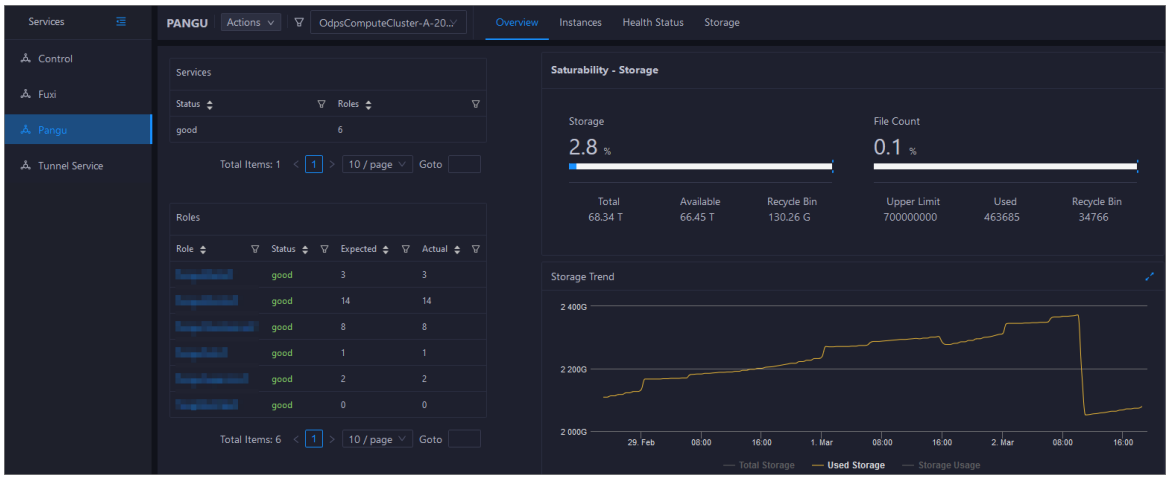

The **Overview** tab shows the key operating information about Apsara Distributed File System. The information includes the service overview, service status, health check result, health check history, storage usage, storage node overview, and the trend charts of storage usage and file count.

Services

This section shows the status of Apsara Distributed File System and the number of server roles.

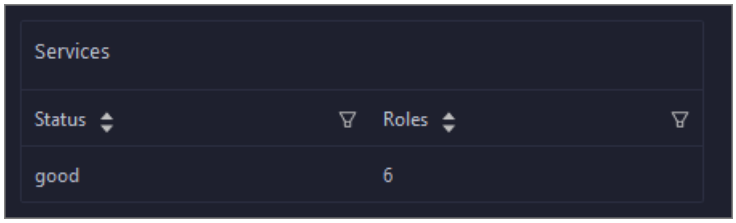

Roles

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.

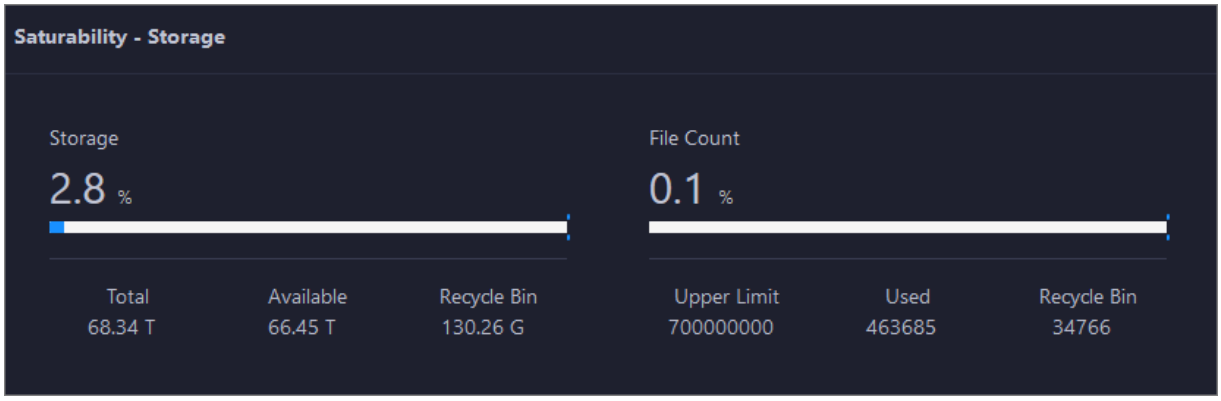

Storage Trend and File Count Trend

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.

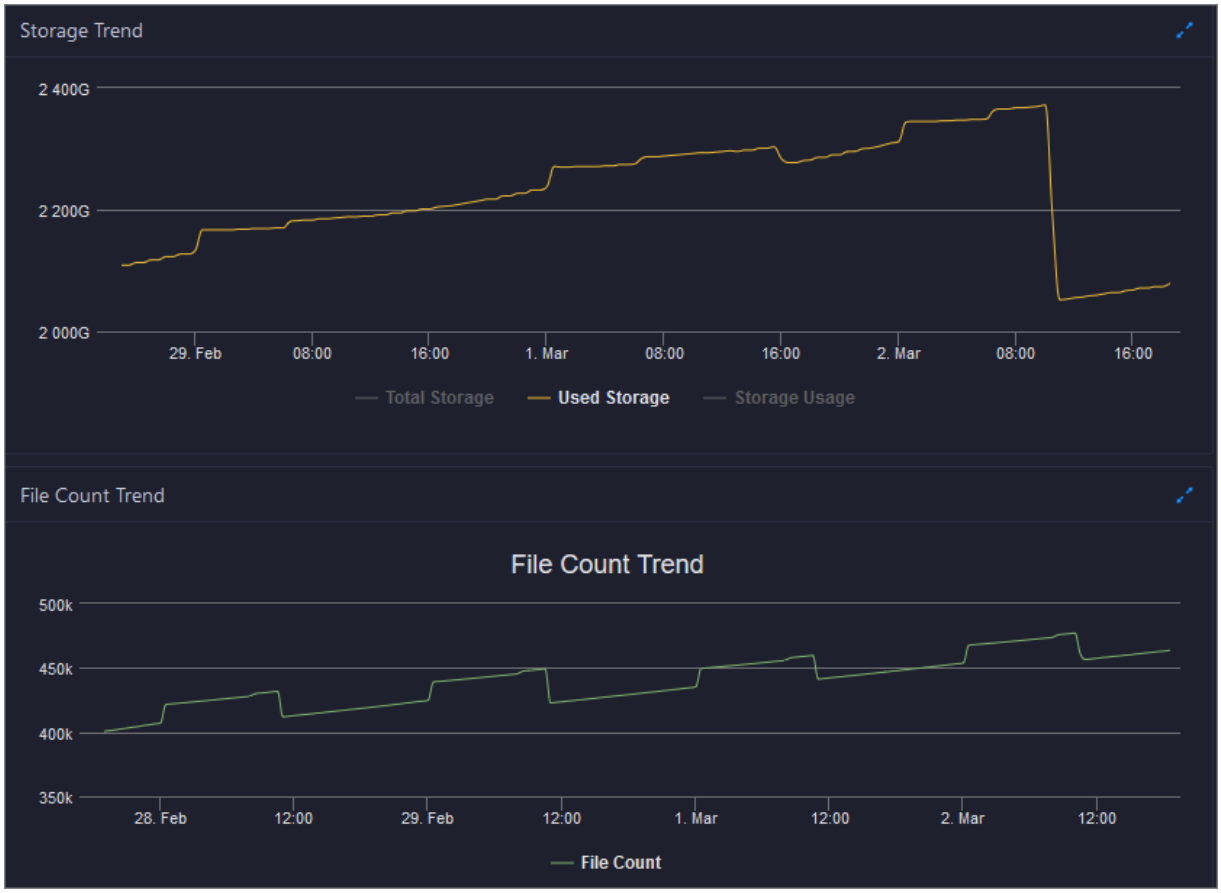

In the upper-right corner of the chart, click the icon to zoom in the chart. The following figure shows an enlarged chart of storage usage.

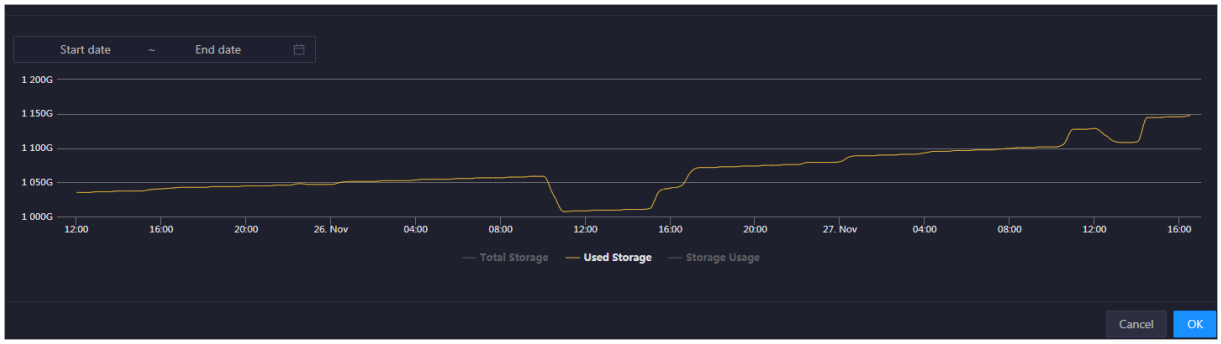

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

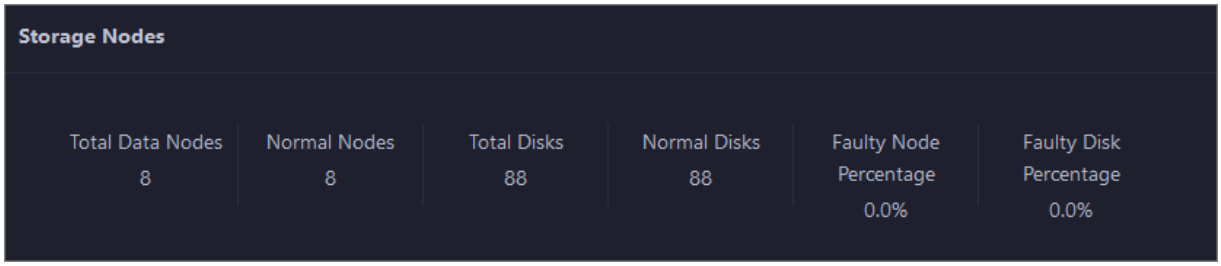

Instances

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.

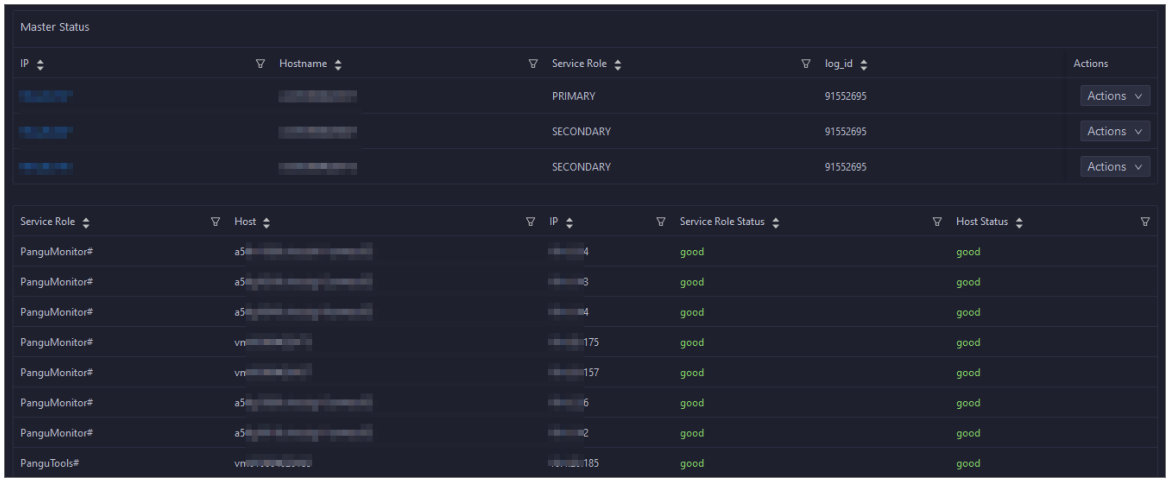

The **Instances** tab shows information about the master nodes and server roles of Apsara Distributed File System. The information about a master node includes the IP address, hostname, server role, and log ID. The information about a server role includes the role name, hostname, role status, and host 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.

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.

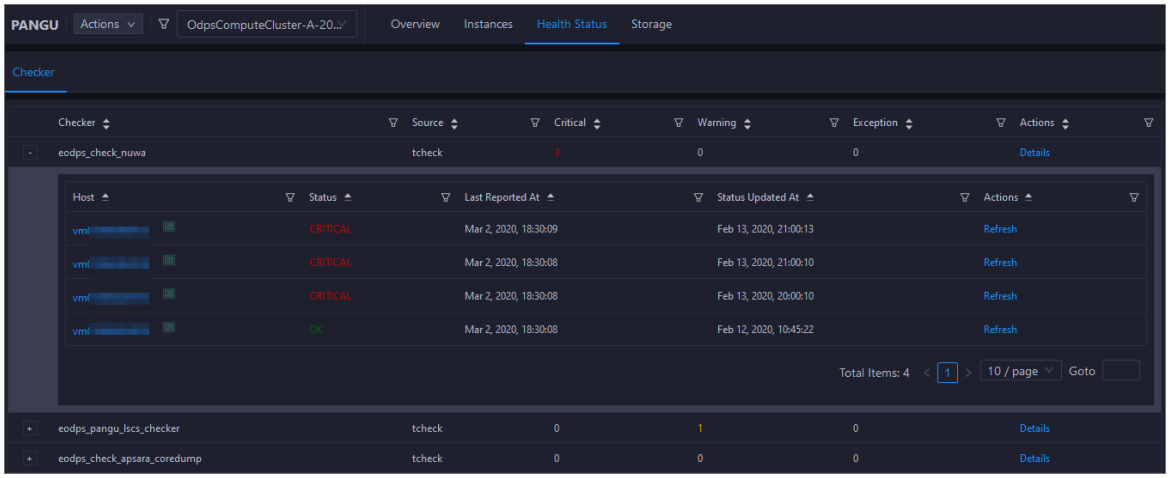

On the **Health Status** page, you can view all checkers of Apsara Distributed File System and the check results for all hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

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.

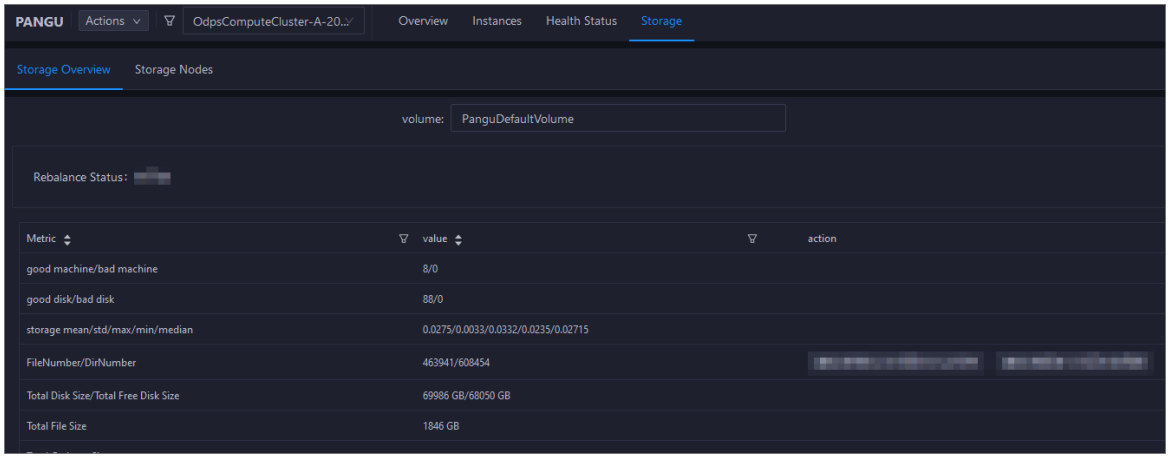

The **Storage Overview** page displays whether data rebalancing is enabled, key metrics and their values, suggestions to handle exceptions, and rack specifications of Apsara Distributed File System.

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.

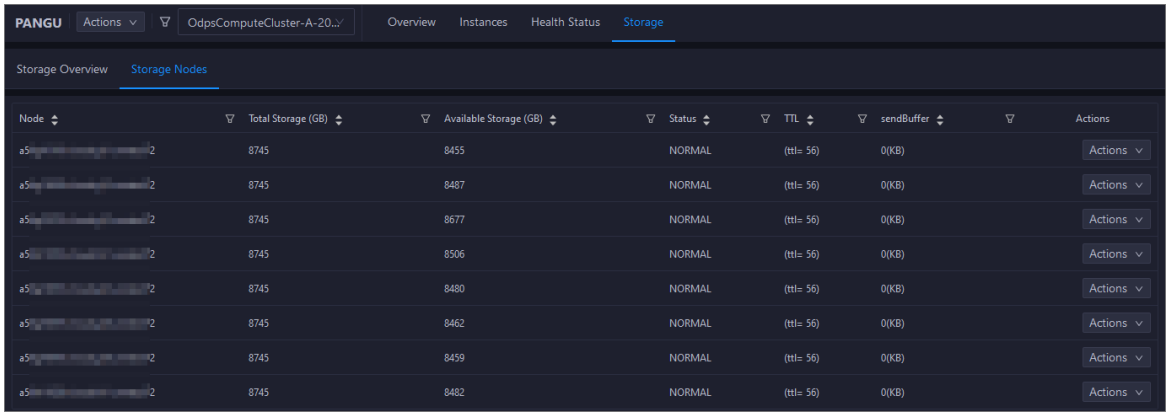

The **Storage Nodes** page displays the information about all storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, TTL, and send buffer size.

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 **Set Node Status to Shutdown** panel, click **Run**. A message appears, indicating that the action has been submitted.

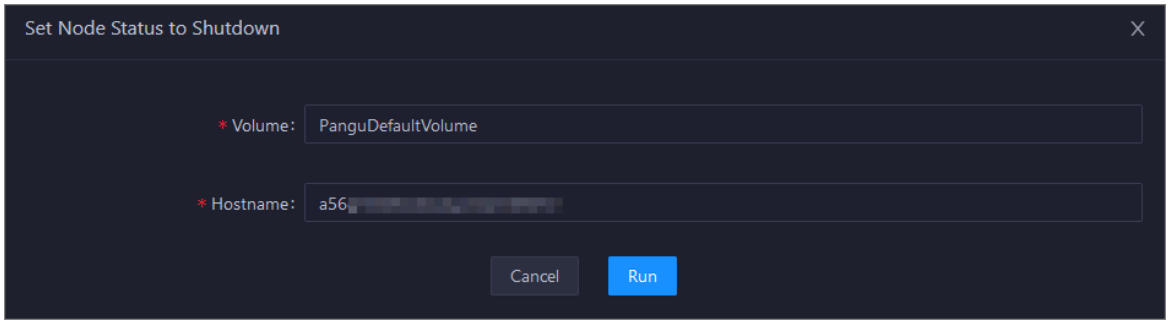

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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 **Set Disk Status to Error** panel, set the **Diskid** parameter.

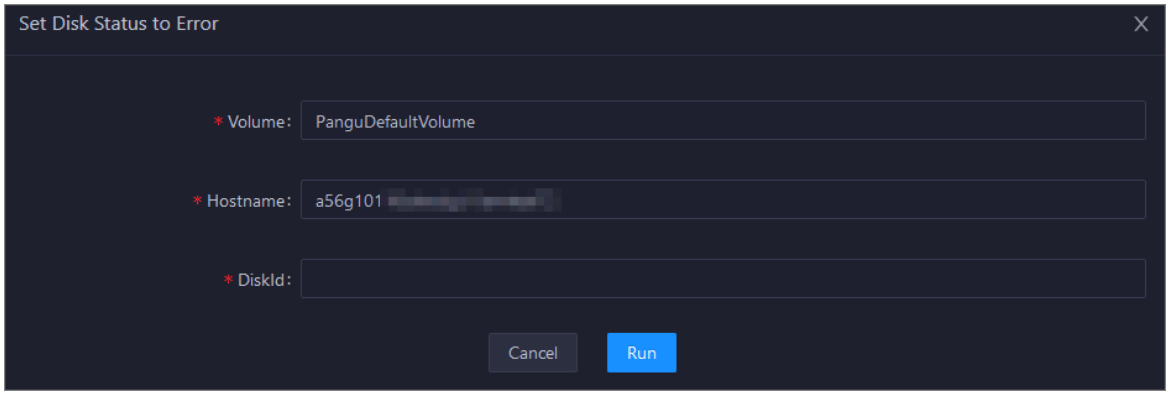

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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

Change the primary master node of Apsara Distributed File System

Apsara Big Data 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 required master node and choose **Actions** > **Change Primary Master Node** in the Actions column. In the Change Primary Master Node panel, specify the required parameters.

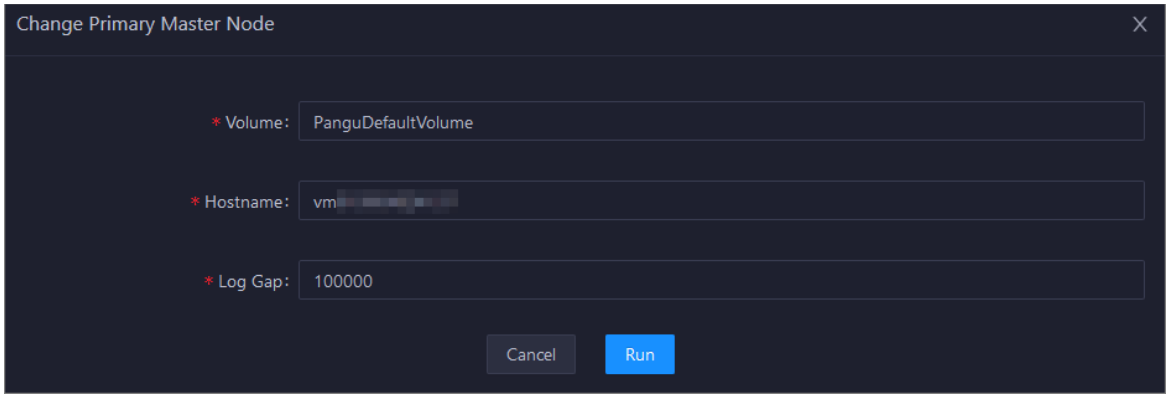

Parameter description:

- **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** panel appears.

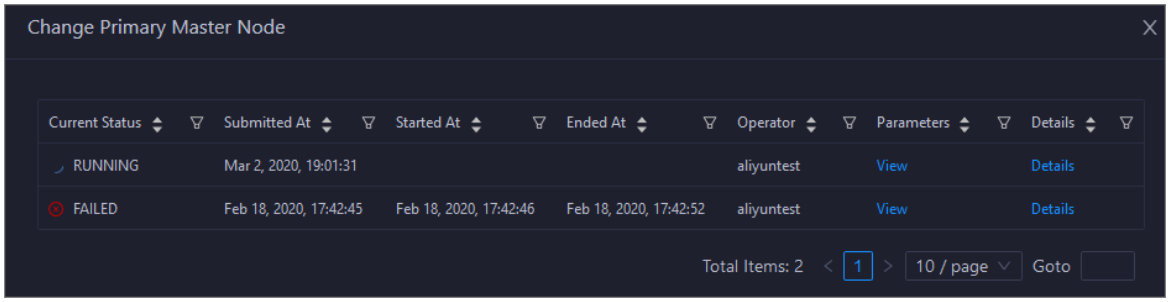

The **Change Primary Master Node** panel shows the switchover history. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

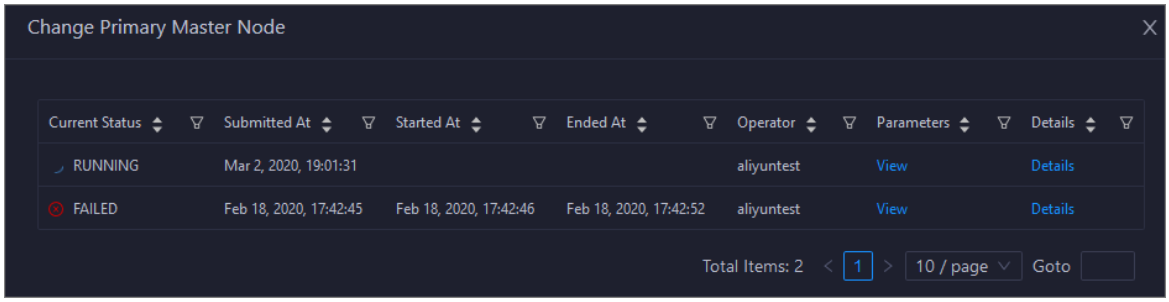

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

Clear the recycle bin of Apsara Distributed File System

Apsara Big Data 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, choose **Actions** > **Empty Recycle Bin**.
3. In the Empty Recycle Bin panel, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

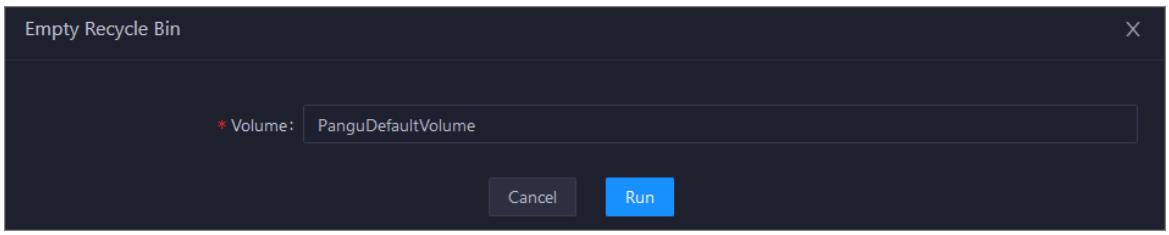

4. Click **Run**.
5. View the execution status.

   In the upper-right corner, click **Actions** and select **Execution History** next to **Empty Recycle Bin** to view the execution history.

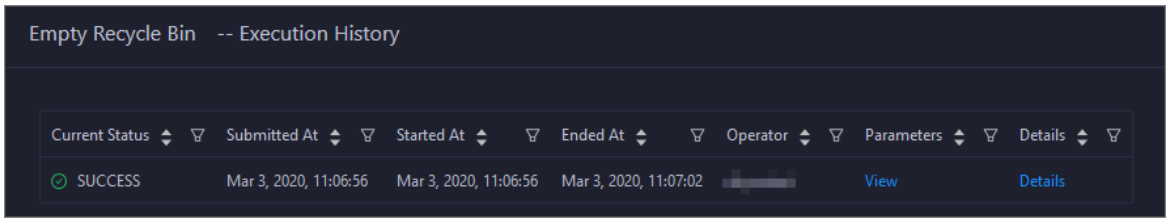

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

6. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

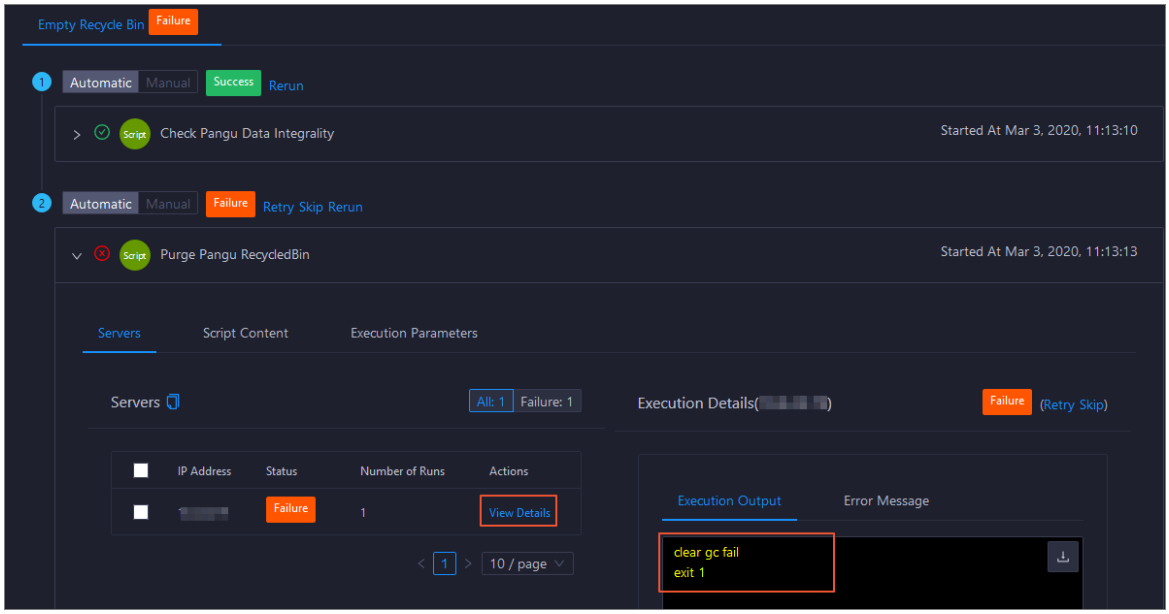

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

Enable or disable data rebalancing for Apsara Distributed File System

Apsara Big Data 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, choose **Actions** > **Disable Data Rebalancing**.
3. In the Disable Data Rebalancing panel, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

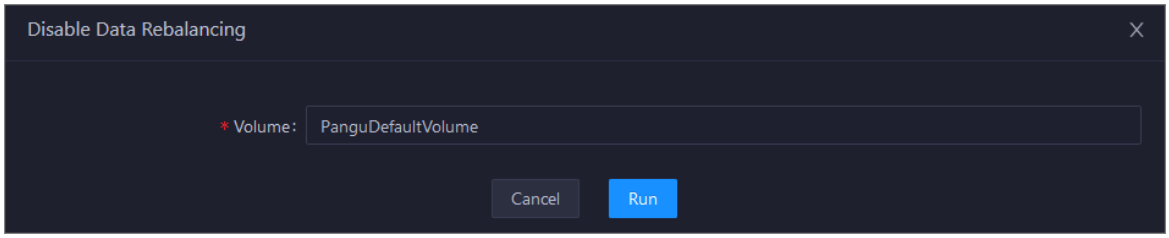

4. Click **Run**.
5. View the execution status.

   Click **Actions** and select **Execution History** next to **Disable Data Rebalancing** to view the execution history.

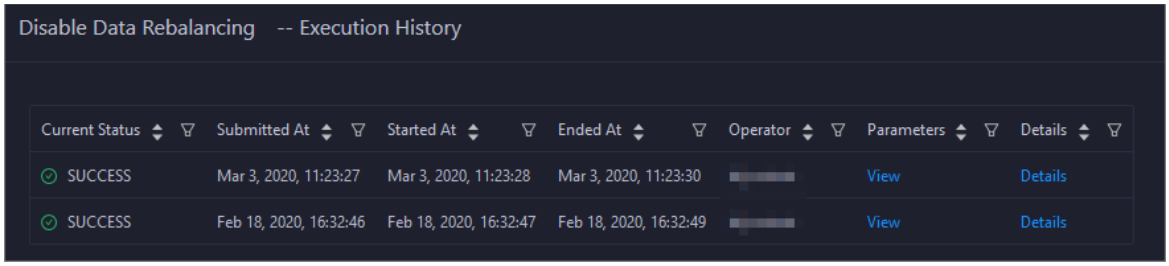

RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

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, choose **Actions** > **Enable Data Rebalancing**.
3. In the Enable Data Rebalancing panel, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

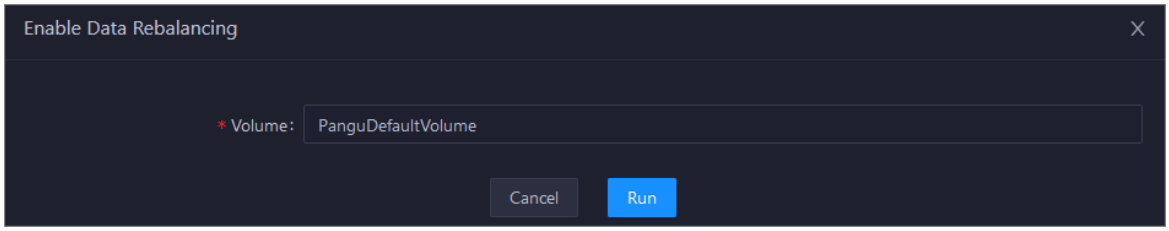

4. Click **Run**.
5. View the execution status.

   Click **Actions** and select **Execution History** next to **Enable Data Rebalancing** to view the execution history.

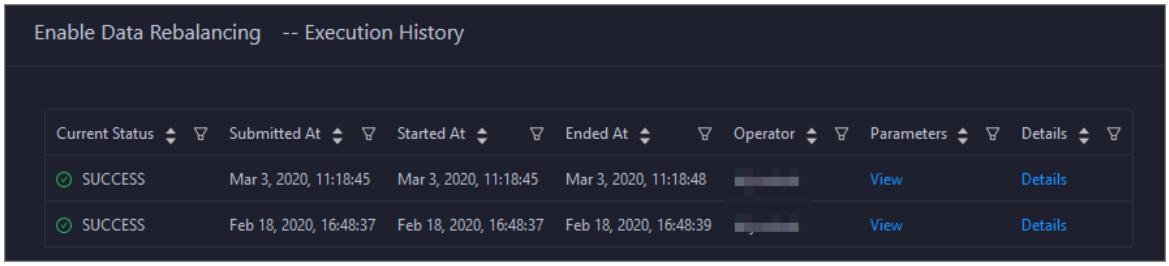

RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

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 panel, click **Details** in the Details column for a failed execution.

2. In the Enable Data Rebalancing panel, 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.

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

Apsara Big Data 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 required master node and choose **Actions** > **Run Checkpoint on Master Node** in the Actions column. In the Run Checkpoint on Master Node panel, set the **Volume** parameter.

   > **Note** The default value of **Volume** is **PanguDefaultVolume**.

6. Click **Run**. The **Run Checkpoint on Master Node** panel appears.

   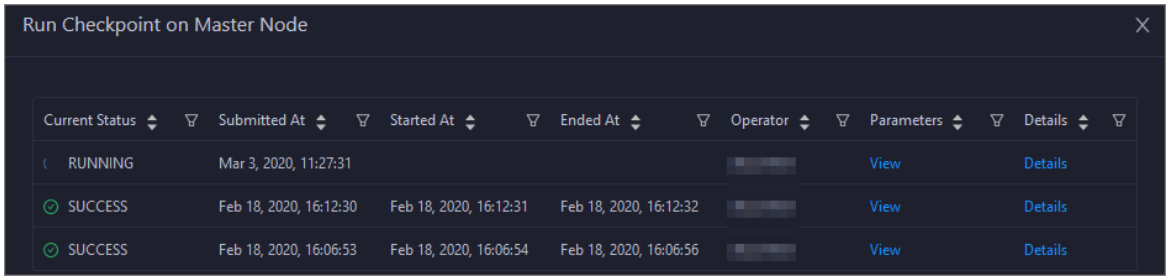

   The **Run Checkpoint on Master Node** panel shows the execution history of the checkpoint on the master node. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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

6.1.1.5.3.4. Tunnel service

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 O&M page of the Tunnel service.

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.

O&M features of the Tunnel service

- 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.
- Traffic Analysis: shows the traffic curves of specific projects in a specific period. The curves show traffic types and the peak throughout in the specified period, which helps you make informed decisions.
- Restart Tunnel Server: allows you to restart one or more Tunnel servers.

Go to the Tunnel Service page

1. Log on to the Apsara Big Data Manager (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 **Tunnel Service**. The **Overview** tab for the Tunnel service appears.

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

The Services section shows the numbers of available services, unavailable services, and services that are being updated.

Roles

The Roles section shows all Tunnel server roles and their status. 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. This trend chart can be automatically or manually refreshed. You can view the trend chart of Tunnel throughput in a specific period.

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.

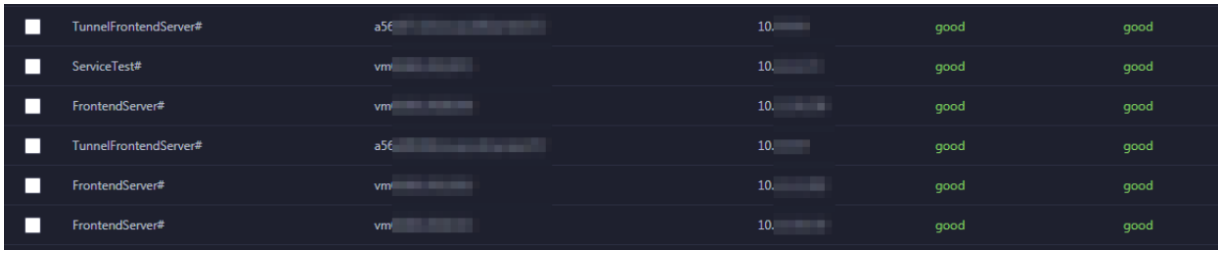

The **Instances** tab shows information about all Tunnel server roles. The information includes the role name, hostname, IP address, role status, and host status. The status can be good, error, or upgrading.

Restart Tunnel servers

Apsara Big Data 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, choose **Actions > Restart Tunnel Server**.
6. In the **Restart Tunnel Server** panel, configure the required parameters.

   The following table describes the required 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 Execution History panel, 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.

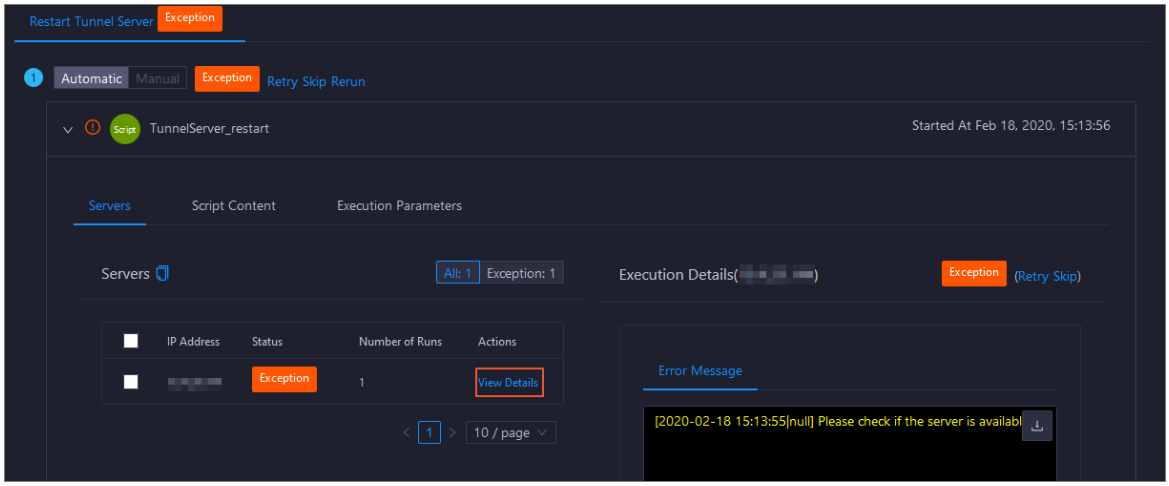

6.1.1.5.4. Cluster O&M

6.1.1.5.4.1. O&M features and entry

This topic describes the O&M features of MaxCompute clusters. It also provides more information about how to go to the MaxCompute cluster O&M page.

Cluster O&M features

O&M features of MaxCompute clusters:

- 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 out Cluster or Scale in Cluster: allows you to add or remove physical hosts to scale out or scale in a MaxCompute cluster.

- 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 a MaxCompute cluster at a time.

Go to the Clusters tab

1. Log on to the Apsara Big Data 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, click a cluster. The **Overview** tab for the selected cluster appears.

6.1.1.5.4.2. 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 Big Data 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 upper-left corner, click the **Login in** icon of the host. The **TerminalService** page appears.

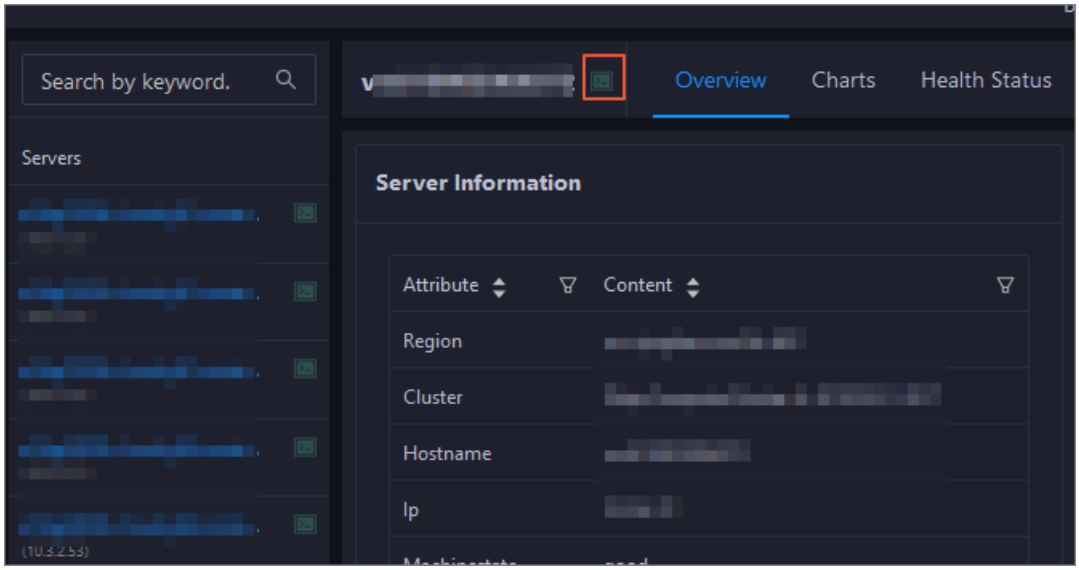

3. In the left-side navigation pane, click the hostname to log on to the host.

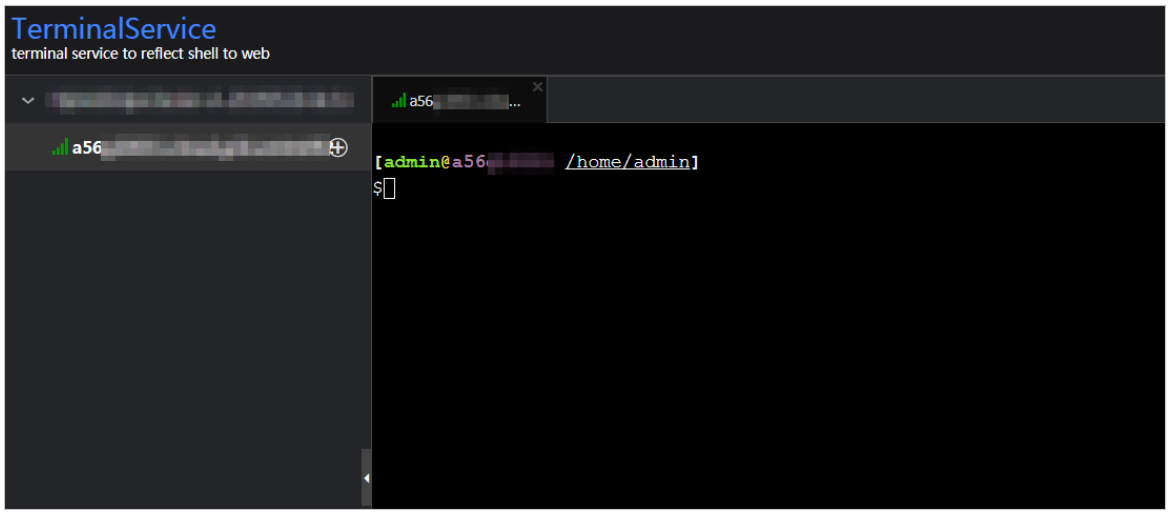

Servers

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 displays all services deployed in the cluster and the respective number 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.

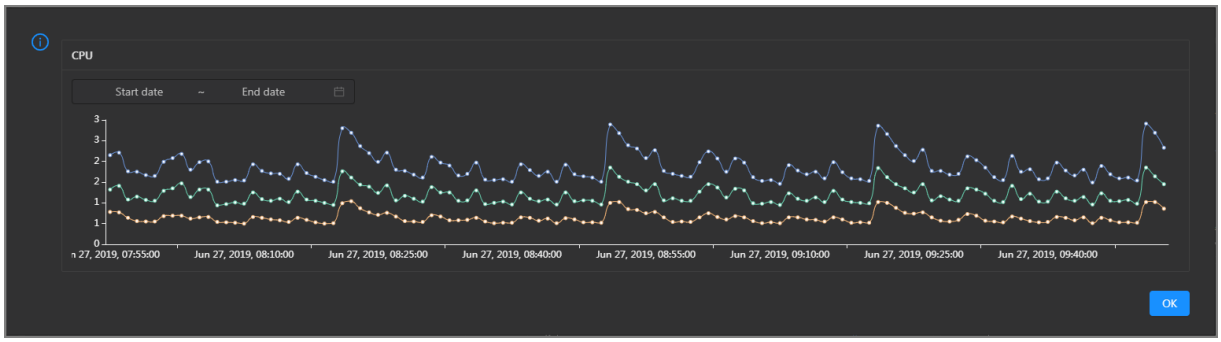

DISK

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

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

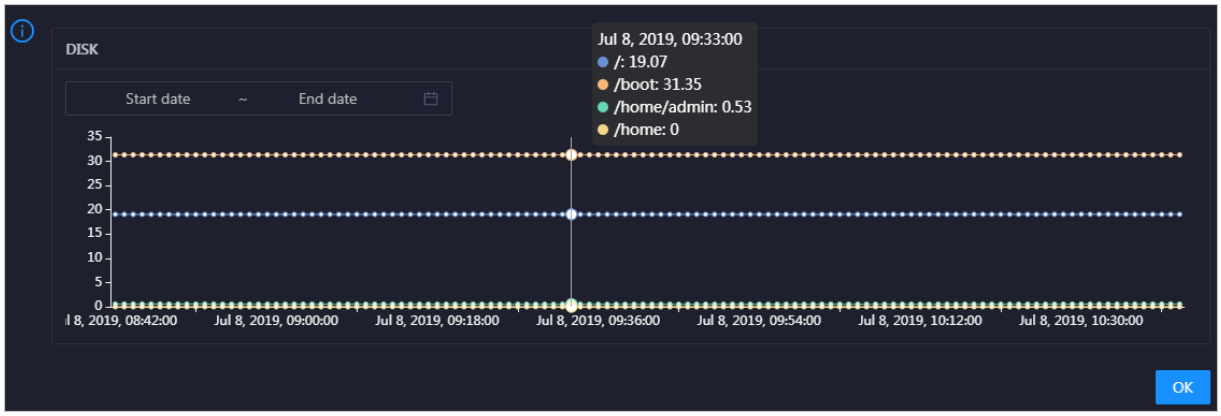

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

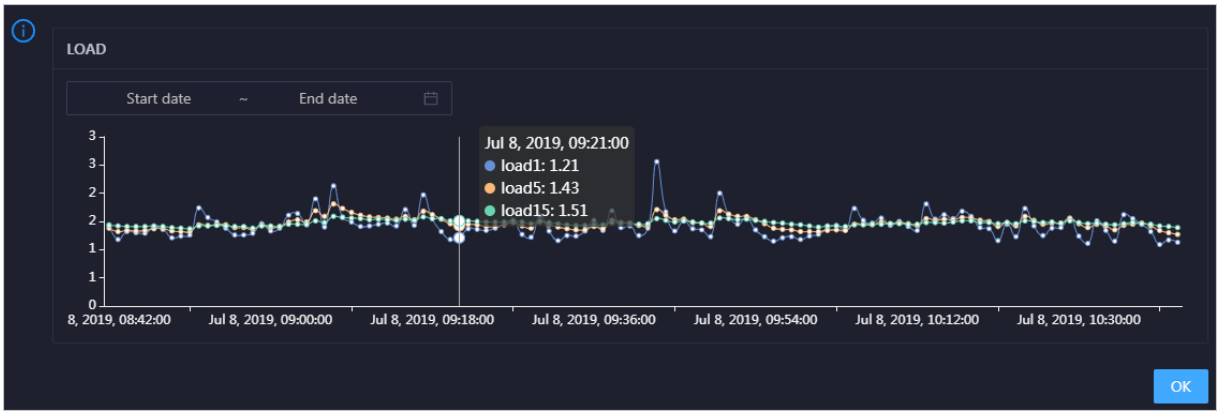

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

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.

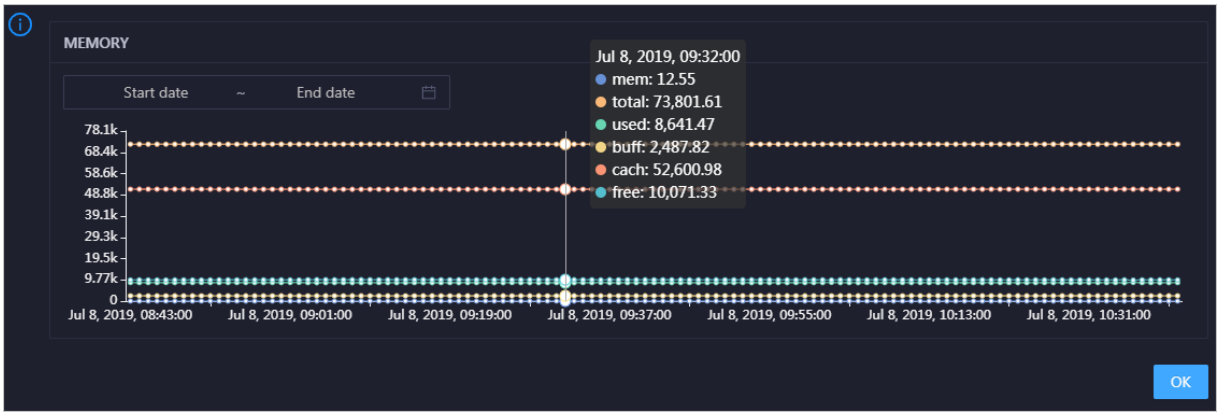

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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.

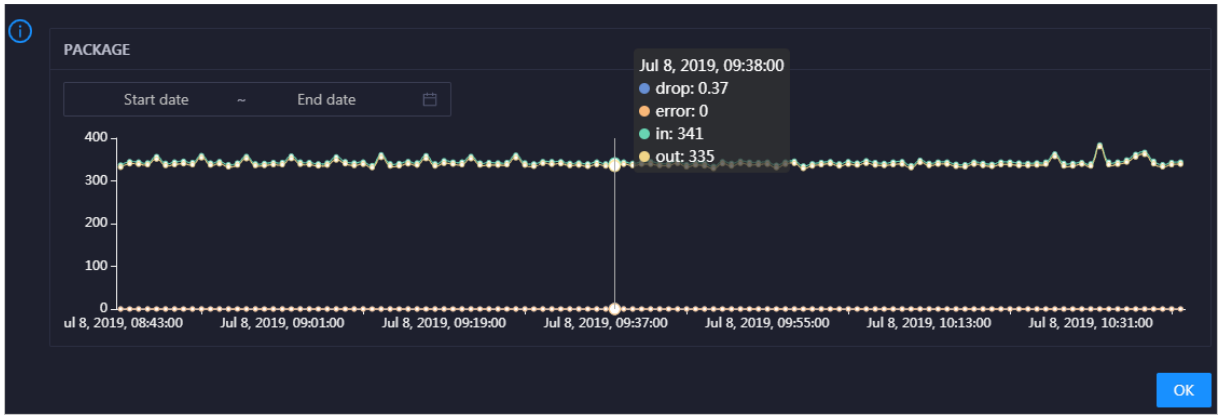

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

Health Check

This section shows the number of checkers for the cluster and the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

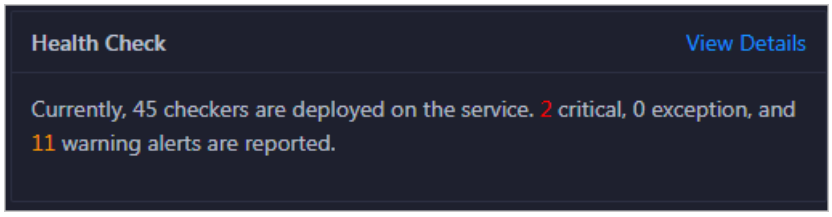

Click **View Details** to go to the Health Status tab. On this tab, you can view health check details. For more information, see Cluster health.

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.

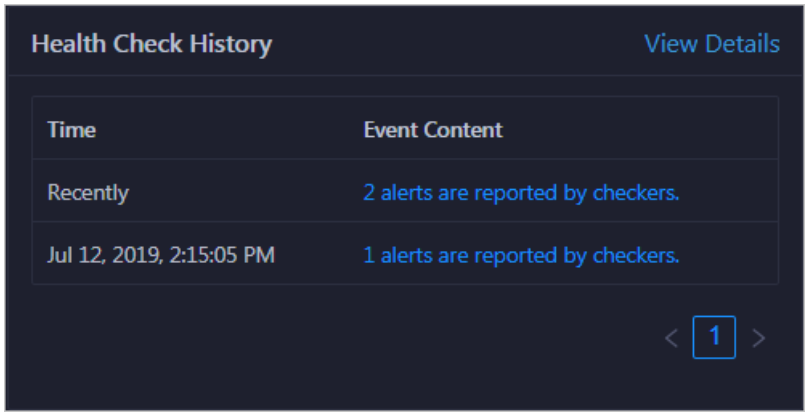

Click **View Details** to go to the Health Status tab. On this tab, you can view health check details. For more information, see Cluster health.

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

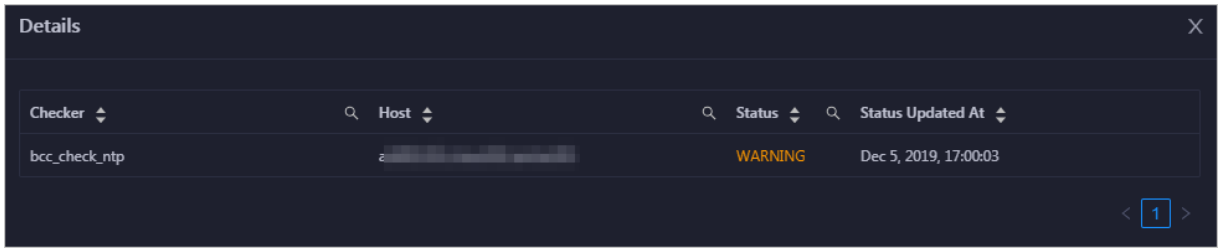

6.1.1.5.4.3. Cluster health

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 Big Data 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.

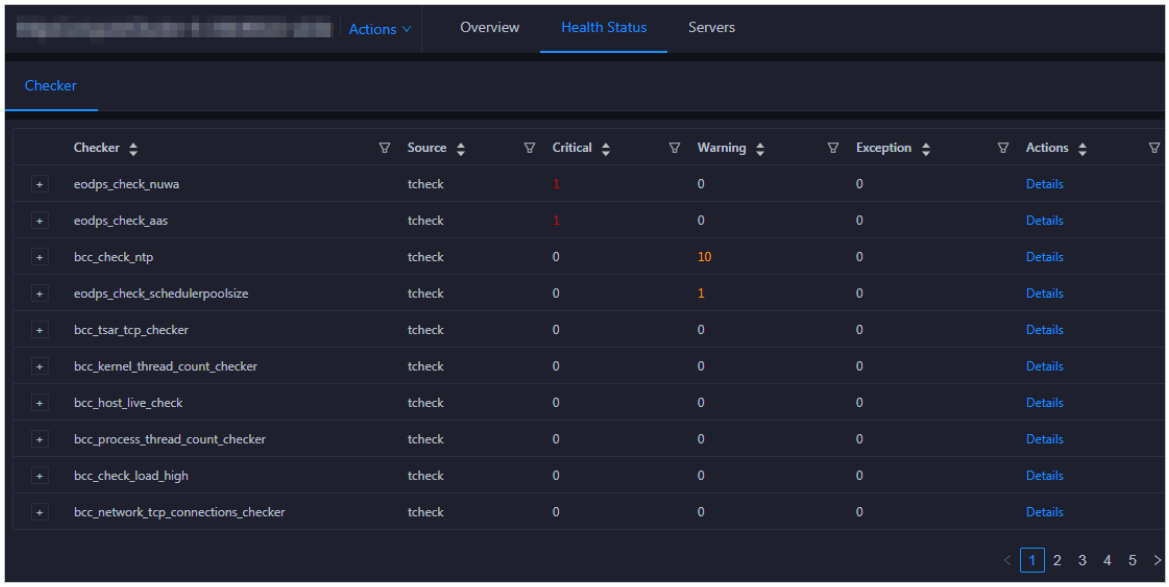

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are represented in different colors. You must handle the alerts in a timely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

View checker details

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

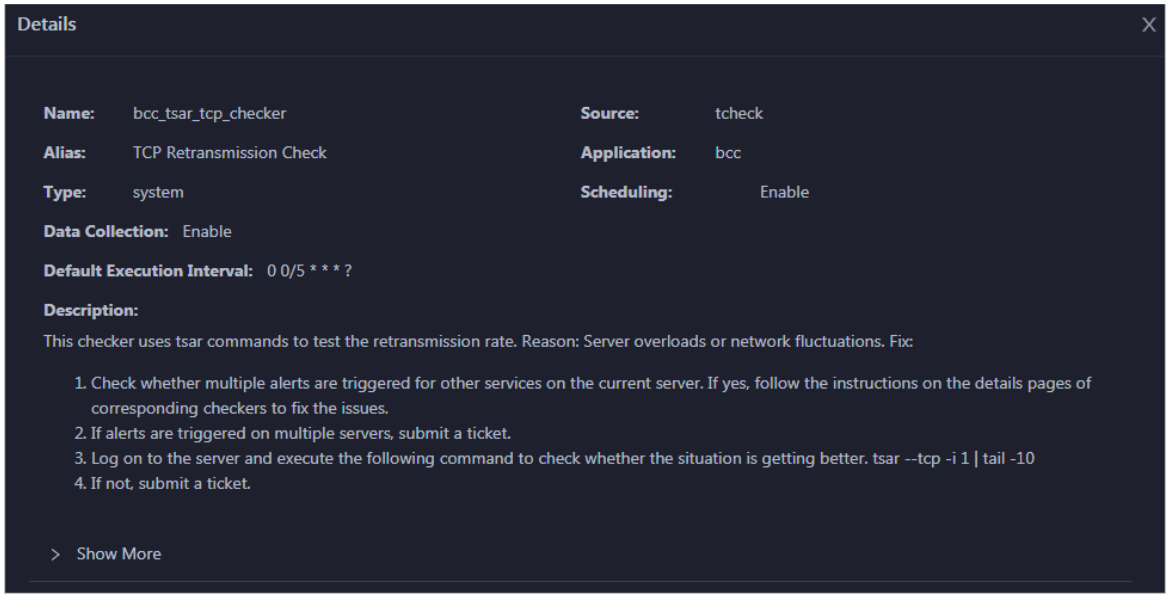

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

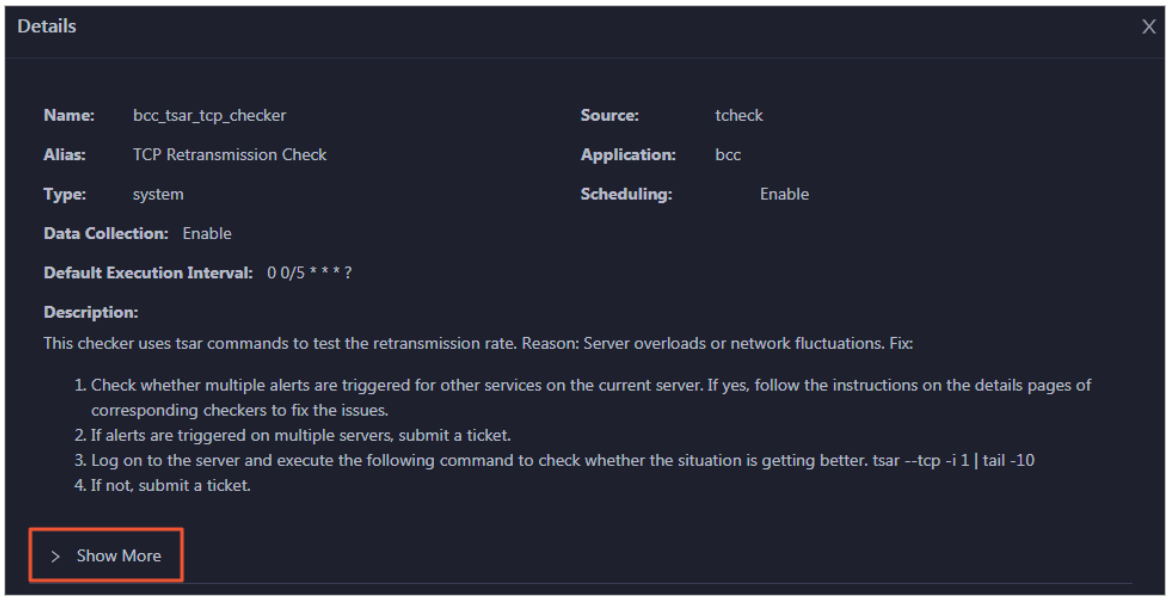

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

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

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

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

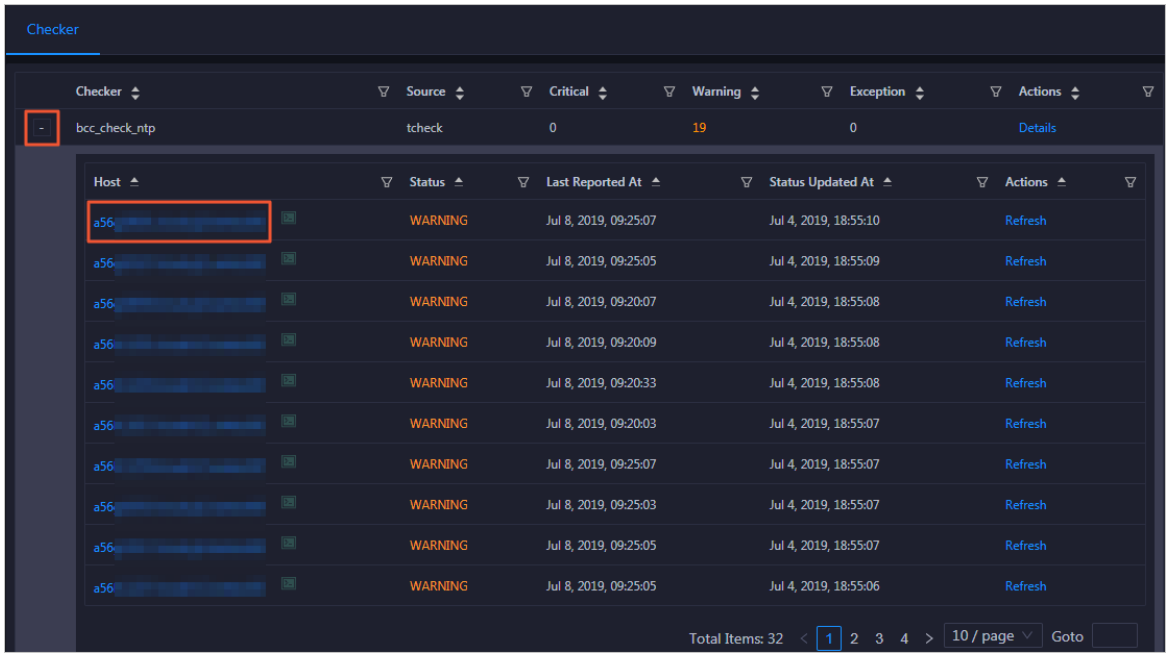

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

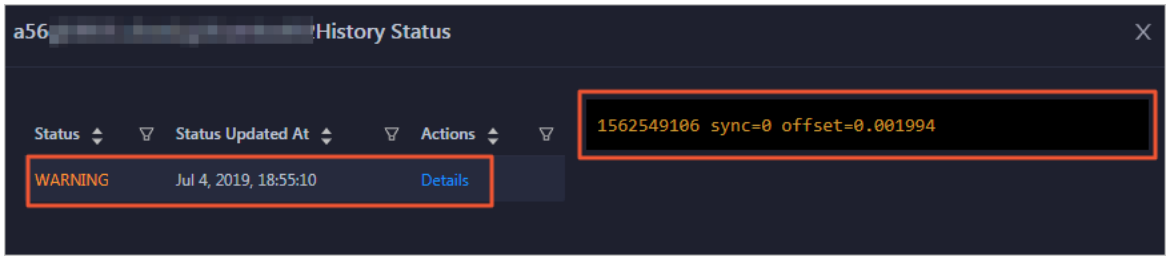

Clear alerts

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

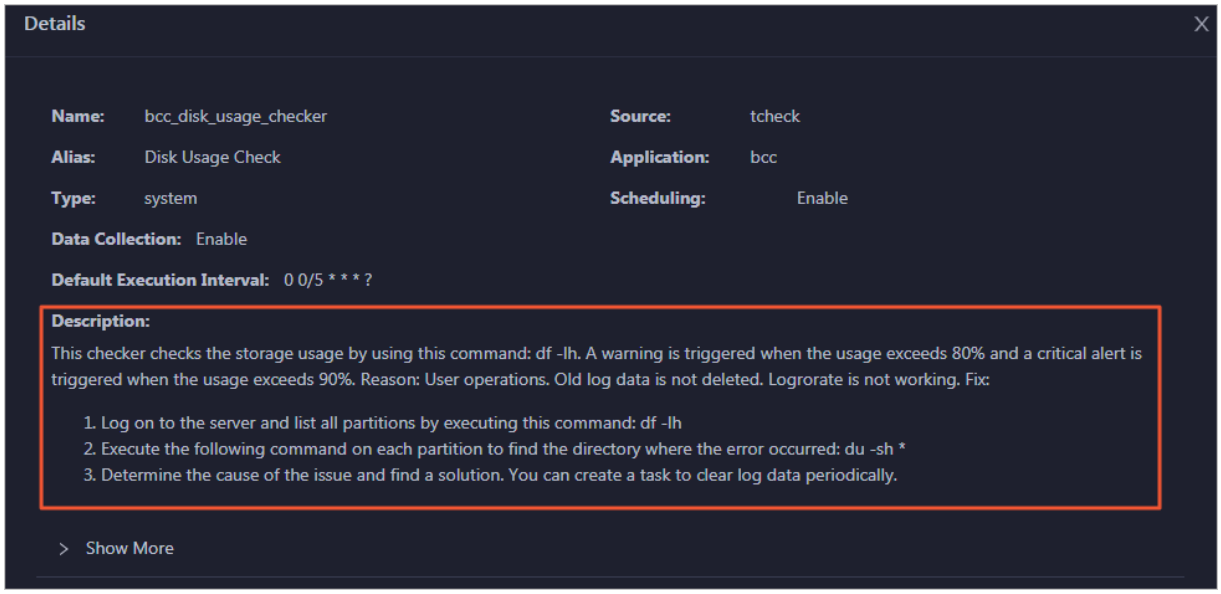

Log on to a host

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

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

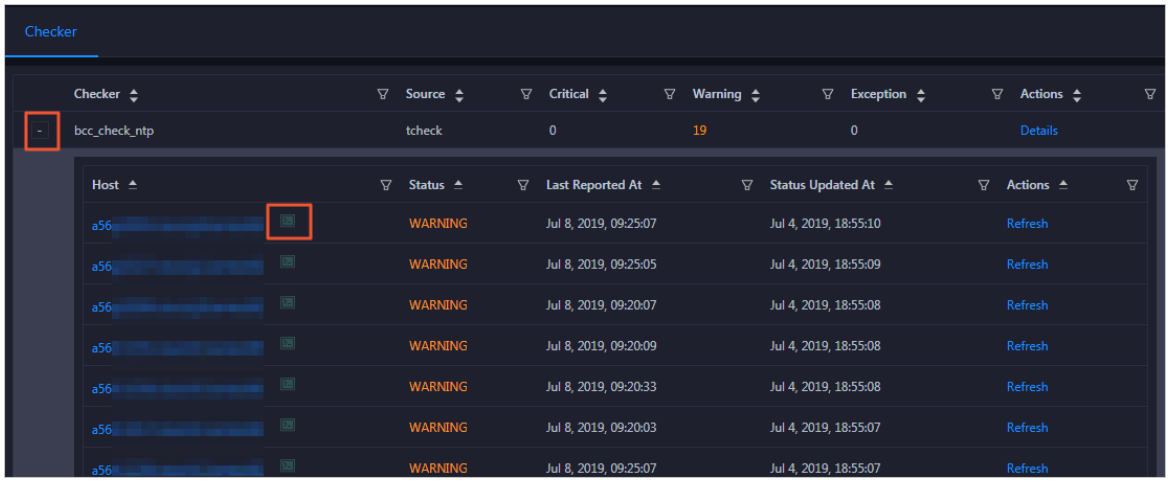

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

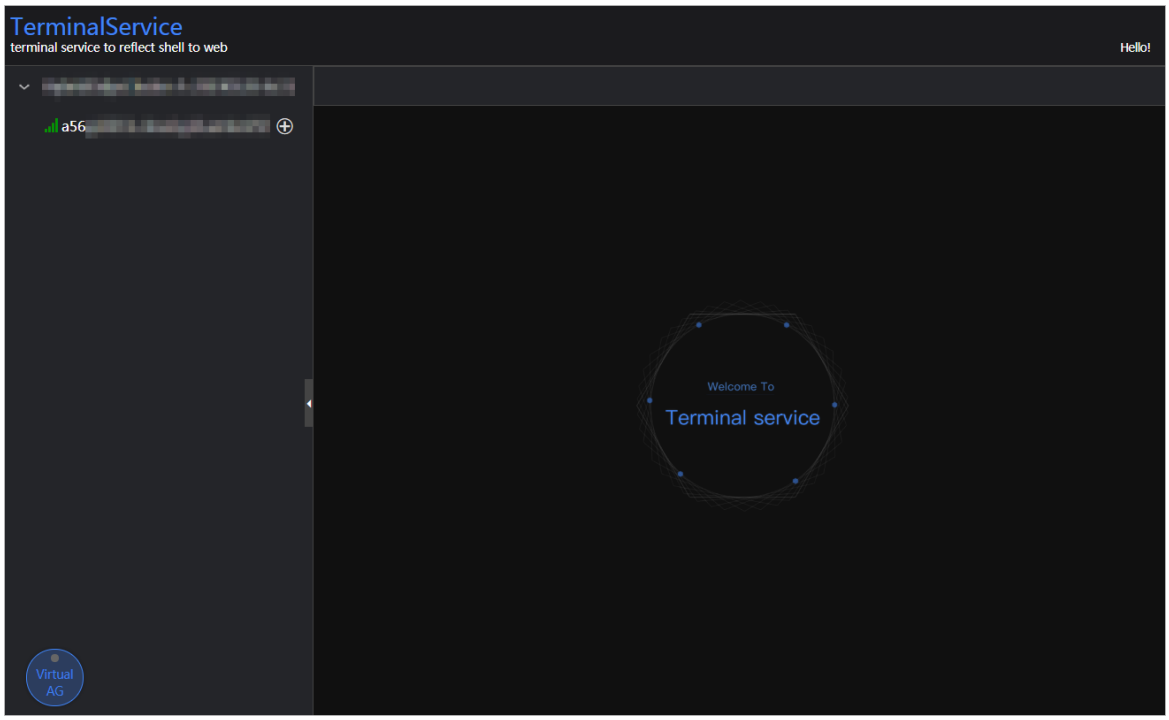

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

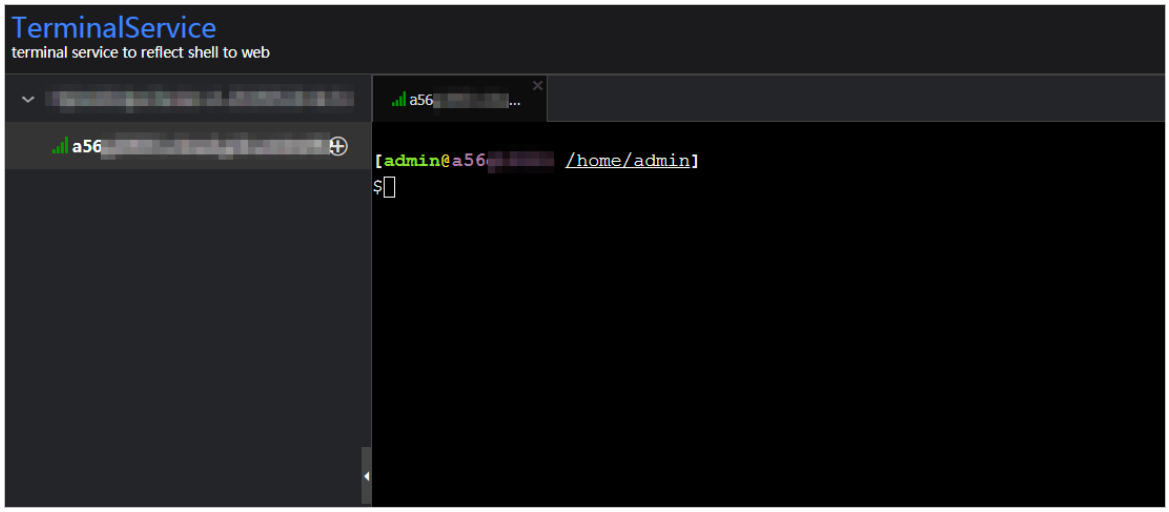

Run a checker again

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

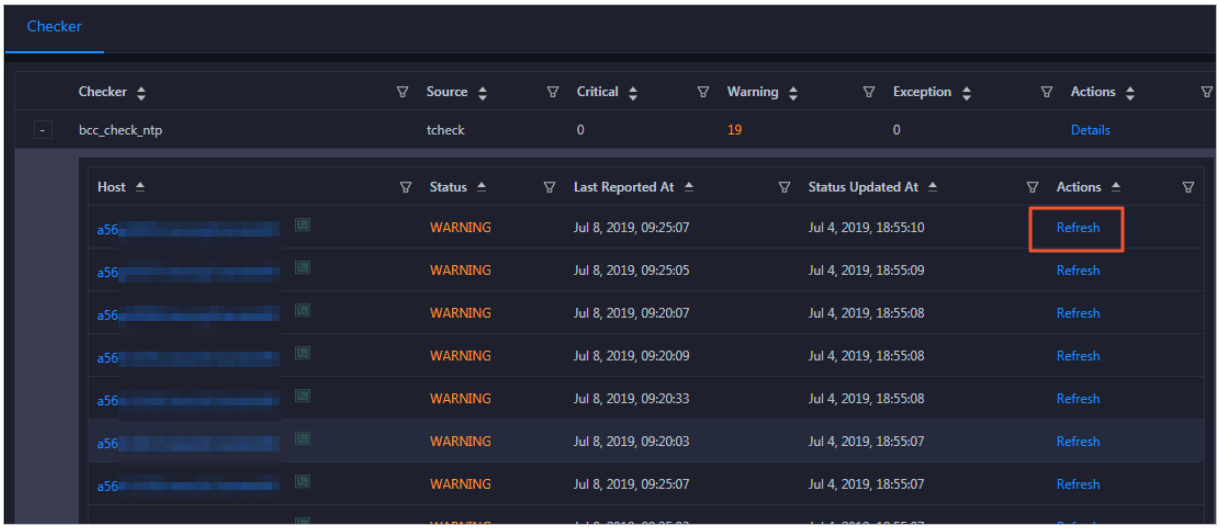

6.1.1.5.4.4. Servers

The Servers tab shows information about hosts. The information includes the hostname, IP address, role, type, CPU utilization, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

In the left-side navigation pane of the **Clusters** tab, click a cluster. Then, click the **Servers** tab. The **Servers** tab for the selected cluster appears.

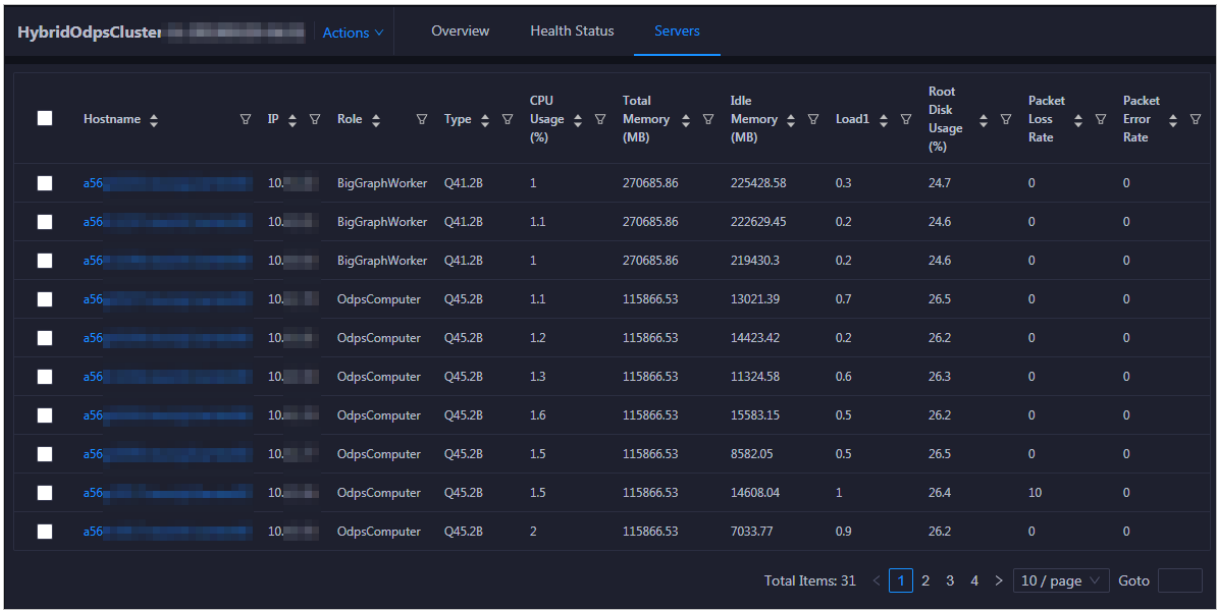

To view more information about a host, click the name of the host. The Hosts tab appears.

6.1.1.5.4.5. Scale in and scale out a MaxCompute cluster

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

Description

In Apsara Stack, scaling out a cluster involves complex operations. You must configure a new physical host on Deployment Planner and Apsara Infrastructure Management Framework 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 out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. If you want to scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

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 `rttrl` command to query and record SInstance hosts. For more information about how to log on to a host, see Log on to a host.

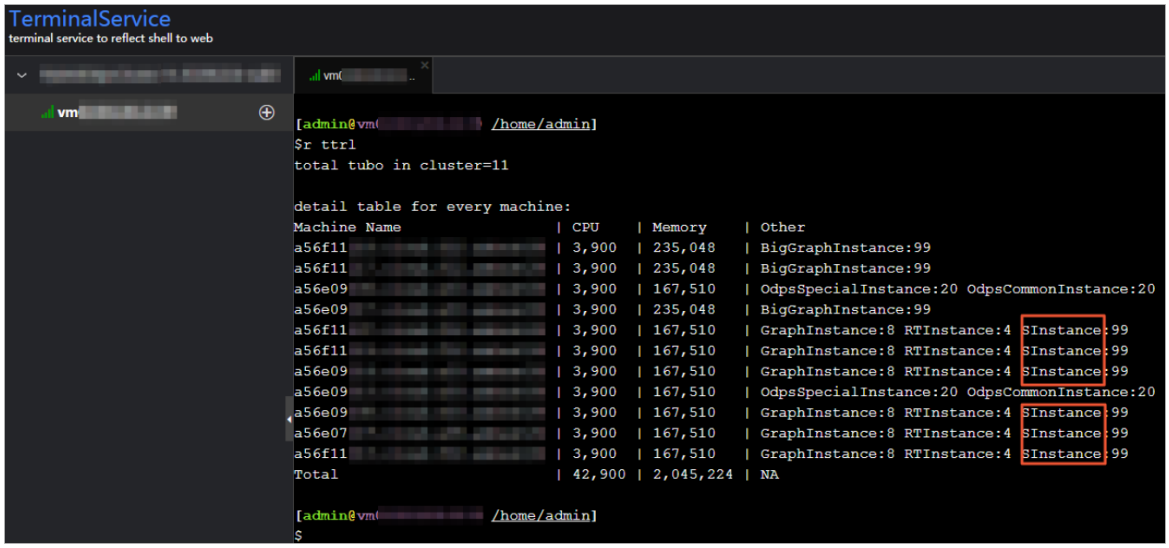

2. In the left-side navigation pane of the **Clusters** tab, click a cluster. Then, click the **Servers** tab. On the tab that appears, select an SInstance host and use it as the template host.

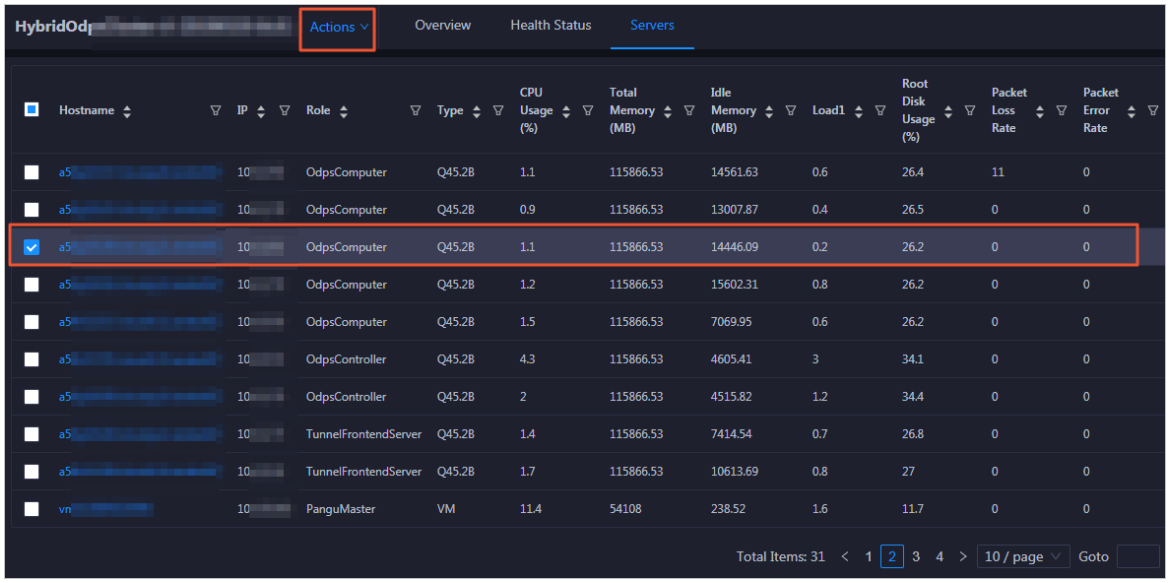

3. In the upper-right corner, choose **Actions > Scale out Cluster**. In the **Scale out Cluster** panel, configure the parameters.

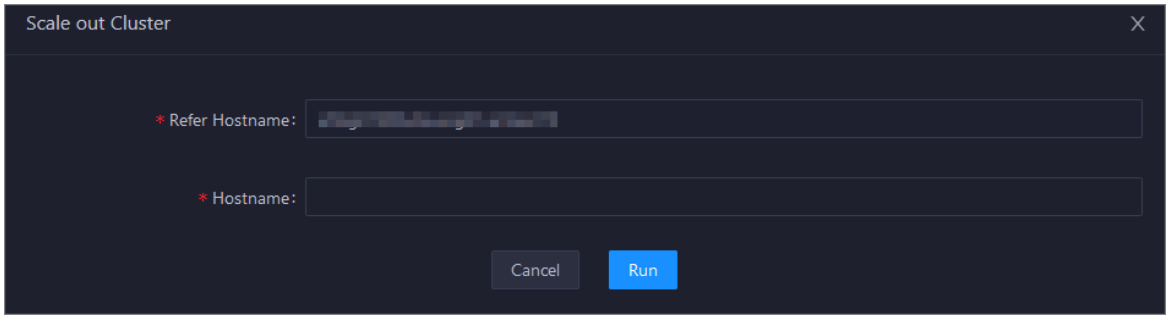

Parameters:

- 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**. A message appears, indicating that the request has been submitted.
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.

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 `rttrl` command to query

and record SInstance hosts. For more information about how to log on to a host, see Log on to a host.

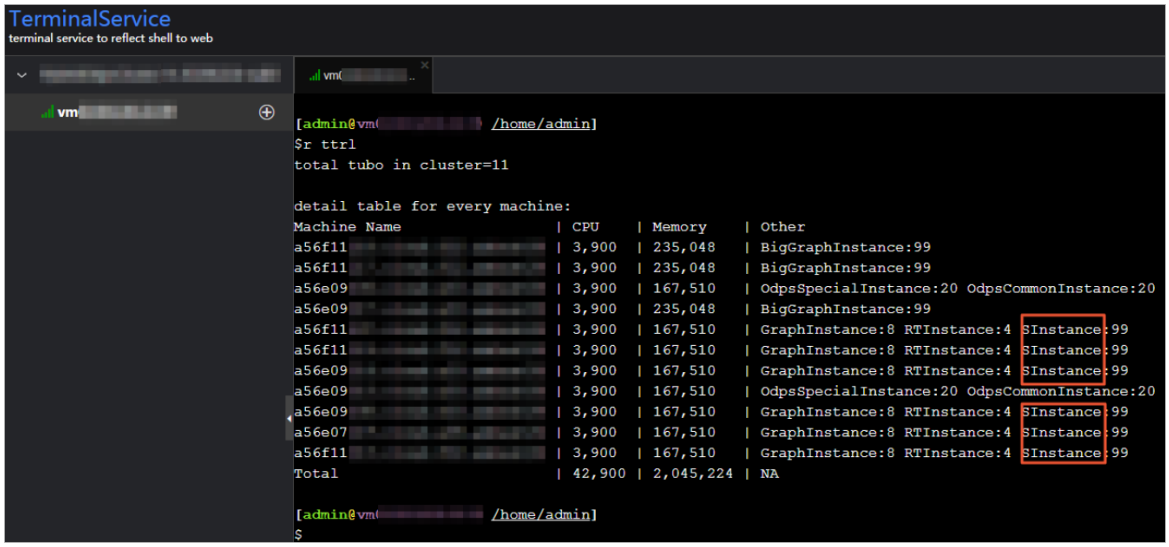

2. In the left-side navigation pane of the **Clusters** tab, click a cluster. Then, click the **Servers** tab. On the tab that appears, select one or more SInstance hosts that you want to remove.

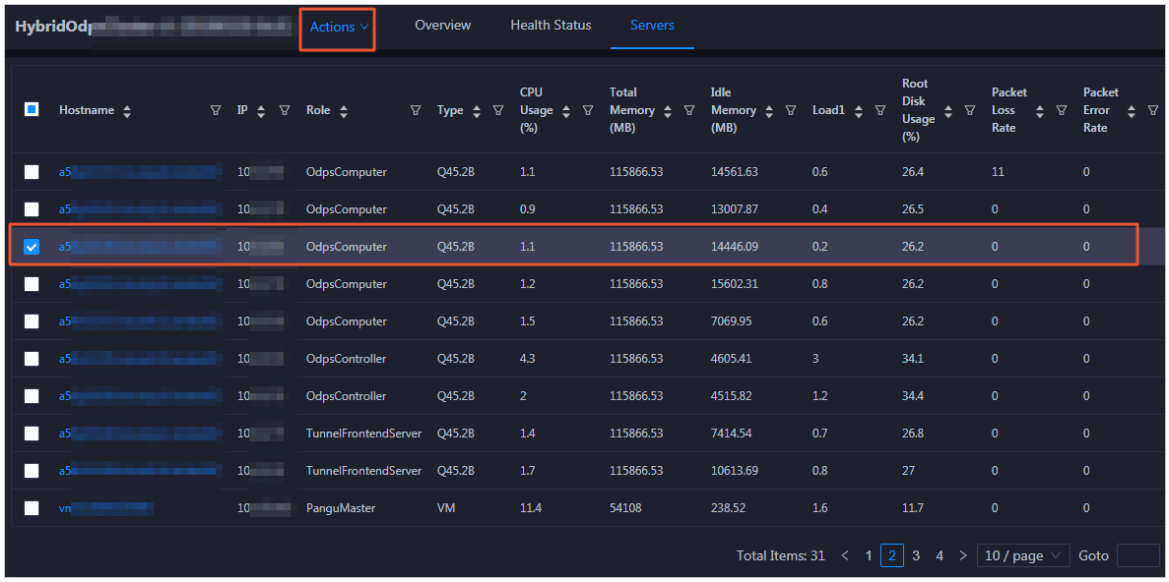

3. In the upper-right corner, choose **Actions > Scale in Cluster**. In the **Scale in Cluster** panel, configure the parameters.

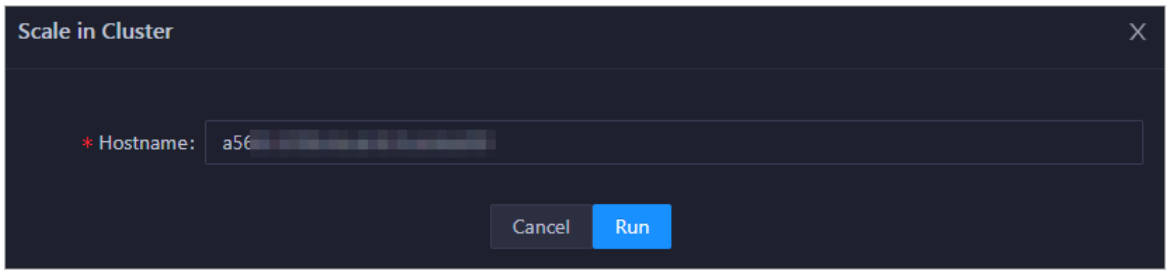

Parameters:

- 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**. A message appears, indicating that the request has been submitted.
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.

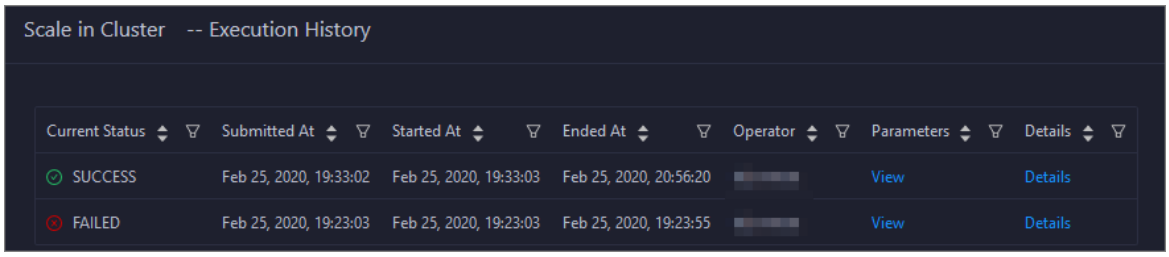

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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.

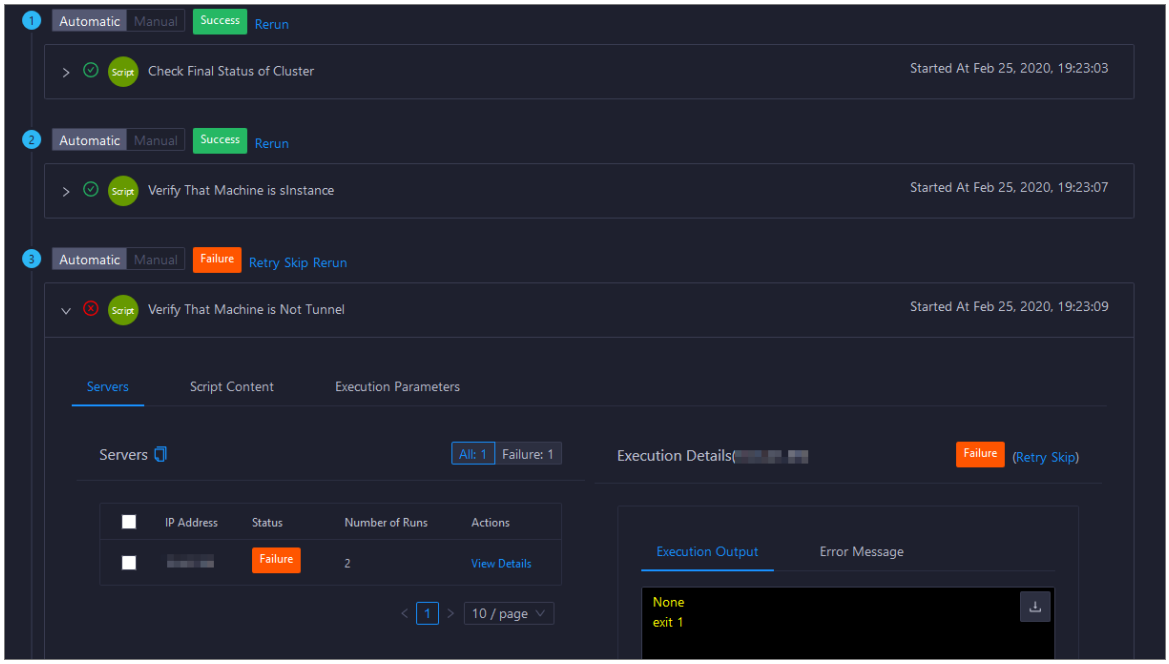

You can view information about parameter settings, host details, scripts, and runtime parameters to identify the cause of the failure.

6.1.1.5.5. Host O&M

6.1.1.5.5.1. 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 Big Data 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.

6.1.1.5.5.2. Host overview

The Overview tab for a host shows brief information about the host in a MaxCompute cluster. On this tab, you can view server information, service role status, 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, and packet transmission for the host.

Go to the Overview tab

In the left-side navigation pane of the **Hosts** tab, click a host. Then, click the **Overview** tab. The **Overview** tab for the host appears.

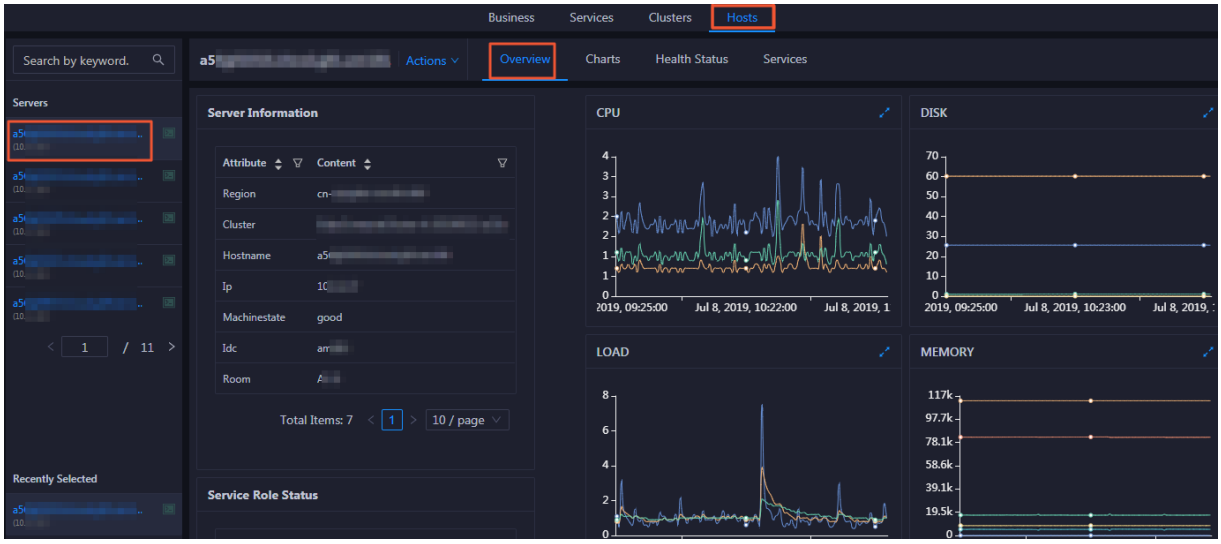

On the **Overview** tab, you can view server information, service role status, health check result, and health check history of the host. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, and packet transmission for the host.

Server Information

The Server Information section shows information about the host. Server information includes the region, cluster, name, IP address, data center, and server room.

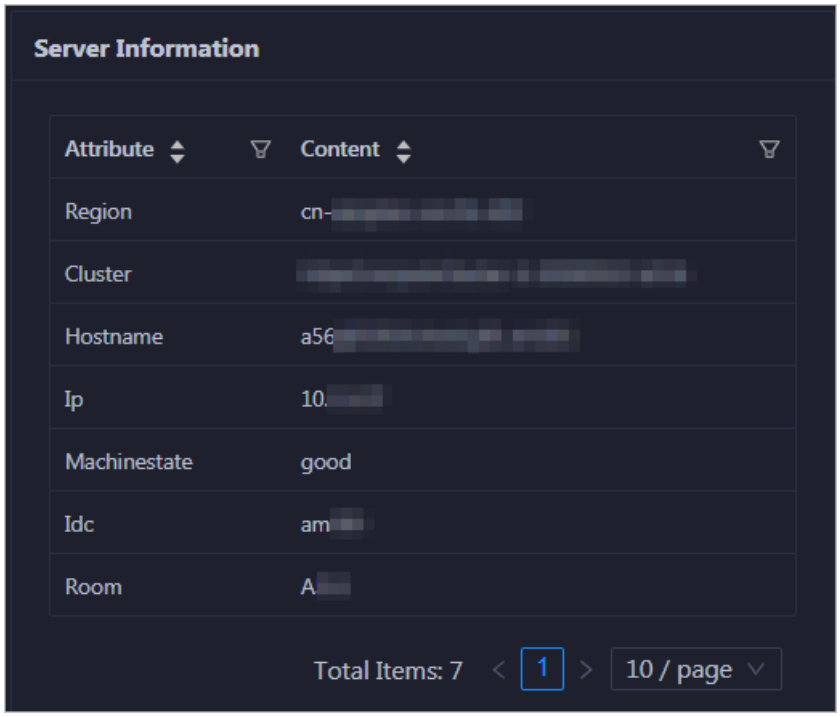

Service Role Status

The Service Role Status section shows information about the services deployed on the host, including the roles, status, and number of services.

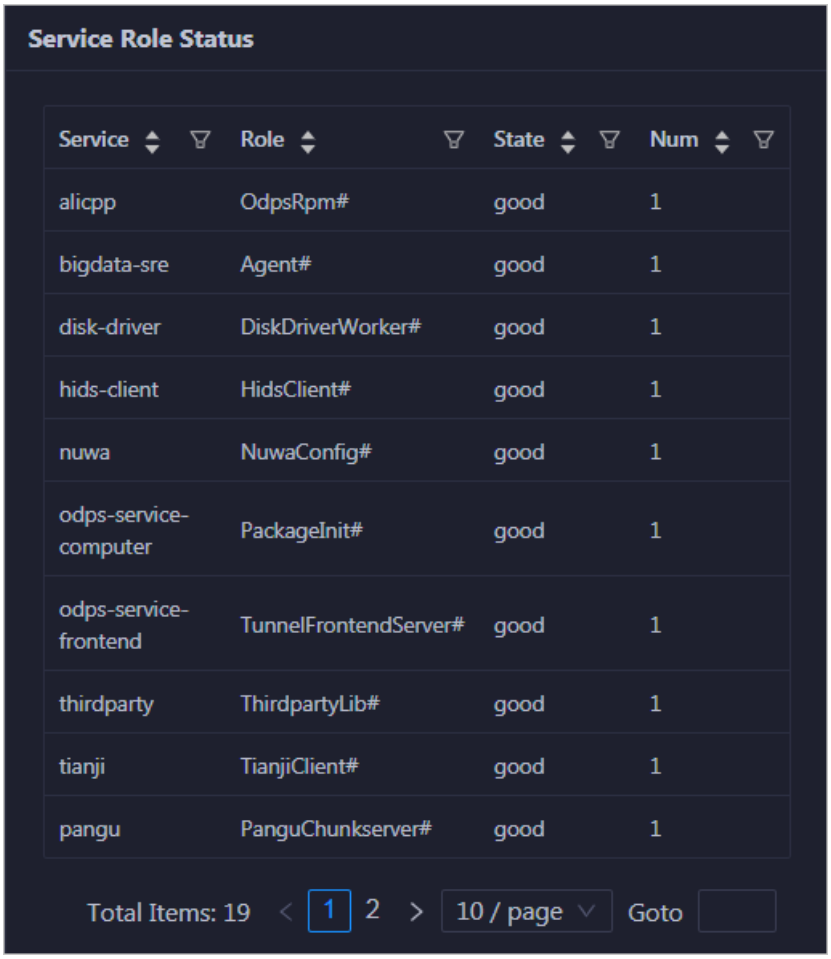

CPU

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

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

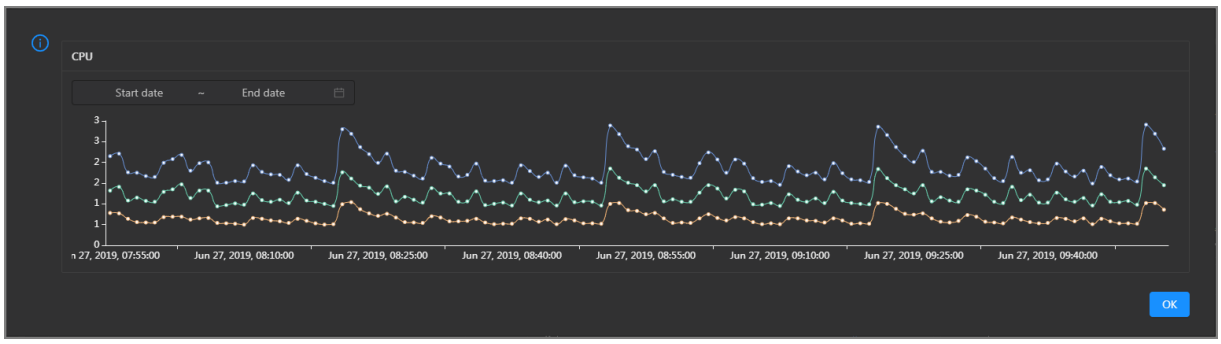

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

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

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

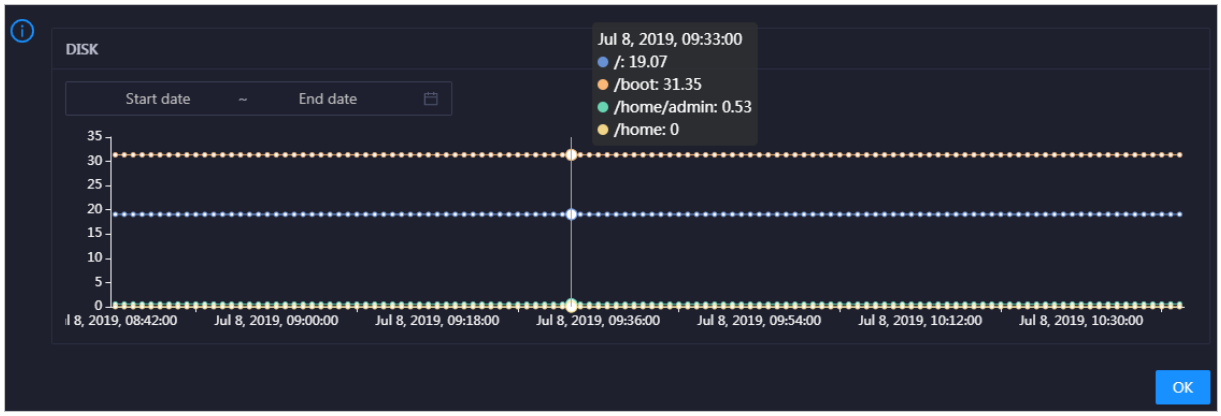

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

LOAD

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

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

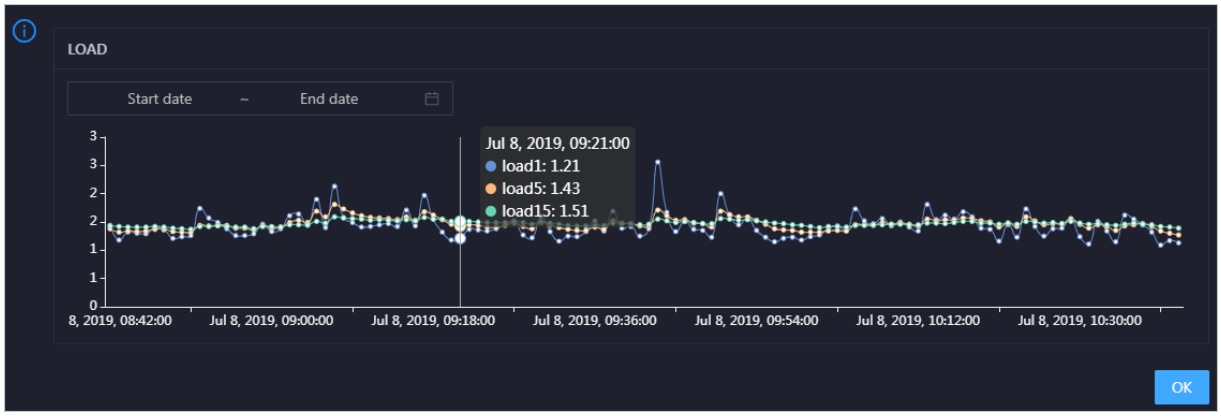

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

MEMORY

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

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

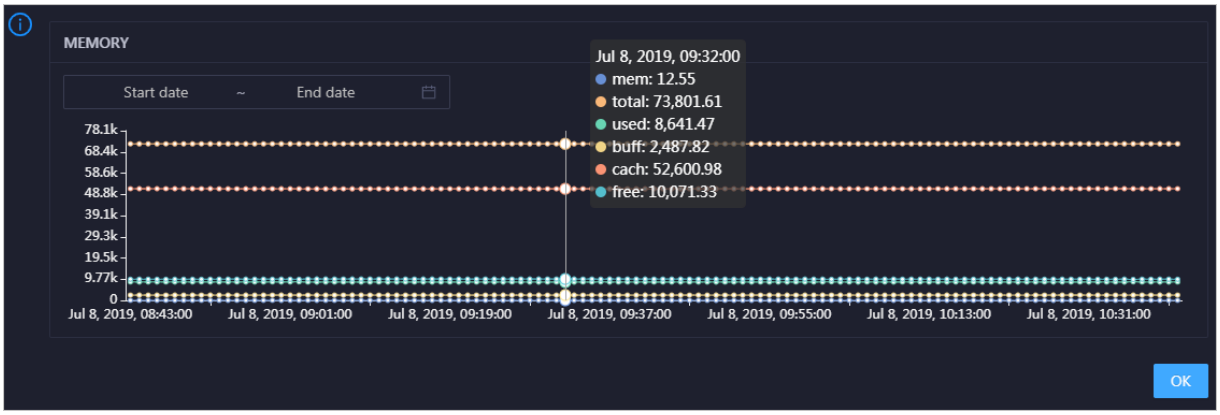

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

PACKAGE

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

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

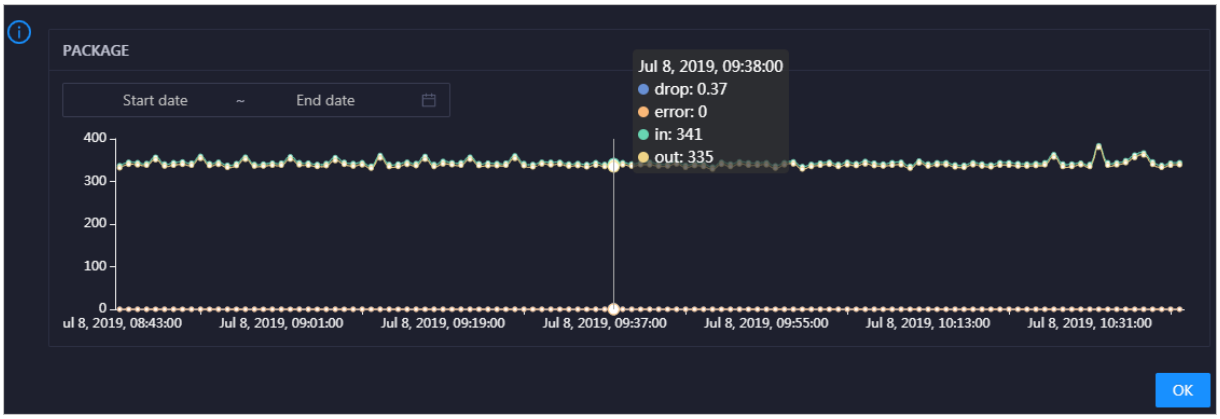

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

The Health Check section shows the number of checkers deployed for the host and the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

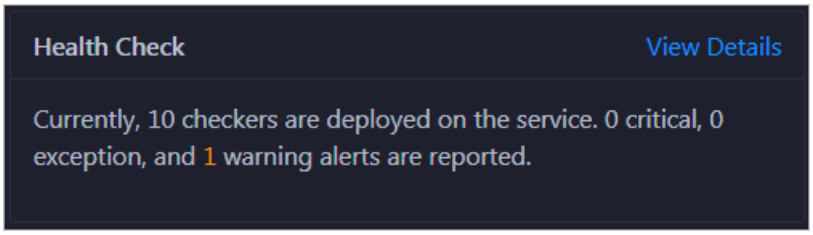

Click **View Details** to go to the Health Status tab. On this tab, you can view the health check details.

Health Check History

The Health Check History section shows the records of the health checks performed on the host.

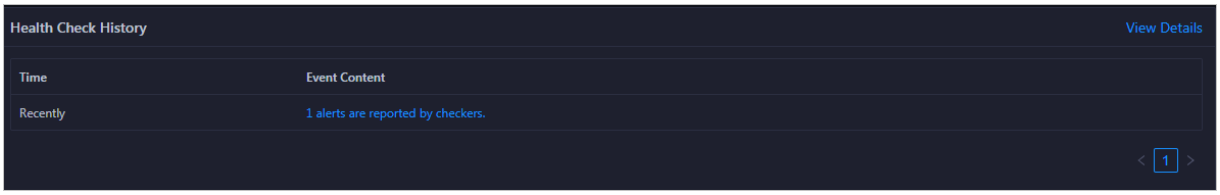

Click **View Details** to go to the Health Status tab. On this tab, you can view the health check details.

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

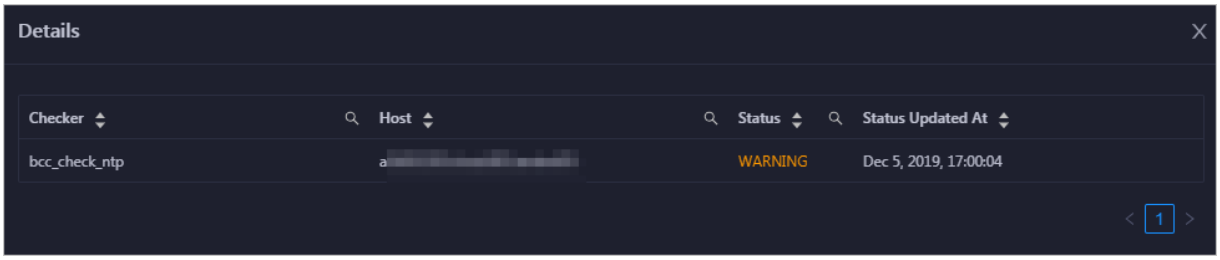

6.1.1.5.5.3. Host charts

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

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

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

6.1.1.5.5.4. Host health

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

Entry

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

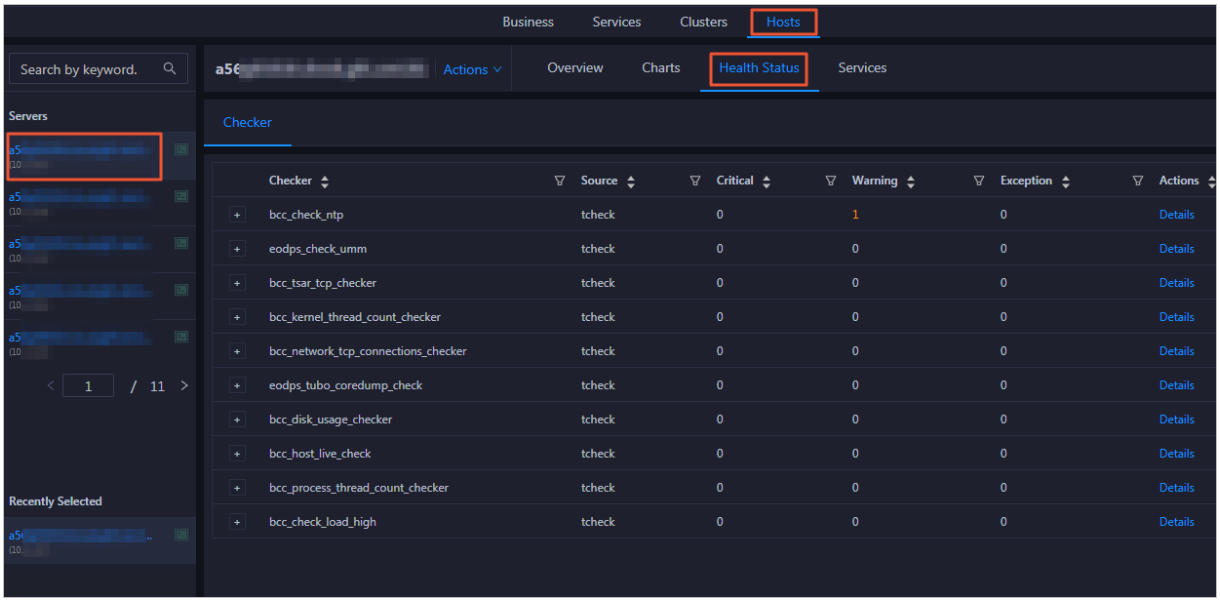

On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

View checker details

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

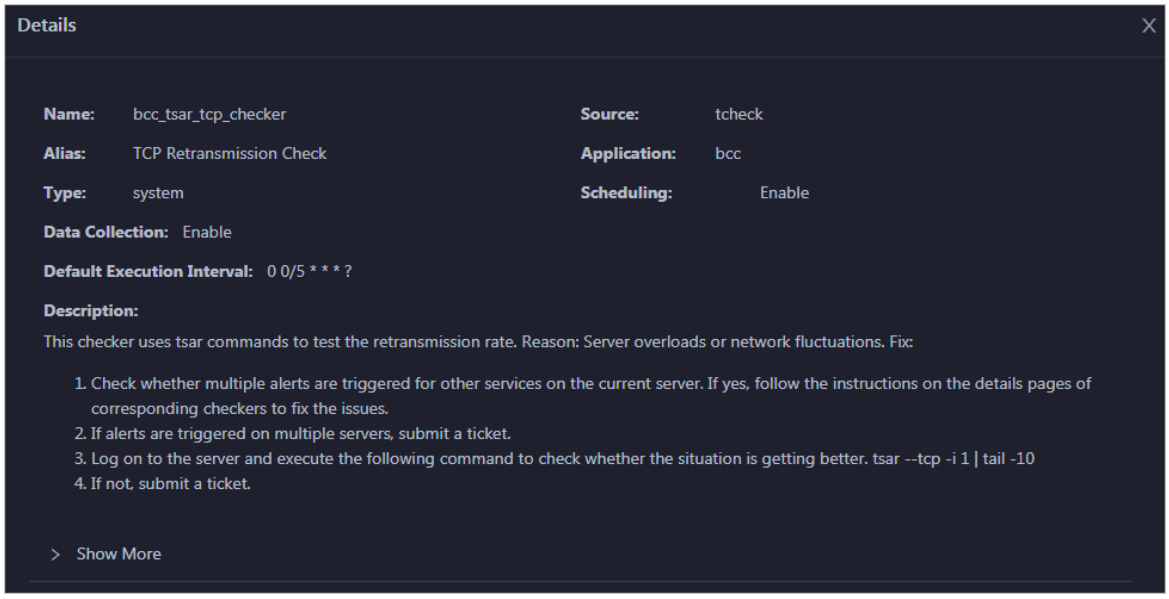

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

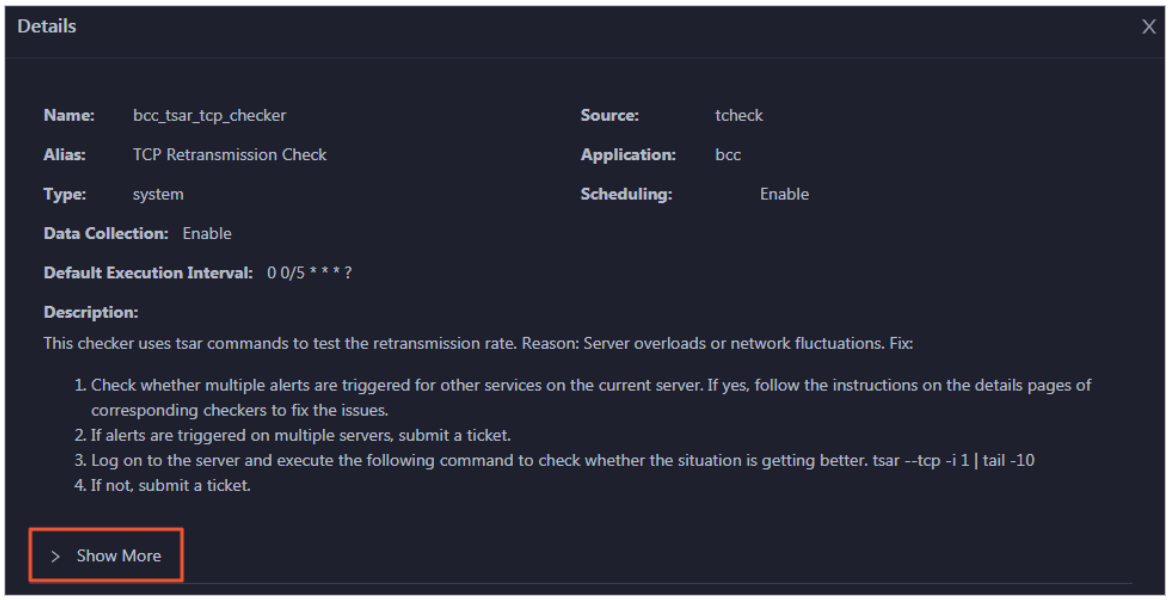

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

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

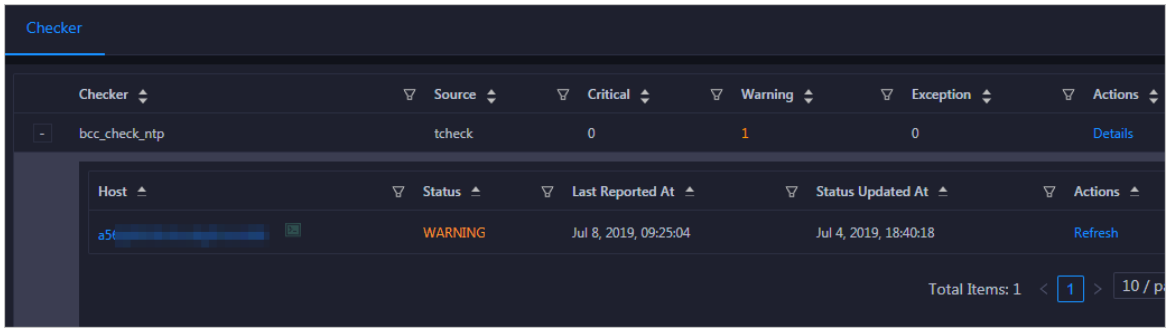

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

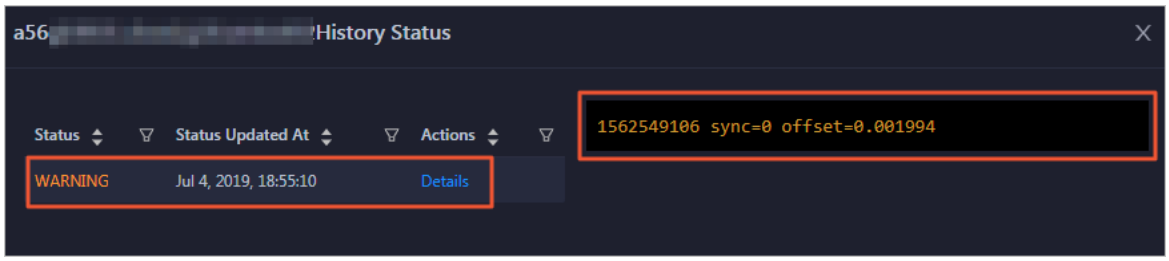

Clear alerts

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

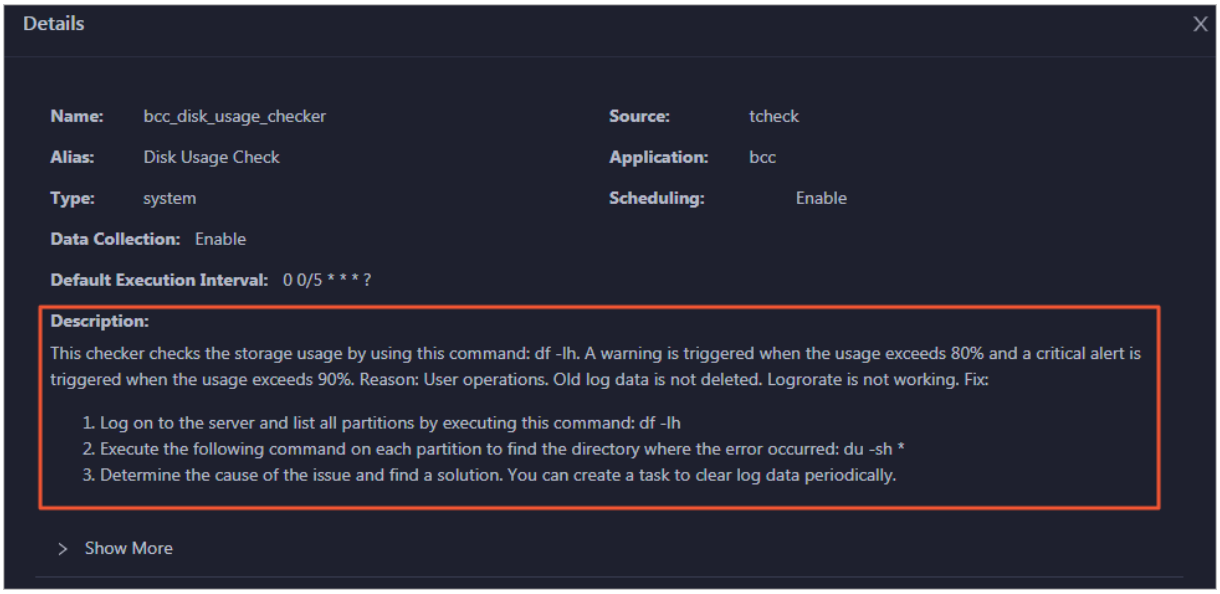

Log on to a host

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

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

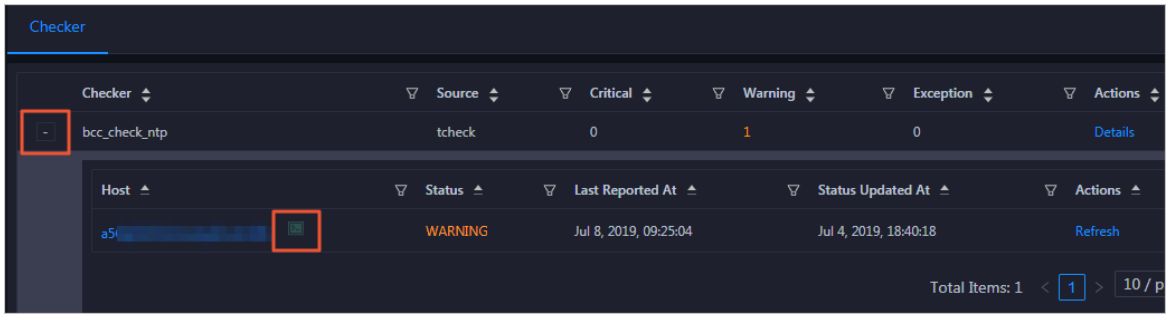

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

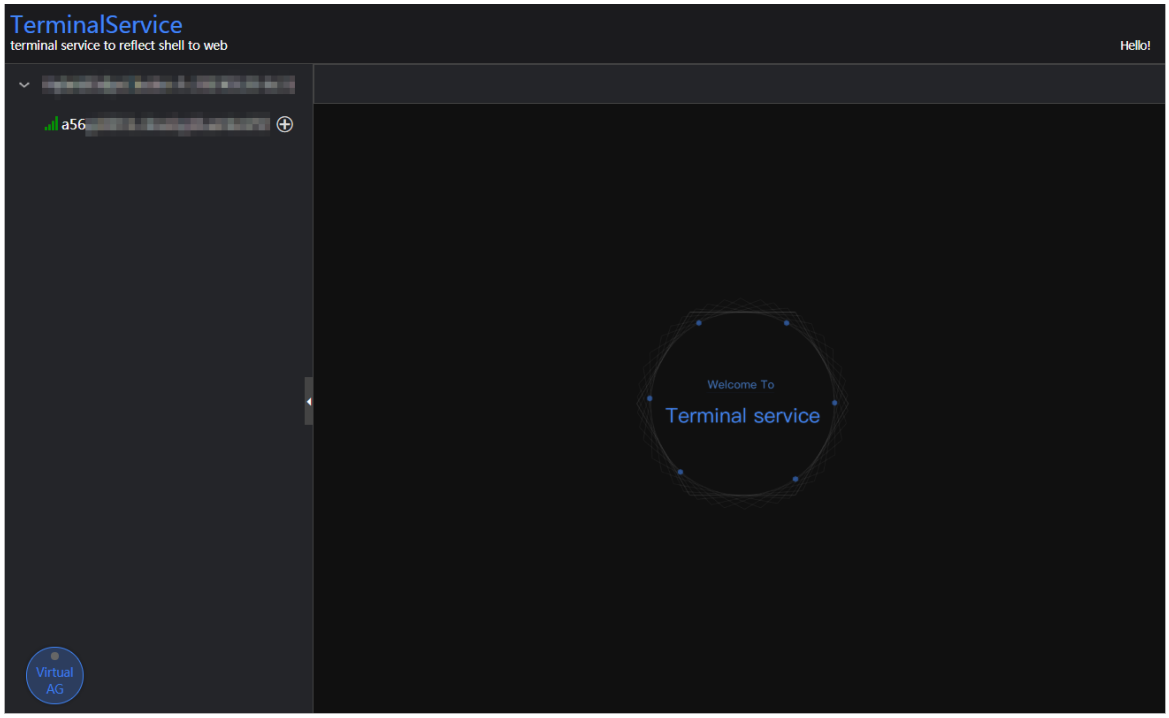

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

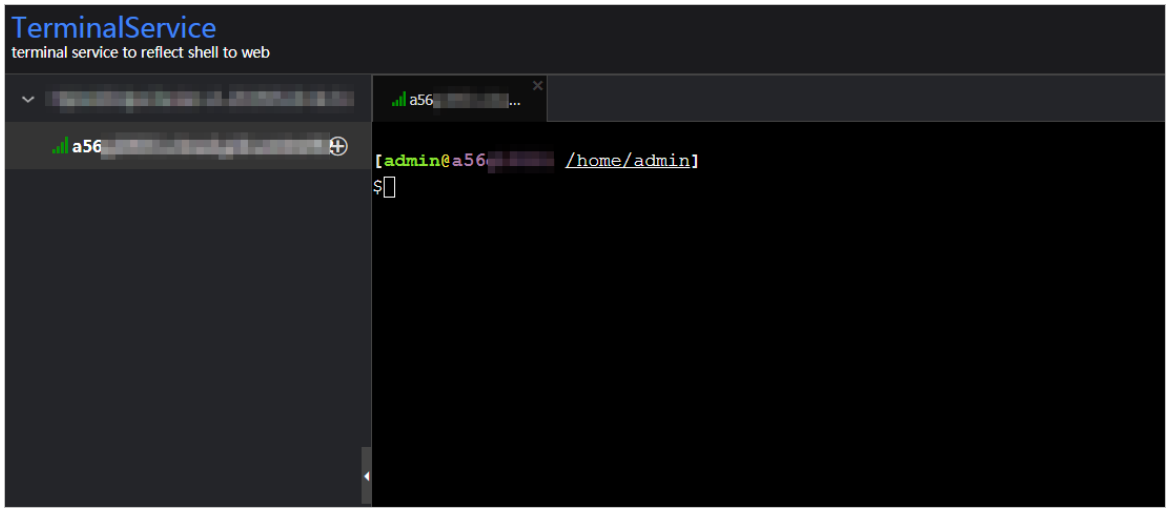

Run a checker again

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

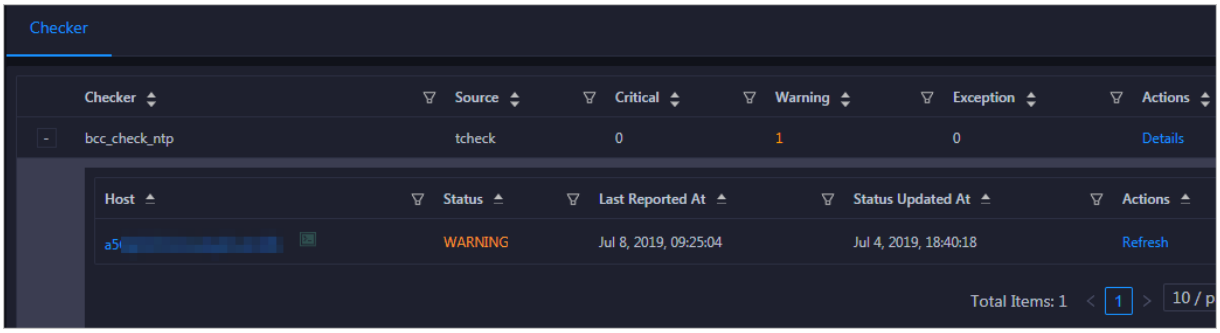

6.1.1.5.5.5. Host services

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

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

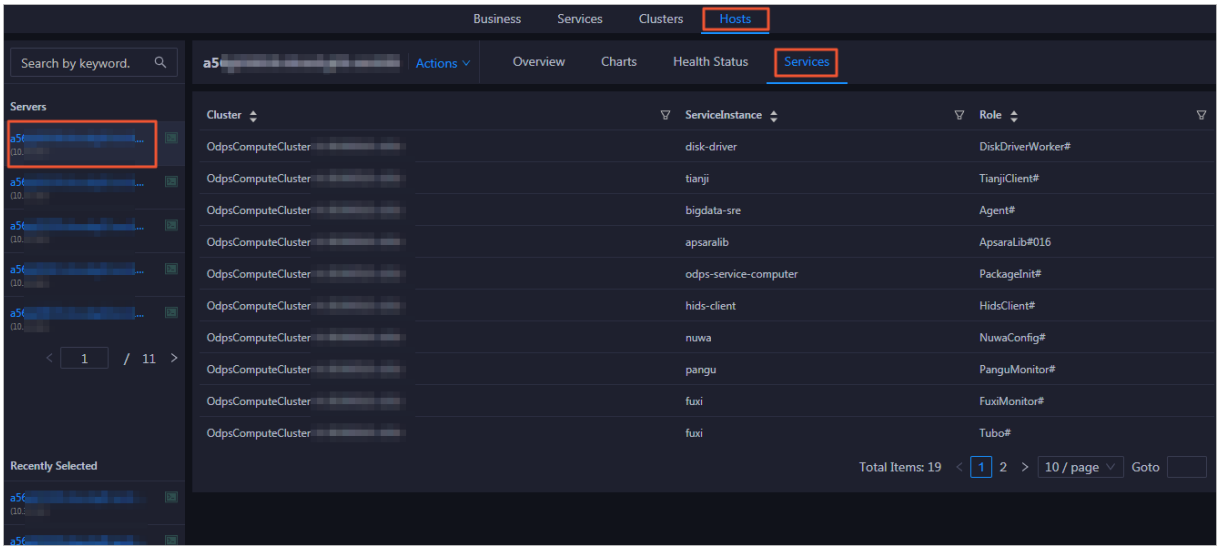

On the **Services** page, you can view the cluster, service instances, and service instance roles of the host.

6.1.1.6. Management

6.1.1.6.1. Overview

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

The management module supports the following features:

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

6.1.1.6.2. Jobs

6.1.1.6.2.1. Overview

This topic describes the Jobs page and concepts related to jobs.

Apsara Big Data Manager (ABM) runs jobs to perform O&M operations on big data services. Jobs, also known as service O&M tasks, are O&M operations performed on physical devices in the cluster. The Jobs page consists of the **Job Execution** and **Job Management** tabs.

Concepts

Concepts related to jobs include:

- Ordinary job: a job that can only be manually run.
- Cron job: a job that is automatically run based on timer settings.
- Scheme: a job template provided by ABM. You can use schemes to generate jobs.
- Atom: a step template provided by ABM. You can use atoms as steps when generating jobs.
- Ordinary step: a step that you need to create when using schemes to generate jobs. Step types include the following: command execution, script execution, file push, API call, and manual step.
- Atomic step: a step that you can directly use when using schemes to generate jobs.

ABM provides common schemes and atoms that support most O&M scenarios.

Job Execution page

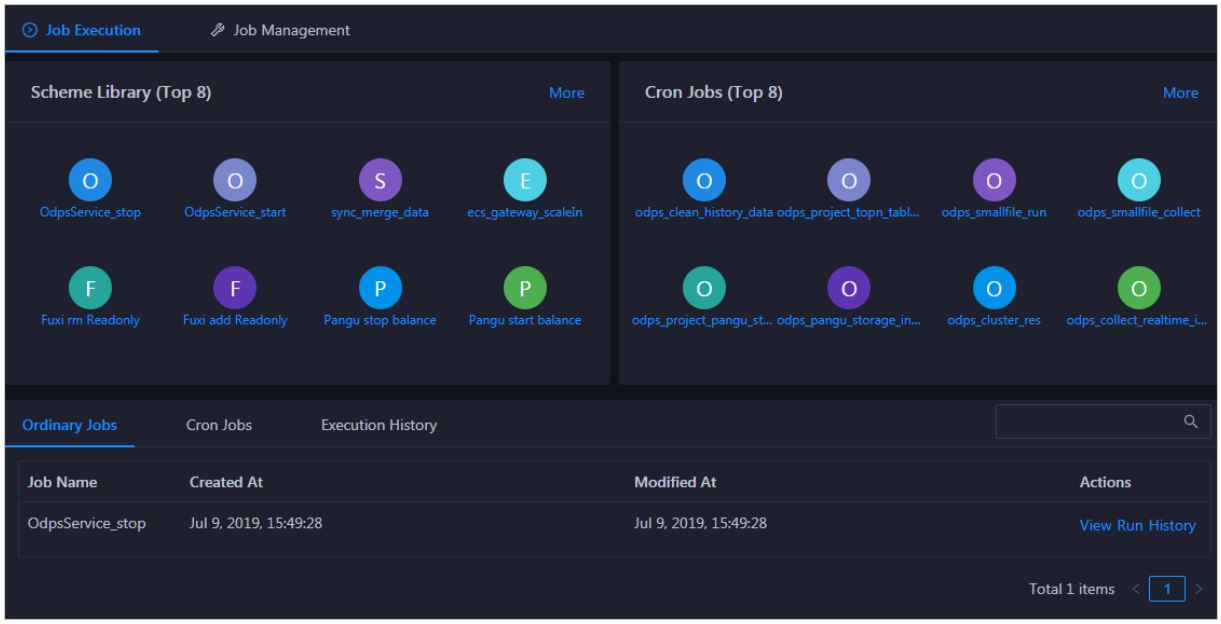

The **Job Execution** page provides the following features:

- Ordinary Jobs:

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

  You can search for a specific ordinary job.

- Cron Jobs:

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

  You can search for a specific cron job.

- Scheme Library (Top 8): dynamically displays the top 8 most used schemes.
- Cron Jobs (Top 8): dynamically displays the top 8 most used cron jobs.
- Execution History:

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

You can search for the execution record of a specific job by multiple conditions.

Job Management page

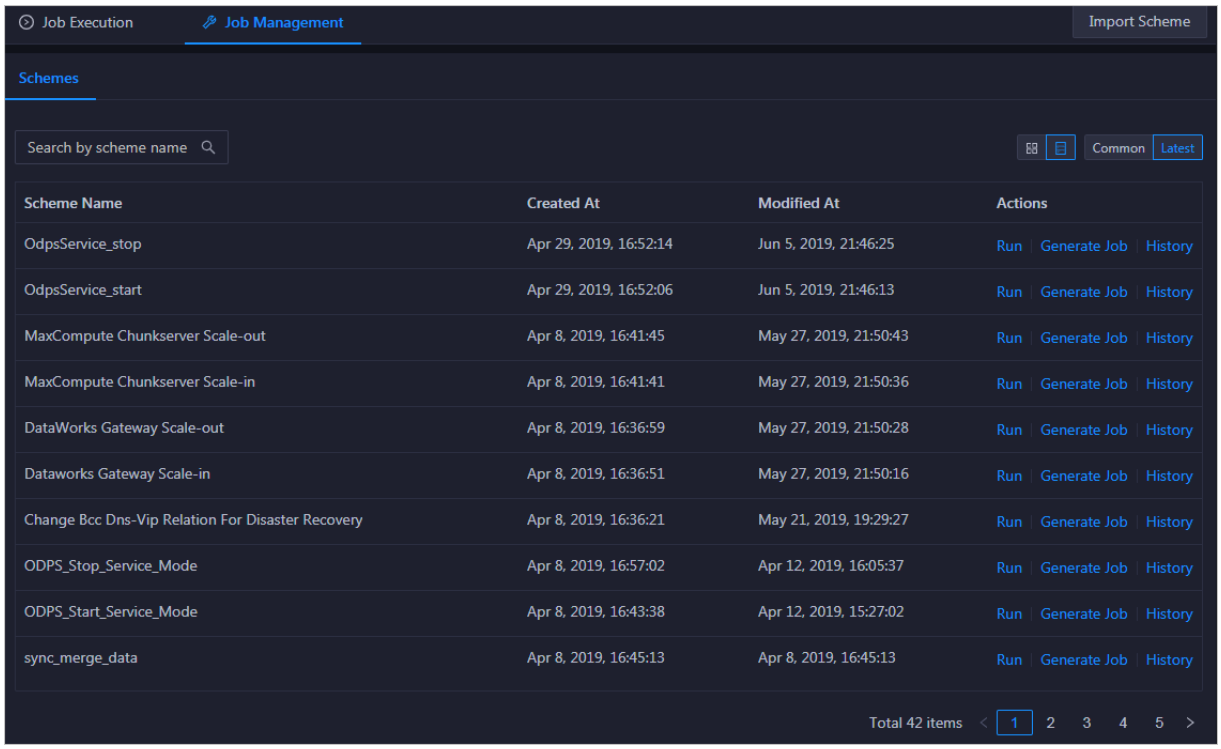

The **Job Management** page provides the following features:

- You can generate and run jobs based on schemes and view the execution history of schemes.
- You can search for a specific scheme.
- You can view schemes in grid or list mode.

6.1.1.6.2.2. Jobs

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.

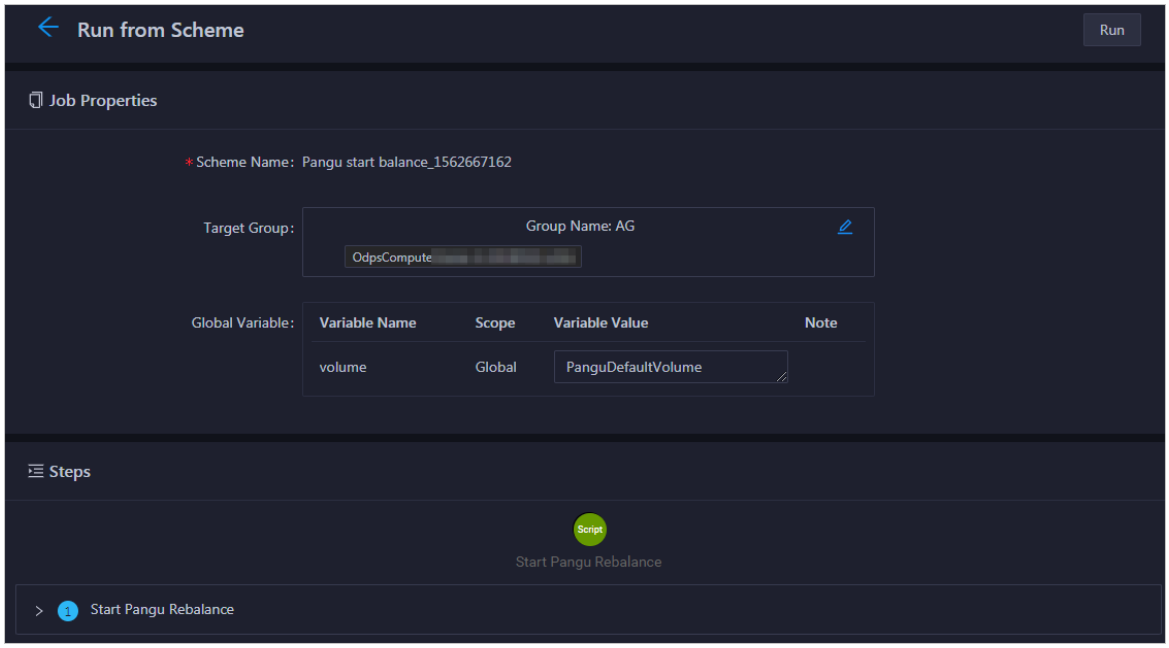

Job parameters

Parameter	Description

Parameter	Description
Target Group	A collection of target nodes on which the operations are performed. After you have added nodes to target groups, you can select a value for **Target Group** based on your needs when you configure the steps. Click next to the target group, and set the nodes to be included in the target group as needed. When you add a node, you can either select the name of the node in **Apsara Infrastructure Management Framework** or enter the IP address of the node under **Servers**.
Global Variable	If global variables are set in the scheme, you need to enter the variable value.

5. After you have configured the preceding parameters, click **Run** in the upper-right corner.
6. Confirm the job risks in the displayed dialog box, and click **Confirm Execution**.

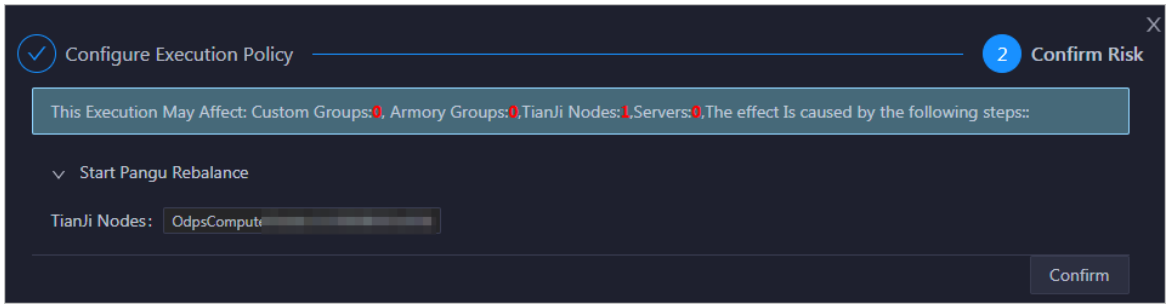

After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see View the execution history.

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

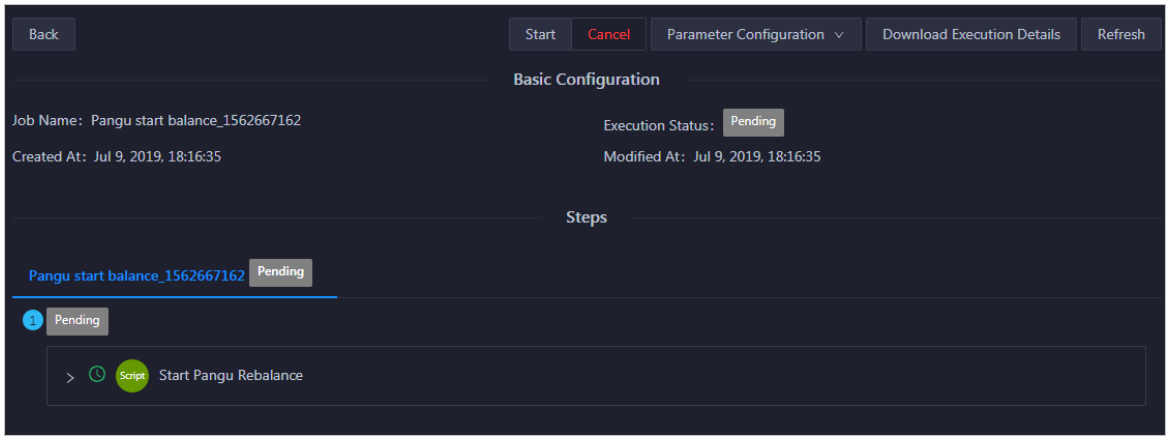

If you do not perform any operation and exit the job execution page, you can find a job record on the **Execution History** page. Click **View** to go to the job execution page again.

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.

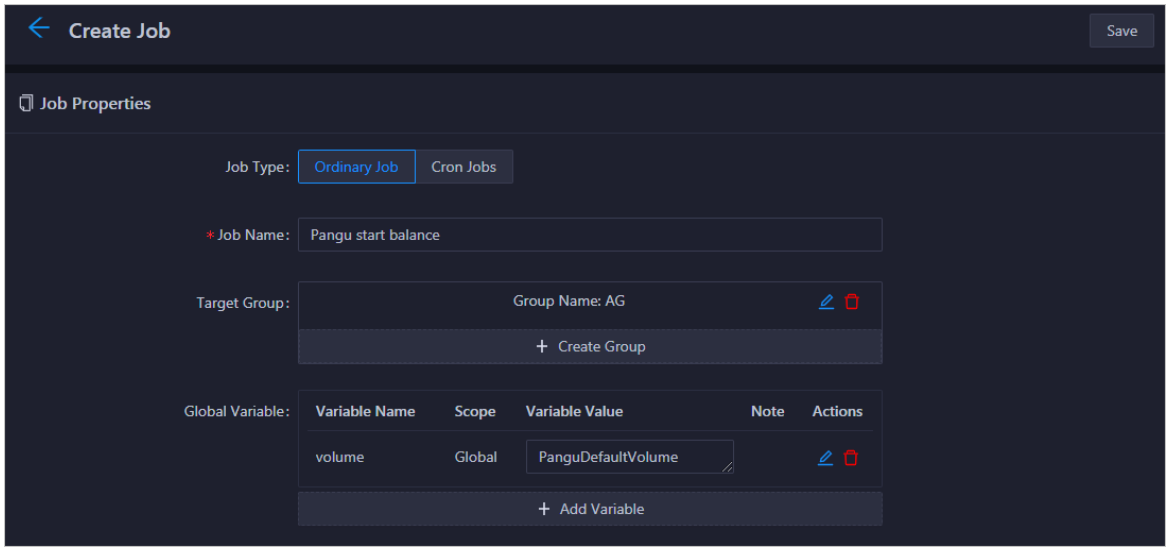

Job properties

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.

Parameter	Description
Job Name	The name of the job. Set the job name based on the functionality of the job to be created so that the user understands what it is and can search for it.
Target Group	A collection of target nodes on which the operations are performed. After you have added nodes to the target groups, you can select the target group based on your needs when you configure the steps. After you have created a group, click to add nodes to the group. When you add a node, you can either select the name of the node in **Tianji** or enter the IP address of the node under **Servers**.
Global Variable	Click **Add Variable** and set the parameters in the dialog box that appears. The **Scope** parameter is used to set the scope of the variable. If it is set **Global**, it is valid for the entire job. If you select a certain step, it is only valid for this step.

5. On the **Create Job** page, add steps as needed.

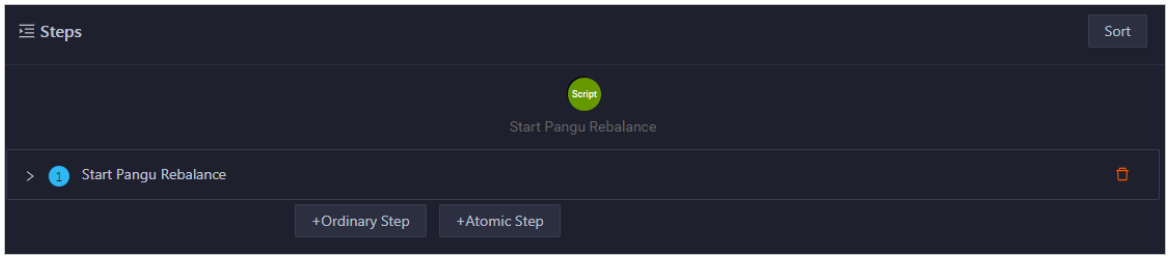

The steps include ordinary steps and atomic steps.

- **+Atomic Step**: a range of built-in steps provided by the system.
- **+Ordinary Step**: Ordinary steps are classified into multiple types. Choose the required type and set the parameters accordingly. 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**.

Section	Parameter	Description
	Description	The description of the step.
Advanced Configuration	**Input Context**	Enable this option if you need to obtain the output of the previous step. When enabled, this step reads the file specified by the *$contextInput* variable to obtain the context.
	Output Context	Enable this option if you need to export the context to the next step. When enabled, this step writes the context to the file specified by the *$contextOutput* variable to export the context.
	Timeout Period	The maximum time period allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out. The default value is **60** seconds.
	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.

Section	Parameter	Description
Advanced Configuration	Input Context	Enable this option if you need to obtain the output of the previous step. When enabled, this step reads the file specified by the *$contextInput* variable to obtain the context.
	Output Context	Enable this option if you need to export the context to the next step. When enabled, this step writes the context to the file specified by the *$contextOutput* variable to export the context.
	Timeout Period	The maximum time period allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out. The default value is **60** seconds.
	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.
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.

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**.

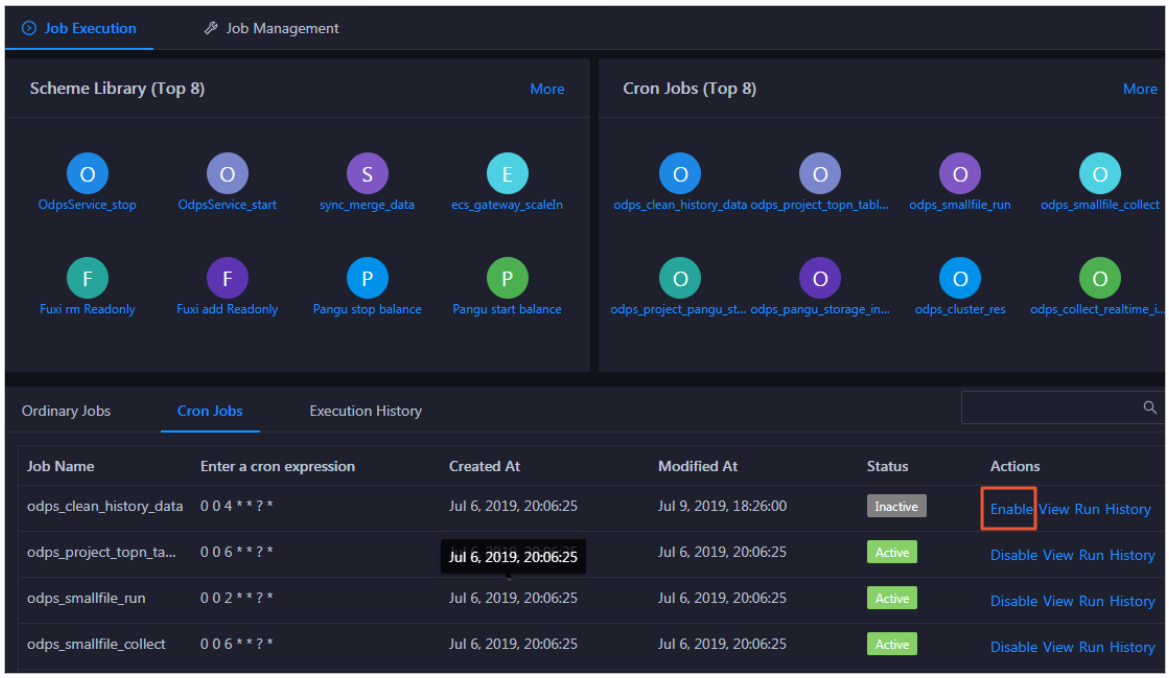

4. On the **Cron Jobs** page, you can enable or disable a cron job.
   - To enable a cron job in the inactive status, click **Enable** in the Actions column of the cron job.

     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**.

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.

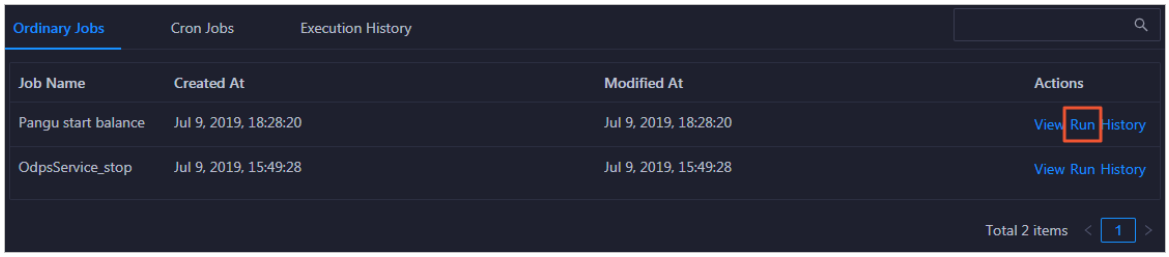

4. In the **Ordinary Jobs** list, click **Run** in the Actions column of a job.
5. Confirm the job risks in the dialog box that appears, and click **Confirm**.

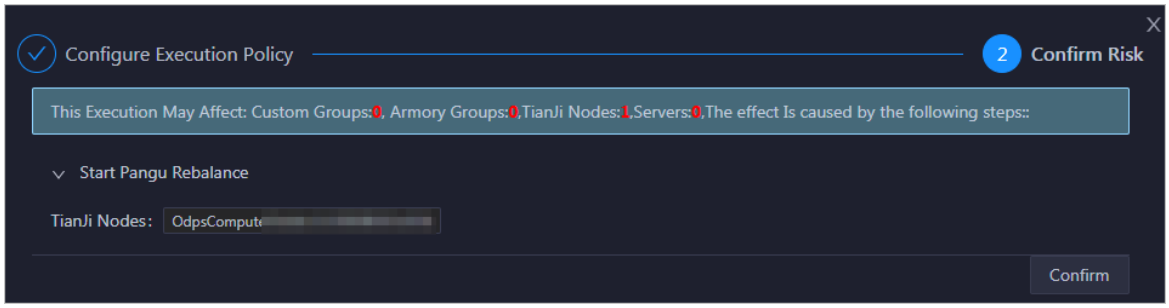

   After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see View the execution history.

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

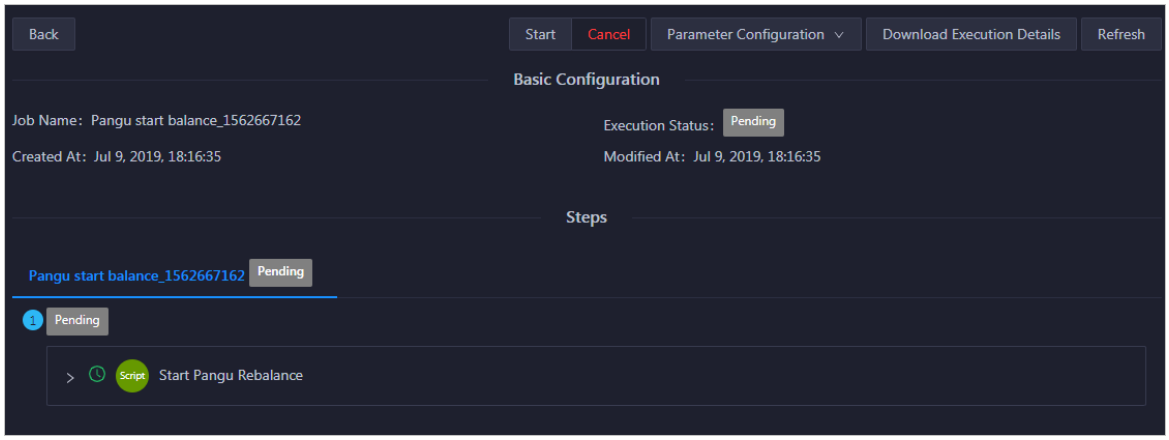

You can find the record about a job on the **Execution History** page, and click **View** to go to the detailed execution page.

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.

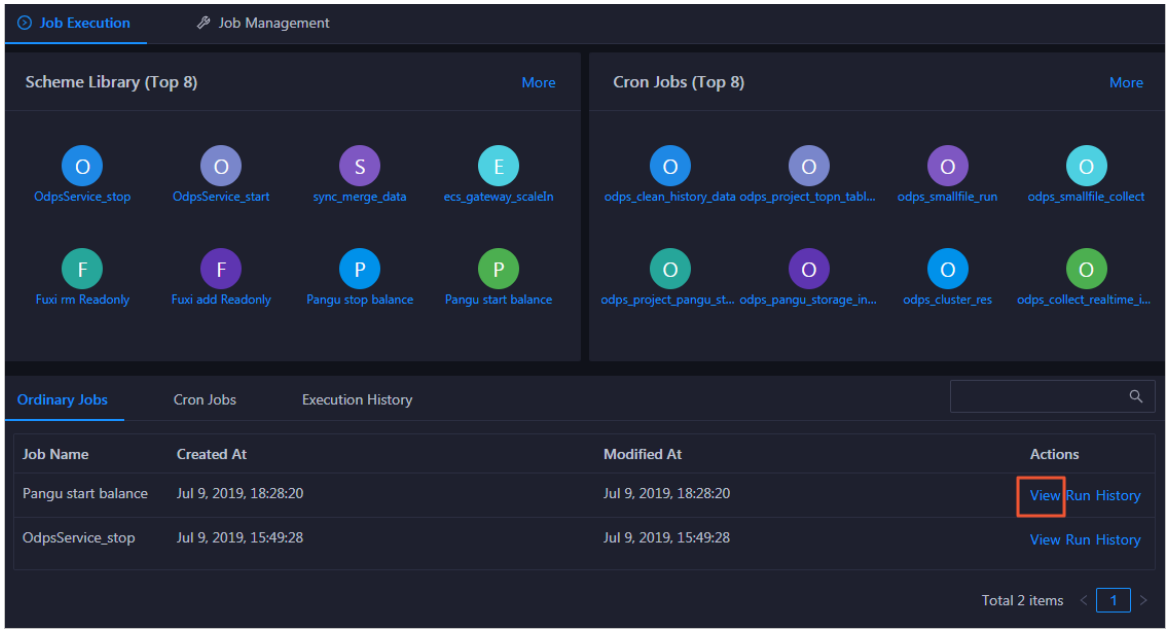

View the execution history of a job

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.

6.1.1.6.2.3. Schemes

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.

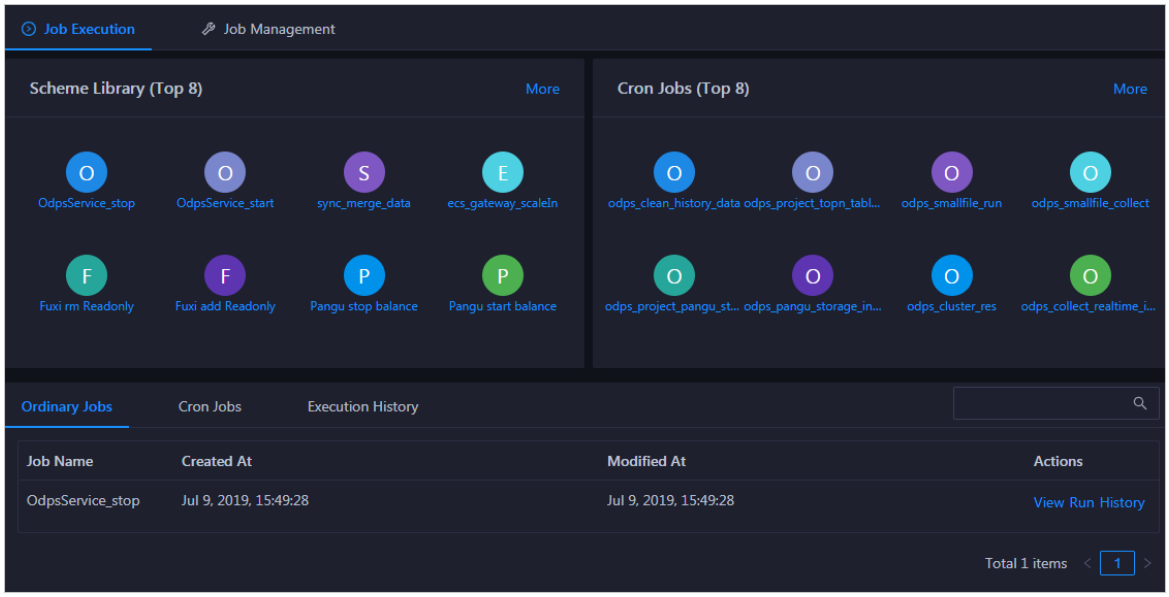

5. On the **Job Details** page, click **Save as Scheme** in the upper-right corner. The system prompts that you have saved the scheme.

Result

The new scheme has the same name as the job from which it was created and is listed on the **Schemes** page.

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**.

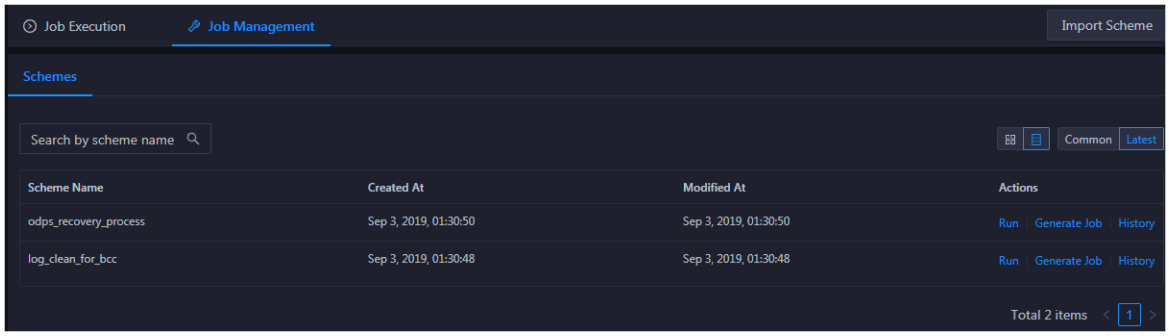

4. If there are too many schemes, you can enter the scheme name in the search bar to search for the required scheme.
5. Change the method for viewing schemes
   - View schemes in list (default): Click in the upper-right corner.
   - View schemes in cards: Click in the upper-right corner.

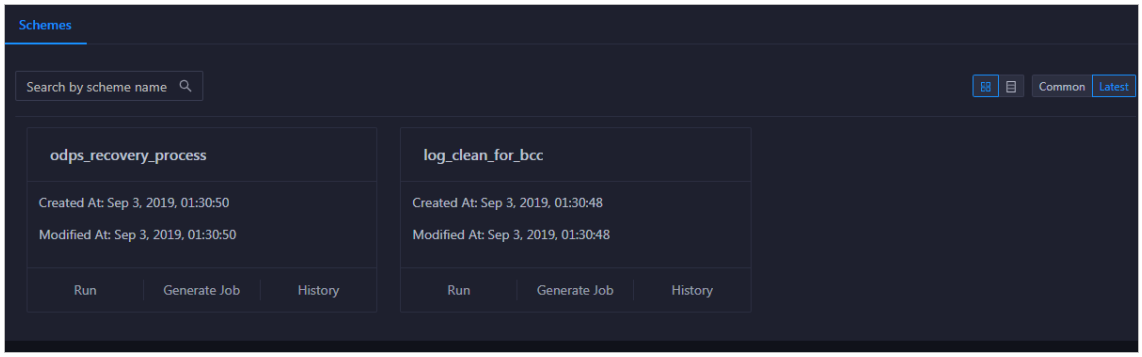

View the execution history of a scheme

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.

6.1.1.6.2.4. 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.

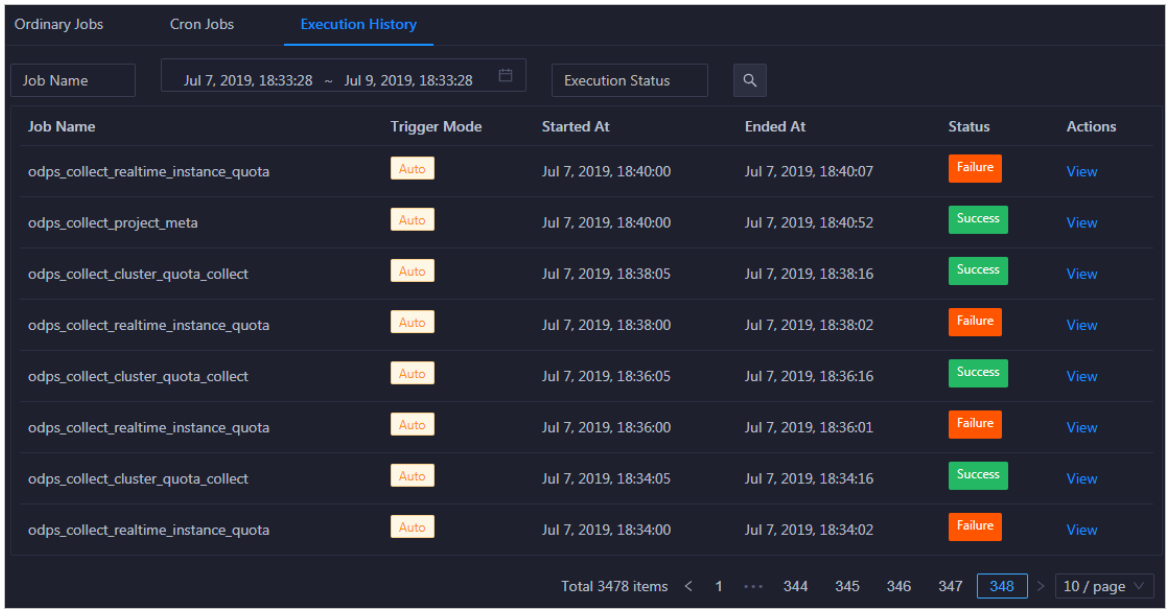

4. If there are too many execution records, filter them by a combination of one or more of the following filter conditions: job name, creator, execution status, and time range. Then, click to search for required records.
5. Click **View** in the Actions column of a record to view the execution details.

   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.

6.1.1.6.3. 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**.

6.1.1.6.4. 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**.

6.1.1.6.5. 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.

6.1.1.6.6. Operation auditing

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

Prerequisites

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

Background

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

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

Procedure

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

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

You can also view the causes of failed steps in detail.

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

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

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

6.1.1.7. Go to other consoles

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

Prerequisites

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

Procedure

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

Result

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

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

6.2. MaxCompute

6.2.1. Operations and Maintenance Guide

6.2.1.1. Concepts and architecture

The topic describes the concepts and architecture of MaxCompute. The architecture and descriptions are for reference only. They are subject to the released product type and supplementary features.

Architecture

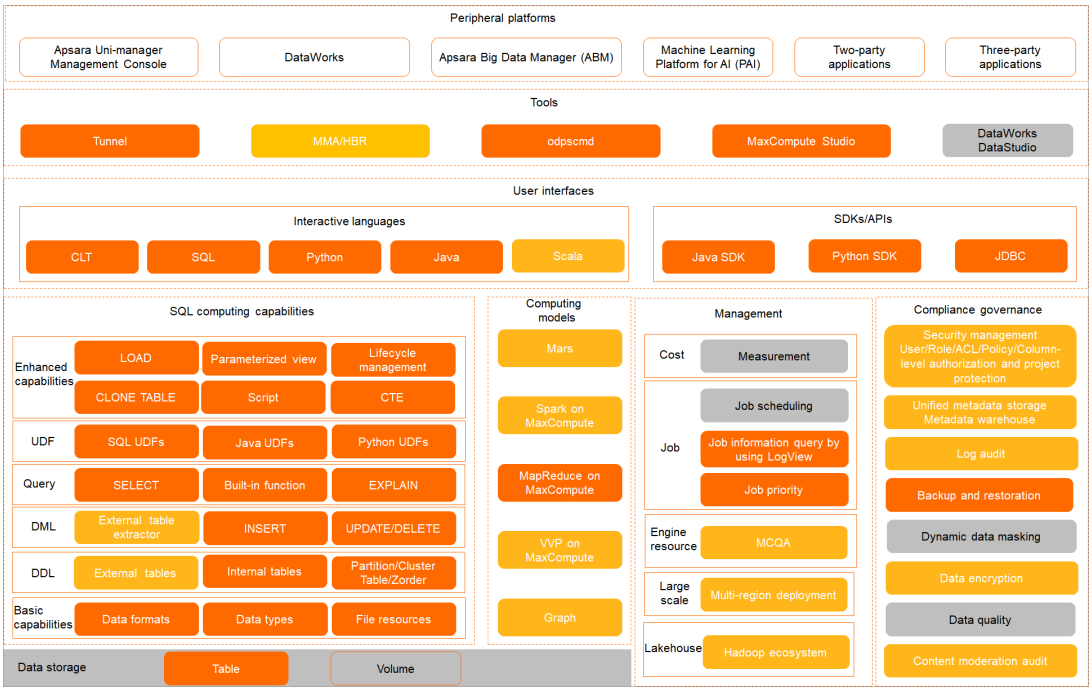

indicates the basic features of MaxCompute. indicates the enhanced features of MaxCompute. indicates the features provided by external systems.

Category	Description
Peripheral platforms	MaxCompute supports the following peripheral platforms: • Apsara Uni-manager Management Console: a unified and intelligent O&M platform. For more information, see *Apsara Uni-manager Management Console User Guide*. • DataWorks: a visualization tool. You can use DataWorks to perform common operations, such as synchronize data, schedule jobs, and generate reports. For more information, see *DataWorks Technical White Paper*. • Apsara Big Data Manager (ABM): provides an easy method for field engineers to manage MaxCompute. For more information, see *Apsara Big Data Manager Technical White Paper*. • Machine Learning Platform for AI (PAI): a machine learning algorithm platform based on MaxCompute. For more information, see *Machine Learning Platform for AI Technical White Paper*. • Two-party applications: other Alibaba Cloud services supported by MaxCompute, such as DataV. • Three-party applications: other services that are compatible with MaxCompute.
Tools	MaxCompute supports the following tools: • Tunnel: a tunnel service. MaxCompute allows you to import heterogeneous data into or export the data from MaxCompute by using Tunnel. For more information, see *Tunnel* in *MaxCompute Product Introduction*. • MaxCompute Migration Assist (MMA): the data migration tool of MaxCompute. If you use MMA, Meta Carrier is used to access your Hive metastore service and capture Hive metadata. Then, MMA uses the Hive metadata to generate data definition language (DDL) statements and SQL statements of Hive user-defined table-valued functions (UDTFs). The DDL statements are used to create MaxCompute tables and their partitions. The SQL statements of Hive UDTFs are used to migrate data. • Hybrid backup recovery (HBR): integrates data backup and migration capabilities of Apsara Stack. • odpscmd: the MaxCompute client. For more information, see *Client* in *MaxCompute User Guide*. • MaxCompute Studio: the big data integrated development environment tool that is provided by MaxCompute. MaxCompute Studio is installed on a developer client. It is a development plug-in that Alibaba Cloud provides for the popular integrated development environment (IDE) IntelliJ IDEA. • DataWorks DataStudio: a visualized development platform provided by DataWorks. For more information, see *DataWorks User Guide*.

Category	Description
User interfaces	MaxCompute supports the following interfaces: • Interactive languages: CLI, SQL, Python, Java, and Scala. • SDKs and APIs: SDK for Java, SDK for Python, and Java Database Connectivity (JDBC). For more information, see *MaxCompute Developer Guide*.
SQL computing capabilities	MaxCompute supports the following SQL computing capabilities: • Enhanced capabilities: support LOAD, parameterized view, lifecycle management, and CLONE TABLE. • User-defined functions (UDFs): include SQL UDFs, Java UDFs, and Python UDFs. • Query: the query operations, such as SELECT and EXPLAIN statements and built-in functions. • Data manipulation language (DML) statements: include INSERT, UPDATE, and DELETE. • DDL statements: allow you to create internal tables, external tables, clustered tables, and partitioned tables. • Basic capabilities: support multiple data types and data formats and allow you to upload resource files. For more information, see *MaxCompute SQL* in *MaxCompute User Guide*.
Computing models	MaxCompute supports the following computing models: • Mars: a tensor-based unified distributed computing framework. Mars can use parallel and distributed computing technologies to accelerate data processing for Python data science stacks. For more information, see *Mars* in *MaxCompute User Guide*. • Spark on MaxCompute: a solution developed by Alibaba Cloud to enable the seamless use of Spark on the MaxCompute platform. It supplements a wide variety of features to MaxCompute. For more information, see *Spark on MaxCompute* in *MaxCompute User Guide*. • MapReduce on MaxCompute: allows you to run MapReduce jobs on MaxCompute. For more information, see *MaxCompute MapReduce* in *MaxCompute User Guide*. • VVP on MaxCompute: encapsulates the features of Realtime Compute for Apache Flink that is developed on the Ververica Platform (VVP) based on MaxCompute resources. You can use the Cupid joint computing platform to complete the operations related to real-time computing by using the underlying storage and computing resources of MaxCompute on the VVP UI. For more information, see *VVP On MaxCompute* in *MaxCompute User Guide*. • Graph: a processing framework designed for iterative graph computing. For more information, see *MaxCompute Graph* in *MaxCompute User Guide*.

Category	Description
Management	MaxCompute can be managed from the following aspects: • Cost: measures resource usage. • Job: provides mechanisms to manage jobs. For example, you can use these mechanisms to schedule jobs, use LogView to view job information, and set job priorities. • Engine resource: supports high-performance MaxCompute Query Acceleration (MCQA). • Large scale: allows you to deploy MaxCompute clusters across regions. • Lakehouse: a data management platform that combines data lakes and data warehouses. It integrates the flexibility and diverse ecosystems of data lakes with the enterprise-class deployment of data warehouses. For more information, see *MaxCompute Operations and Maintenance Guide*.
Compliance governance	MaxCompute allows you to use the following methods for compliance governance: • Security management: allows you to control the permissions of users and roles, and supports multiple authorization methods, such as ACL-based, policy-based, and column-level authorization. • Unified metadata storage: stores metadata in a centralized manner. • Log audit: audits different log data of different users. • Backup and restoration: allows you to back up and restore data from a storage system. • Dynamic data masking: allows you to query data masking rules in DataWorks. • Data encryption: uses Key Management Service (KMS) to encrypt data for storage. This way, MaxCompute can provide static data protection to meet the requirements of enterprise governance and security compliance. • Data quality: DataWorks provides an end-to-end platform that supports quality verification, notification, and management services for various heterogeneous data sources. • Content moderation audit: uses the content moderation engine to identify and audit pornographic, violent, and illegal content. For more information about security, see *MaxCompute Security White Paper*.
Data storage	MaxCompute stores data as tables or volumes.

The following figure shows how a MaxCompute job is run.

Procedure to run a MaxCompute job

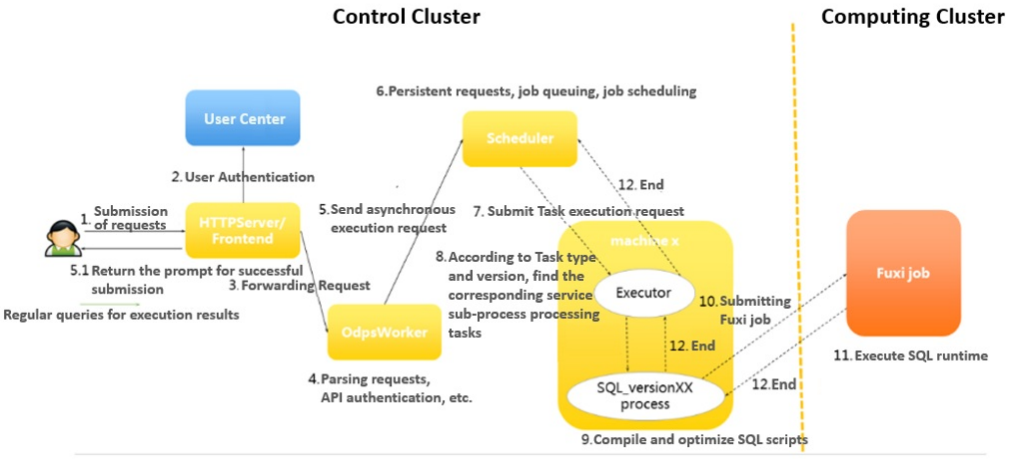

The following concepts are involved in the procedure to run a MaxCompute job.

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 in sequence. This application is rarely implemented because MaxCompute jobs are not commonly used. In most cases, an instance contains only one task.
2. MaxCompute task: a specific task in MaxCompute. Almost 20 task types, such as SQL, MapReduce, Admin, Lot, and Xlib, are supported. The execution logic varies greatly based on the task type. Different tasks in an instance are differentiated by their task name. MaxCompute tasks 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 directed acyclic graph (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 tasks are available. This reduces the time required for starting and stopping a job.
   - You can submit a MaxCompute task to multiple compute clusters. The primary key name of a Fuxi job is in the format of cluster name + 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. When a task is executed, it 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.

6.2.1.2. O&M commands and tools

6.2.1.2.1. 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.

6.2.1.2.2. 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

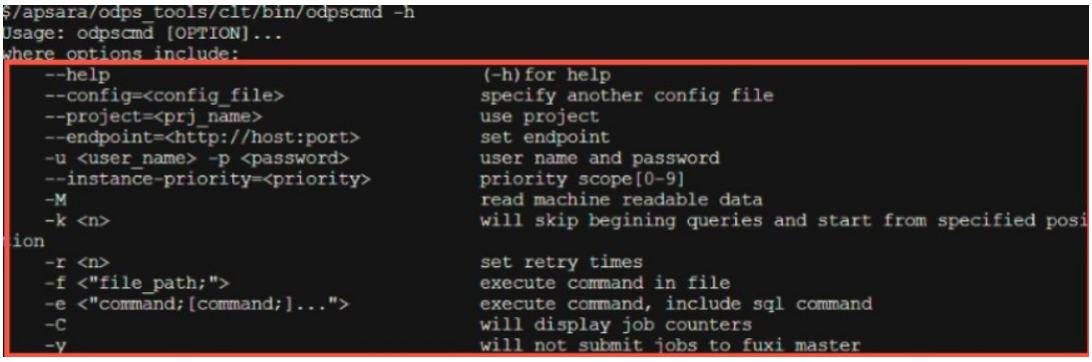

- **-e**: The MaxCompute client does not execute SQL statements in interactive mode.
- **--project, -u, and -p**: The client directly uses the specified values for the project, user, and pass parameters. If you do not specify a parameter, the client uses the corresponding value configured in the conf 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.
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.

Command	Description
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.

6.2.1.2.3. 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.

Command	Description
tunnel resume	If an error occurs because of network or Tunnel service faults, you can resume file or directory transmission after interruption. This command only allows you to resume the previous data upload. Every data upload or download operation is called a session. Run the resume command and specify the ID of the session to be resumed.
tunnel show	Allows you to view historical task information.
tunnel purge	Purges the session directory. Sessions from the last three days are purged by default.

Tunnel commands allow you to view help information by using the Help sub-command on the client. The sub-commands of each Tunnel command are described as follows:

Upload

Imports data of a local file into a MaxCompute table. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help upload;
usage: tunnel upload [options] <path> <[project.]table[/partition]>
       upload data from local file
 -acp,-auto-create-partition <ARG>  auto create target partition if not
                    exists, default false
 -bs,-block-size <ARG>          block size in MiB, default 100
 -c,-charset <ARG>              specify file charset, default ignore.
                    set ignore to download raw data
 -cp,-compress <ARG>            compress, default true
 -dbr,-discard-bad-records <ARG>   specify discard bad records
                    action(true|false), default false
 -dfp,-date-format-pattern <ARG>   specify date format pattern, default
                    yyyy-MM-dd HH:mm:ss
 -fd,-field-delimiter <ARG>       specify field delimiter, support
                    unicode, eg \u0001. default ","
 -h,-header <ARG>               if local file should have table
                    header, default false
 -mbr,-max-bad-records <ARG>      max bad records, default 1000
 -ni,-null-indicator <ARG>        specify null indicator string,
                    default ""(empty string)
 -rd,-record-delimiter <ARG>      specify record delimiter, support
                    unicode, eg \u0001. default "\r\n"
 -s,-scan <ARG>                 specify scan file
                    action(true|false|only), default true
 -sd,-session-dir <ARG>         set session dir, default
                    D:\software\odpscmd_public\plugins\ds
                    hip
 -ss,-strict-schema <ARG>         specify strict schema mode. If false,
                    extra data will be abandoned and
                    insufficient field will be filled
                    with null. Default true
 -te,-tunnel_endpoint <ARG>       tunnel endpoint
  -threads <ARG>                number of threads, default 1
 -tz,-time-zone <ARG>           time zone, default local timezone:
                    Asia/Shanghai
Example:
  tunnel upload log.txt test_project.test_table/p1="b1",p2="b2"
```

Parameters:

- -acp: indicates whether to automatically create the destination partition if it does not exist. No destination partition is created by default.
- -bs: specifies the size of each data block uploaded with Tunnel. Default value: 100 MiB (MiB = 1024 * 1024B).
- -c: specifies the local data file encoding format. Default value: UTF-8. If this parameter is not set, the encoding format of the downloaded source data is used by default.
- -cp: indicates whether to compress the local data file before it is uploaded to reduce network traffic. By default, the local data file is compressed before it is uploaded.
- -dbr: indicates whether to ignore dirty data (such as additional columns, missing columns, and columns with mismatched data types).
  - If this parameter is set to true, all data that does not comply with table definitions is ignored.

  - If this parameter is set to false, an error is returned when dirty data is found, so that raw data in the destination table is not contaminated.

- -dfp: specifies the DateTime format. Default value: yyyy-MM-dd HH:mm:ss.
- -fd: specifies the column delimiter used in the local data file. Default value: comma (,).
- -h: indicates whether the data file contains the header. If this parameter is set to true, Dship skips the header row and starts uploading data from the second row.
- -mbr: terminates any attempts to upload more than 1,000 rows of dirty data. This parameter allows you to adjust the maximum allowable volume of dirty data.
- -ni: specifies the NULL data identifier. Default value: an empty string ("").
- -rd: specifies the row delimiter used in the local data file. Default value: \r\n.
- -s: indicates whether to scan the local data file. Default value: false.
  - If this parameter is set to true, the system scans the source data first, and then imports the data if the format is correct.
  - If this parameter is set to false, the system imports data directly without scanning.
  - If this parameter is set to only, the system only scans the source data, and does not import the data after scanning.
- -sd: sets the session directory.
- -te: specifies the Tunnel endpoint.
- -threads: specifies the number of threads. Default value: 1.
- -tz: specifies the time zone. Default value: Asia/Shanghai.

Show

Displays historical records. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help show;
usage: tunnel show history [options]
        show session information
 -n,-number <ARG>  lines
Example:
  tunnel show history -n 5
  tunnel show log
```

Parameters:

-n: specifies the number of rows to be displayed.

Resume

Resumes the execution of historical operations (only applicable to data upload). The following example shows how to use the sub-commands:

```
odps@  project_name>tunnel help resume;
usage: tunnel resume [session_id] [-force]
        resume an upload session
 -f,-force  force resume
Example:
  tunnel resume
```

Download

The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help download;
usage: tunnel download [options] <[project.]table[/partition]> <path>
       download data to local file
 -c,-charset <ARG>            specify file charset, default ignore.
                              set ignore to download raw data
 -ci,-columns-index <ARG>     specify the columns index(starts from
                              0) to download, use comma to split each
                              index
 -cn,-columns-name <ARG>      specify the columns name to download,
                              use comma to split each name
 -cp,-compress <ARG>          compress, default true
 -dfp,-date-format-pattern <ARG>  specify date format pattern, default
                              yyyy-MM-dd HH:mm:ss
 -e,-exponential <ARG>        When download double values, use
                              exponential express if necessary.
                              Otherwise at most 20 digits will be
                              reserved. Default false
 -fd,-field-delimiter <ARG>   specify field delimiter, support
                              unicode, eg \u0001. default ","
 -h,-header <ARG>             if local file should have table header,
                              default false
  -limit <ARG>                specify the number of records to
                              download
 -ni,-null-indicator <ARG>    specify null indicator string, default
                              ""(empty string)
 -rd,-record-delimiter <ARG>  specify record delimiter, support
                              unicode, eg \u0001. default "\r\n"
 -sd,-session-dir <ARG>       set session dir, default
                              D:\software\odpscmd_public\plugins\dshi
                              p
 -te,-tunnel_endpoint <ARG>   tunnel endpoint
  -threads <ARG>              number of threads, default 1
 -tz,-time-zone <ARG>         time zone, default local timezone:
                              Asia/Shanghai
usage: tunnel download [options] instance://<[project/]instance_id> <path>
       download instance result to local file
 -c,-charset <ARG>            specify file charset, default ignore.
                              set ignore to download raw data
 -ci,-columns-index <ARG>     specify the columns index(starts from
                              0) to download, use comma to split each
                              index
 -cn,-columns-name <ARG>      specify the columns name to download,
                              use comma to split each name
 -cp,-compress <ARG>          compress, default true
 -dfp,-date-format-pattern <ARG>  specify date format pattern, default
                              yyyy-MM-dd HH:mm:ss
 -e,-exponential <ARG>        When download double values, use
                              exponential express if necessary.
                              Otherwise at most 20 digits will be
                              reserved. Default false
 -fd,-field-delimiter <ARG>   specify field delimiter, support
```

```
-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
```

6.2.1.2.4. LogView tool

6.2.1.2.4.1. 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.

6.2.1.2.4.2. 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

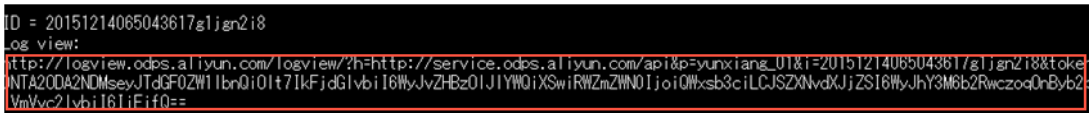

Enter the string with all carriage return and line feed characters removed in the address bar of the browser.

Composition of a LogView string

A LogView string consists of five parts, as shown in the following figure.

Composition of a LogView string

6.2.1.2.4.3. 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

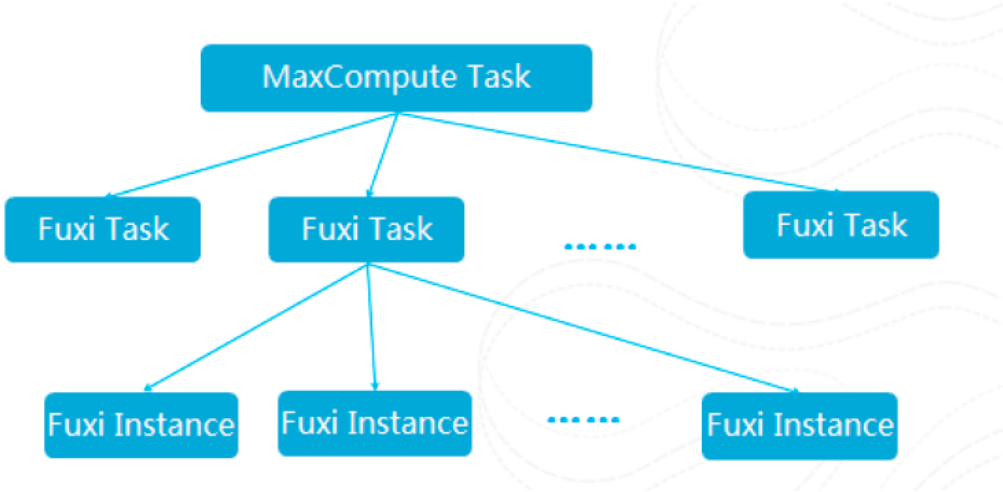

The following figures show the relevant information in LogView.

MaxCompute Instance

MaxCompute Instance

MaxCompute Instance

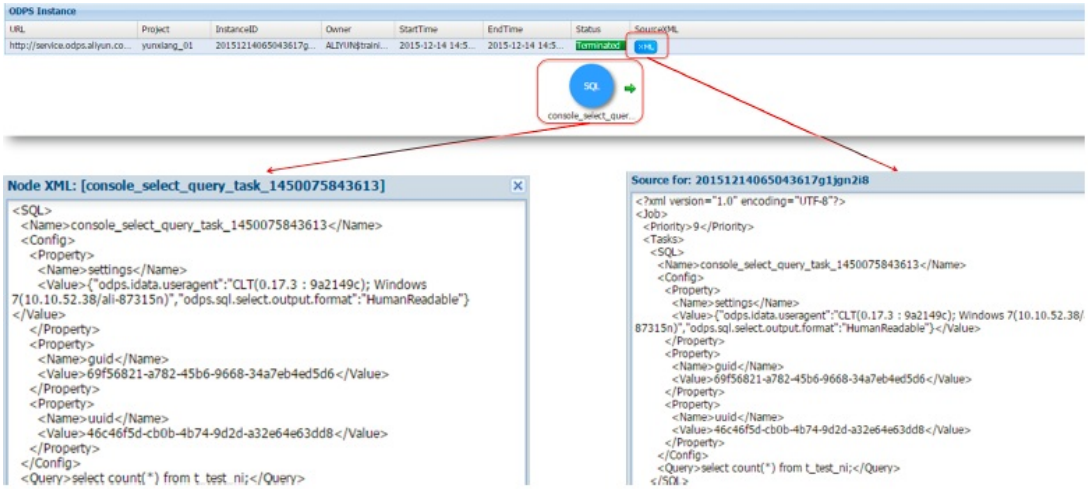

MaxCompute Task

MaxCompute Task

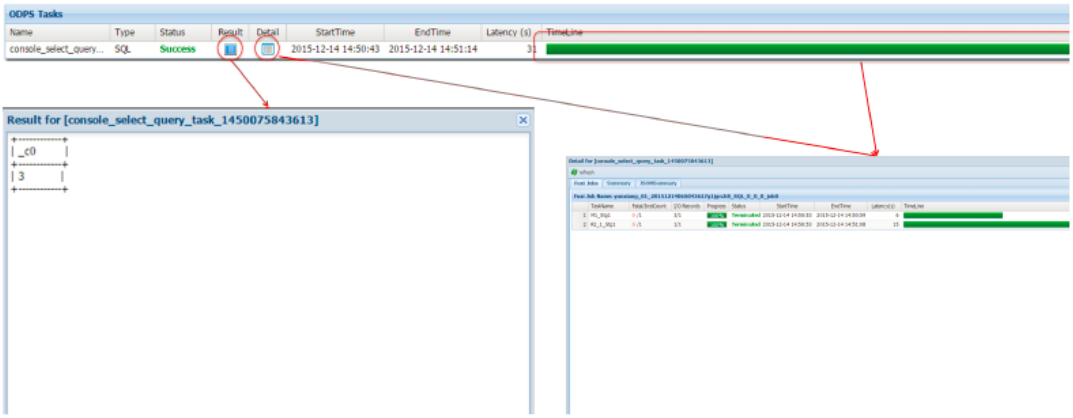

Task Detail - Fuxi Job

Task Detail - Fuxi Job(1)

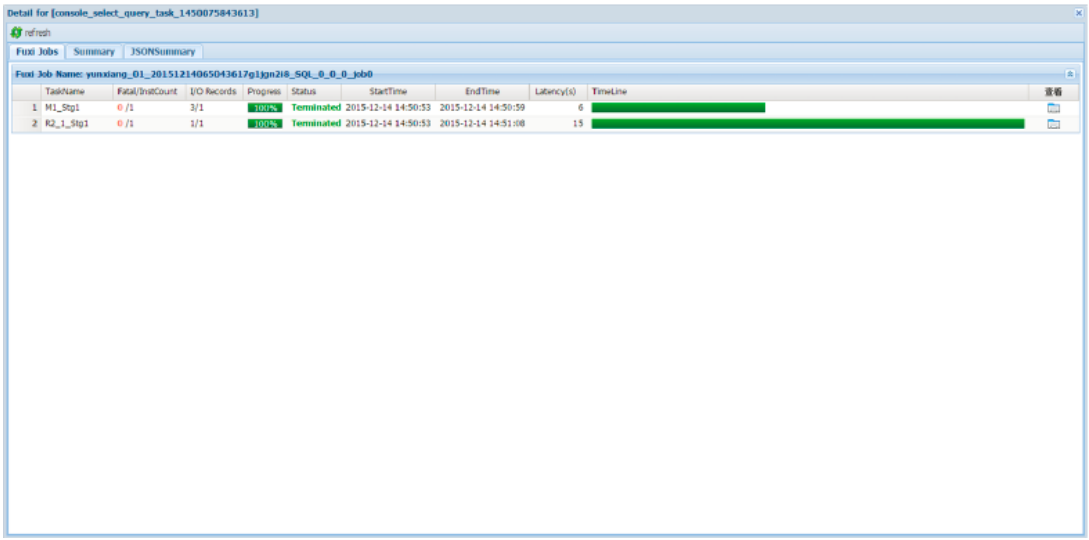

Task Detail - Fuxi Job(2)

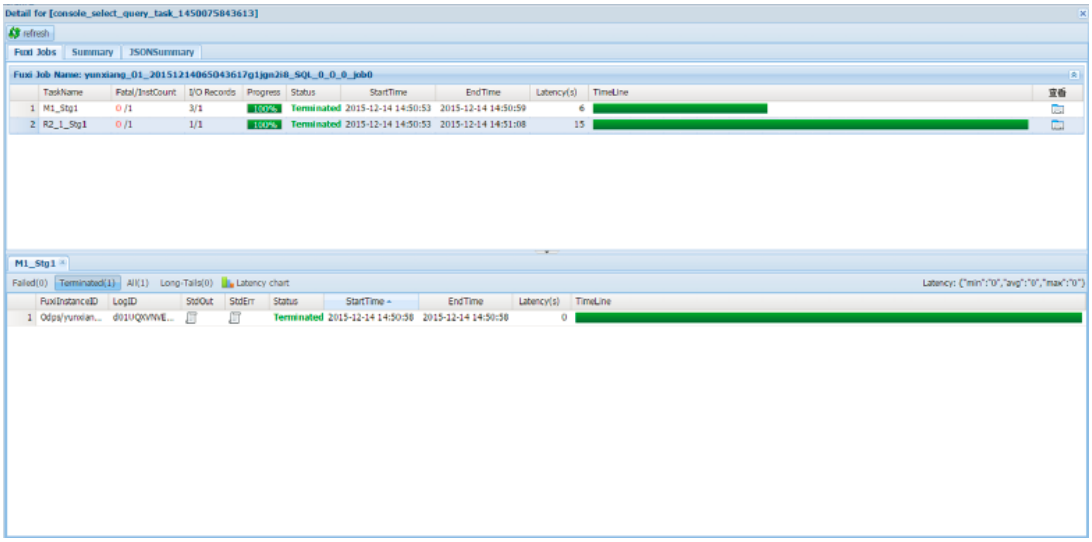

Task Detail - Summary

Task Detail - Summary

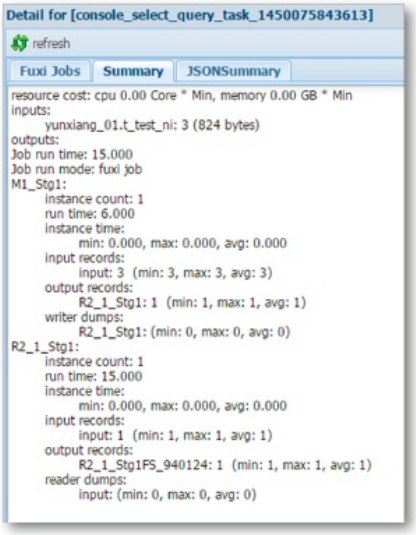

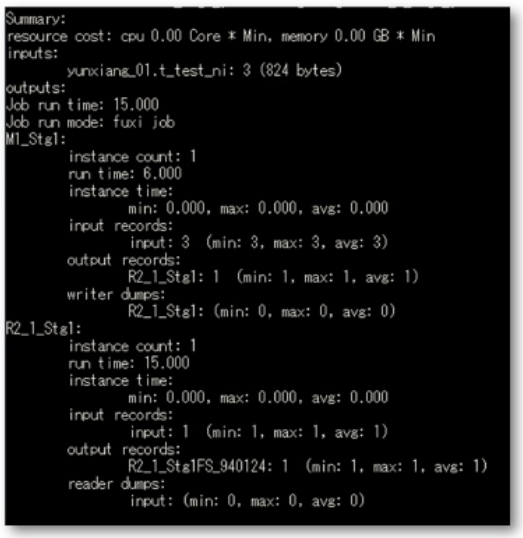

Task Detail - JSONSummary

Task Detail - JSONSummary

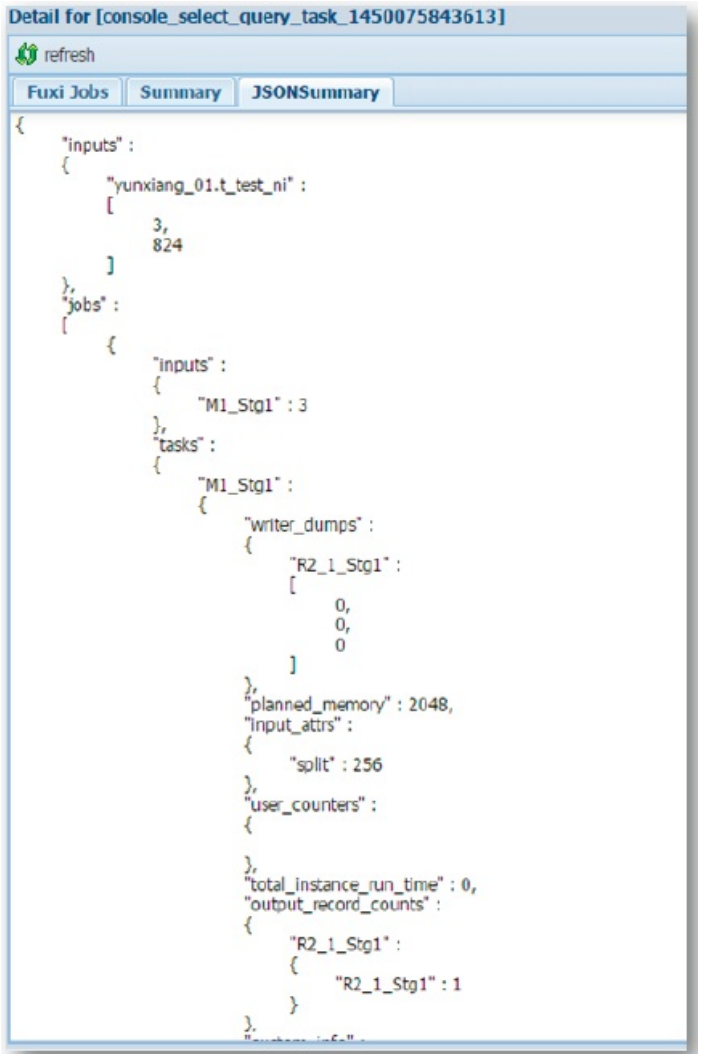

6.2.1.2.4.4. Basic operations and examples

View each point in time in the life cycle of a job.

View each point in time in the life cycle of a job

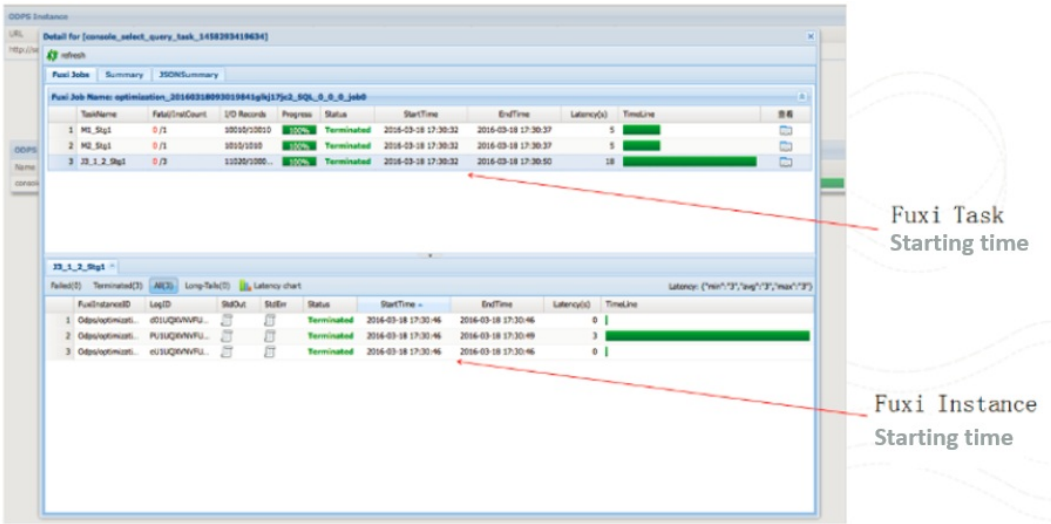

View the time it takes for Job Scheduler to schedule an instance.

View the time it takes for Job Scheduler to schedule an instance

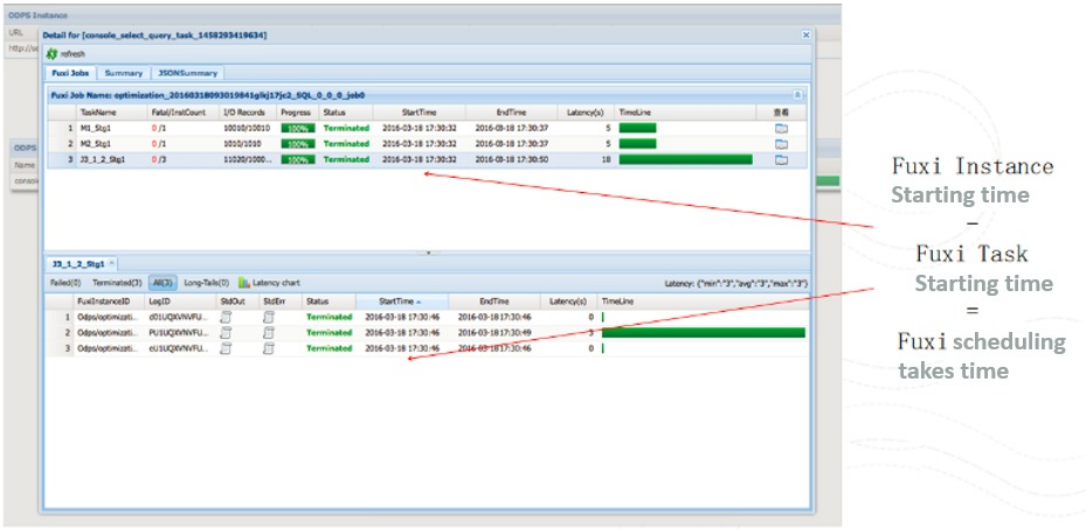

View the polling interval.

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
lbnQiOlt7IkFjdGlvbiI6WyJvZHBzOlJlYWQiXSwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2RwczoqOnBybZplY3RzLZ9wdGlt
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

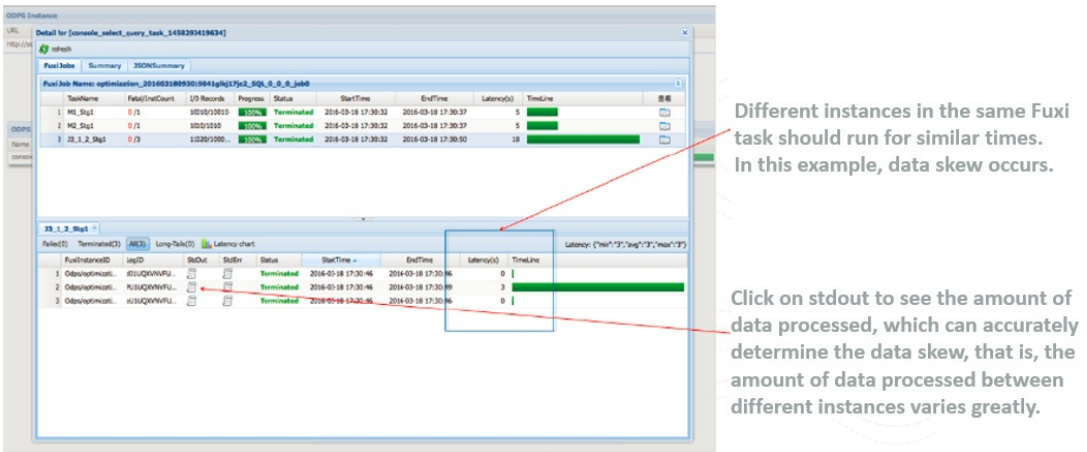

View the UDF and MR debugging information

View the UDF and MR debugging information

View the UDF and MR debugging information

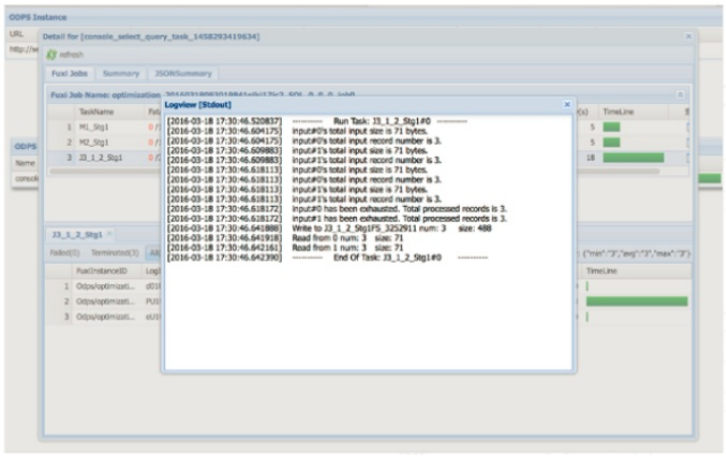

View debugging information in Fuxi Instance Stuout and Stderr

View the task status - Terminated

View the task status - Terminated

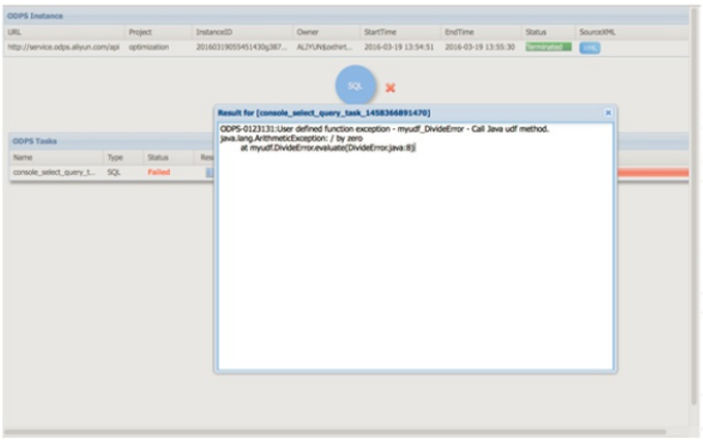

Error messages can be seen from the results of the job

You can also click Detail to go into details to see what went wrong.

6.2.1.2.4.5. 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

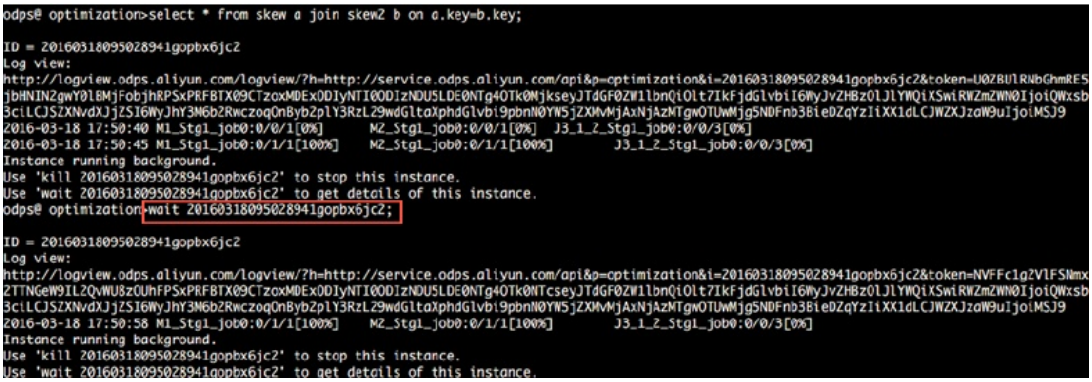

Locate running tasks

After you exit the control window, you can run the `show p;` command to locate currently running tasks and historical tasks.

Locate running tasks

6.2.1.2.5. Apsara Big Data Manager

Apsara Big Data Manager (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 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 *MaxCompute O&M*.

6.2.1.3. Routine O&M

6.2.1.3.1. 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.

6.2.1.3.2. 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 and click the search icon to search for the expected cluster.
   ii. Check whether all machines have reached the desired state based on the information in the Status, Machine Status, and Server Role Status columns. 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=aEVmNTF1dm5GMnFOVlBSWjViZE0rOWRERnZFPSxPRFBTX09CTzox
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          |
```

The following figure shows that fuxi job is running. The command output indicates that fuxi job 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       101h09  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 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
```

Examples:

```
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 Big Data Manager (ABM) to check disk usage.

6.2.1.3.3. 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
```

6.2.1.3.4. 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.

6.2.1.3.5. 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

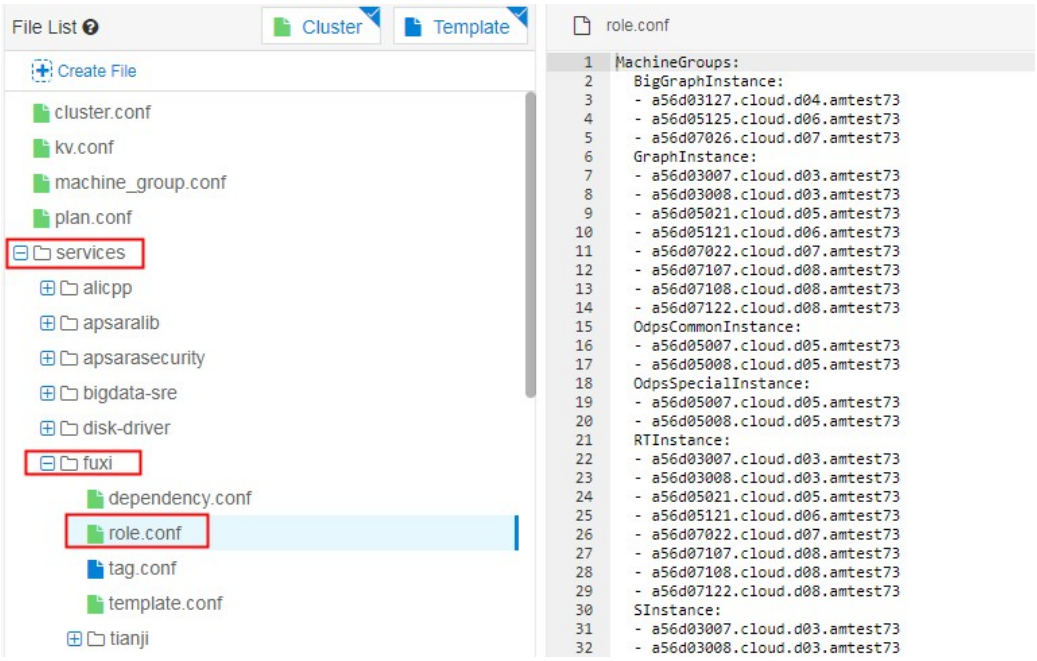

3. Adjust the machine tags on the right and click **Preview and Submit**.

   Adjust machine tags

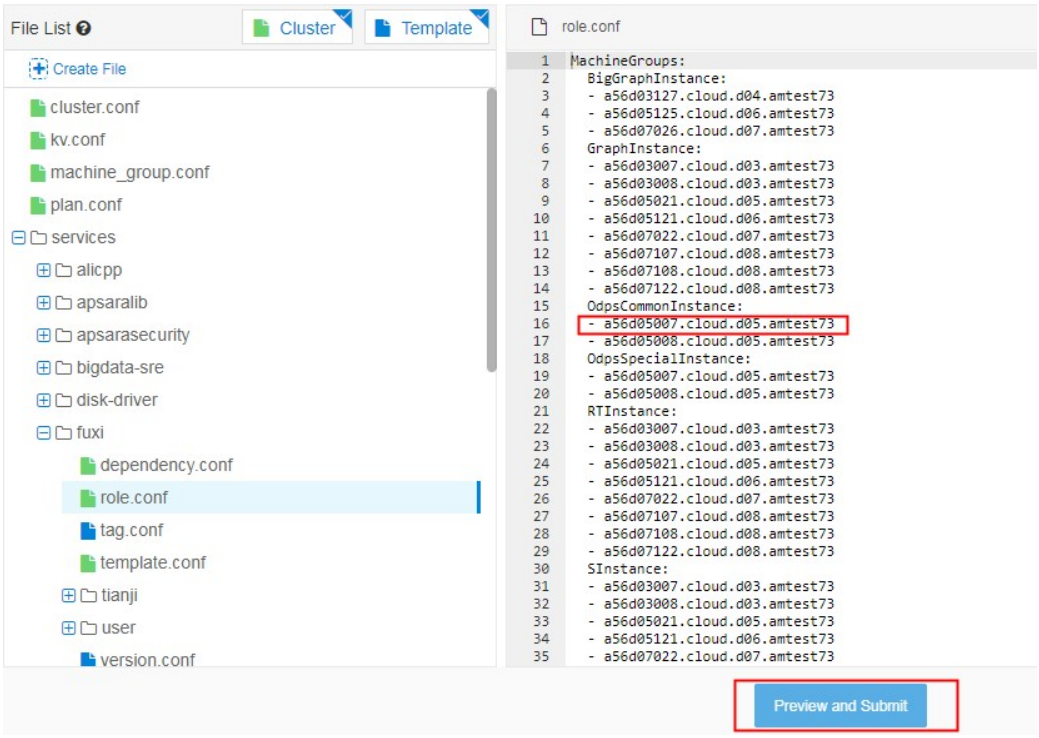

4. In the **Confirm and Submit** dialog box that appears, enter the change description and click **Submit**.

   Submit

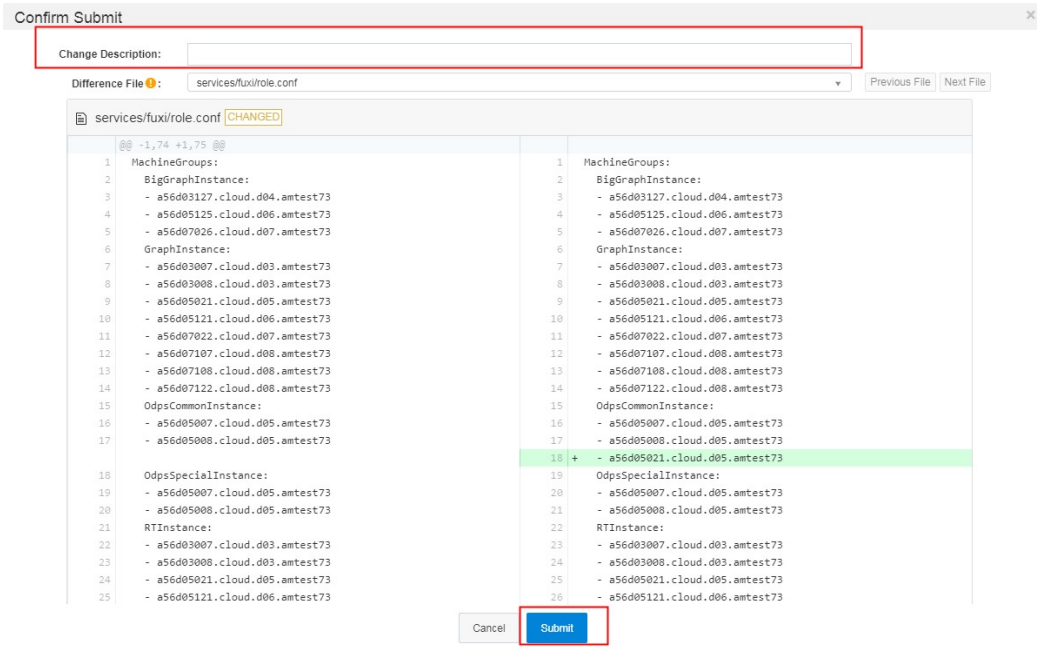

5. The cluster starts rolling and the changes start to take effect.

   > **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.

6.2.1.3.6. 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.
```

6.2.1.4. Common issues and solutions

6.2.1.4.1. 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

  ```
  The pangu disk status:
  Total Disk Size:681225 GB
  Total Free Disk Size:635921 GB
  Total File Size:997 GB
  Total UnReserved Disk Space4Piops:0 GB
  Total Disk Space4Piops:0 GB
  Total UnReserved Disk Iops4Piops:0
  Total Disk Iops4Piops:0
  ```

  > **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

Numbers of files that can be stored for different Pangu-cluster memory capacities

48G memory	Upper limit of total number of files : 87.5 million
96G memory	Upper limit of total number of files : 175 million
128G memory	Upper limit of total number of files : 233 million

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

```
$pu quota
quota under pangu://localcluster/
EntryNumber Limit:unlimited
 Used:16632877
 Used(excluding hardlink):16632712
FileNumber Limit:unlimited
 Used:8594596
 Used(excluding hardlink):8594431
FilePhysicalLength Limit:unlimited
 Used:14151159608 95
 Used(excluding hardlink):1414395196936
FileLogicalLength Limit:unlimited
 Used:467814050981
 Used(excluding hardlink):467573796328
```

- 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

```
$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:
  >
  > - 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

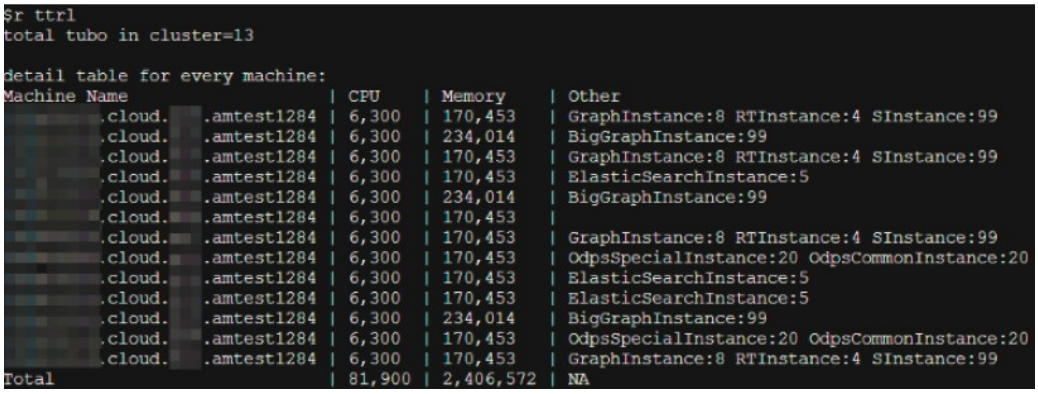

```
$r ttrl
total tubo in cluster=13

detail table for every machine:
Machine Name                    | CPU    | Memory     | Other
          .cloud.    .amtest1284 | 6,300  | 170,453    | GraphInstance:8 RTInstance:4 SInstance:99
          .cloud.    .amtest1284 | 6,300  | 234,014    | BigGraphInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453    | GraphInstance:8 RTInstance:4 SInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453    | ElasticSearchInstance:5
          .cloud.    .amtest1284 | 6,300  | 234,014    | BigGraphInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453    |
          .cloud.    .amtest1284 | 6,300  | 170,453    | GraphInstance:8 RTInstance:4 SInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453    | OdpsSpecialInstance:20 OdpsCommonInstance:20
          .cloud.    .amtest1284 | 6,300  | 170,453    | ElasticSearchInstance:5
          .cloud.    .amtest1284 | 6,300  | 170,453    | ElasticSearchInstance:5
          .cloud.    .amtest1284 | 6,300  | 234,014    | BigGraphInstance:99
          .cloud.    .amtest1284 | 6,300  | 170,453    | OdpsSpecialInstance:20 OdpsCommonInstance:20
          .cloud.    .amtest1284 | 6,300  | 170,453    | GraphInstance:8 RTInstance:4 SInstance:99
Total                             | 81,900 | 2,406,572  | NA
```

  > **Note** In the command output, the domain name, total CPU capacity (Unit: U. 100 U = 1 core), and total memory (Unit: MB) of each Tubo machine, as well as the role of each Tubo machine in Job Scheduling System are listed in four columns.

- Run the `rtfrl` command on the cluster AG to view the remaining computing resources.

  Remaining computing resources

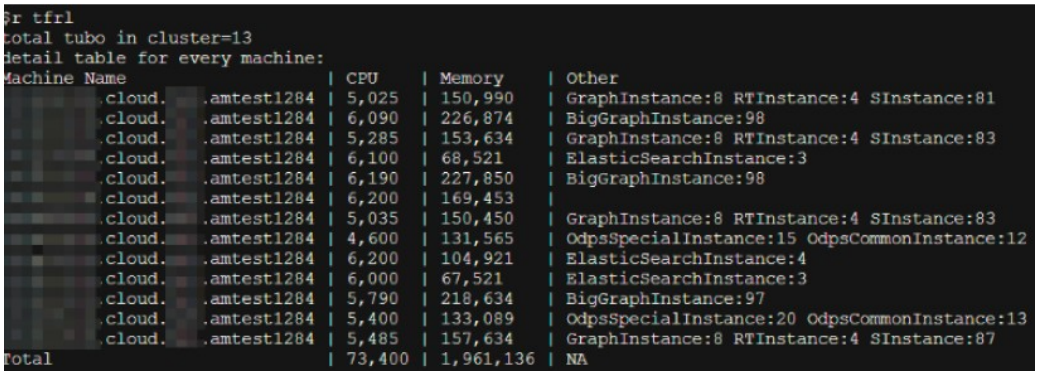

```
$r tfrl
total tubo in cluster=13
detail table for every machine:
Machine Name                    | CPU    | Memory     | Other
          .cloud.    .amtest1284 | 5,025  | 150,990    | GraphInstance:8 RTInstance:4 SInstance:81
          .cloud.    .amtest1284 | 6,090  | 226,874    | BigGraphInstance:98
          .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 `r cru` command on the cluster AG to view the resources used by all running jobs in MaxCompute.

  Resources used by all running jobs

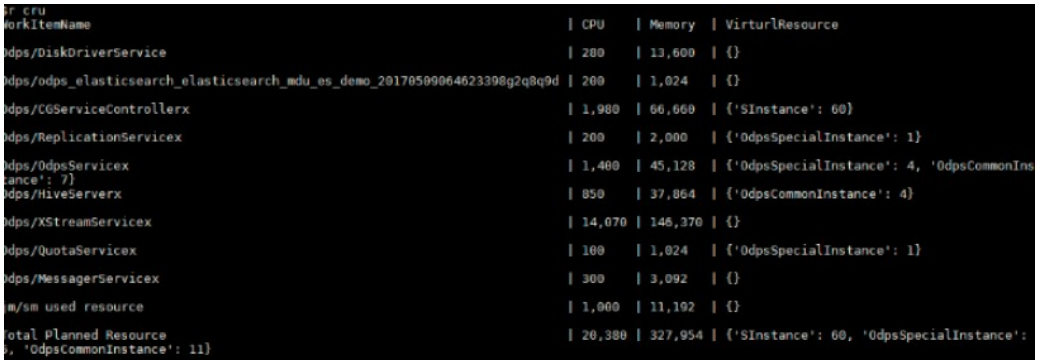

> **Note** The name, total CPU capacity, total memory of each job, as well as the number of Fuxi instances started in the role of each job in Job Scheduling System are listed in four columns.

How to allocate project resources in a MaxCompute cluster

- Storage resource allocation: Based on the characteristics of a project, the space size and file size limit are configured when you create the project.

  If the following error messages are displayed, the file size limit of the project has been exceeded. In this case, you must organize the data in the project by deleting unnecessary table data or increasing the storage resource quota.

  Error messages

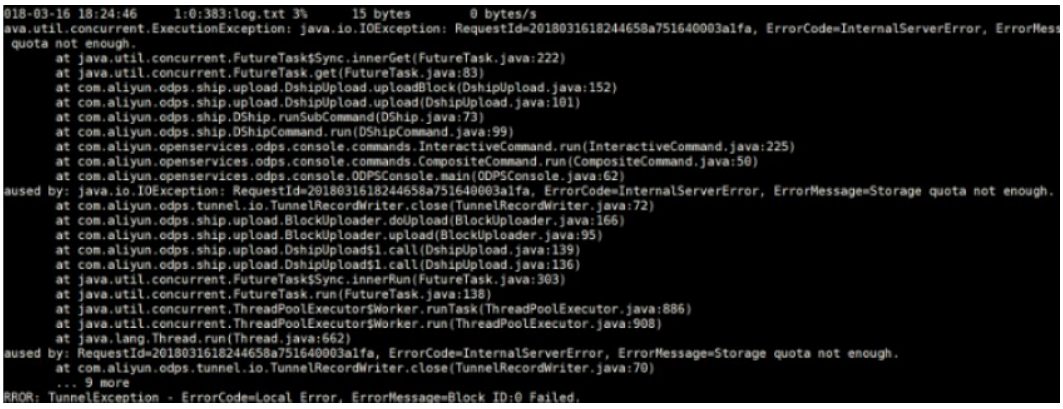

> **Notice** The sum of the storage capacity of all projects cannot exceed the total allowable storage capacity of a service. Similarly, the total file size of all projects cannot exceed the total allowable file size. Therefore, you must properly allocate the storage space and file size limit by project and make timely adjustment based on your business requirements.

- Computing resource allocation: division of quota groups.
  - What is a quota group?

A MaxCompute cluster allows you to divide computing resources into different quota groups, and schedule them as required. A quota group represents a certain amount of CPU and memory resources. MinQuota and MaxQuota are used for CPU and memory configurations. MinQuota is the minimum quota allowed for the quota group, and MaxQuota is the maximum quota allowed for the quota group. For example, MinCPU=500 indicates that the quota group has been assigned at least 500/100=5 cores. MaxCPU=2000 indicates that the quota group has been assigned at least 2000/100=20 cores.

MaxCompute uses a FAIR scheduling policy and a first-in-first-out (FIFO) scheduling policy by default. The difference between the FAIR and FIFO scheduling polices lies in the keys by which tasks in waiting queues are sorted. If each schedule unit has its own priority, both FAIR and FIFO scheduling policies allocate high-priority schedule units first. If all schedule units share the same priority, the FIFO scheduling policy sorts the schedule units by the time when they are submitted. The earlier they are submitted, the higher priority they have. The FAIR scheduling policy sorts the scheduling units by the slotNum allocated to them. The smaller the slotNum is, the higher priority they have. For the FAIR policy group, this can basically ensure that the same amount of resources are assigned to schedule units with the same priority.

You can run the `r quota` command on the cluster AG to view quota group settings.

View quota group settings

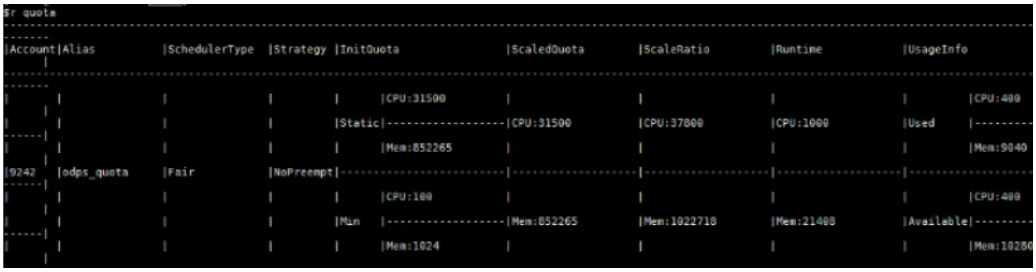

You can run the following command on the cluster AG to create and modify a quota as needed:

```
sh /apsara/deploy/rpc_wrapper/rpc.sh setquota -i $QUOTAID -a $QUOTANAME -t fair -s $max_cpu_quota $max_mem_quota -m $min_cpu_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

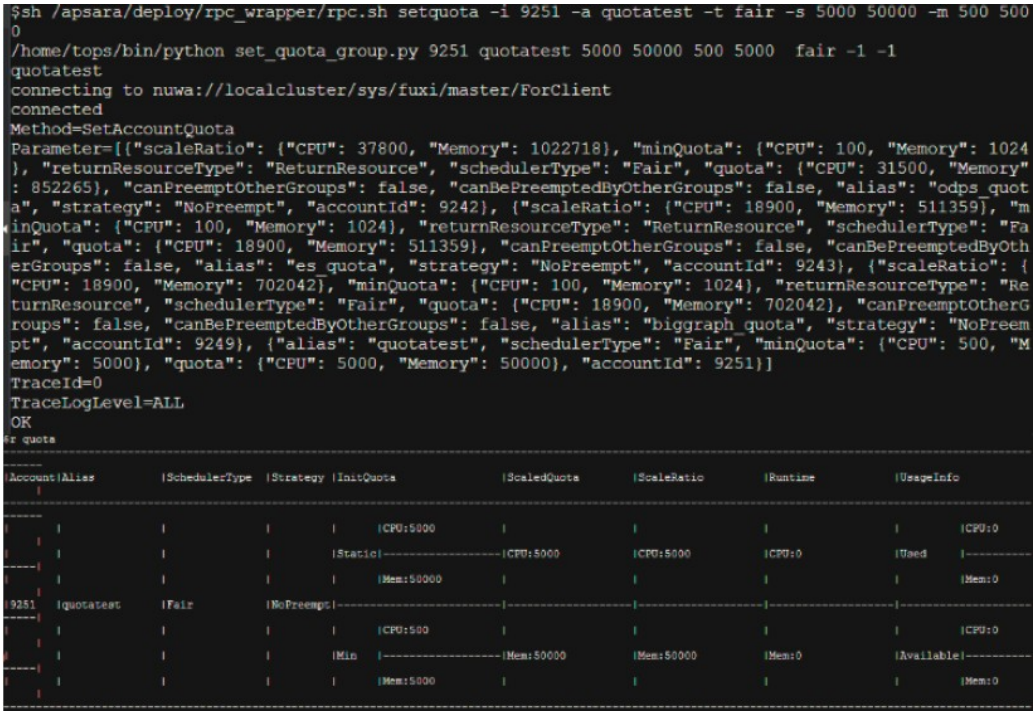

Modify a quota

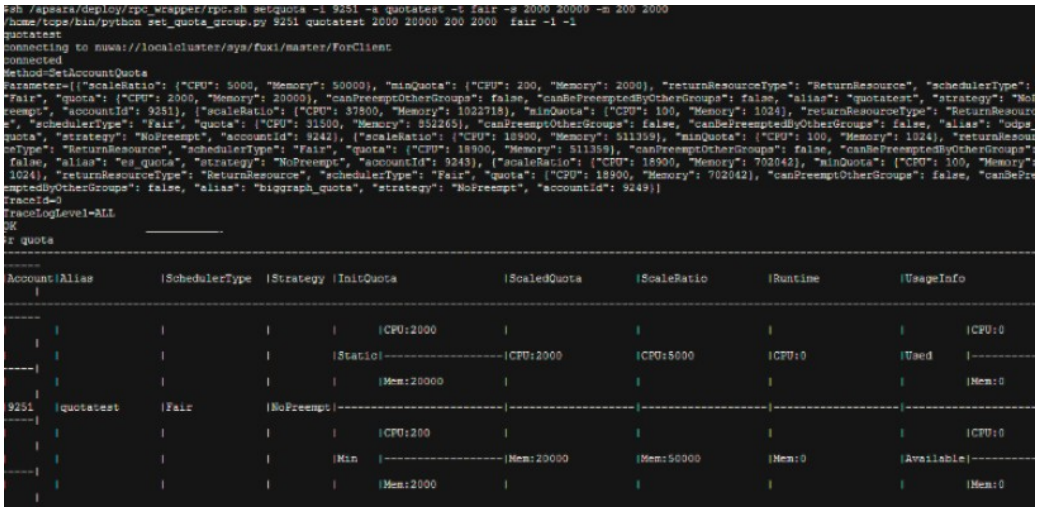

- How to divide quota groups

  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

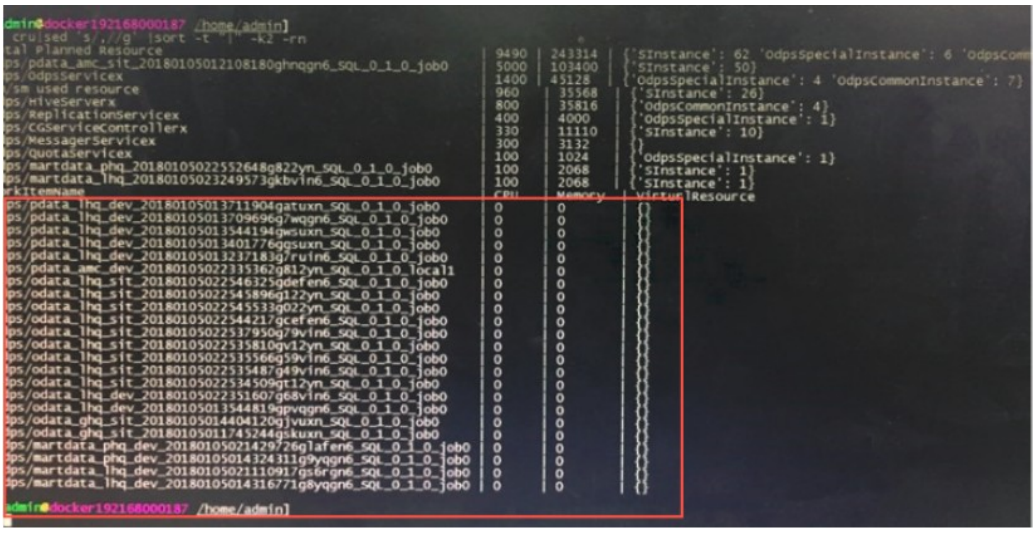

Quota used up

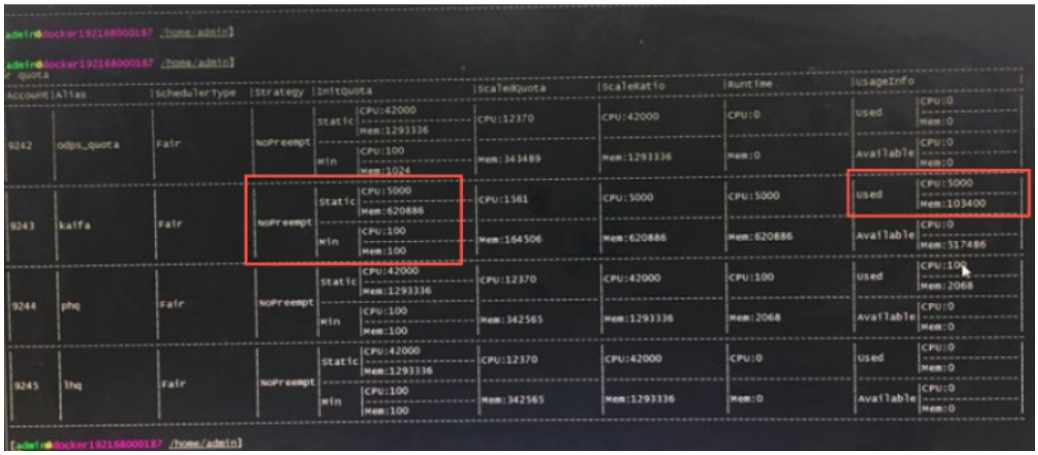

You must divide quota groups based on the following general principles:

- 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.

6.2.1.4.2. 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

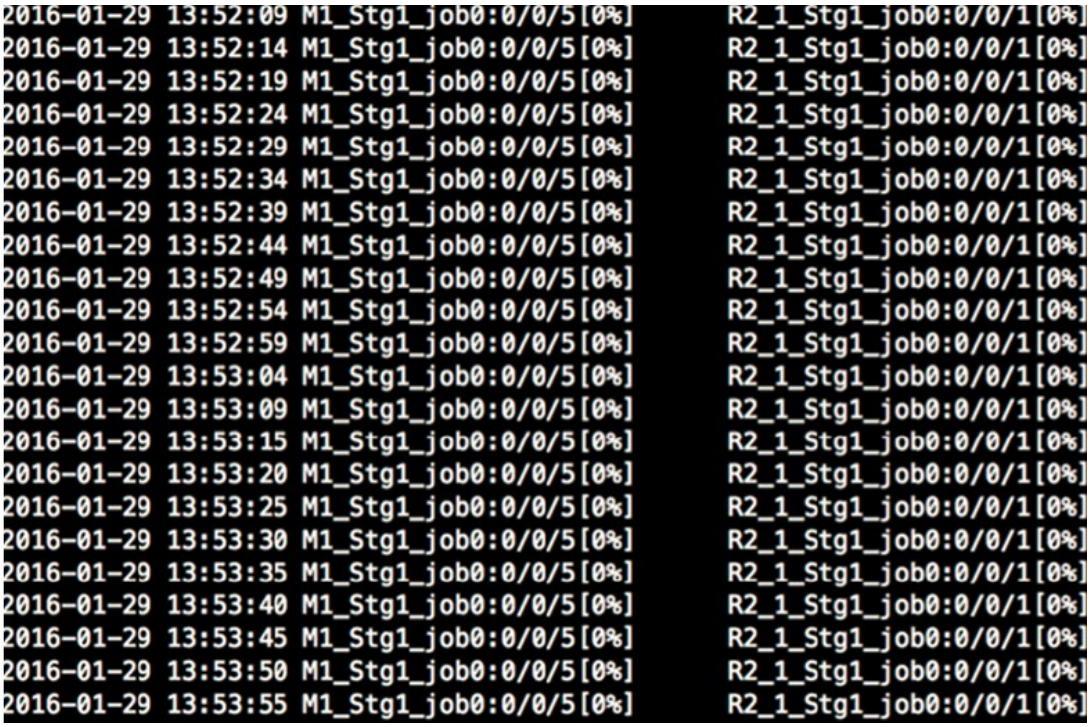

Cause: The issue is typically caused by insufficient resources. You can use LogView to determine the status of job resources (task instance status).

- 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.

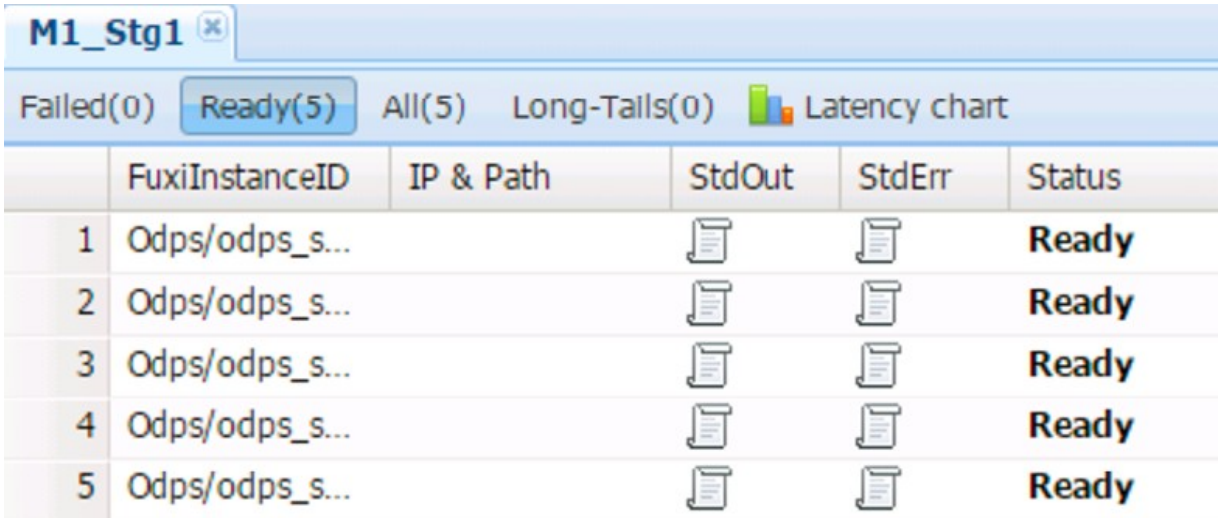

Solution:

- 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

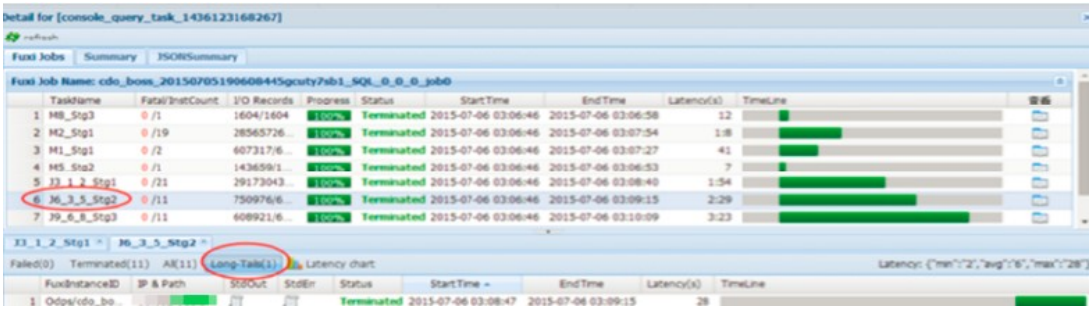

Further analysis: Analyze the job summary in LogView, and calculate the difference between the max and avg values of input and output records of a slow task. If the max and avg values differ by several orders of magnitude, it can be initially determined that the job data is skewed.

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 c2, b.key c3, b.value c4 from src a join src_skewjoin1 b on a.key = b.key;
```

Scenario 5: how to configure common SQL parameters

Map settings

```
set odps.sql.mapper.cpu=100
```

Description: allows you to set the number of CPUs used by each instance in a Map task. Default value: 100. Valid values: 50 to 800.

```
set odps.sql.mapper.memory=1024
```

Description: allows you to set the memory size of each instance in a Map task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

```
set odps.sql.mapper.merge.limit.size=64
```

Description: allows you to set the maximum size of control files to be merged. Unit: MB. Default value: 64. You can set this variable to control the inputs of mappers. Valid values: 0 to Integer.MAX_VALUE.

```
set odps.sql.mapper.split.size=256
```

Description: allows you to set the maximum data input volume for a Map task. Unit: MB. Default value: 256. You can set this variable to control the inputs of mappers. Valid values: 1 to Integer.MAX_VALUE.

Join settings

```
set odps.sql.joiner.instances=-1
```

Description: allows you to set the number of instances in a Join task. Default value: -1. Valid values: 0 to 2000.

```
set odps.sql.joiner.cpu=100
```

Description: allows you to set the number of CPUs used by each instance in a Join task. Default value: 100. Valid values: 50 to 800.

```
set odps.sql.joiner.memory=1024
```

Description: allows you to set the memory size of each instance in a Join task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

Reduce settings

```
set odps.sql.reducer.instances=-1
```

Description: allows you to set the number of instances in a Reduce task. Default value: -1. Valid values: 0 to 2000.

```
set odps.sql.reducer.cpu=100
```

Description: allows you to set the number of CPUs used by each instance in a Reduce task. Default value: 100. Valid values: 50 to 800.

```
set odps.sql.reducer.memory=1024
```

Description: allows you to set the memory size of each instance in a Reduce task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

UDF settings

```
set odps.sql.udf.jvm.memory=1024
```

Description: allows you to set the maximum memory size used by the UDF JVM heap. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

```
set odps.sql.udf.timeout=600
```

Description: allows you to set the timeout period of a UDF. Unit: seconds. Default value: 600. Valid values: 0 to 3600.

```
set odps.sql.udf.python.memory=256
```

Description: allows you to set the maximum memory size used by the UDF Python API. Unit: MB. Default value: 256. Valid values: 64 to 3072.

```
set odps.sql.udf.optimize.reuse=true/false
```

Description: When this parameter is set to true, each UDF function expression can only be calculated once, improving performance. Default value: true.

```
set odps.sql.udf.strict.mode=false/true
```

Description: allows you to control whether functions return NULL or an error if dirty data is found. If the parameter is set to true, an error is returned. Otherwise, NULL is returned.

MapJoin settings

```
set odps.sql.mapjoin.memory.max=512
```

Description: allows you to set the maximum memory size for a small table when running MapJoin. Unit: MB. Default value: 512. Valid values: 128 to 2048.

```
set odps.sql.reshuffle.dynamicpt=true/false
```

Description:

- Dynamic partitioning scenarios are time-consuming. Disabling dynamic partitioning can accelerate SQL.
- If there are few dynamic partitions, disabling dynamic partitioning can prevent data skew.

Scenario 6: how to check the storage usage of a single project

Launch the MaxCompute console as a project owner and run the `desc project <project_name>-extended;` command to view the following information.

Storage information

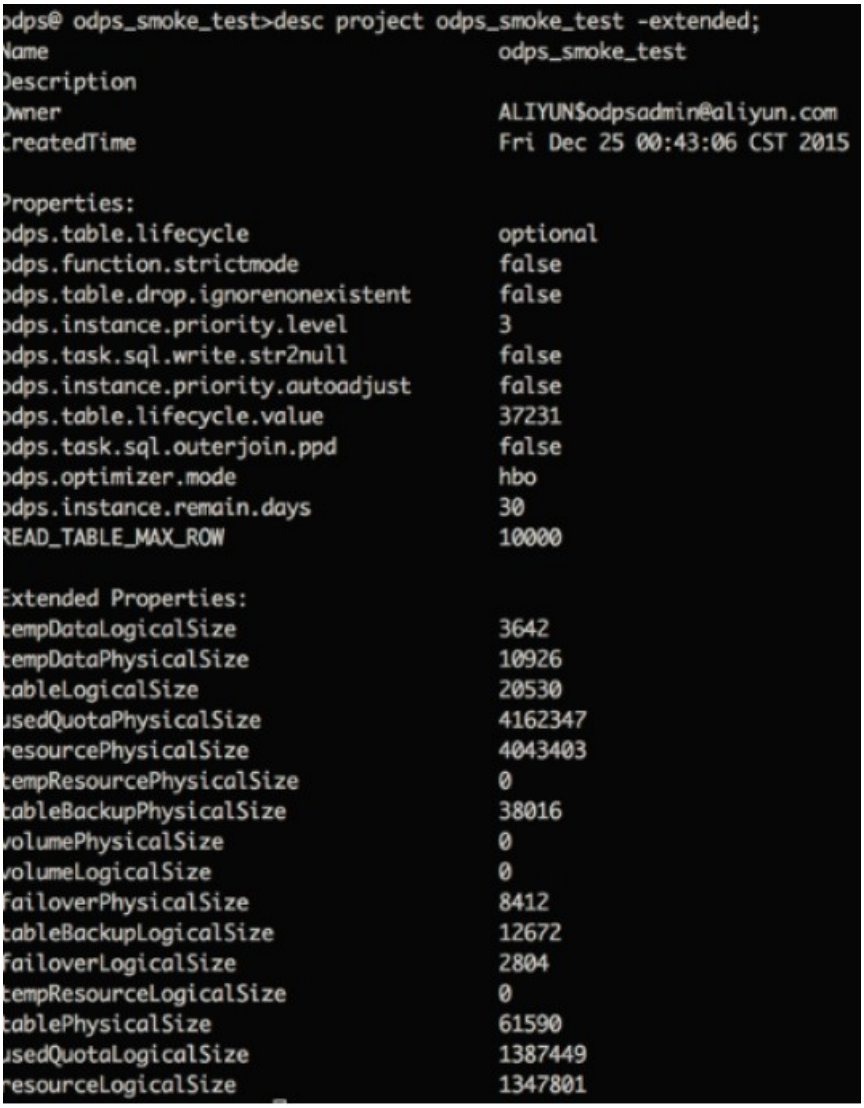

```
odps@ odps_smoke_test>desc project odps_smoke_test -extended;
Name                                    odps_smoke_test
Description
Owner                                   ALIYUN$odpsadmin@aliyun.com
CreatedTime                             Fri Dec 25 00:43:06 CST 2015

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.

6.2.1.5. MaxCompute O&M

6.2.1.5.1. Log on to the ABM console

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

Prerequisites

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

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

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

Procedure

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

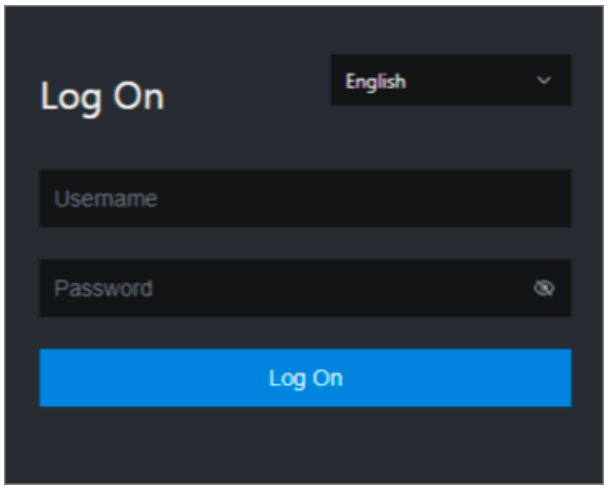

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains the following special characters: ! @ # $ %
- The password must be 10 to 20 characters in length.

4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Big Data Services** section, choose **General-Purpose O&M > Apsara Big Data Manager**.

6.2.1.5.2. Business O&M

6.2.1.5.2.1. 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 search for and filter 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.
- Quota Groups: shows the quota groups of all projects in a MaxCompute cluster. It allows you to create and modify quota groups. You can also view details about quota groups and enable period management for quota groups.
- Jobs: shows information about jobs in a MaxCompute cluster. You can search for and filter jobs. You can also view the operational logs, terminate running jobs, 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 Apsara Big Data 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 **Business** tab. In the left-side navigation pane, choose **Projects > Project List**.

6.2.1.5.2.2. Project management

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 set 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.

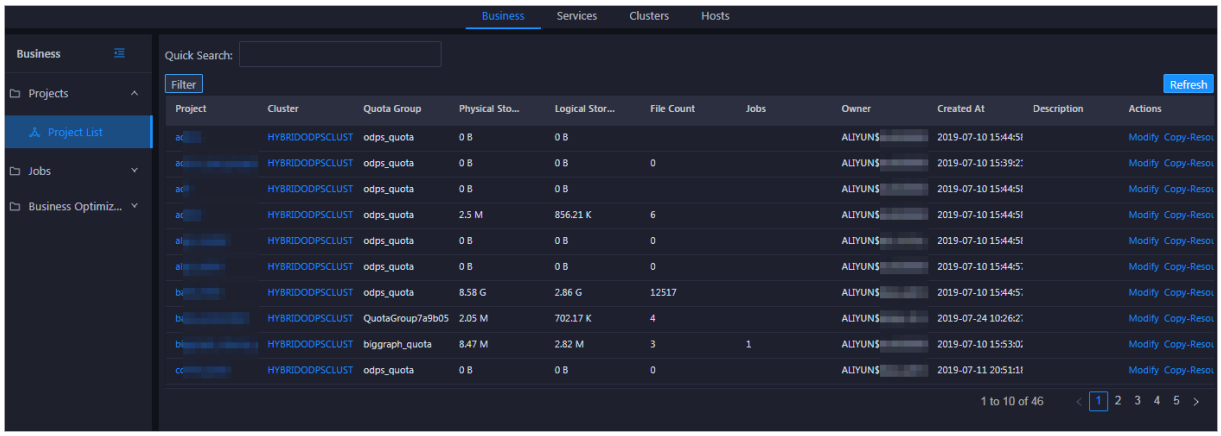

The **Project List** page shows the detailed information about all projects in a cluster. You can view the name, cluster, used storage, storage quota, storage usage, number of files, owner, and creation time of a project.

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 project for which you want to change the quota group, click **Actions** in the Actions column, and select **Change Default Quota Group**. In the **Change Default Quota Group** pane, configure parameters.

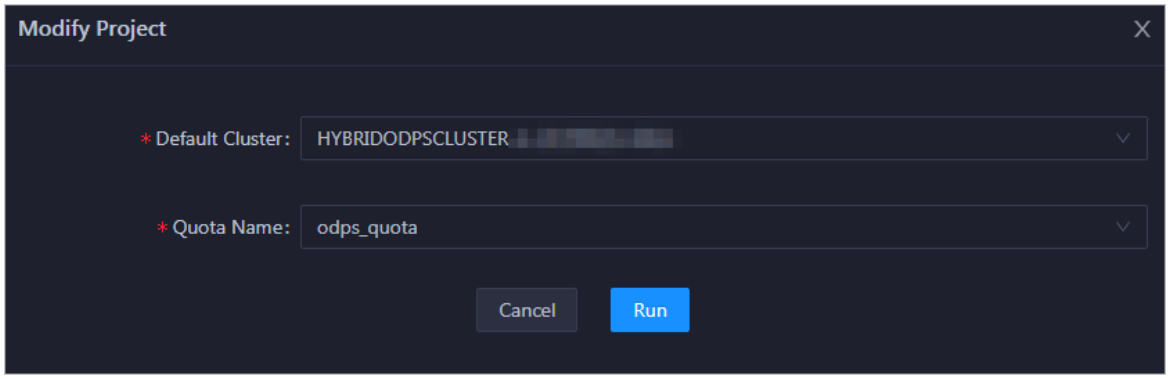

Parameters:

- **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 configure 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 project for which you want to modify the storage quota, click **Actions** in the Actions column, and select **Modify Storage Quota**. In the **Change Storage Quota** pane, configure 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 reason for the modification

2. After you configure 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 project for which you want to configure resource replication, click **Actions** in the Actions column, and select **Resource Replication**. In the **Copy Resource** pane, configure parameters.

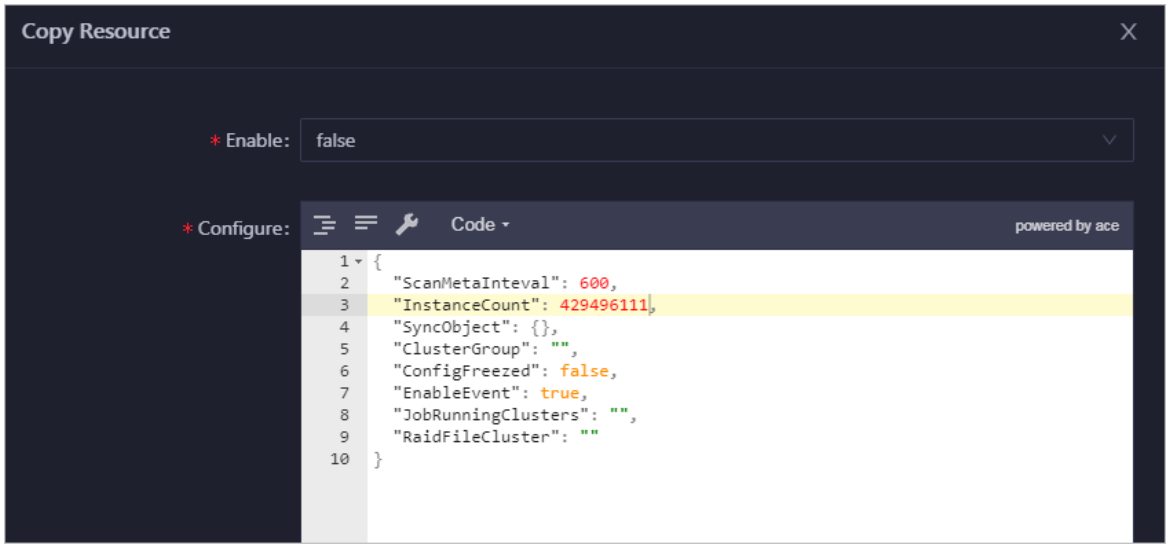

Parameters:

- **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 for a project. In most cases, the default settings are used. If you want to modify the settings, consult O&M engineers.

2. After you modify code in the **Configure** field, click **Compare Versions** to view the differences, which are highlighted.

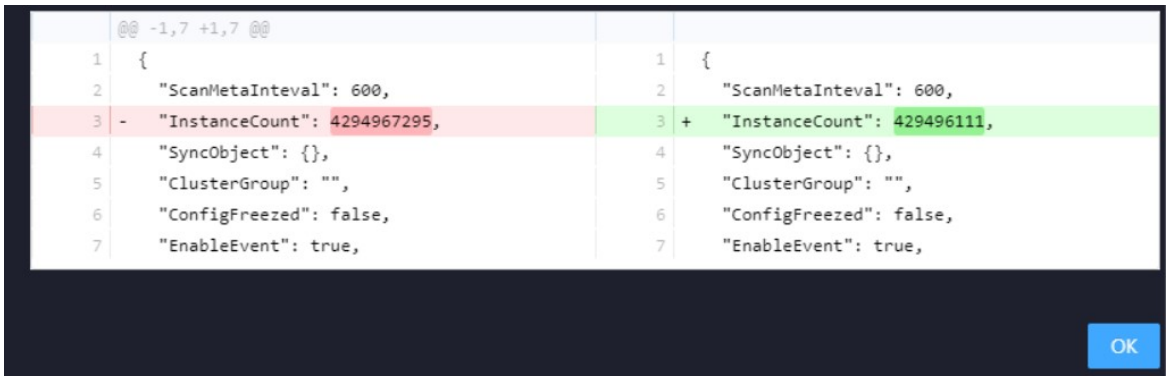

3. Click **Run**.

Project details

The Apsara Big Data Manager (ABM) console shows your MaxCompute projects and project details. You can view the project overview, jobs, storage, configuration, quota group, and tunnel, as well as information about resource analysis, storage encryption, and cross-cluster replication.

Go to the project details page

1. Log on to the Apsara Big Data Manager 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. Choose **Projects** > **Project List**. Click the name of a project to view its details.

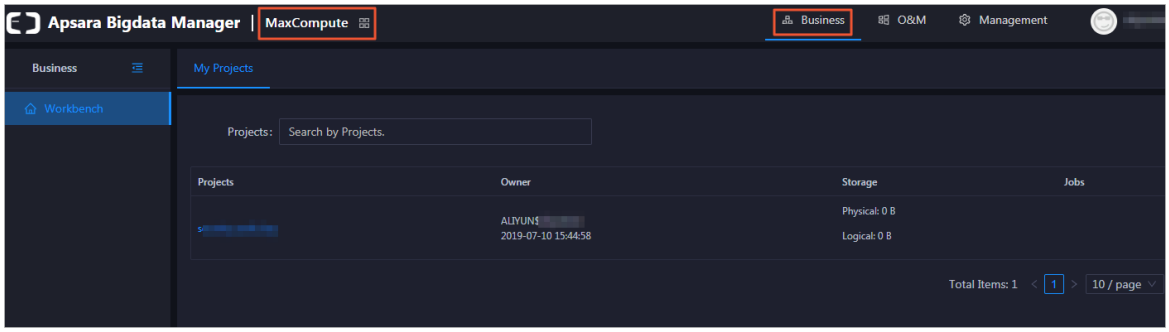

Overview

On the **Overview** tab, you can view the following information about the 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, status, start time, priority, and type. You can also view the operational logs of a job to locate errors during job running.

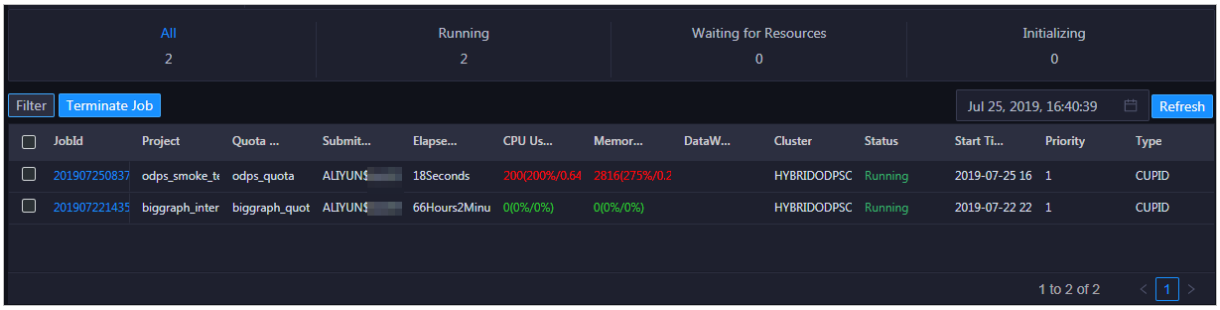

You can perform the following operations on the Jobs tab:

- 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**.

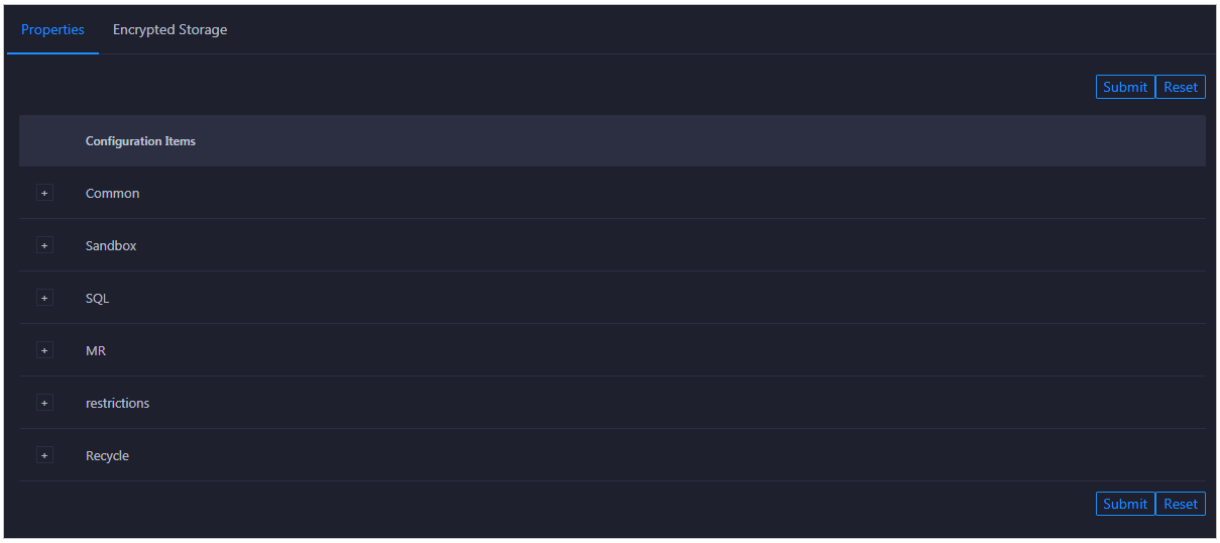

On the **Authorize Package for Metadata Repository** tab, you can install the package and perform package-based authorization.

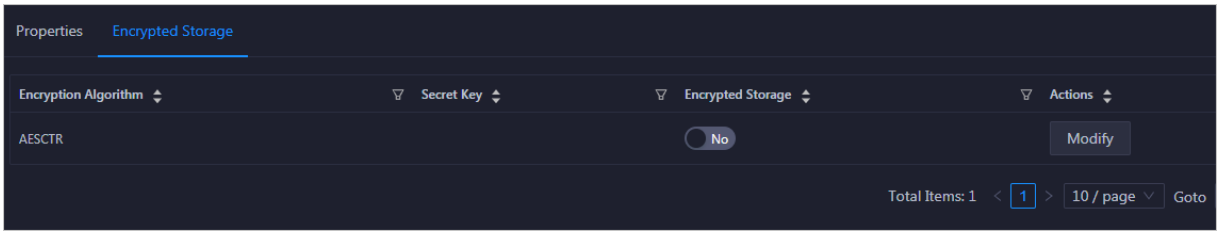

Quota Groups

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

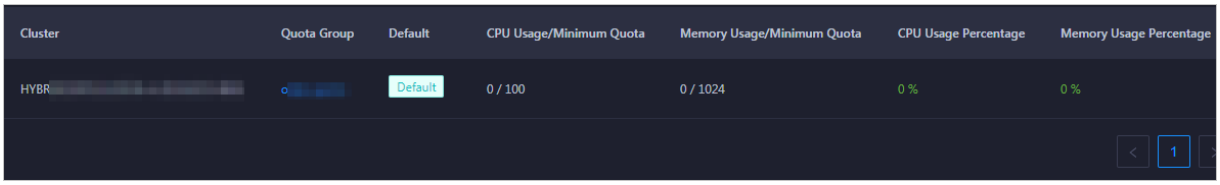

To view details about a quota group, click the quota group name in the **Quota** column.

> **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.

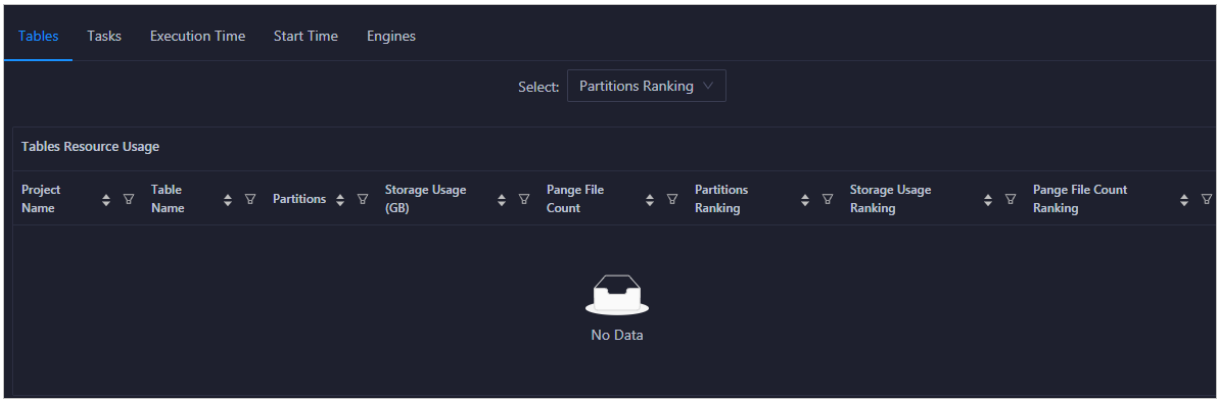

Encryption at Rest

On the **Encryption at Rest** tab, you can encrypt data by using the following encryption algorithms: AES-CTR, AES256, RC4, and SM4.

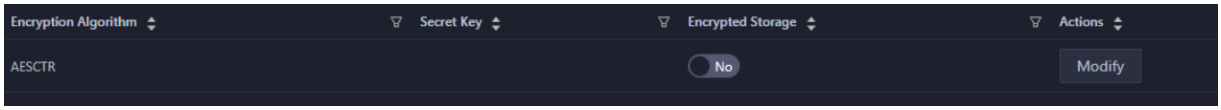

Cross-cluster Replication

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.

For example, 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** sub-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.

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 Big Data 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 required 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** panel, specify **Encryption Algorithm**, **region**, and **project**.

   > **Note** AES-CTR, AES256, RC4, and SM4 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 Create Rule message, 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** panel, 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.

Parameter	Description
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 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	{
	"odps.merge.cross.paths": "true",
	"odps.idata.useragent": "odps encrypt key rotate via force mergeTask",
	"odps.merge.max.filenumber.per.job": "10000000",
	"odps.merge.max.filenumber.per.instance": "10000",
	"odps.merge.failure.handling": "any",
	"odps.merge.maintain.order.flag": "true",
	"odps.merge.smallfile.filesize.threshold": "4096",
	"odps.merge.quickmerge.flag": "true",
	"odps.merge.maxmerged.filesize.threshold": "4096",
	"odps.merge.force.rewrite": "true",
	"odps.merge.restructure.action": "hardlink"
	}

8. (Optional)View the history of 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.

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 Big Data 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. For more information about how to create a workspace, see *Create a workspace* in *DataWorks User Guide*.

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.

Procedure

1. Log on to the Apsara Big Data Manager 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. Choose **Projects > Project List**.
4. Click the name of a 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** message, click **Run**. A message appears, indicating that the permissions are granted.
7. Click **Install** in the Actions column. In the **Install Package** message, 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.

Perform disaster recovery

When a primary MaxCompute cluster fails, you can perform a primary/secondary switchover in the Apsara Big Data 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, choose **Projects > 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** in the upper-right corner. 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_recoverv* directorv of the Docker container in bcc-saa.AG# of the secondary ABM cluster. Then, run the `/home/tops/bin/python change_dns_vip.pv` script in the directorv. If job_success appears, the execution succeeds. Then, run 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 run, you can go to the secondary cluster page.

  > **Note** If an exception occurs when you run 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

Pay attention to the following points for a disaster recovery switchover:

- By default, the logon to Apsara Big Data Manager depends on the Apsara Uni-manager Operations Console. If the Apsara Uni-manager Operations 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, run `change_login_by_bcc.sh` to switch the logon mode to the mode that is independent of the Apsara Uni-manager Operations Console. After the Apsara Uni-manager Operations Console has reached the desired state, run `change_login_by_aso.sh` to switch the logon mode back to the mode that depends on the Apsara Uni-manager Operations 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, choose **Projects > 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 complete 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 run 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.

Migrate projects

Apsara Big Data 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, choose **Projects > Project Migration**.
2. In the upper part of the **Migration Mission** page, select the region where the project resides.

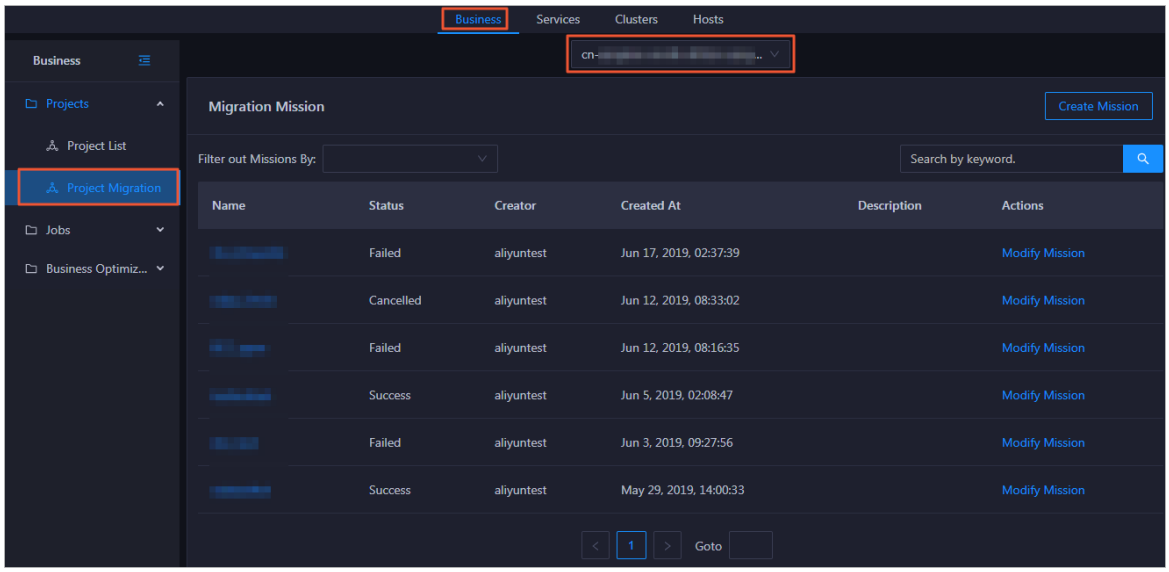

3. In the upper-right corner, click **Create Mission**. On the page that appears, specify the parameters in the **General**, **Source**, **Target Selection**, and **Cluster for Mission Execution** sections as prompted.

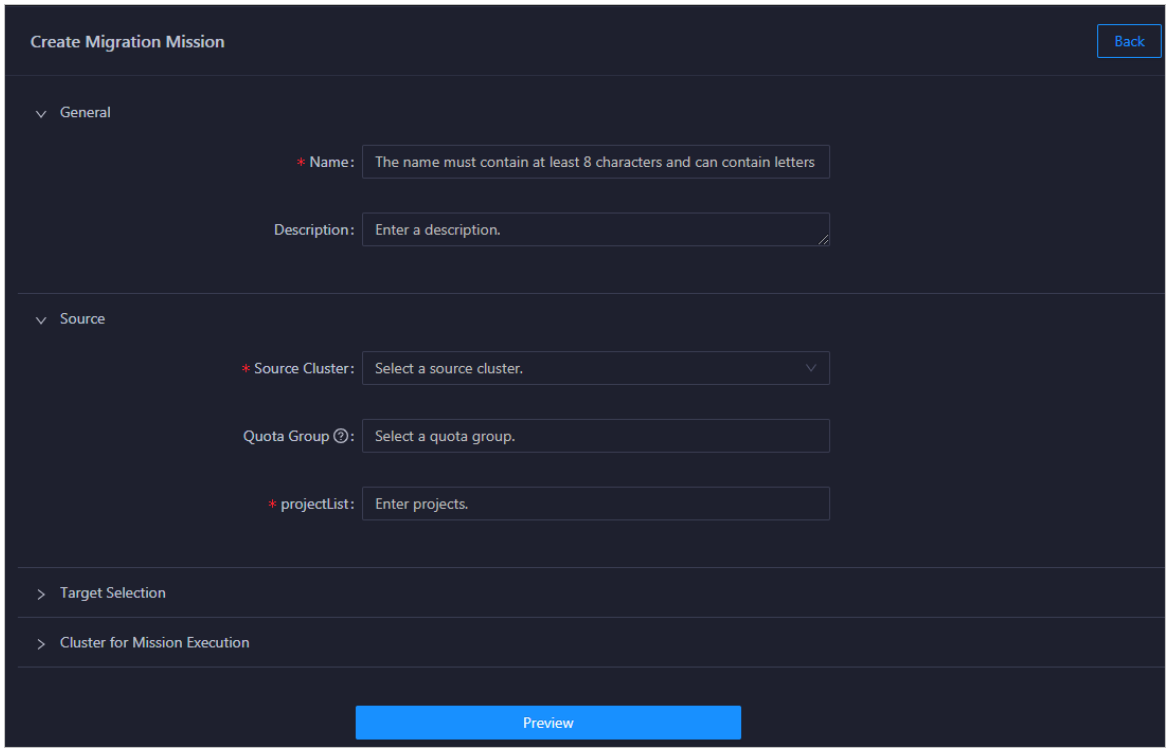

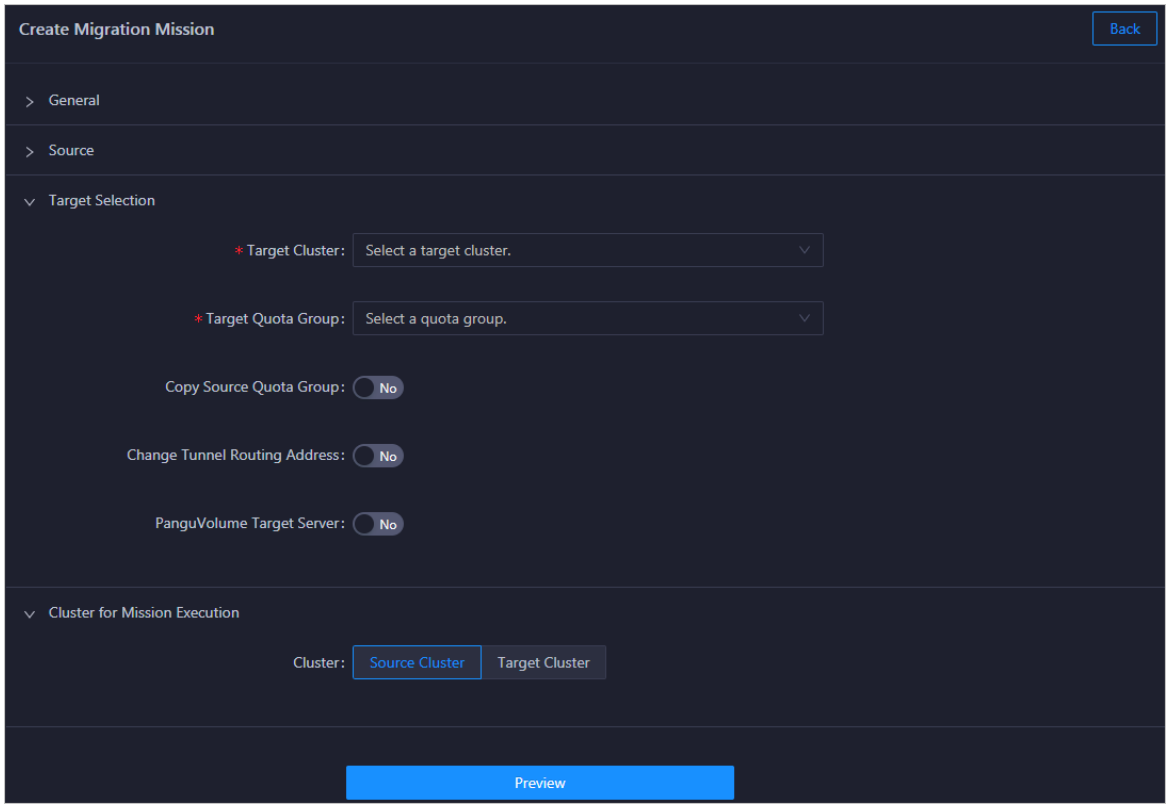

The following table describes the required parameters.

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 destination 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 destination 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.

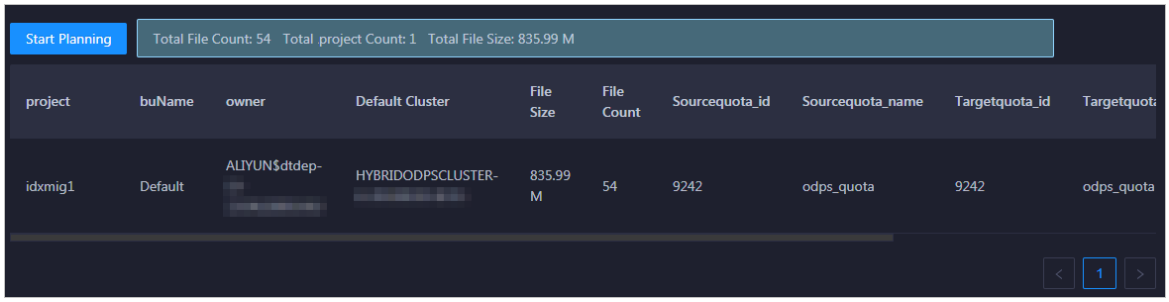

5. After you confirm the configuration, click **Start Planning** in the upper-left corner. A project migration task is generated. The migration details appear.

   It requires some time to generate the task.

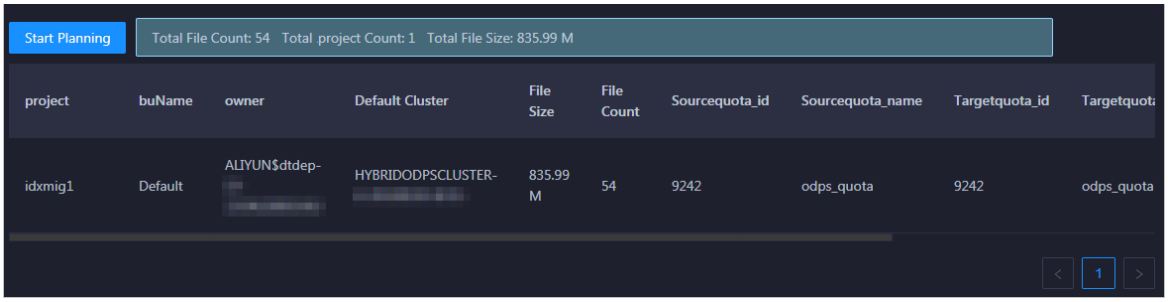

A standard project migration task generally includes five steps:

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 one 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 perform 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.

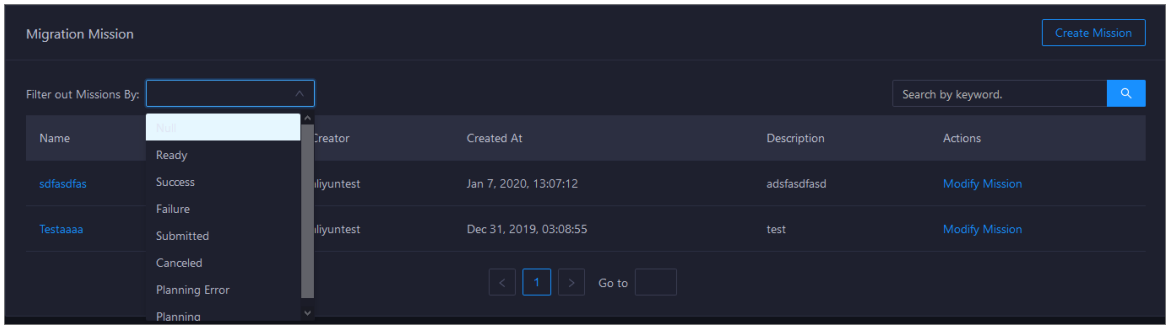

2. Click the name of a task. On the **Migration Details** page, view the details of the task.

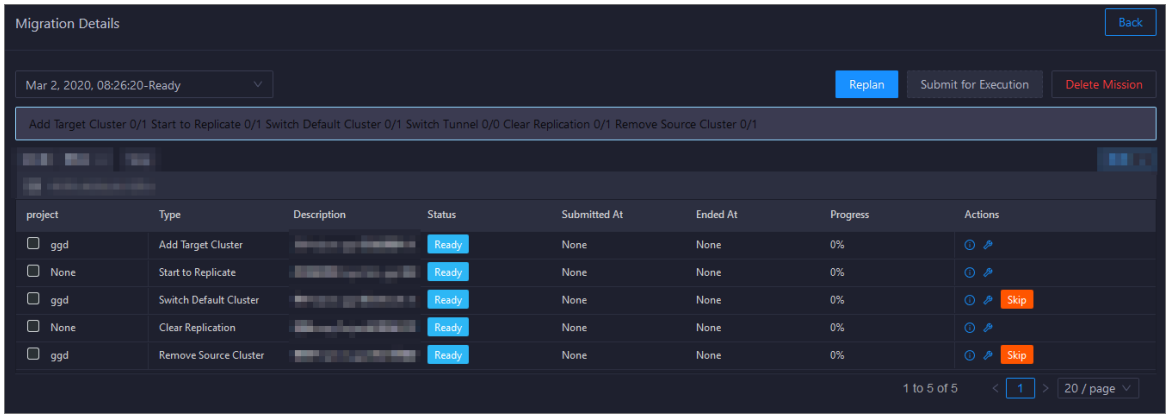

3. If a step fails, click the **Details** or **Debugging** icon in the Actions column to view the details or debugging information of the step. This allows you to identify the cause of the failure.
4. Perform other required operations.

   Click **Menu** in the upper-right corner. You can export the step list, change the column width to automatically fit the content, 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.

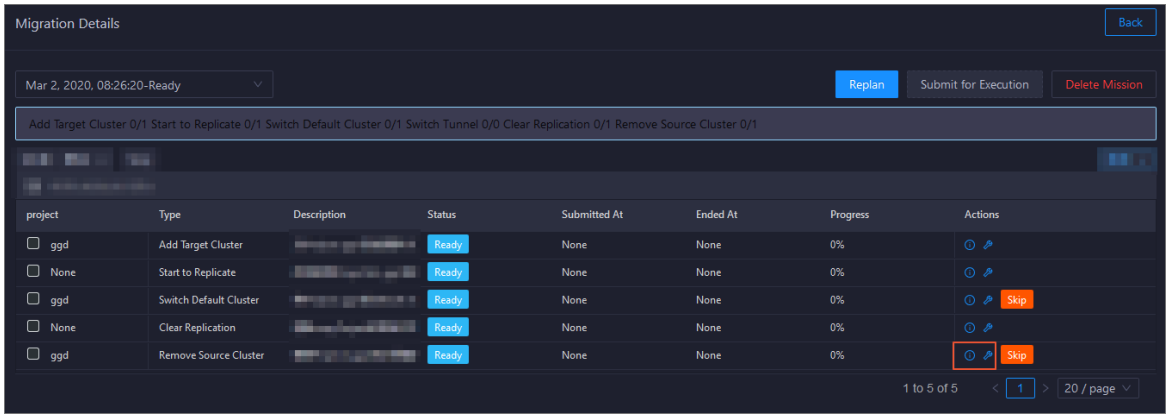

2. Click the **Details** icon in the Actions column to view the details of the step.

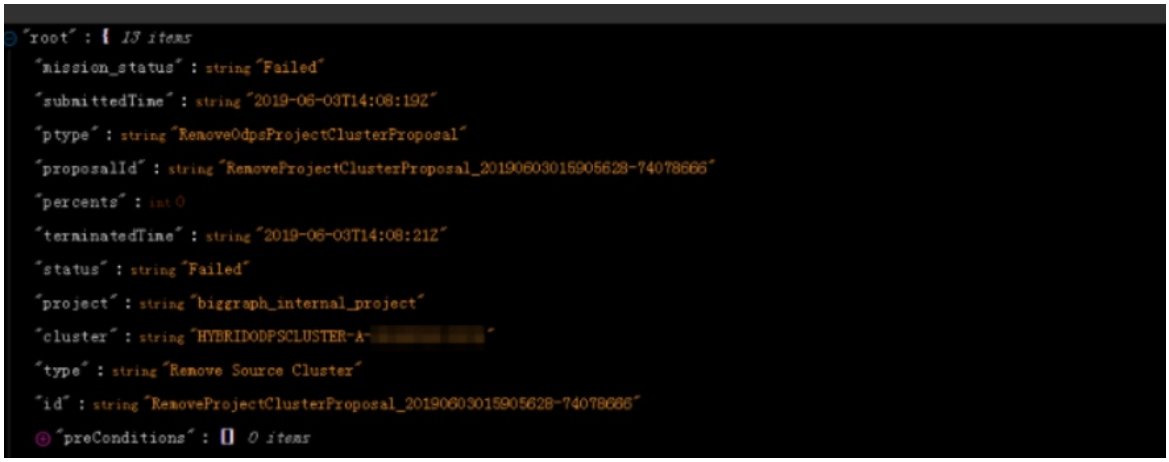

3. Click the **Debugging** icon in the Actions column to view the debugging information of the step.

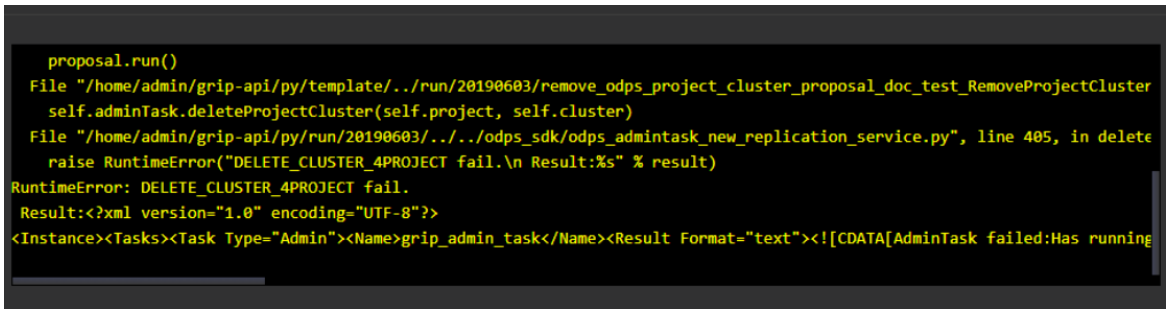

Modify a project migration task

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 required task, click **Modify Mission** in the Actions column, or click **Replan** on the **Migration Details** page.

6.2.1.5.2.3. Manage quota groups

Apsara Big Data Manager (ABM) shows the quota groups of all projects in a MaxCompute cluster. It allows you to create and modify quota groups. You can also view details about quota groups and enable period management for 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 of the tab that appears, click **Quota Groups**. Then, click **Quota Groups** or **Periods** as required.

Create a quota group

In the upper-right corner of the **Quota Groups** page, click **Create Quota Group**. In the panel that appears, configure the parameters and click **Run**.

Parameter	Description
Cluster	The cluster of the quota group that you want to create.
Quota Group	The name of the quota group that you want to create.
Preemption Policy	The preemption policy of the quota group. Valid values: No Preemption and Preemption. Default value: No Preemption.
Scheduling Type	The type of resource scheduling. Valid values: First In, First Out and Average. Default value: First In, First Out.
Minimum CUs	The minimum number of compute units (CUs) that are provided by the quota group.
Maximum CUs	The maximum number of CUs that are provided by the quota group.
CPU-to-Memory Ratio	The ratio of CPUs to memory of hosts in the quota group.

Modify a quota group

On the **Quota Groups** page, find the quota group that you want to modify and click **Modify** in the Actions column. In the panel that appears, modify the settings and click **Run**.

> **Note** If period management has been enabled for the quota group you want to modify, first modify the period management configuration.

View details about a quota group

On the **Quota Groups** page, find the quota group whose details you want to view and click **Details** in the Actions column. Then, you can view information about the resource usage, resource analysis, and period management of the quota group.

Enable period management for a quota group

1. On the **Periods** page, find the quota group for which you want to enable period management and click **Period Management** in the Actions column.
2. On the **Period Management** tab, click **Set Periods**. In the dialog box that appears, set Period and click **Enable Period Management**.

   > **Note**
   > - You can click **Add** to specify more than one period and **Delete** to delete a period.
   > - For the quota group that has period management enabled, click **Edit** in the Actions column. In the **Modify Period Configuration** panel, you can modify the parameters of the quota group within the specified period.

3. To disable period management for a quota group, click **Set Periods** again. In the dialog box that appears, click **Disable Period Management**.

6.2.1.5.2.4. Job management

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 Big Data Manager (ABM). You can also view Logview information about jobs, terminate jobs, and collect job logs.

View job snapshots

You can view job snapshots by day in the last week. The information about a job snapshot includes the job ID, project, quota group, submitter, running duration, minimum CPU utilization, and maximum CPU utilization. It also includes the 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, choose **Jobs > Job Snapshots**. The **Job Snapshots** page appears.

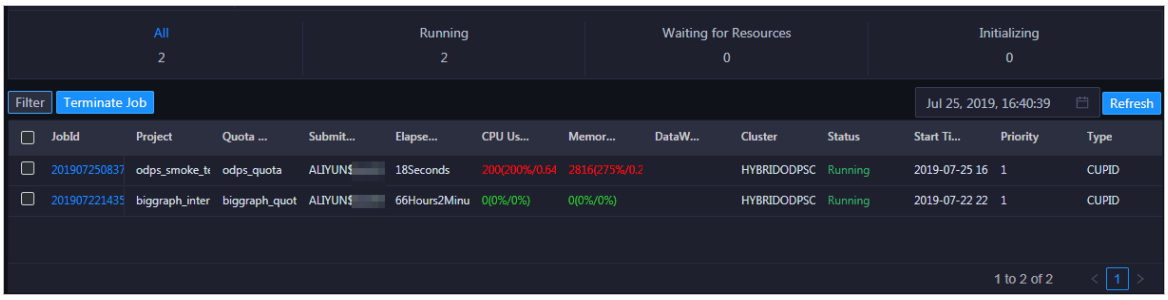

2. In the upper-right corner, select the date and time to view job snapshots by day.

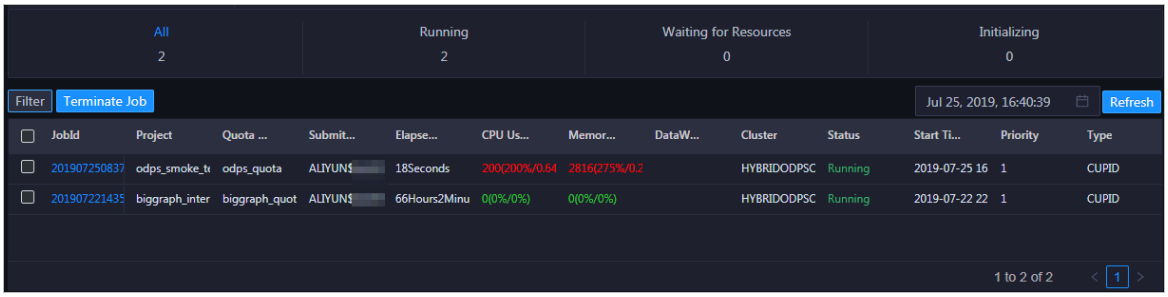

3. Click **All**, **Running**, **Waiting for Resources**, or **Initializing** to view job snapshots on the specified date.
4. Find the required snapshot and click **Logview** in the Actions column. In the dialog box that appears, click **Run** to view Logview information about the job.

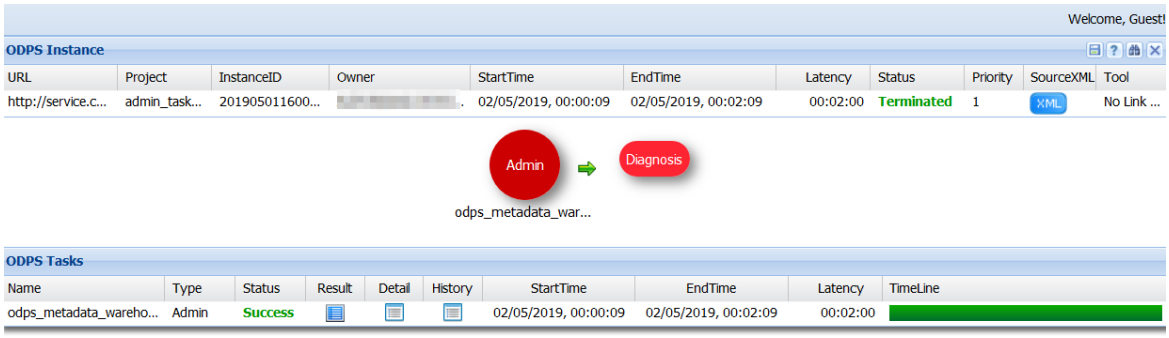

Terminate jobs

1. In the left-side navigation pane of the **Business** tab, choose **Jobs > Job Snapshots**. The **Job Snapshots** page appears.

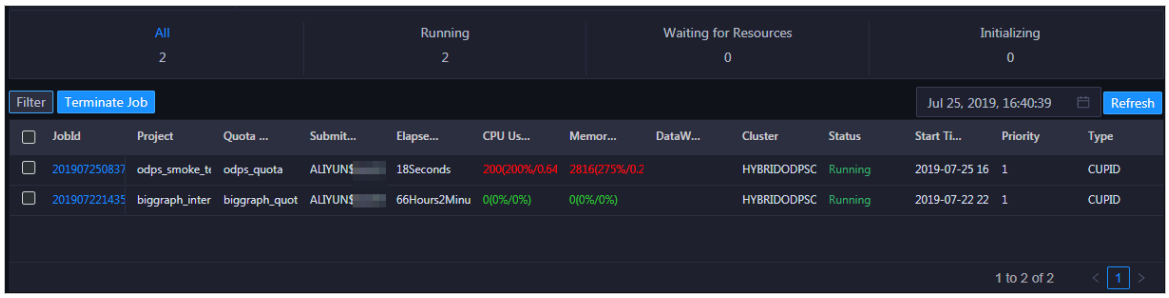

2. Select one or more jobs and click **Terminate Job** above the snapshot list. In the panel that appears, view information about the job or jobs that you want to terminate.

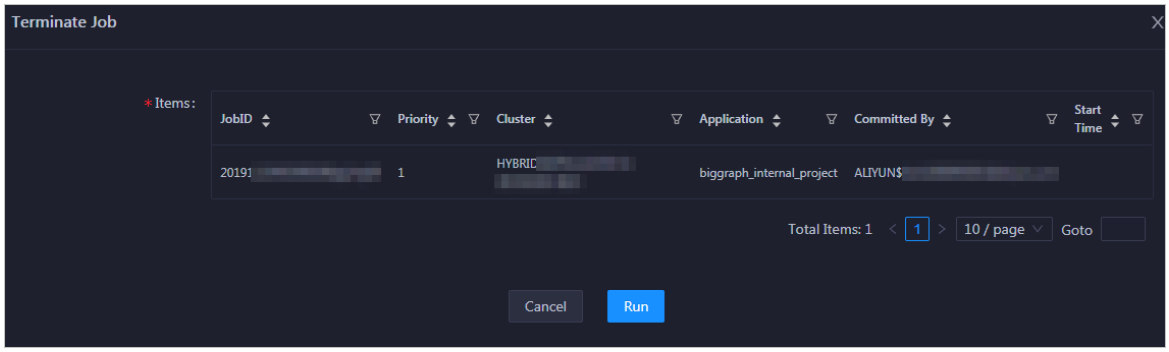

3. Click **Run**. A message appears, indicating the running result.

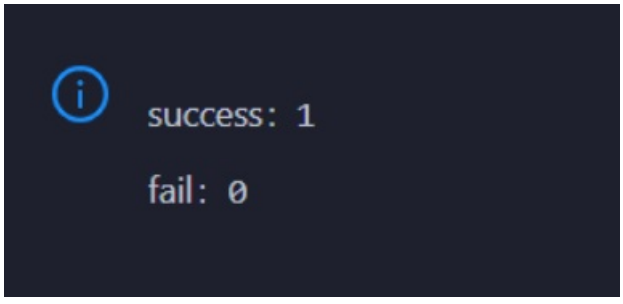

Collect job logs

If 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, choose **Jobs > Job Snapshots**. The **Job**

Snapshots page appears.

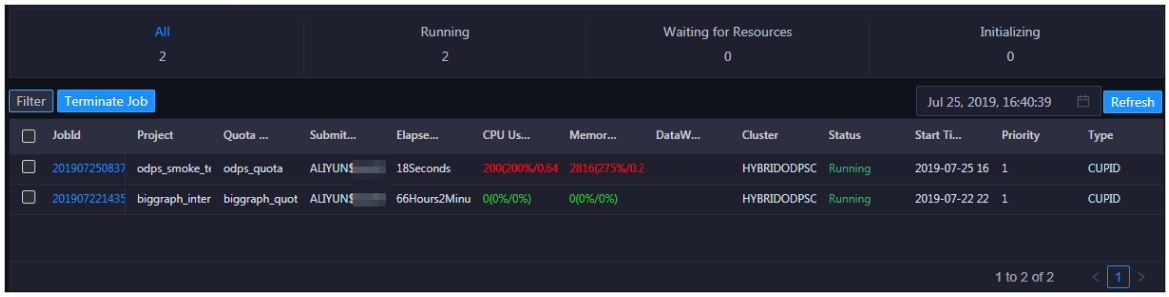

2. In the upper-right corner of the **Job Snapshots** page, choose **Actions > Collect Job Logs**.
3. In the **Collect Job Logs** panel, configure the parameters.

   The following table describes the parameters.

Parameter	Description
Target Service	The service from which you want to collect job logs.
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 **Execution History** panel, 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** panel, 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.

6.2.1.5.2.5. Business optimization

Merge small files

Excessive small files in a MaxCompute cluster occupy a lot of memory resources. Apsara Big Data 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, choose **Business Optimization > File Merging**. The **Merge Tasks** tab appears.

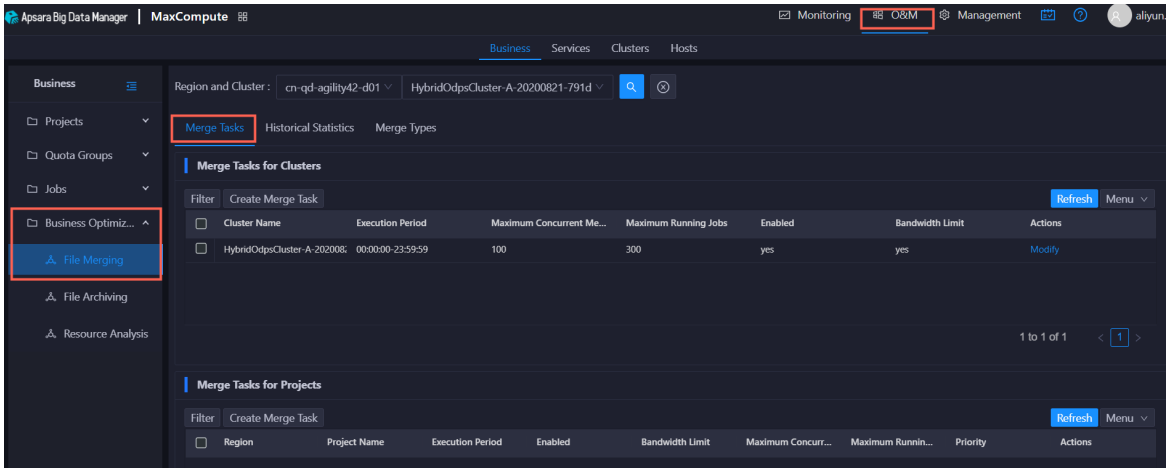

2. In the **Merge Tasks for Clusters** section, click **Create Merge Task**. In the Modify Merge Task for Cluster panel, specify the required parameters.

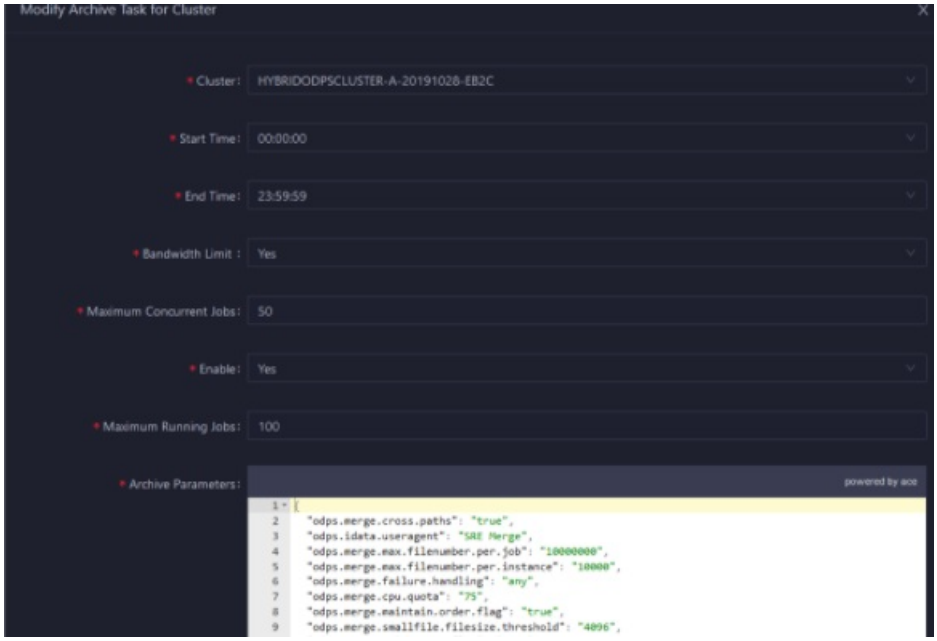

The following table describes the parameters.

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, choose **Business Optimization > File Merging**. The **Merge Tasks** tab appears.

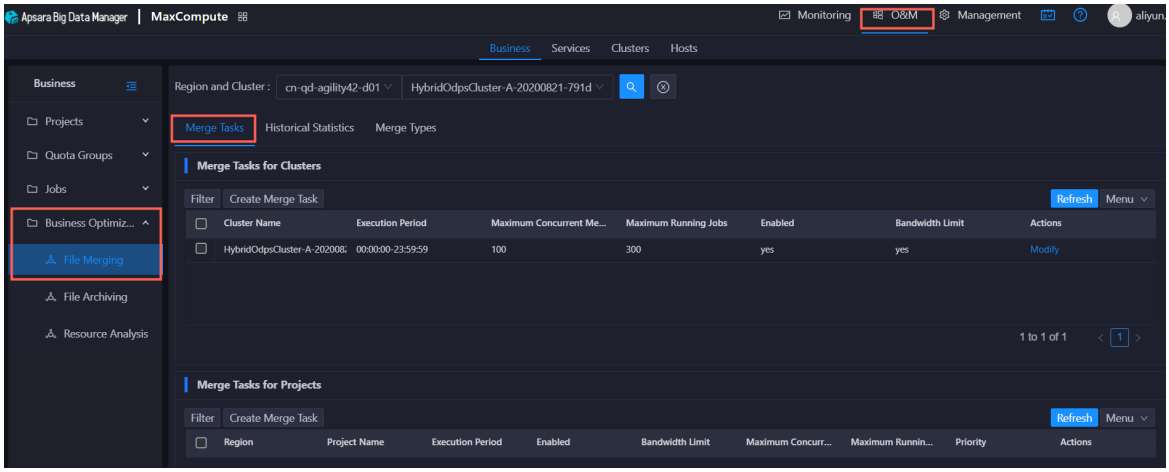

2. In the **Merge Tasks for Projects** section, click **Create Merge Task**. In the Modify Merge Task for Project panel, specify the required parameters.

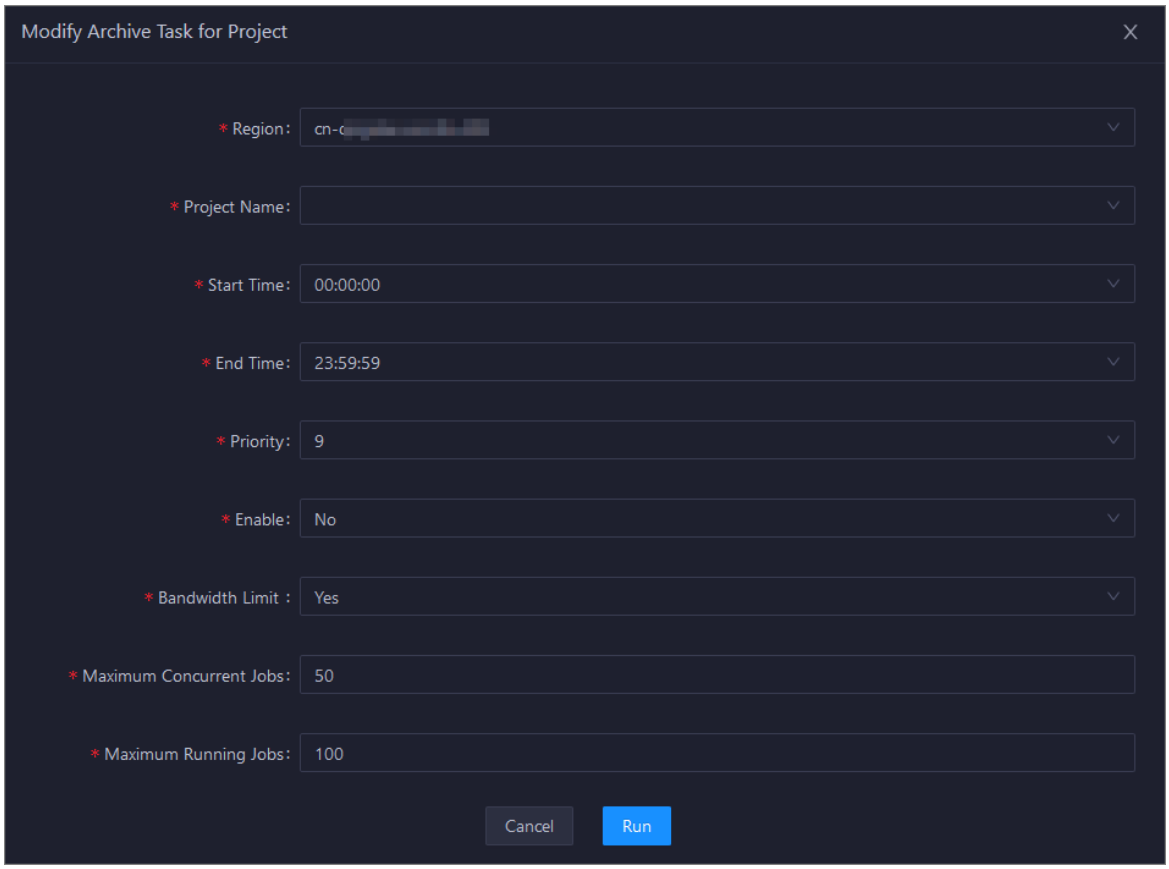

The following table describes the parameters.

Parameter	Description

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 where the selected project resides 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 where the selected project resides at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the cluster where the selected project resides, 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, choose **Business Optimization > 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.

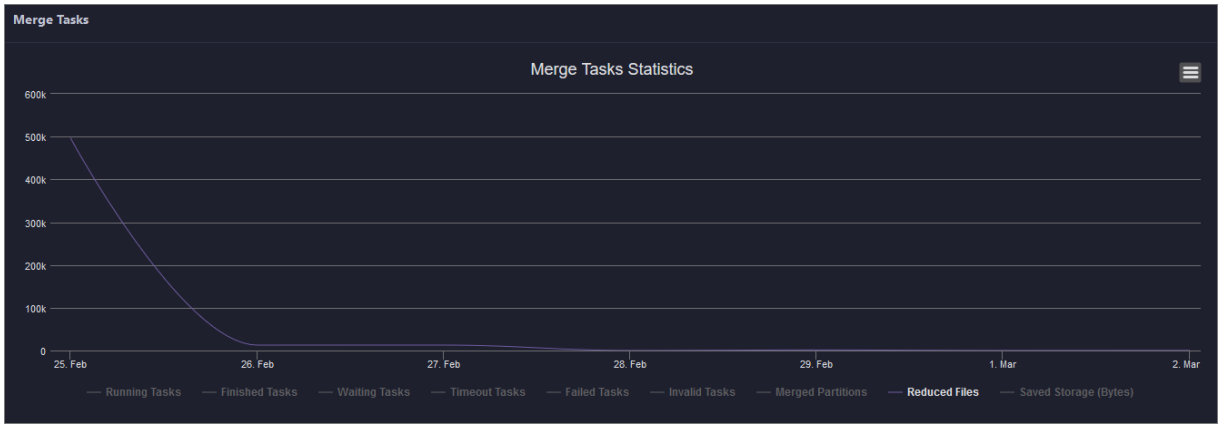

Merge Tasks for Clusters and Merge Tasks for Projects

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.

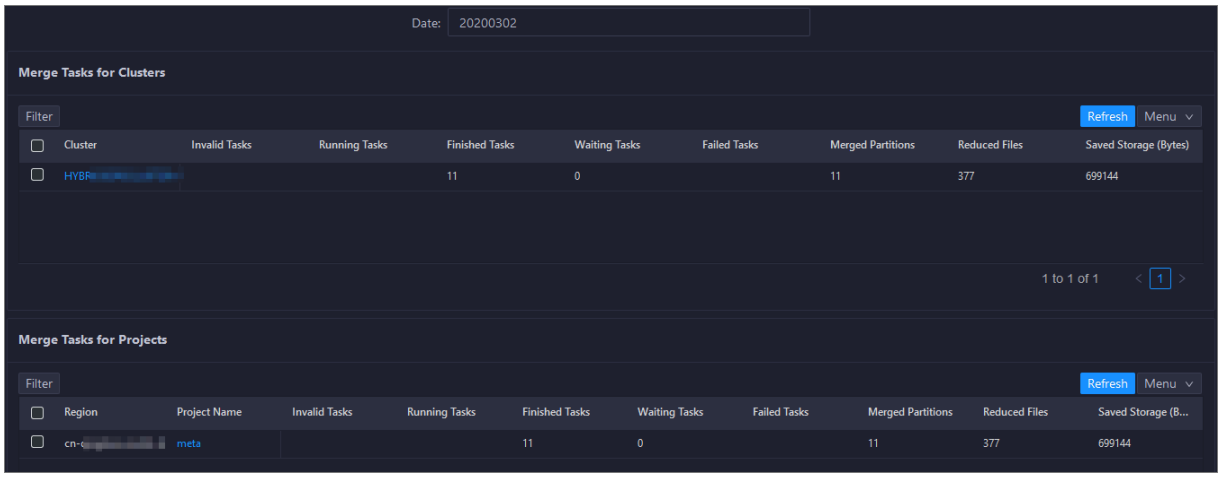

Manage merge types

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > 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 Modify Merge Type panel, specify the required parameters.

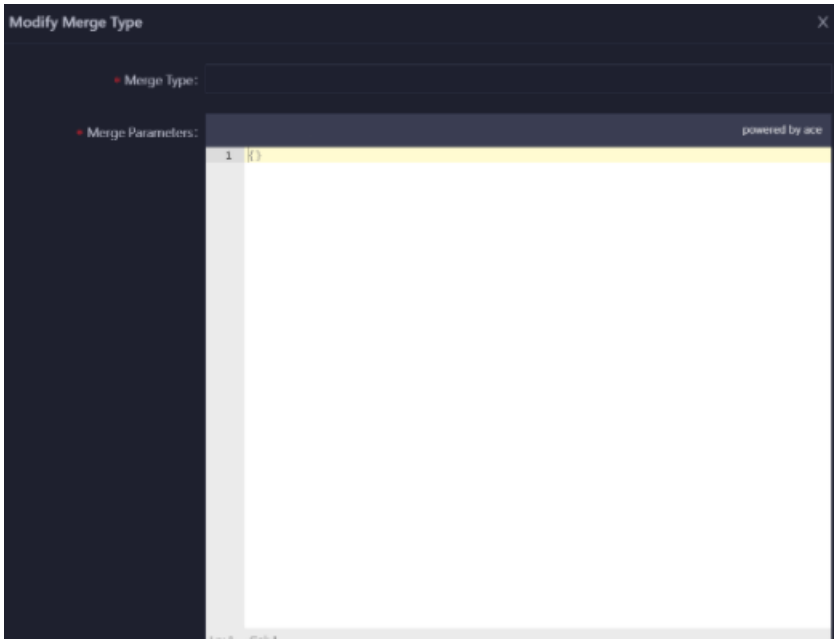

The following table describes the parameters.

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.

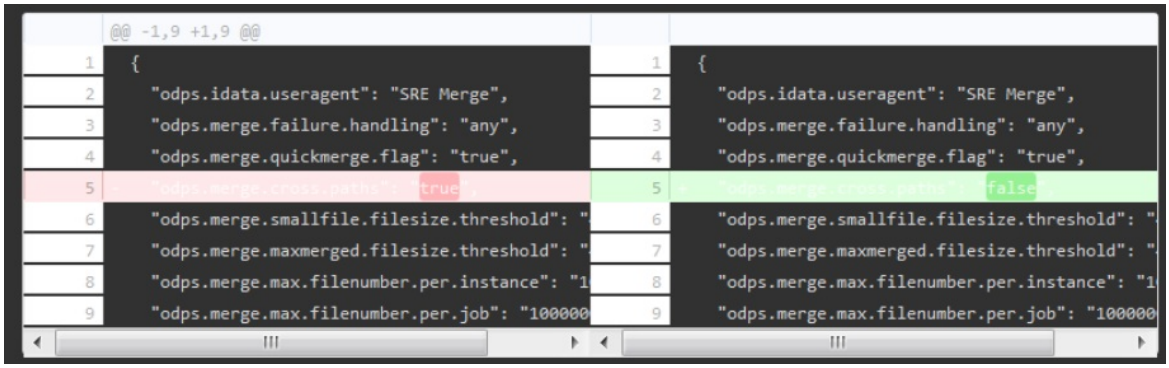

3. Click **Run**.

   The newly created merge type appears in the list of merge types.

Compress idle files

Apsara Big Data 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, choose **Business optimization > File Archiving**. The **Archive Tasks** tab appears.

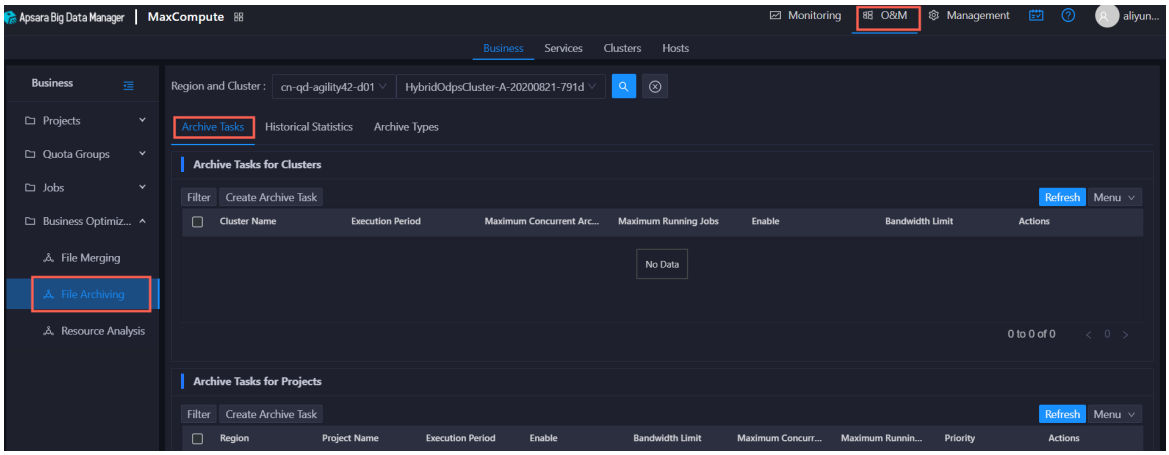

2. In the **Archive Tasks for Clusters** section, click **Create Archive Task**. In the Modify Archive Task for Cluster panel, specify the required parameters.

   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**.
Enable	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, choose **Business Optimization > File Archiving**. The **Archive Tasks** tab appears.

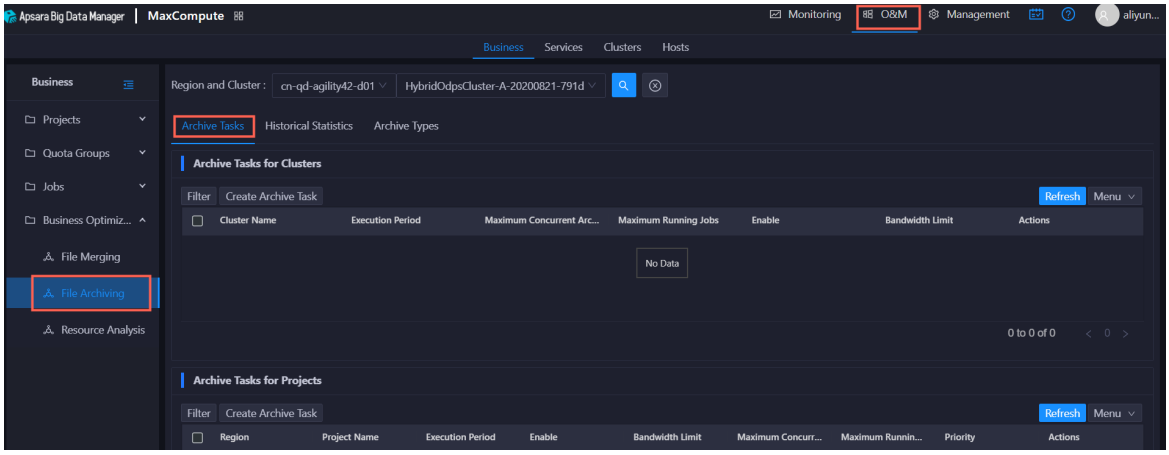

2. In the **Archive Tasks for Projects** section, click **Create Archive Task**. In the Modify Archive Task for Project panel, specify the required parameters.

   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 where the selected project resides at the same time. This parameter is valid only when **Bandwidth Limit** is set to **No**.
Enable	Specifies whether the task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run for the cluster where the selected project resides at the same time. This parameter is a global parameter. The jobs refer to all types of jobs in the cluster where the selected project resides, 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, choose **Business Optimization > 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, choose **Business Optimization > 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 Modify Archive Type panel, specify the required parameters.

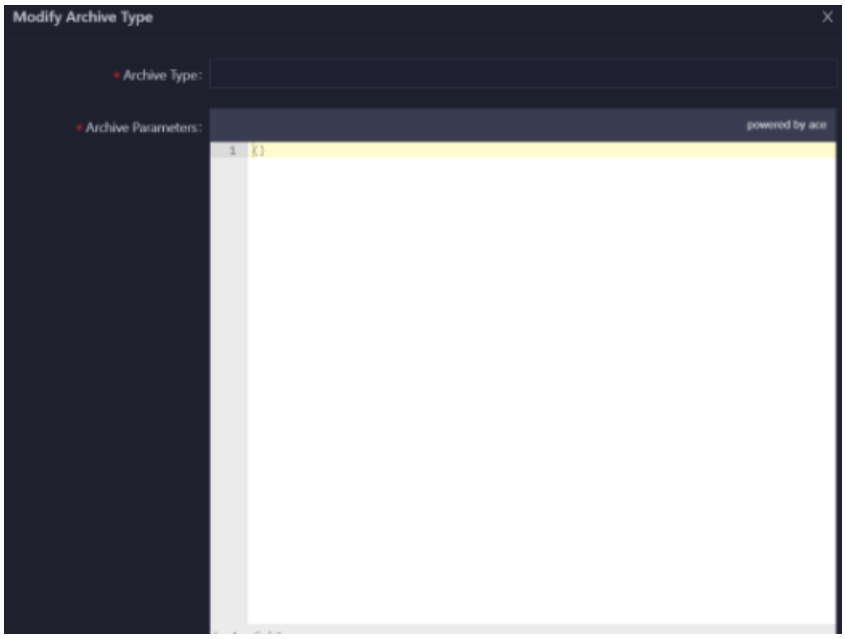

The following table describes the parameters.

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.

Analyze resources

Apsara Big Data Manager (ABM) allows you to analyze the resources for MaxCompute clusters on different tabs in the ABM console. This way, you can better understand the data storage in MaxCompute. The tabs include Tables, Projects, Tasks, Execution Time, Start Time, and Engines.

Tables

On the Tables tab, you can view the detailed information about all tables in each project, including Partitions, Storage Usage (GB), Pangu File Count, Partitions Ranking, Storage Usage Ranking, and Pangu File Count Ranking. You can sort tables by partition quantity, physical storage usage, and file quantity of Apsara Distributed File System.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. The **Tables** tab appears.

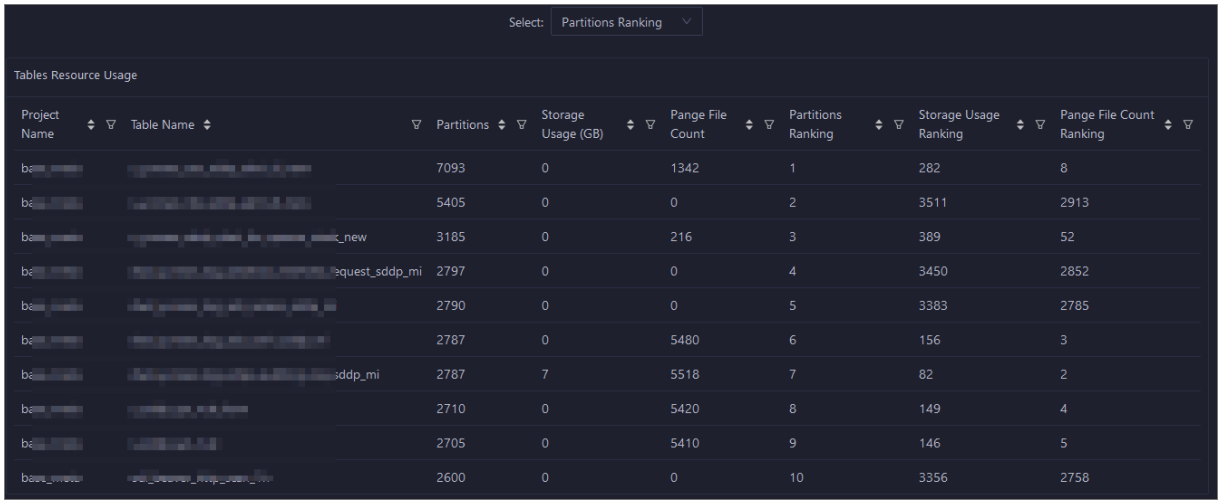

Projects

On the Projects tab, you can view the detailed information about storage for each project, including Pangu File Count, Storage Usage (GB), CU Usage, Total Memory Usage, Tasks, Tables, Idle Storage, and daily and weekly increases in percentage of these items.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Projects** tab.

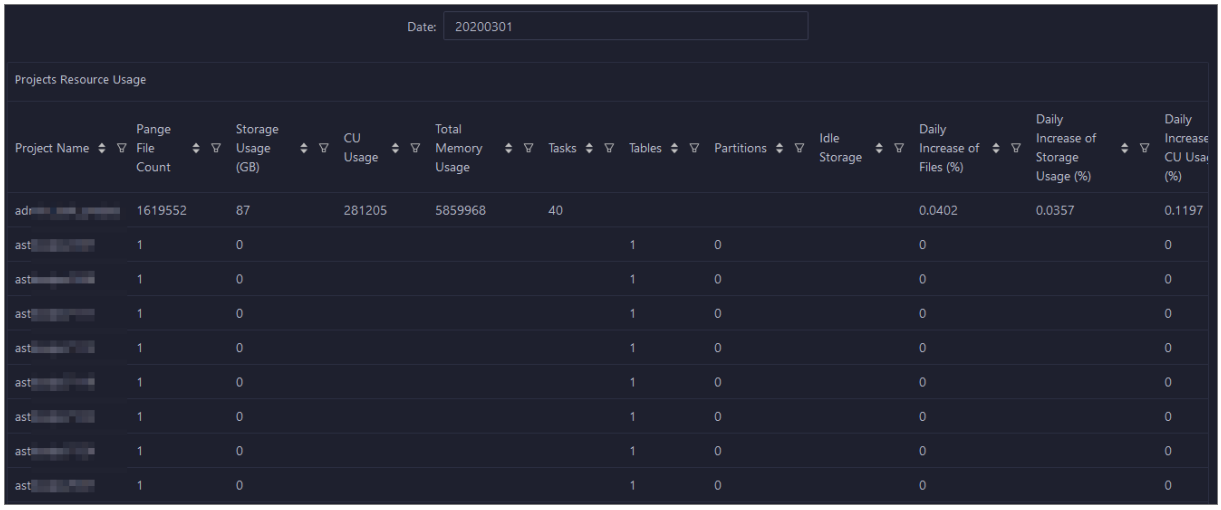

Tasks

On the Tasks tab, you can view the detailed information about all tasks in each project, including instanceid, Status, CU Usage, Start Time, End Time, Execution Time (s), CU Usage Ranking, and SQL Statements.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Tasks** tab.

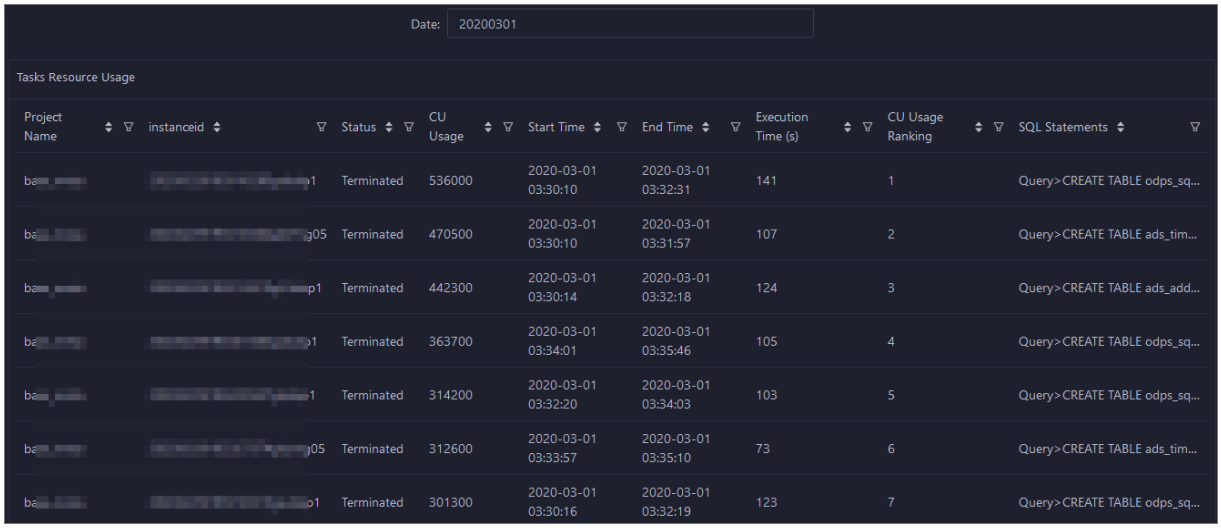

Execution Time

On the Execution Time tab, you can view the numbers of tasks whose execution time is within different time ranges in each project. The metrics include Less than 5 Minutes, Less than 15 Minutes, Less than 30 Minutes, Less than 60 Minutes, and More than 60 Minutes. The Execution Time chart displays the trend lines of task quantity in different colors by day.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Execution Time** tab.

Start Time

On the Start Time tab, you can view the numbers of tasks started in different time periods for each project. The time interval is 30 minutes. The Tasks chart displays the trend line of the number of tasks started in a specified time period by day.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Start Time** tab.

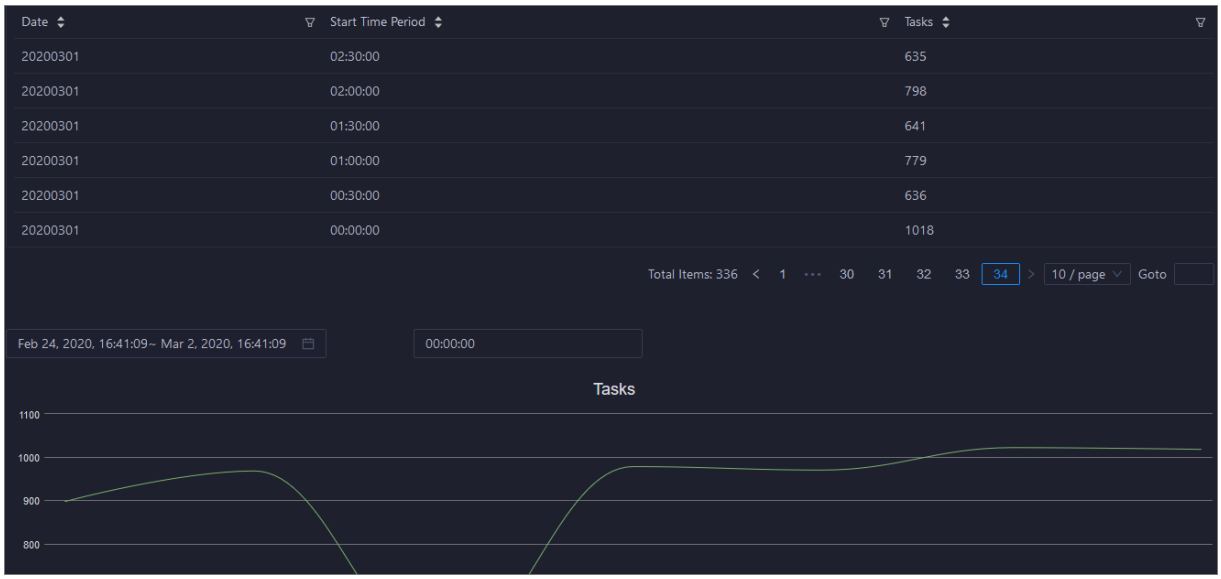

Engines

On the Engines tab, you can view the trend lines of performance statistics of tasks in each project in the Task Performance Analysis chart. The performance metrics include cost_cpu, cost_mem, cost_time, input_bytes, input_bytes_per_cu, input_records, input_records_per_cu, output_bytes, output_bytes_per_cu, output_records, and output_records_per_cu.

In the left-side navigation pane of the **Business** tab, choose **Business Optimization > Resource Analysis**. Click the **Engines** tab.

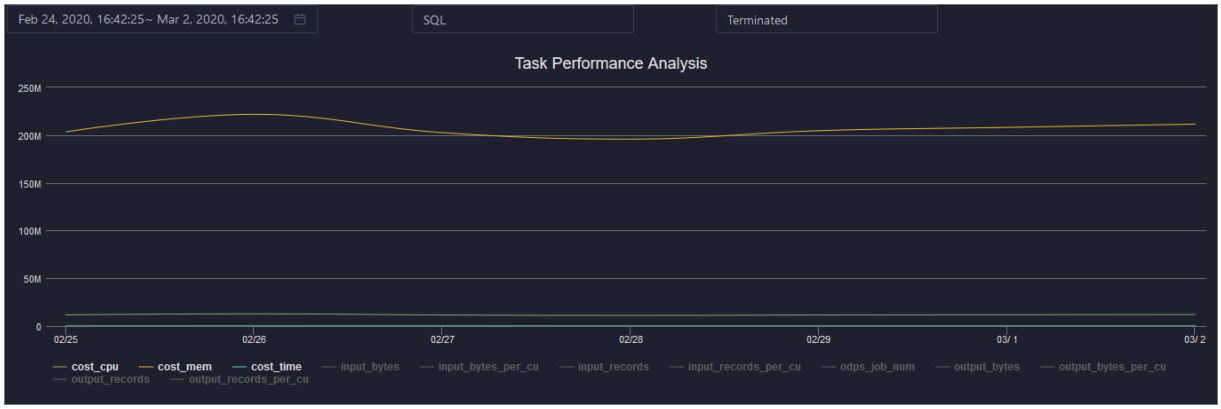

6.2.1.5.3. Service O&M

6.2.1.5.3.1. Control service O&M

O&M features 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 cause of 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 Big Data 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.

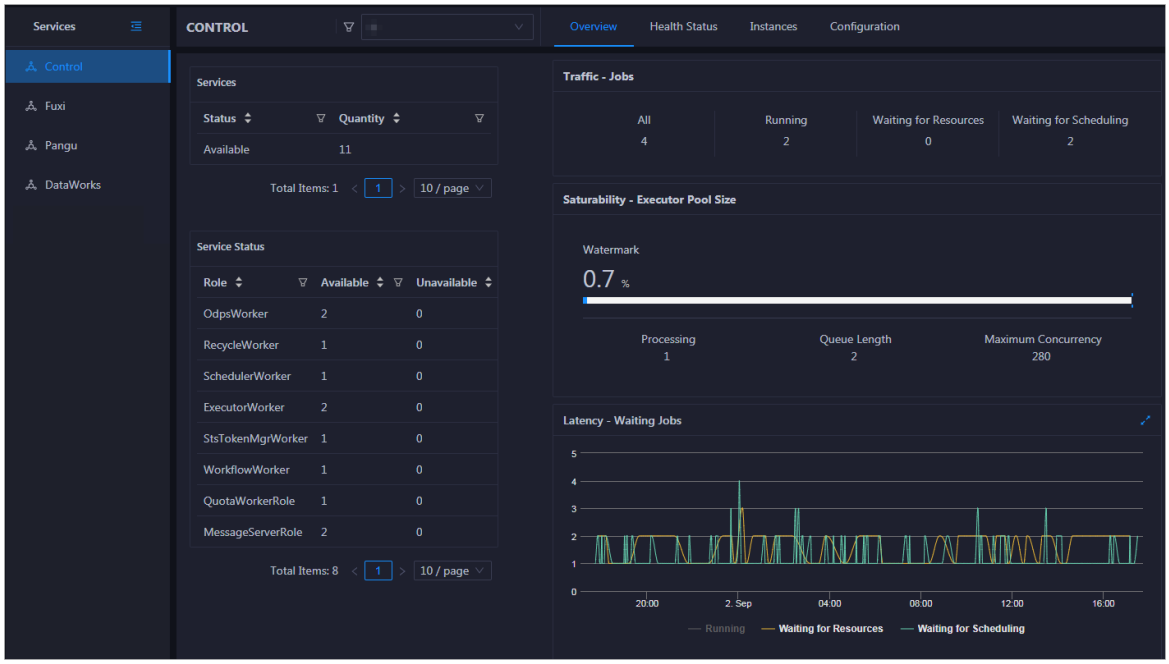

Control service overview

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.

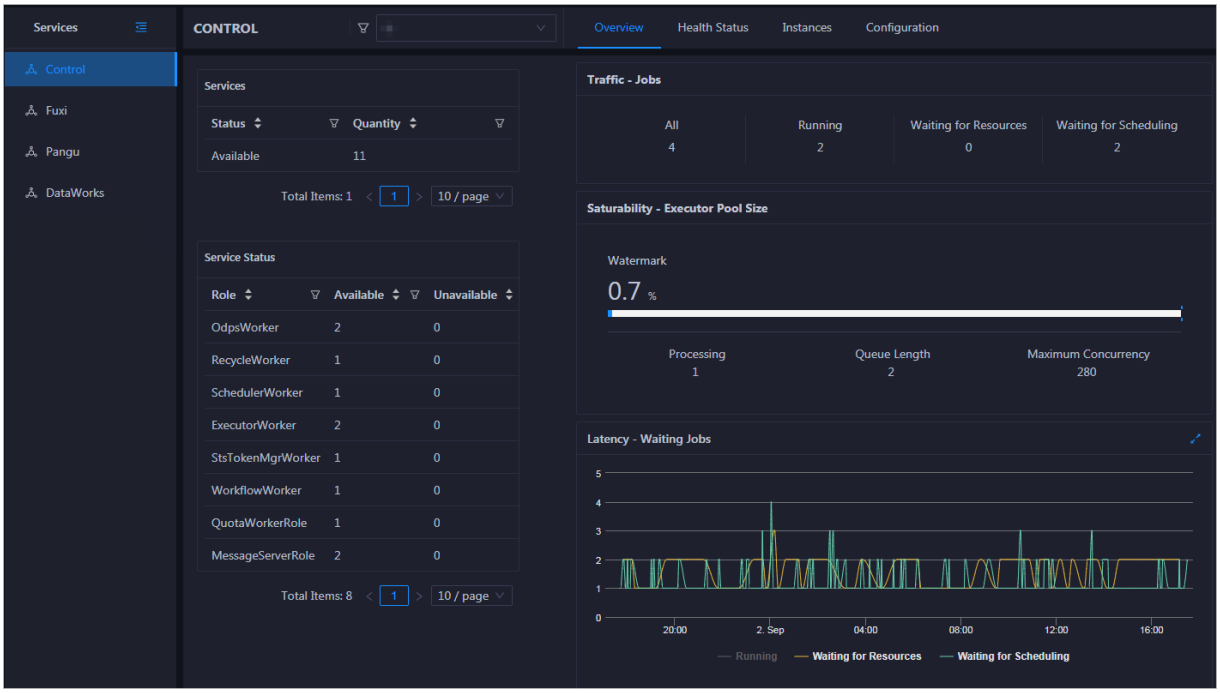

On the **Overview** page, you can view the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

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.

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.

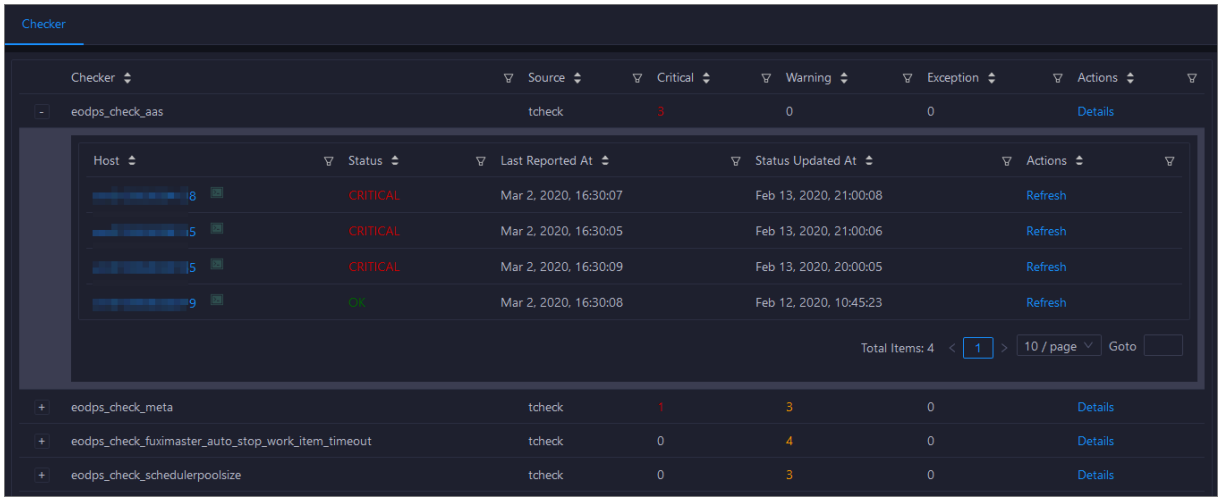

On the **Health Status** page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical**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

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.

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.

Metadata warehouse for the control service

This topic describes how 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.

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

In the left-side navigation pane of the **Services** tab, click **Control**. On the page that appears, click the **Metadata Repository** tab.

The **Metadata Repository** tab displays the completion time of the output tasks of the metadata warehouse and the trend chart of the consumed time for running tasks. The time displayed in the **Completed At** column indicates the time when an output task is complete. The time displayed in the **Collected At** column indicates the last time at which ABM obtains the status of output tasks.

Stop or start a server role

Apsara Big Data 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, choose **Actions** > **Stop Service Role**.
5. In the Stop Service Role panel, select a server role that 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 panel 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, choose **Actions** > **Start Service Role**.
5. In the Start Service Role panel, select a server role that 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 panel 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 panel, click **Details** in the Details column of the task to view the details.
2. In the Start Service Role panel, 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.

Start AdminConsole

AdminConsole is a management platform of MaxCompute. It is disabled by default. Apsara Big Data 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 page that appears, choose **Actions** > **Start Admin Console**.
6. In the **Start Admin Console** panel, 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** 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 **Start Admin Console** in the upper-right corner to view the execution history.
2. In the Execution History panel, 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.

Collect service logs

Apsara Big Data 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 page that appears, choose **Actions > Collect Service Logs**.
6. In the **Collect Service Logs** panel, specify the required parameters.

   The following table describes the parameters.

Parameter	Description
Target Service	The service from which you want to collect service logs. Select a 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 panel, 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.

6.2.1.5.3.2. Job Scheduler O&M

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: displays 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: displays 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: displays information about the master nodes and server roles of Job Scheduler and allows you to restart the master nodes.
- Compute Nodes: displays 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 Big Data 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.

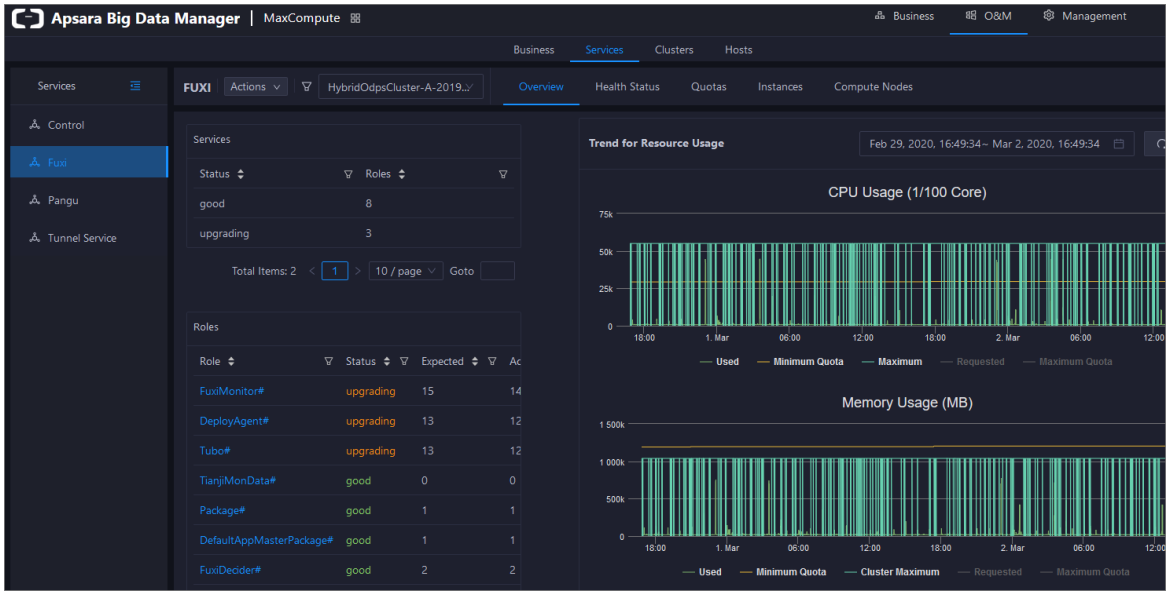

Overview

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.

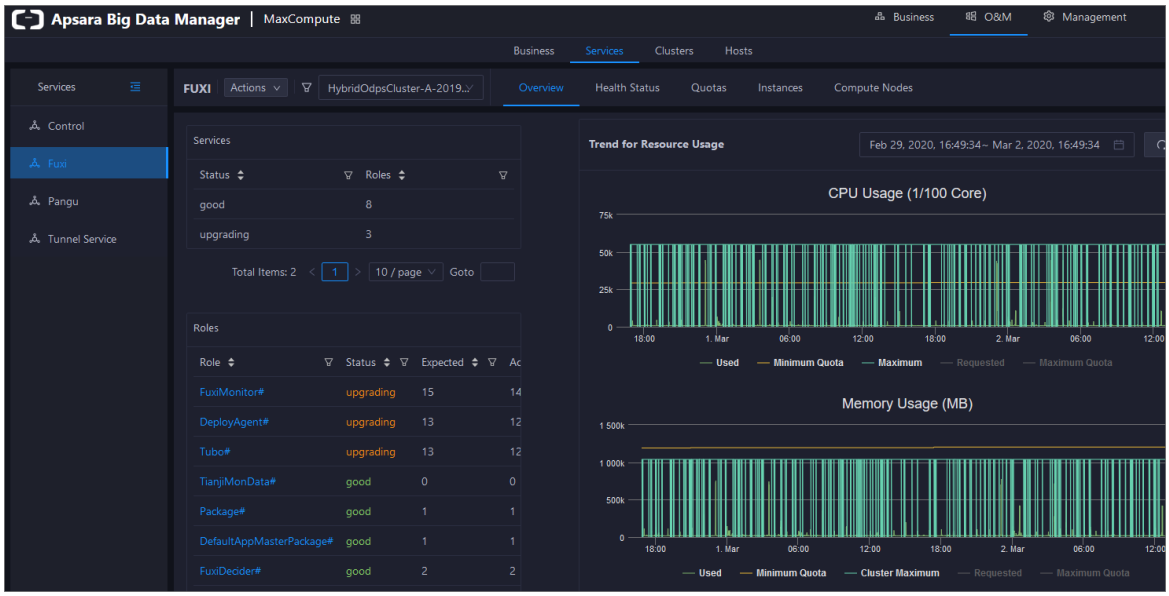

The **Overview** tab shows the key operating information of Job Scheduler. The information includes the service overview, service status, resource usage, compute node overview, and the trend charts of CPU utilization and memory usage.

Services

This section shows the numbers of available services, unavailable services, and services that are being updated.

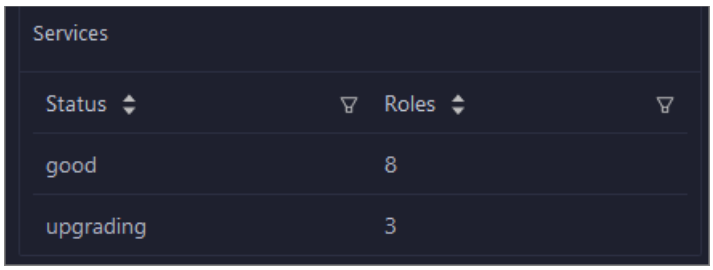

Roles

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.

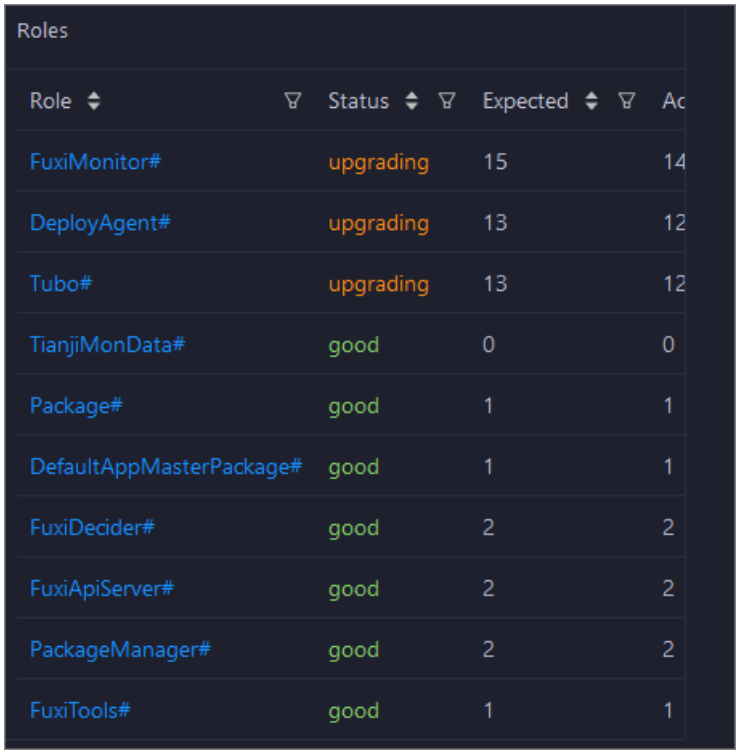

Click the name of a server role to go to the Apsara Infrastructure Management Framework console and view its details.

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.

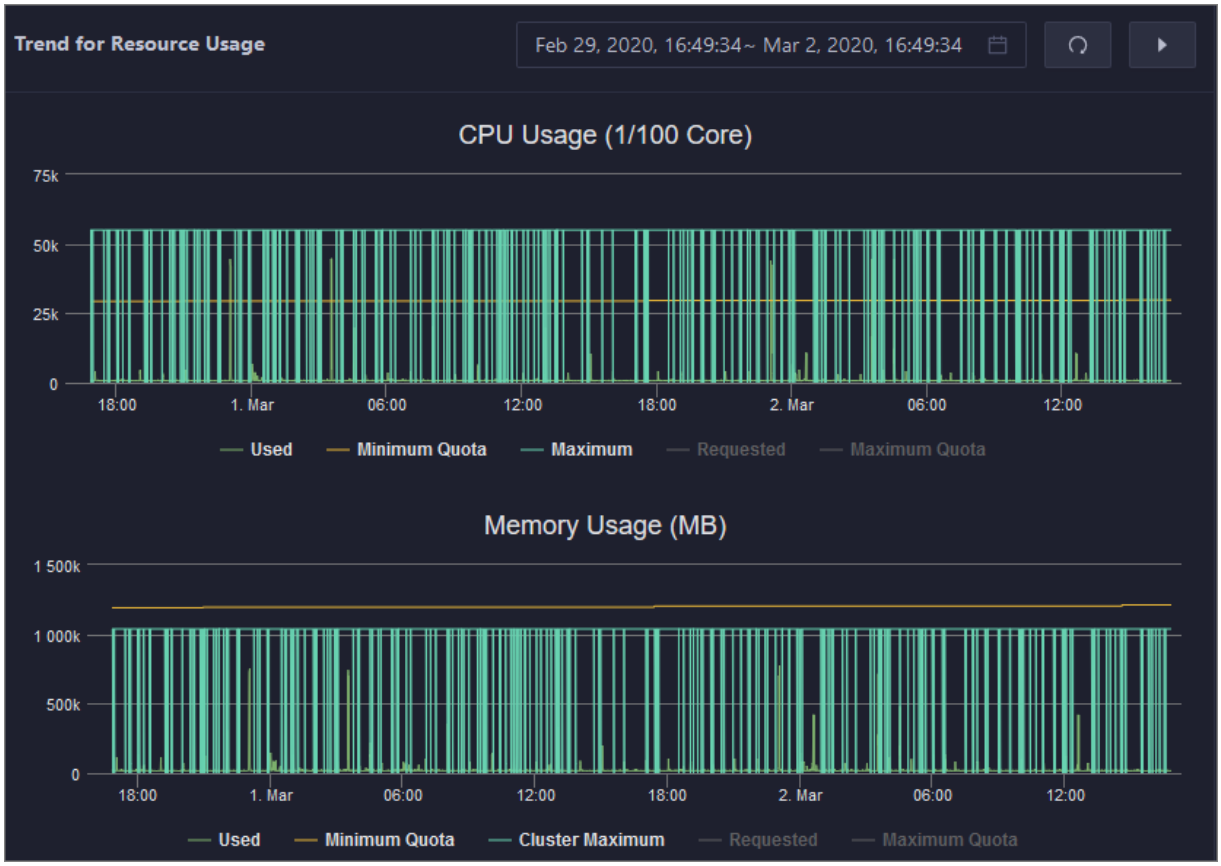

Saturability - Resource Usage

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.

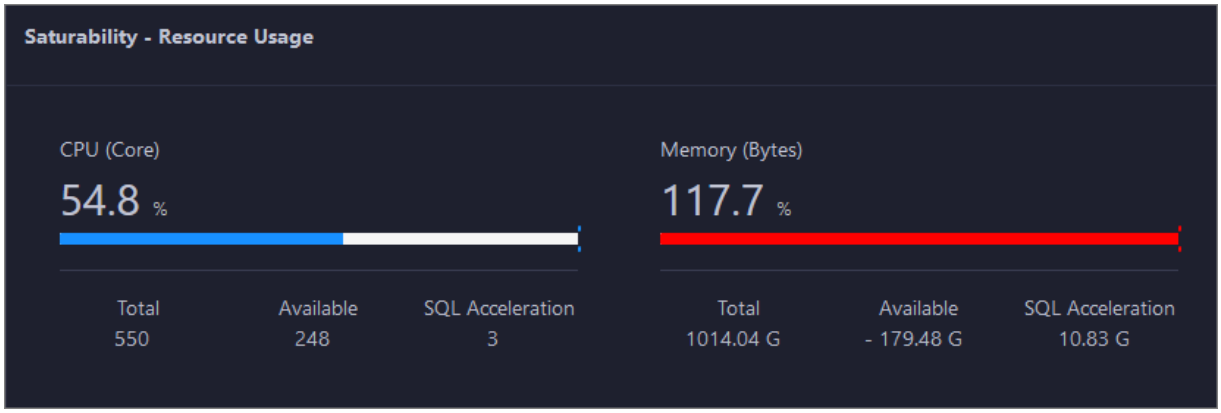

Compute Nodes

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.

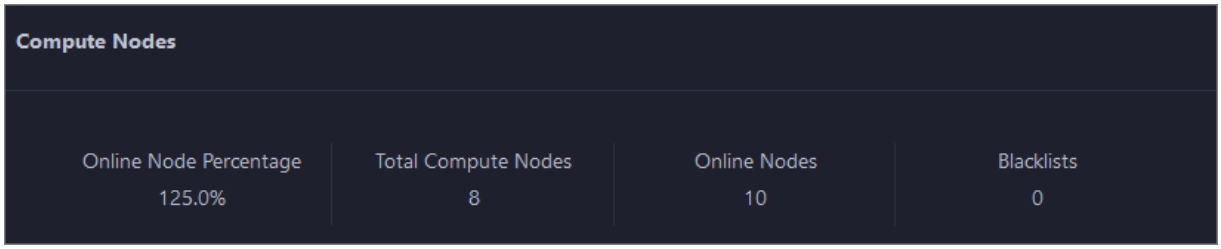

Job Scheduler health

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.

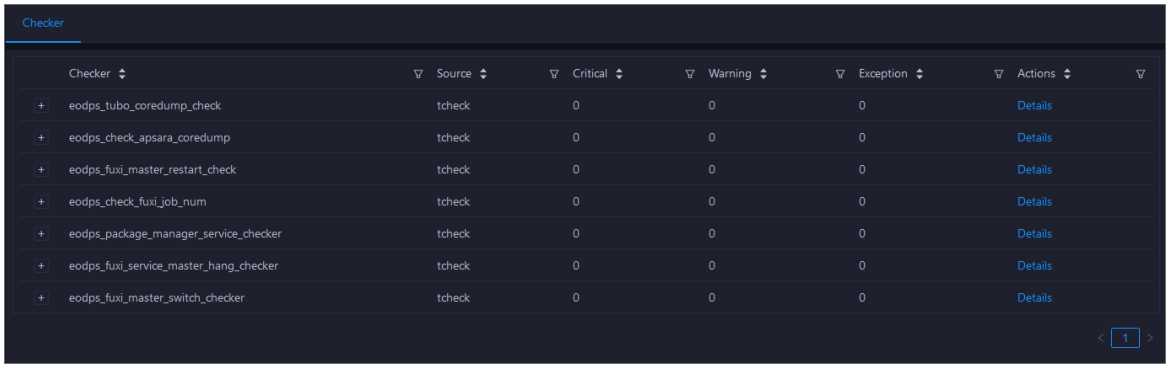

On the **Health Status** page, you can view all checkers of the Job Scheduler service and the check results for all hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

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.

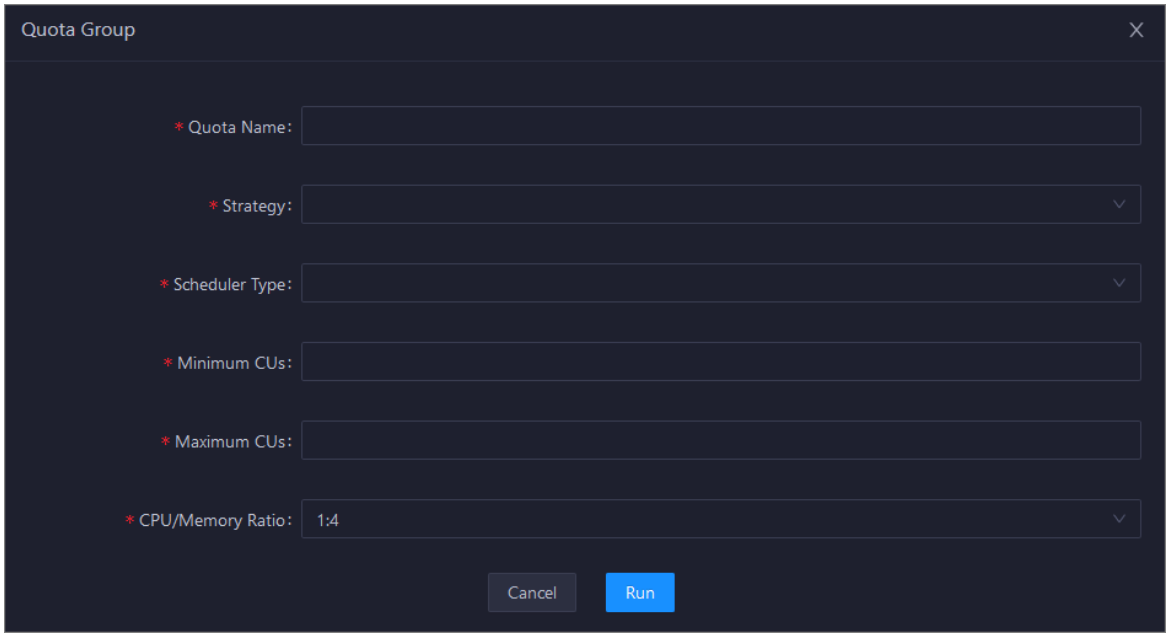

3. Click **Run**.

   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

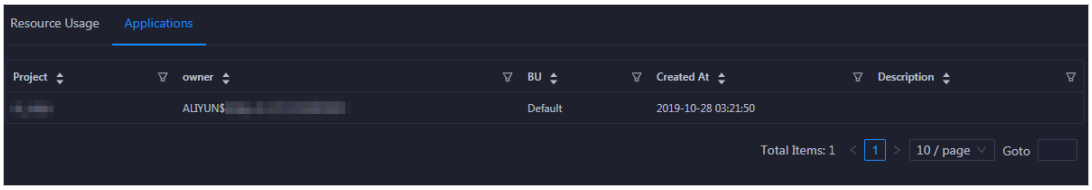

Modify a quota group

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.

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.

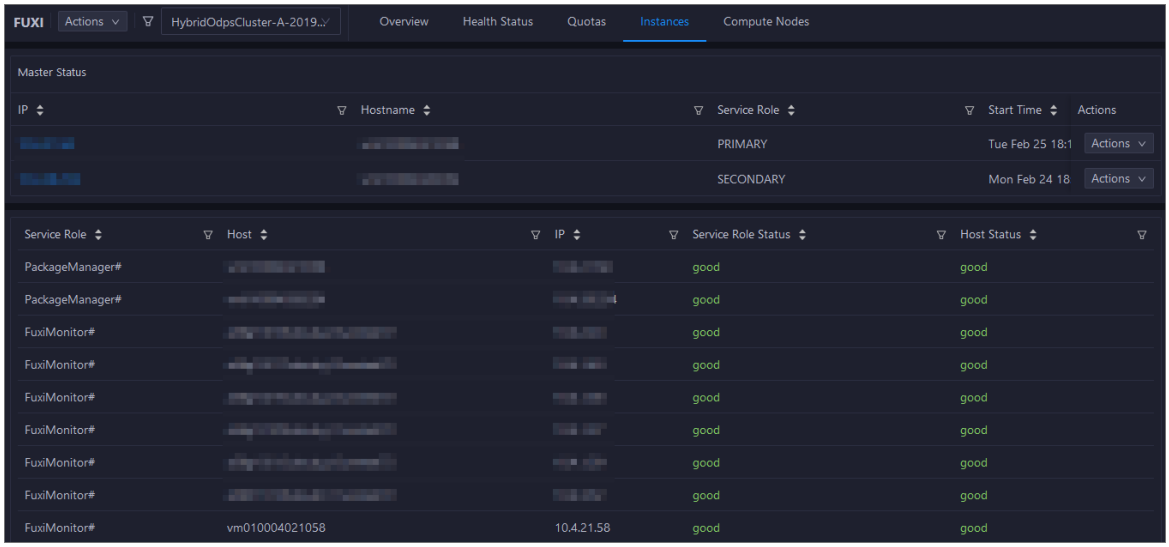

The **Instances** tab shows information about the master nodes and server roles of Job Scheduler. The information about the master nodes includes the IP address, hostname, server role, and start time. The information about a server role includes the role name, hostname, role status, and host status.

Supported operations

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

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.

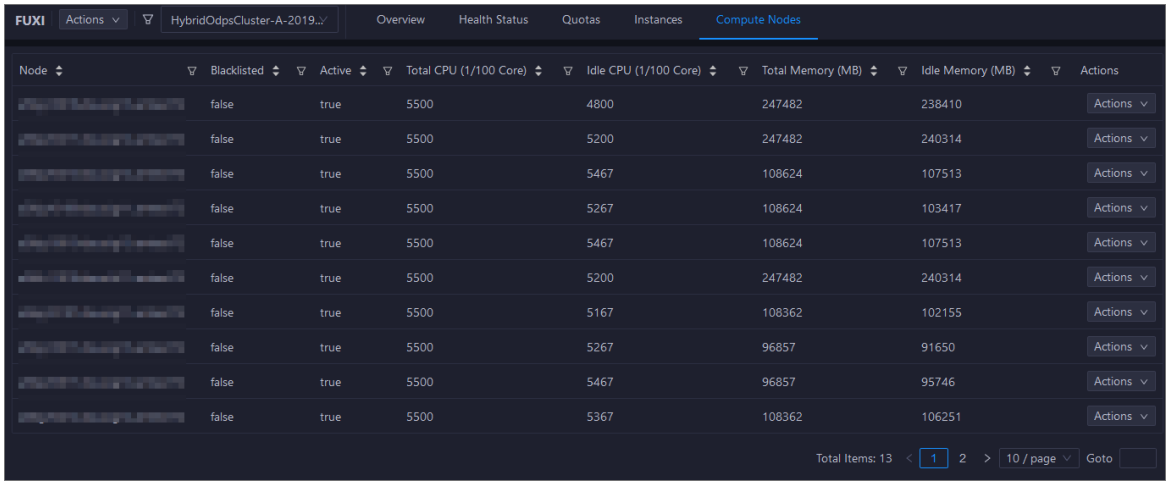

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including the total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active.

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.

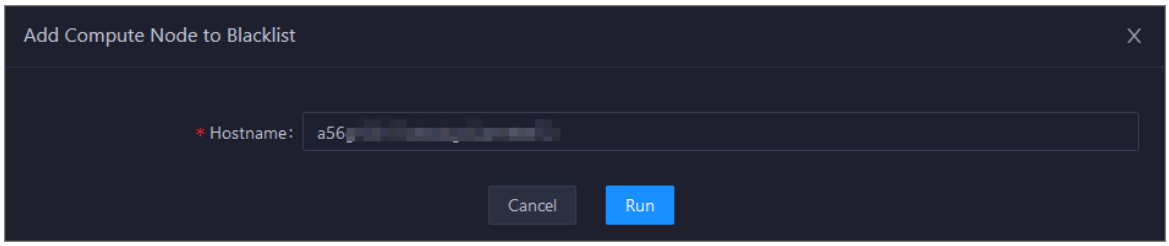

The value of the **Hostname** parameter is automatically filled. You do not need to specify a value for this parameter.

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

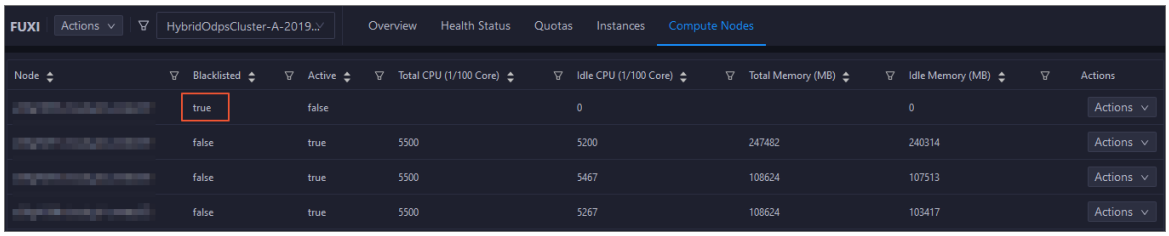

Enable and disable SQL acceleration

You can enable or disable SQL acceleration for Job Scheduler in the Apsara Big Data 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, choose **Actions** > **Enable SQL Acceleration**.
3. In the Enable SQL Acceleration panel, set the **WorkerSpans** parameter.

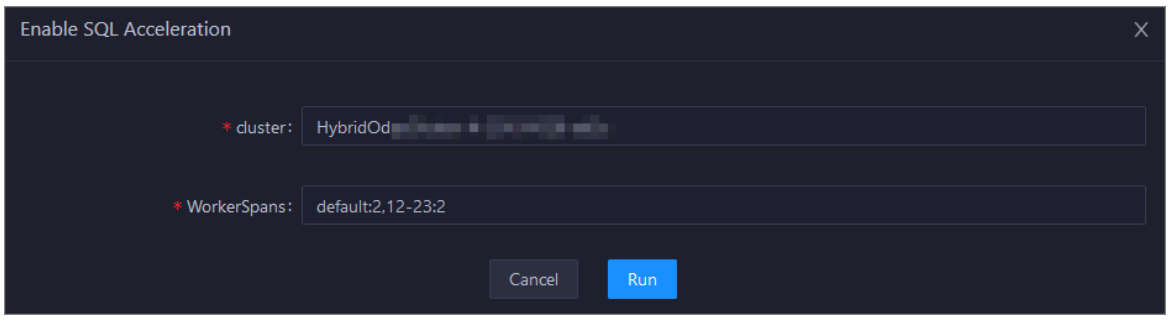

WorkerSpans: the default resource quota of the cluster and the resource quota for a specific period. Default value: **default:2,12-23:2**.

> **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, choose **Actions** > **Disable SQL Acceleration**.
3. In the Disable SQL Acceleration panel, 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 Execution History panel, view the execution history of enabling SQL acceleration.

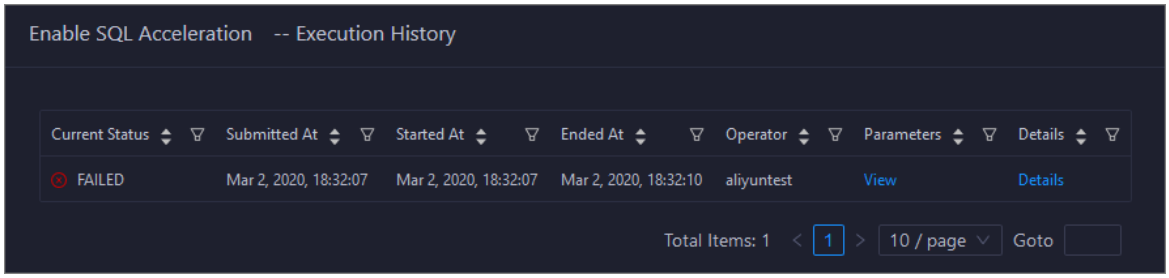

The execution history shows the current status, submission time, start time, end time, and operator of each execution.

4. If the execution fails, click **Details** in the Details column to identify the cause of the failure.

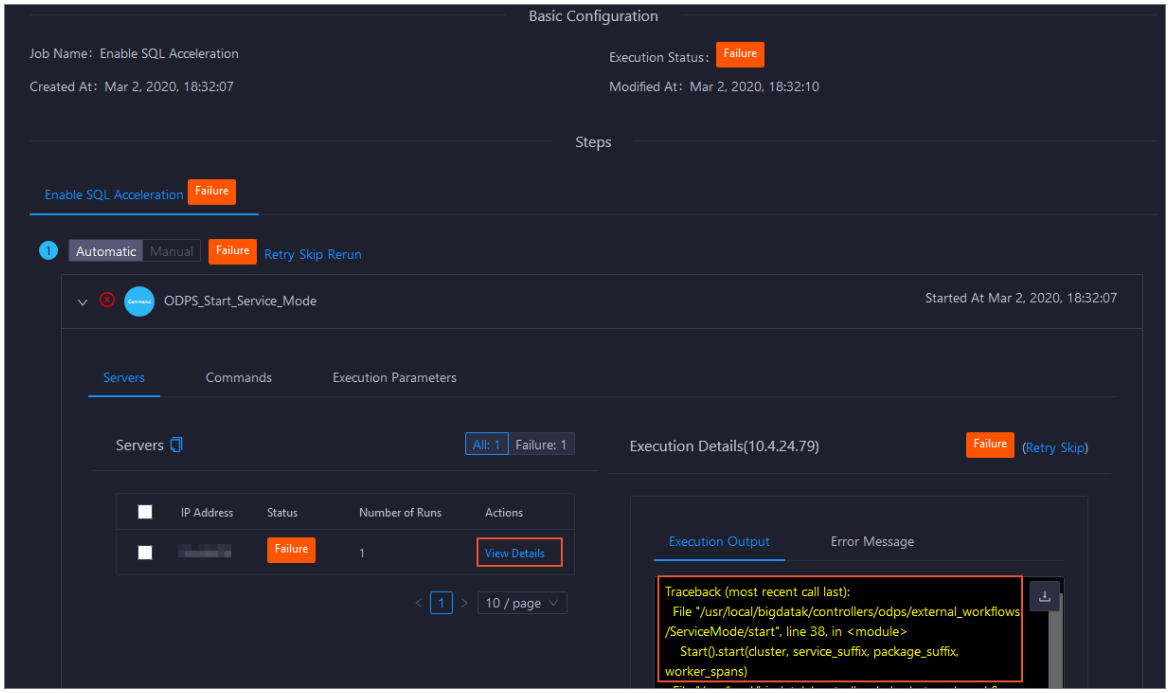

Restart a master node of Job Scheduler

Job Scheduler is the resource management and task scheduling system of the Apsara distributed operating system. Apsara Big Data 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-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**. Then, click the **Instances** tab.
5. On the **Instances** tab, choose **Actions** > **Restart Fuxi Master Node** in the Actions column of a primary or secondary master node.
6. In the **Restart Fuxi Master Node** panel, click **Run**. The **Restart Fuxi Master Node** panel appears.

Step 2: View the execution status or progress

1. In the **Restart Fuxi Master Node** panel, check the execution history of restarting master nodes.

   The **Restart Fuxi Master Node** panel displays 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) 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. In the **Restart Fuxi Master Node** panel, 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.

6.2.1.5.3.3. Apsara Distribute File System O&M

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 Big Data 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.

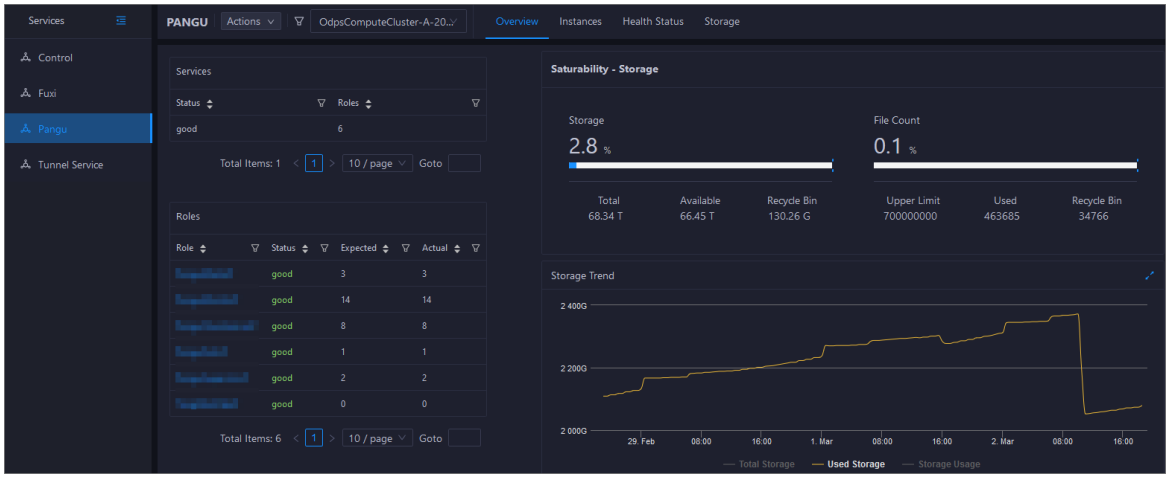

Overview

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.

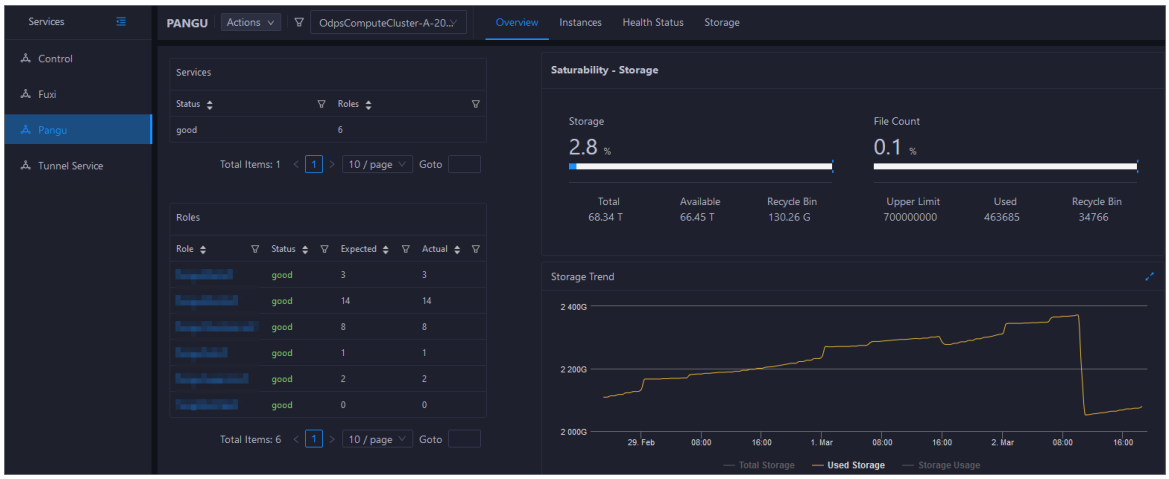

The **Overview** tab shows the key operating information about Apsara Distributed File System. The information includes the service overview, service status, health check result, health check history, storage usage, storage node overview, and the trend charts of storage usage and file count.

Services

This section shows the status of Apsara Distributed File System and the number of server roles.

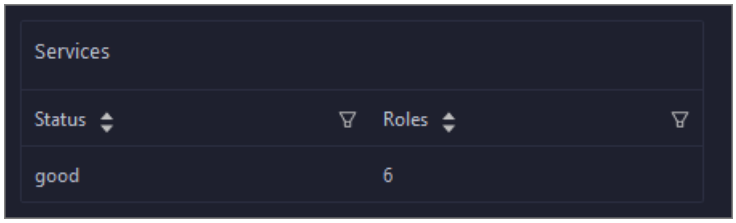

Roles

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.

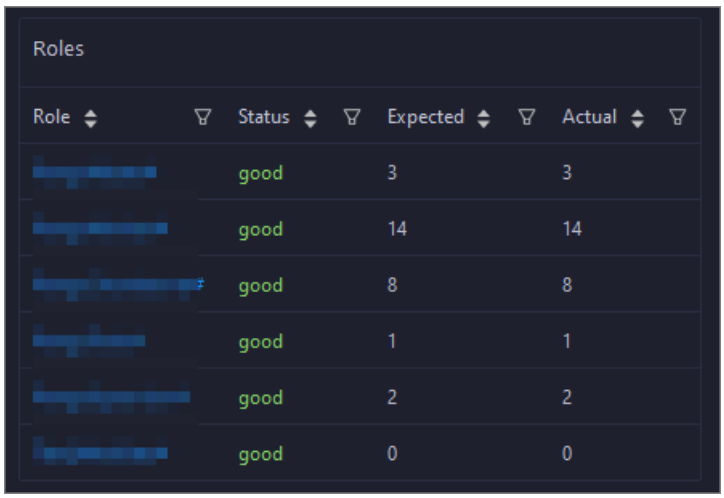

Saturability - Storage

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.

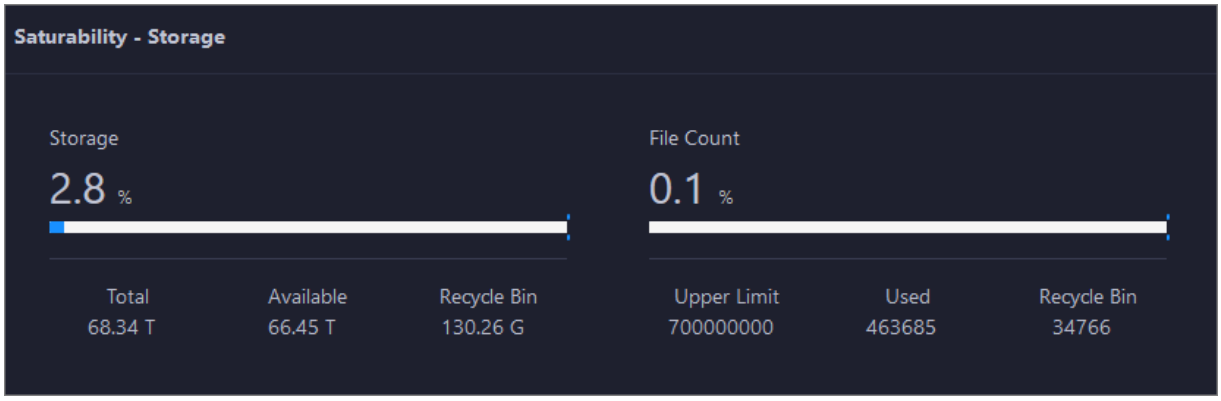

Storage Trend and File Count Trend

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.

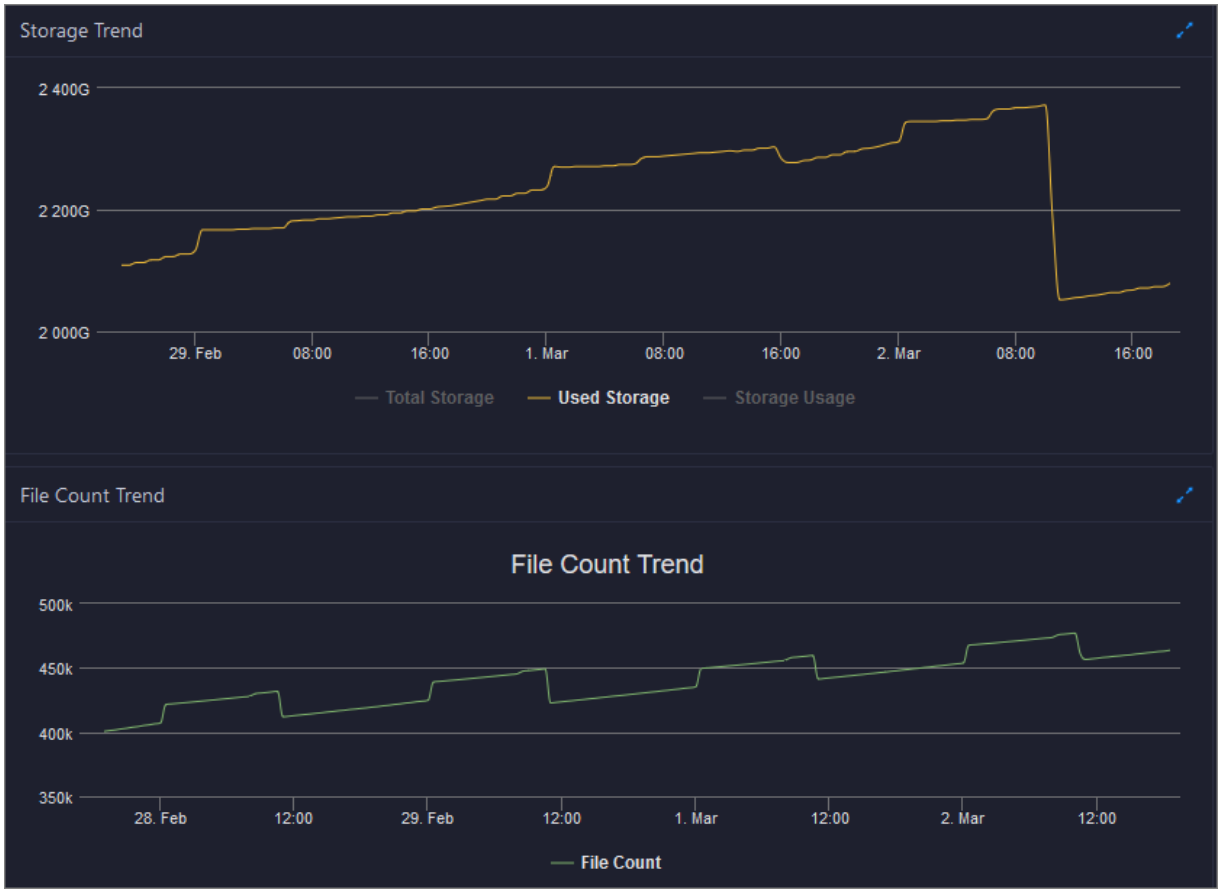

In the upper-right corner of the chart, click the icon to zoom in the chart. The following figure shows an enlarged chart of storage usage.

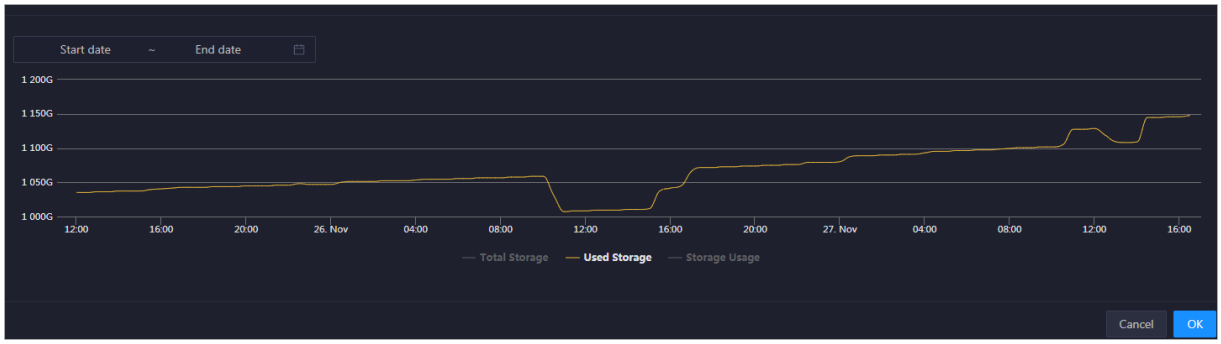

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

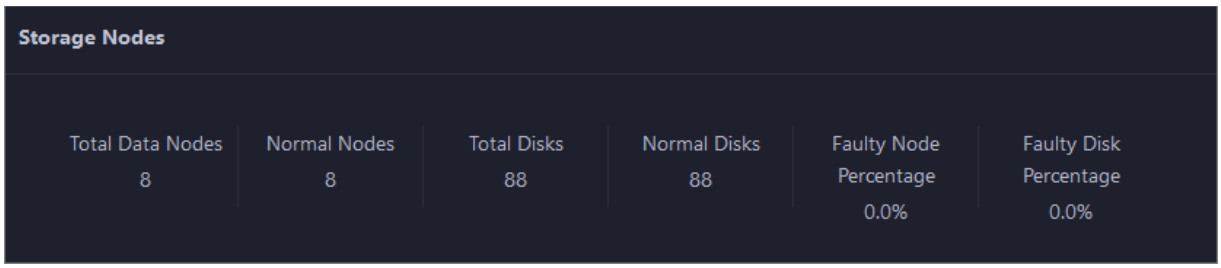

Instances

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.

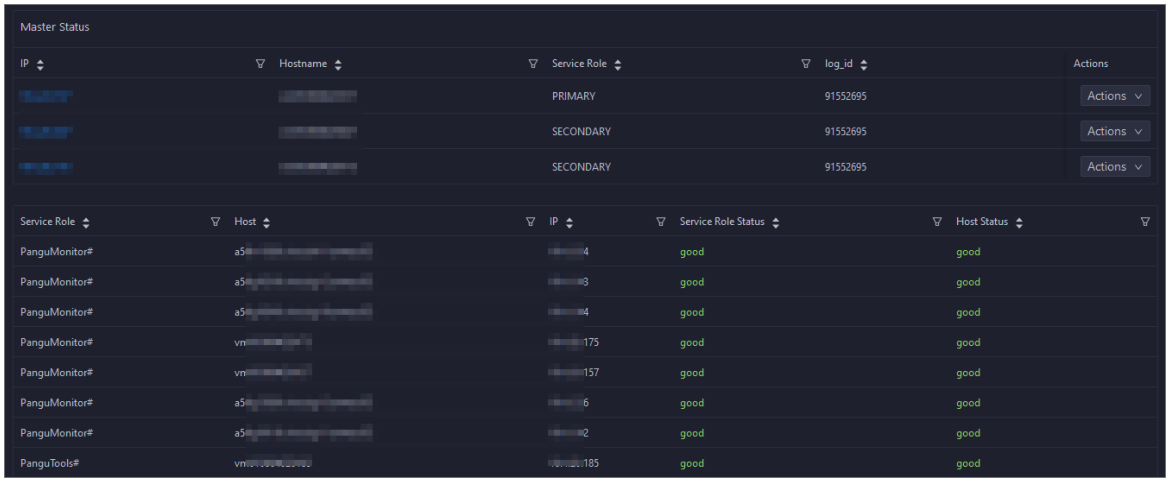

The **Instances** tab shows information about the master nodes and server roles of Apsara Distributed File System. The information about a master node includes the IP address, hostname, server role, and log ID. The information about a server role includes the role name, hostname, role status, and host 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.

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.

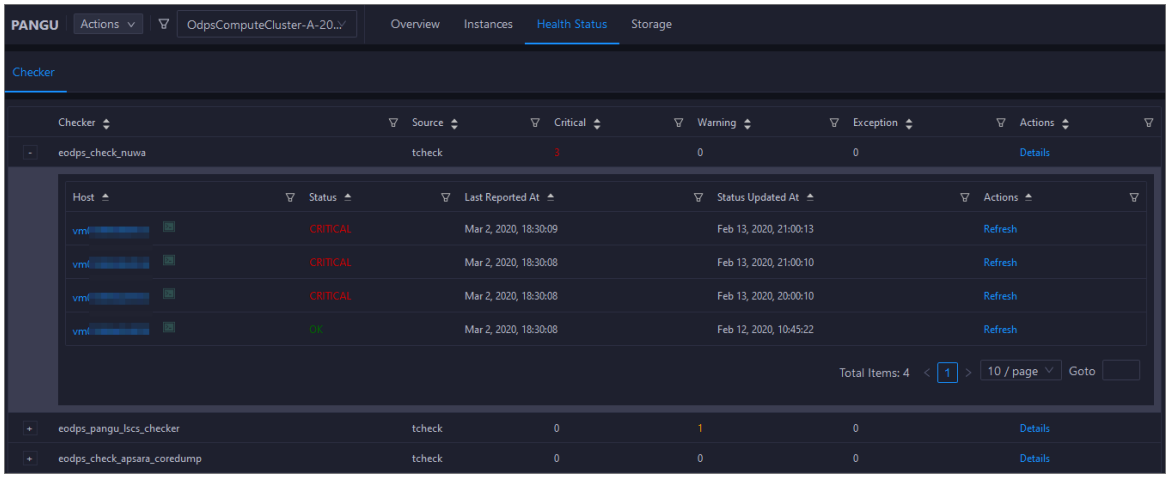

On the **Health Status** page, you can view all checkers of Apsara Distributed File System and the check results for all hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

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.

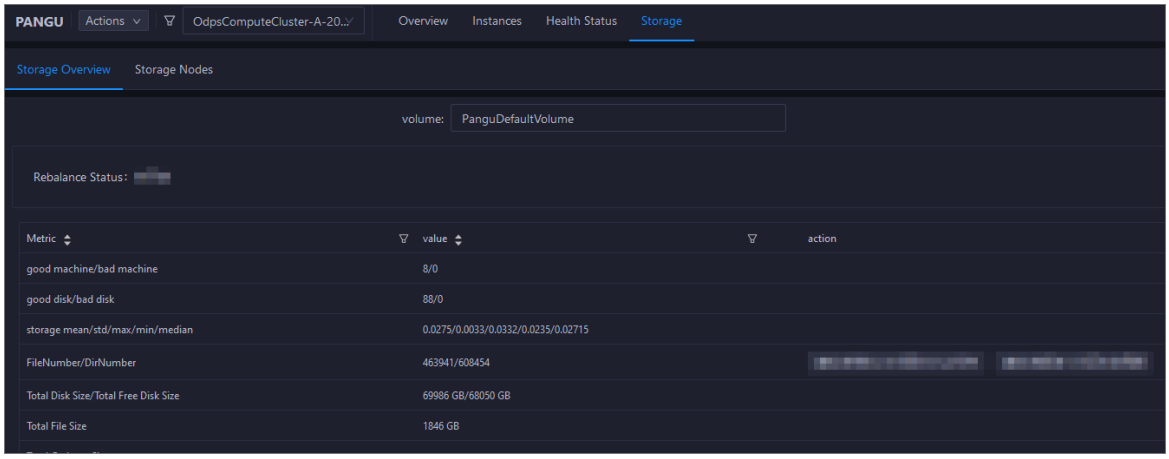

The **Storage Overview** page displays whether data rebalancing is enabled, key metrics and their values, suggestions to handle exceptions, and rack specifications of Apsara Distributed File System.

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.

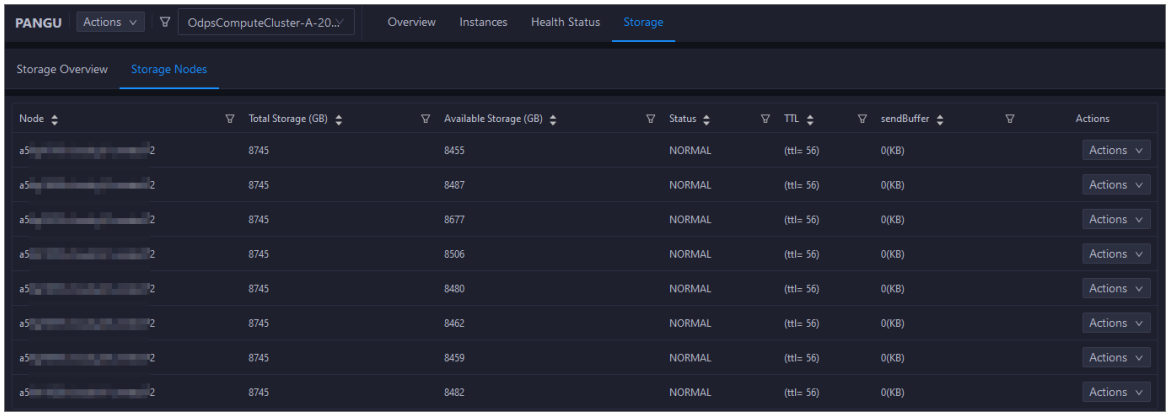

The **Storage Nodes** page displays the information about all storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, TTL, and send buffer size.

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 **Set Node Status to Shutdown** panel, click **Run**. A message appears, indicating that the action has been submitted.

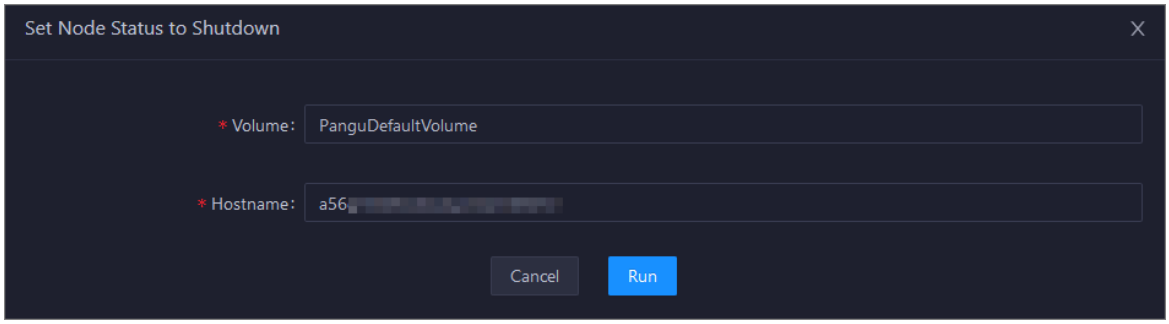

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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 **Set Disk Status to Error** panel, set the **Diskid** parameter.

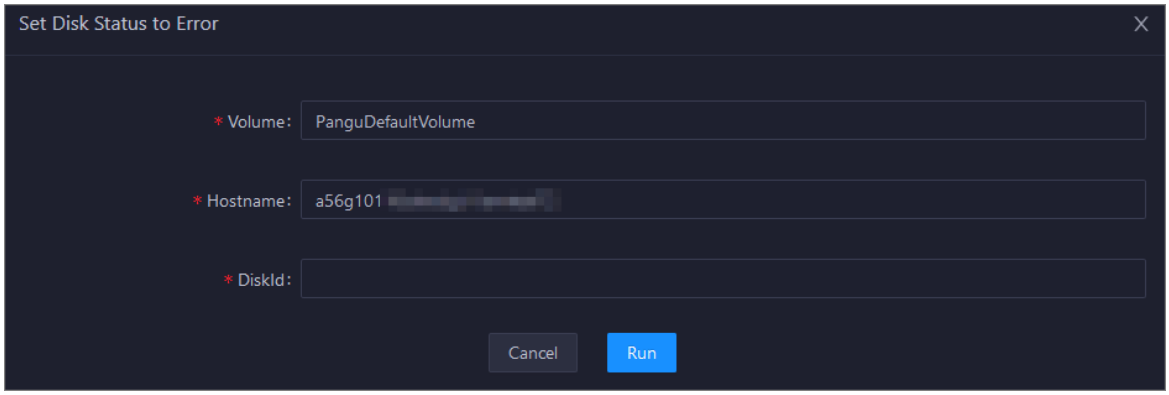

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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

Change the primary master node of Apsara Distributed File System

Apsara Big Data 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 required master node and choose **Actions** > **Change Primary Master Node** in the Actions column. In the Change Primary Master Node panel, specify the required parameters.

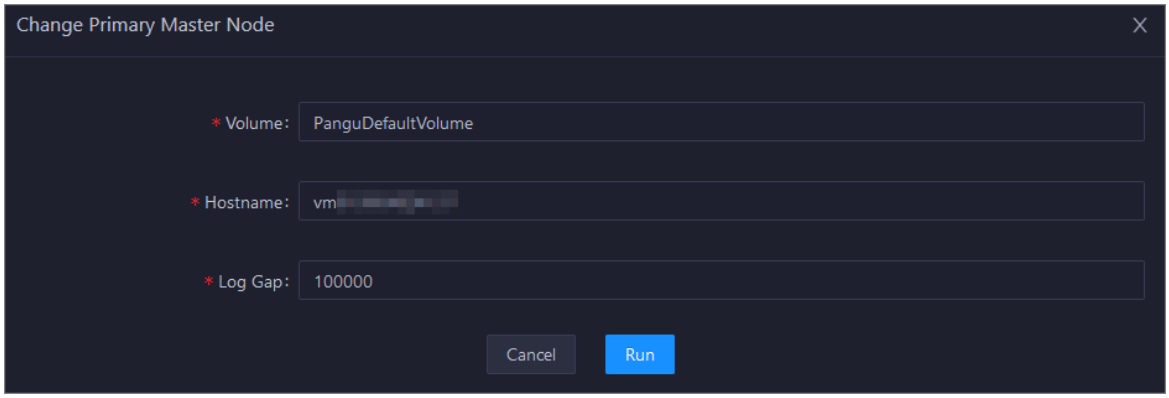

Parameter description:

- **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** panel appears.

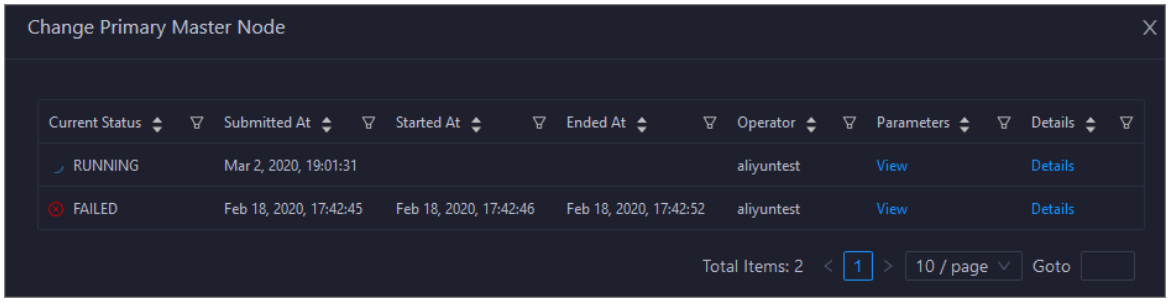

The **Change Primary Master Node** panel shows the switchover history. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

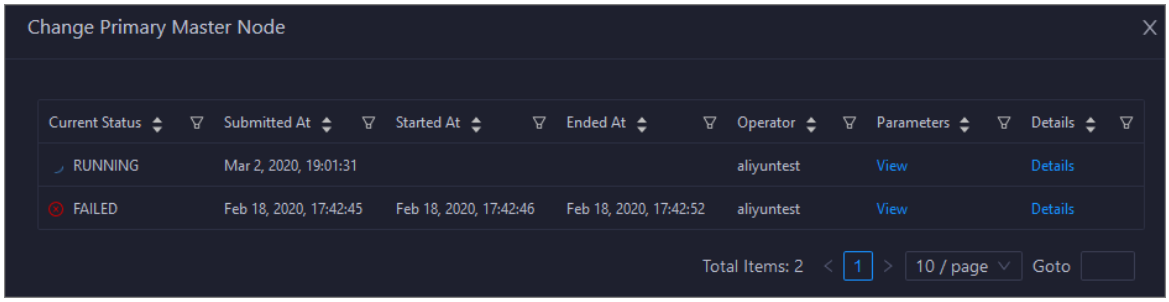

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

Clear the recycle bin of Apsara Distributed File System

Apsara Big Data 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, choose **Actions** > **Empty Recycle Bin**.
3. In the Empty Recycle Bin panel, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

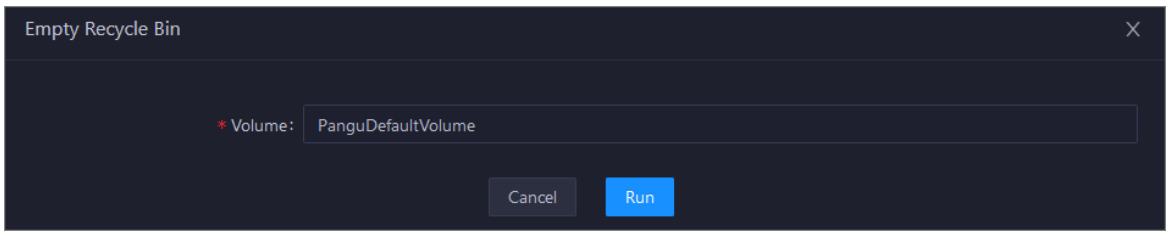

4. Click **Run**.
5. View the execution status.

   In the upper-right corner, click **Actions** and select **Execution History** next to **Empty Recycle Bin** to view the execution history.

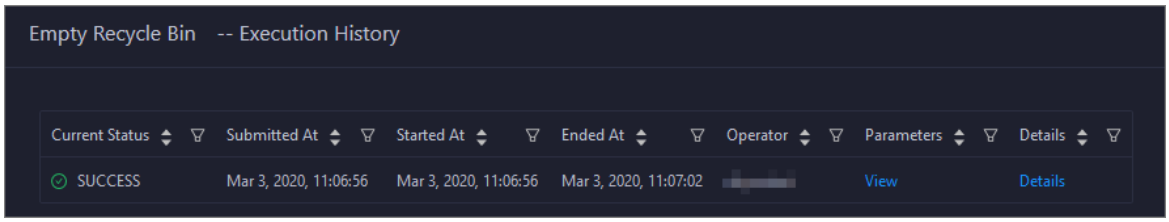

   RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

6. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

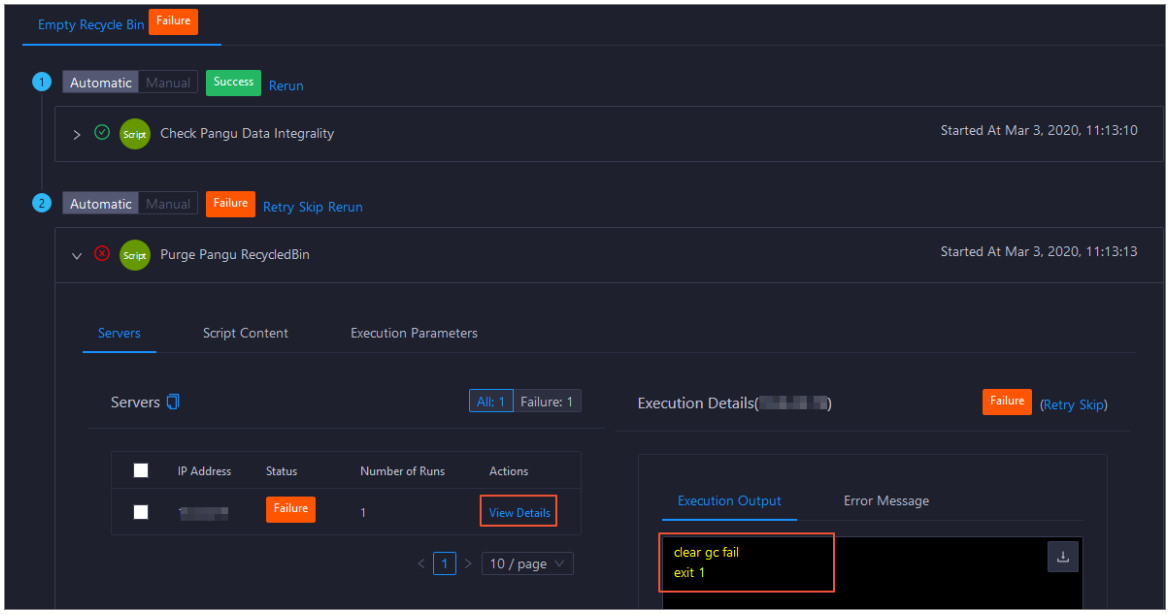

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

Enable or disable data rebalancing for Apsara Distributed File System

Apsara Big Data 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, choose **Actions** > **Disable Data Rebalancing**.
3. In the Disable Data Rebalancing panel, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

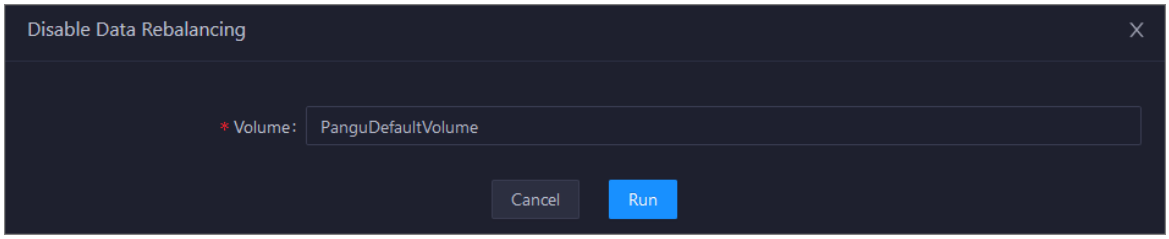

4. Click **Run**.
5. View the execution status.

   Click **Actions** and select **Execution History** next to **Disable Data Rebalancing** to view the execution history.

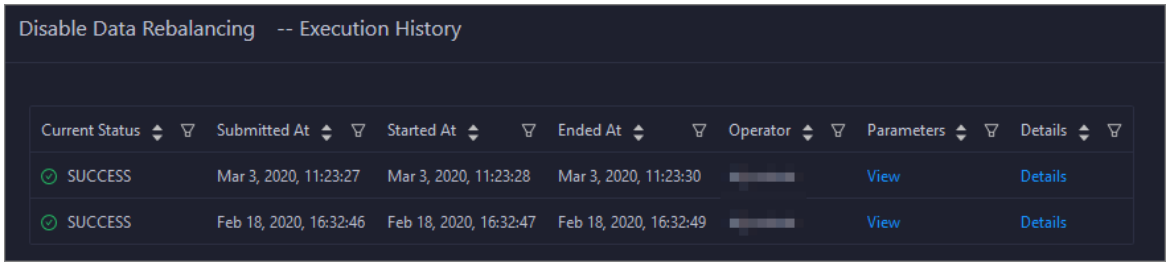

RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

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, choose **Actions** > **Enable Data Rebalancing**.
3. In the Enable Data Rebalancing panel, set the **Volume** parameter. The default value is **PanguDefaultVolume**.

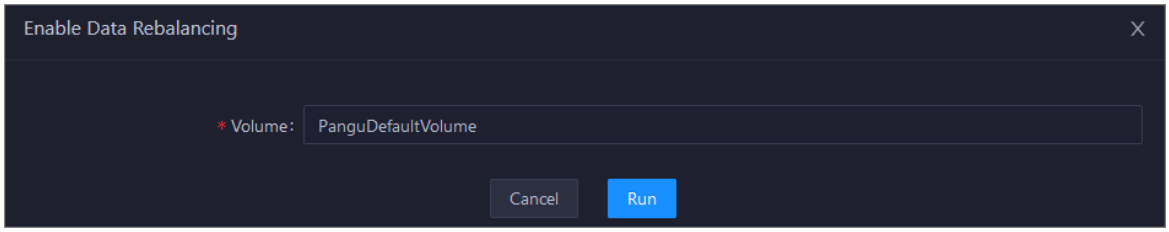

4. Click **Run**.
5. View the execution status.

   Click **Actions** and select **Execution History** next to **Enable Data Rebalancing** to view the execution history.

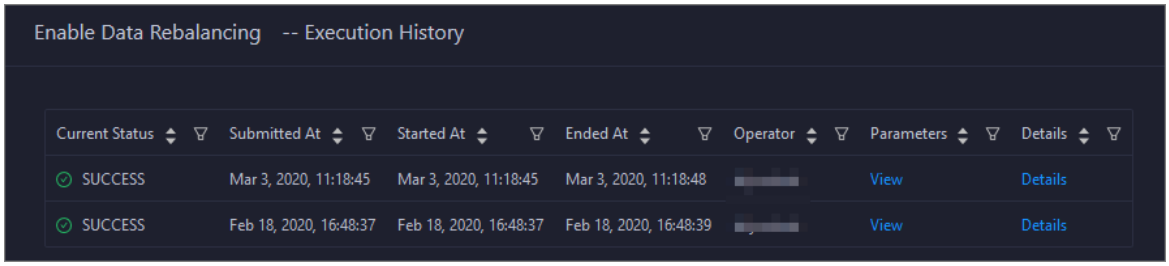

RUNNING indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

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 panel, click **Details** in the Details column for a failed execution.

2. In the Enable Data Rebalancing panel, 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.

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

Apsara Big Data 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 required master node and choose **Actions** > **Run Checkpoint on Master Node** in the Actions column. In the Run Checkpoint on Master Node panel, set the **Volume** parameter.

   > **Note** The default value of **Volume** is **PanguDefaultVolume**.

6. Click **Run**. The **Run Checkpoint on Master Node** panel appears.

   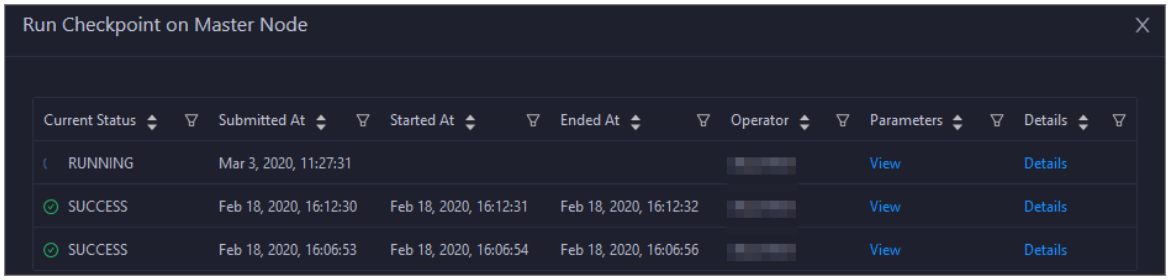

   The **Run Checkpoint on Master Node** panel shows the execution history of the checkpoint on the master node. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeds. **FAILED** indicates that the execution fails.

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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

6.2.1.5.3.4. Tunnel service

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 O&M page of the Tunnel service.

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.

O&M features of the Tunnel service

- 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.
- Traffic Analysis: shows the traffic curves of specific projects in a specific period. The curves show traffic types and the peak throughout in the specified period, which helps you make informed decisions.
- Restart Tunnel Server: allows you to restart one or more Tunnel servers.

Go to the Tunnel Service page

1. Log on to the Apsara Big Data Manager (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 **Tunnel Service**. The **Overview** tab for the Tunnel service appears.

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

The Services section shows the numbers of available services, unavailable services, and services that are being updated.

Roles

The Roles section shows all Tunnel server roles and their status. 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. This trend chart can be automatically or manually refreshed. You can view the trend chart of Tunnel throughput in a specific period.

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.

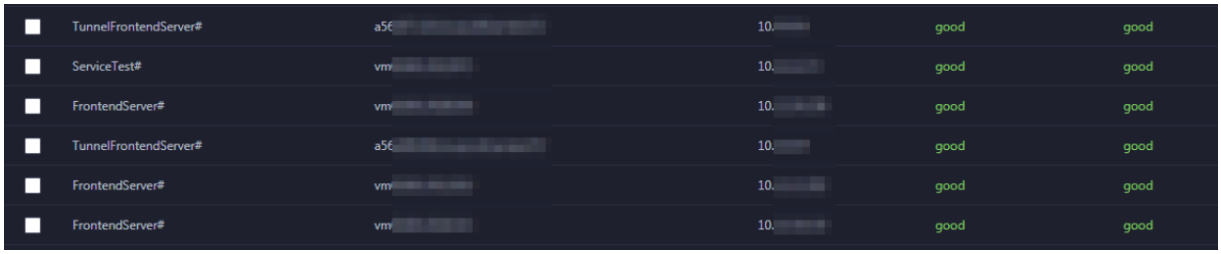

The **Instances** tab shows information about all Tunnel server roles. The information includes the role name, hostname, IP address, role status, and host status. The status can be good, error, or upgrading.

Traffic analysis

The Traffic Analysis tab displays the traffic curves of specific projects in a specific period. The curves show traffic types and the peak throughout in the specified period, which helps you make informed decisions.

Go to the Traffic Analysis tab

In the left-side navigation pane of the **Services** tab, click **Tunnel Service**. Then, click the **Traffic Analysis** tab. The **Traffic Analysis** tab for the Tunnel service appears.

After you specify a period and the project for traffic analysis, click the icon. Then, you can view the upstream and downstream throughput curves of Tunnel traffic for traffic analysis.

> **Note**
> - The traffic data comes from Monitoring System. Make sure that this system is normal.
> - By default, the top five projects that have the most traffic are selected. You can also filter projects based on your business requirements.
> - By default, the beginning of the period is two days before the current time, and the end of the period is one day before the current time. You can also specify the period based on your business requirements.

Restart Tunnel servers

Apsara Big Data 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, choose **Actions > Restart Tunnel Server**.
6. In the **Restart Tunnel Server** panel, configure the required parameters.

   The following table describes the required 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 Execution History panel, 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.

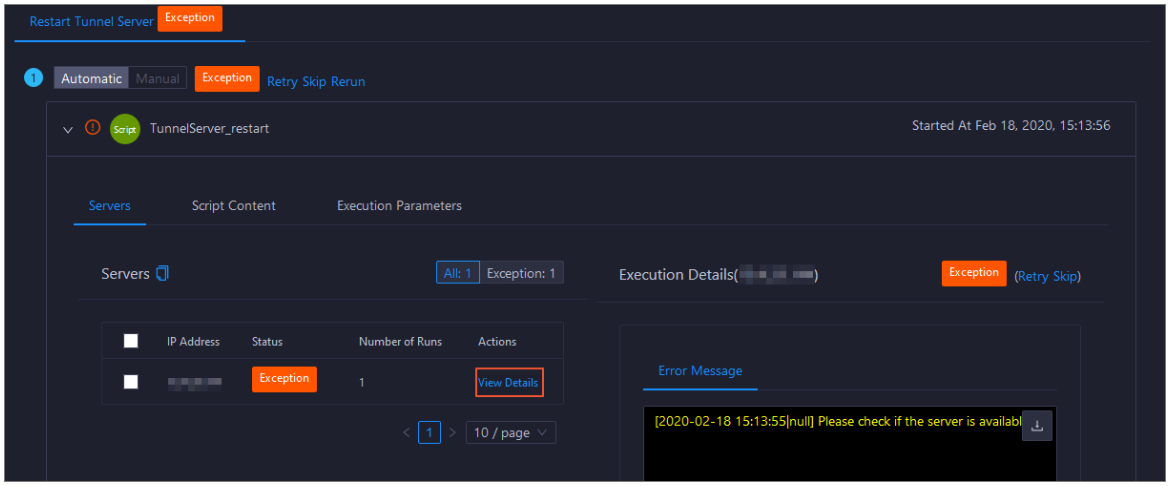

6.2.1.5.4. Cluster O&M

6.2.1.5.4.1. O&M features and entry

This topic describes the O&M features of MaxCompute clusters. It also provides more information about how to go to the MaxCompute cluster O&M page.

Cluster O&M features

O&M features of MaxCompute clusters:

- 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 out Cluster or Scale in Cluster: allows you to add or remove physical hosts to scale out or scale in a MaxCompute cluster.
- 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 a MaxCompute cluster at a time.

Go to the Clusters tab

1. Log on to the Apsara Big Data 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, click a cluster. The **Overview** tab for the selected cluster appears.

6.2.1.5.4.2. 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 Big Data 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.

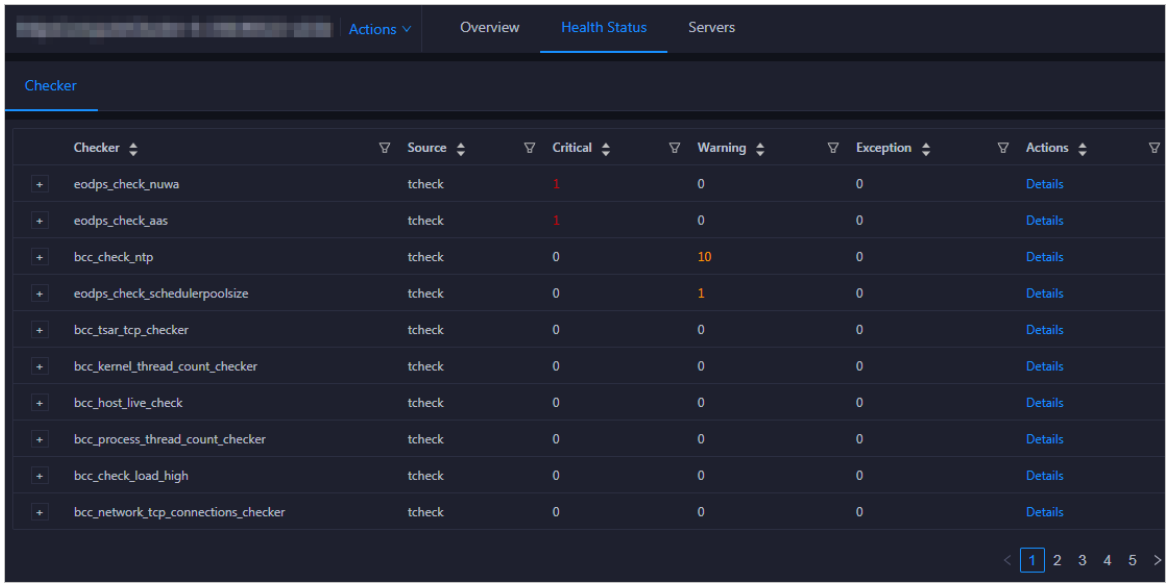

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are represented in different colors. You must handle the alerts in a timely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

View checker details

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

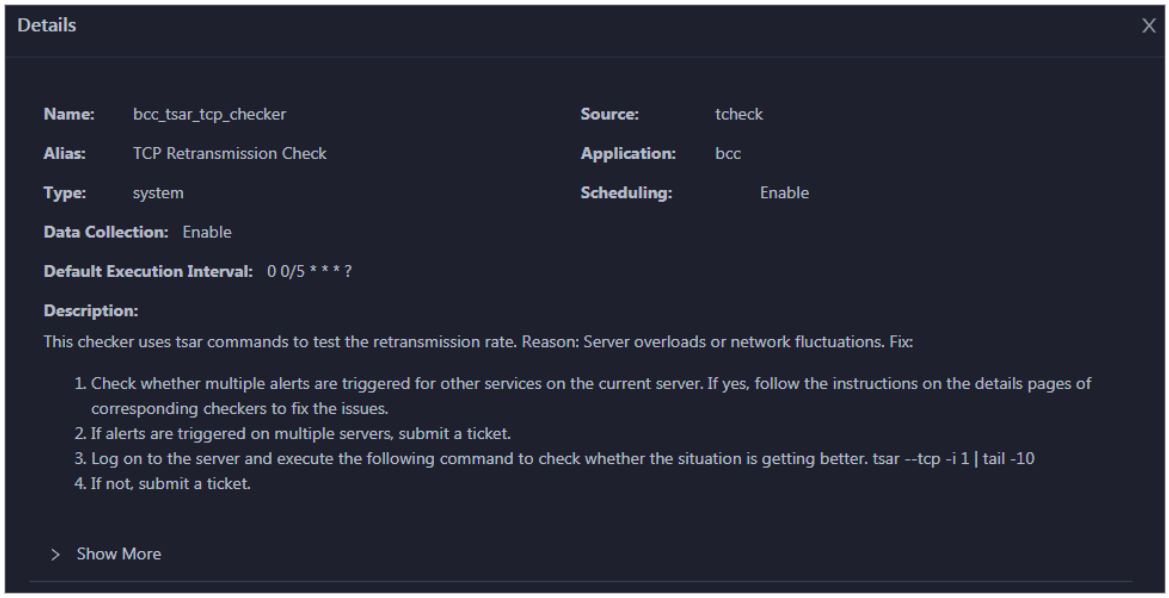

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

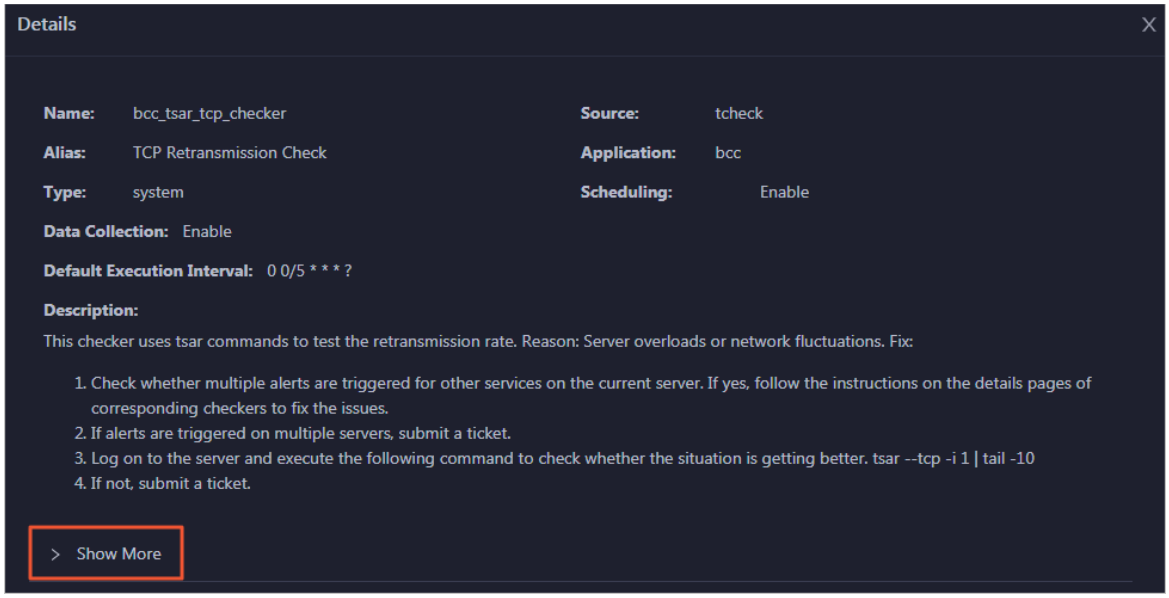

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

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

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

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

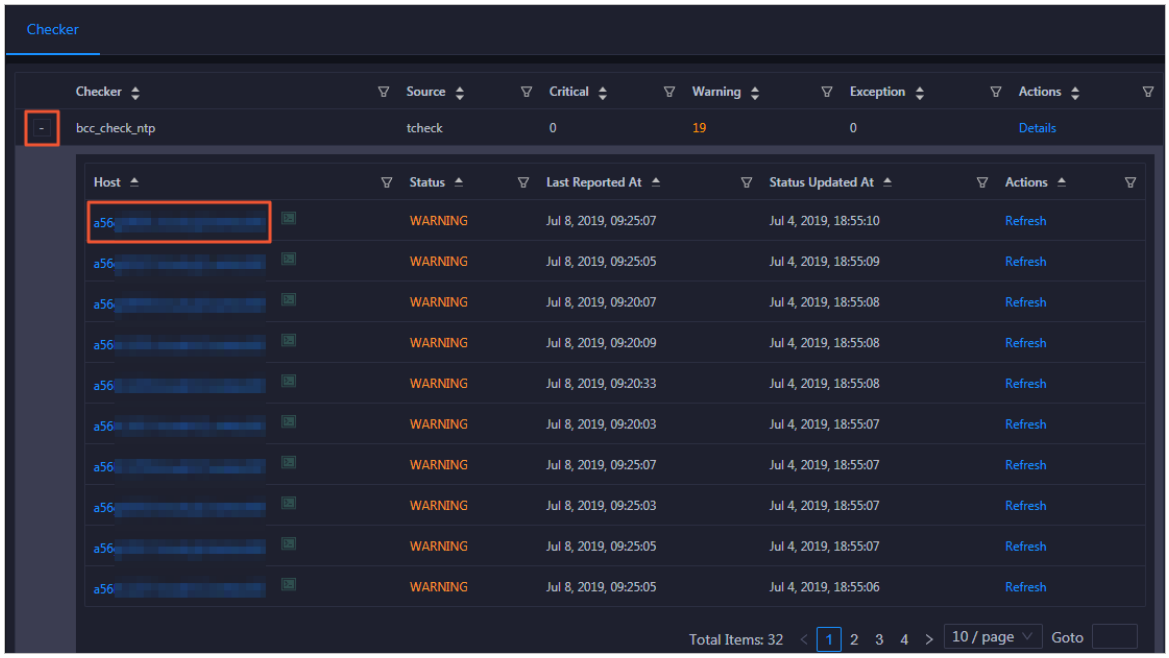

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

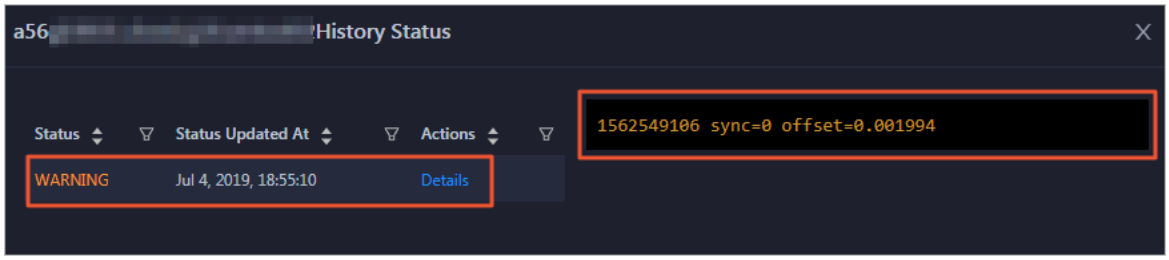

Clear alerts

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

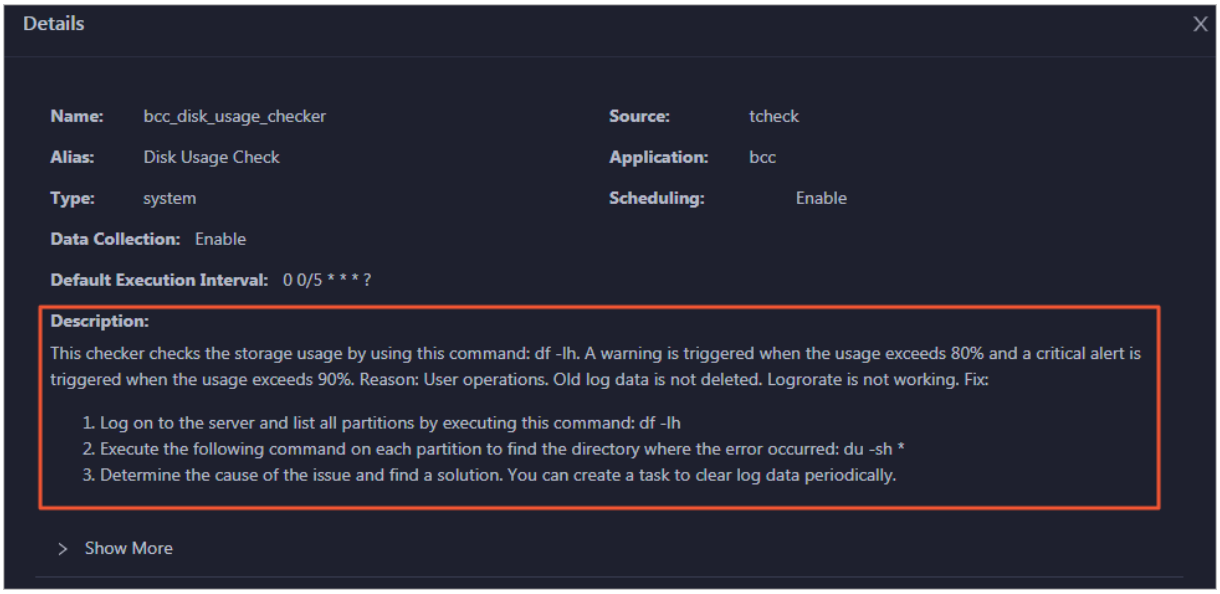

Log on to a host

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

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

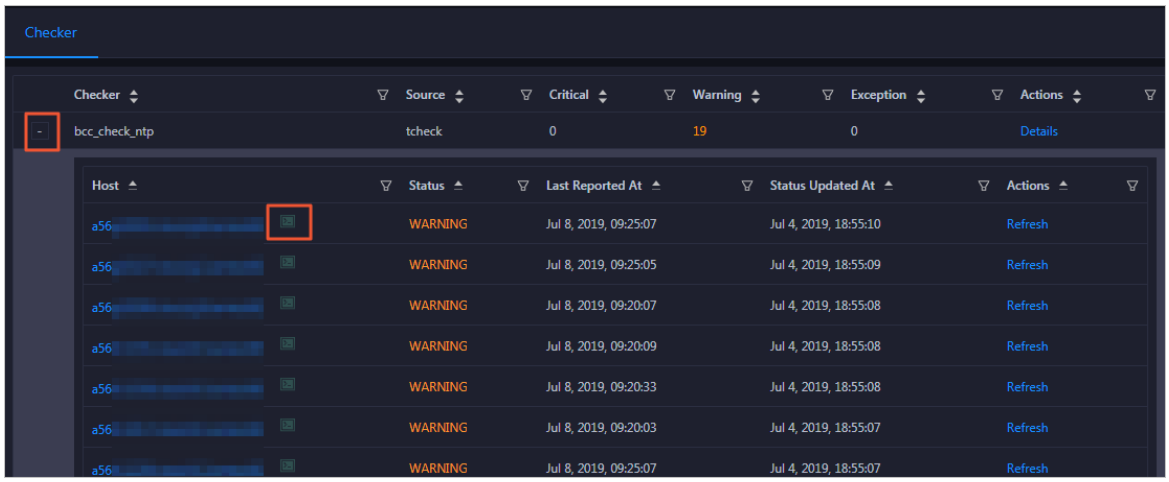

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

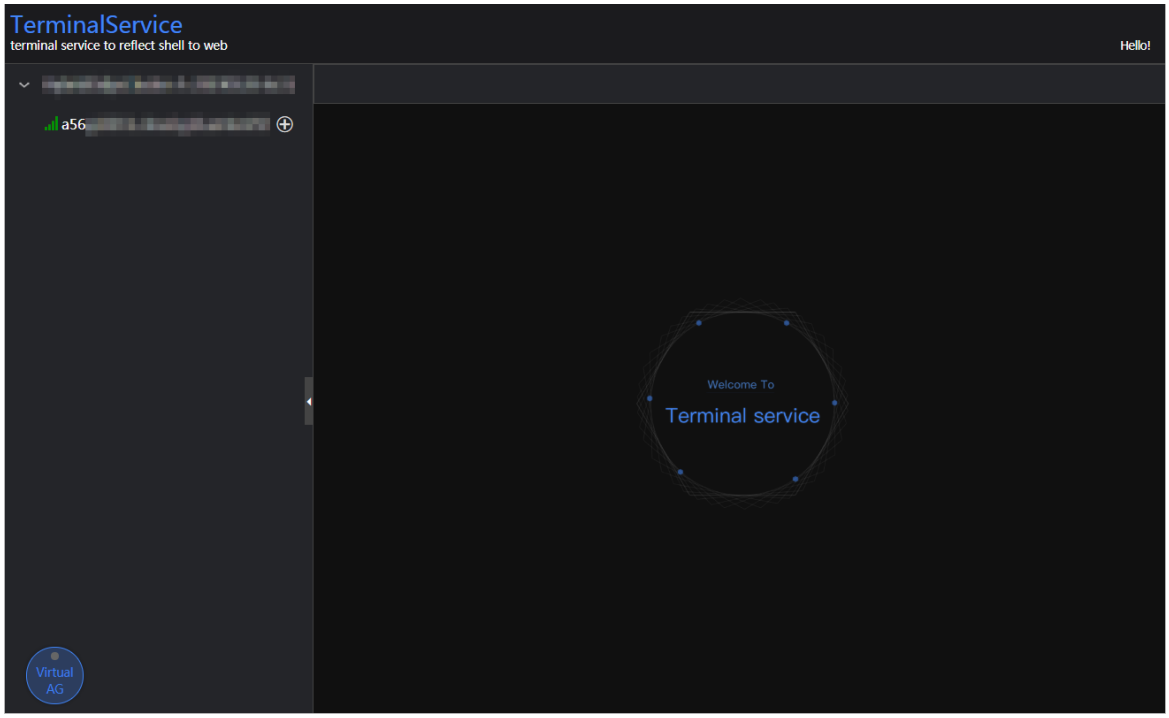

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

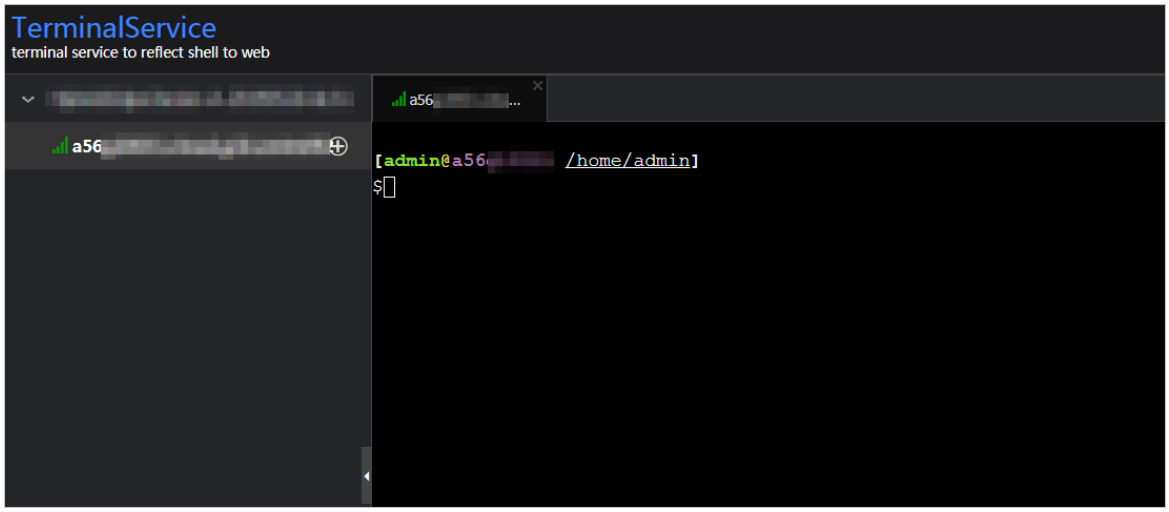

Run a checker again

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

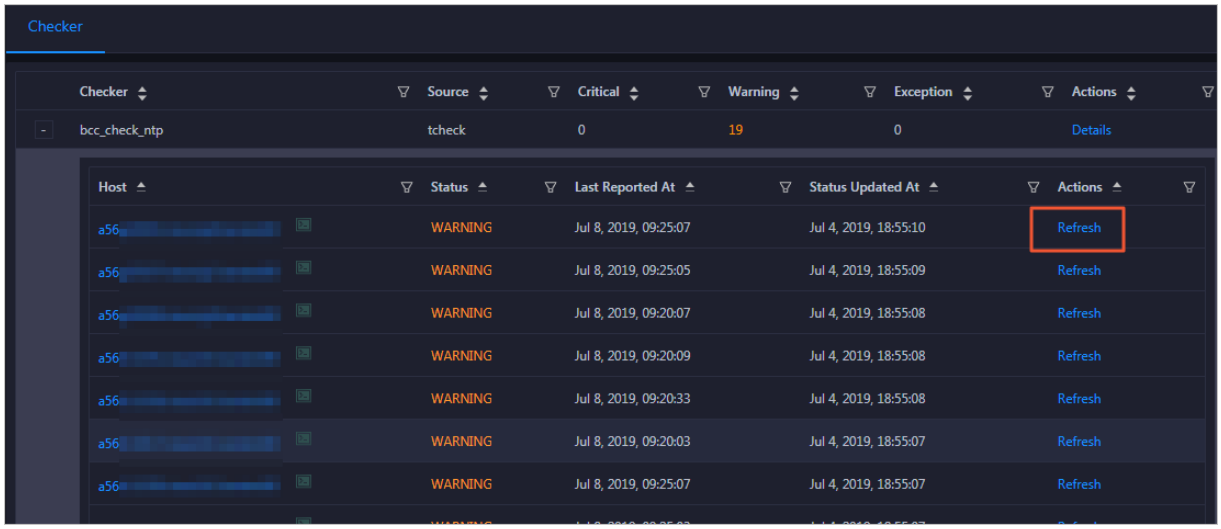

6.2.1.5.4.3. Overview

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 Big Data 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 upper-left corner, click the **Login in** icon of the host. The **TerminalService** page appears.

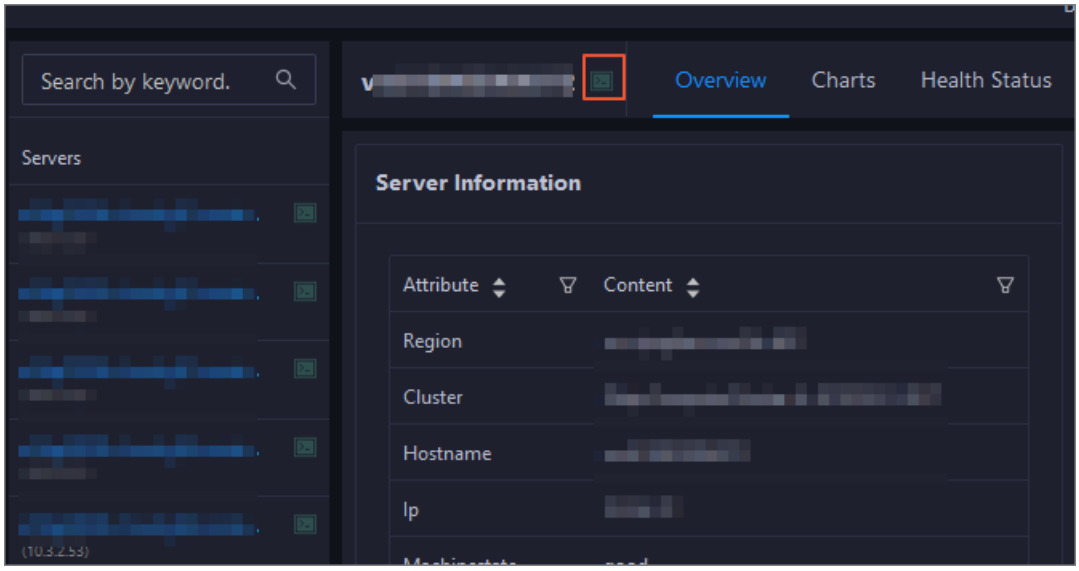

3. In the left-side navigation pane, click the hostname to log on to the host.

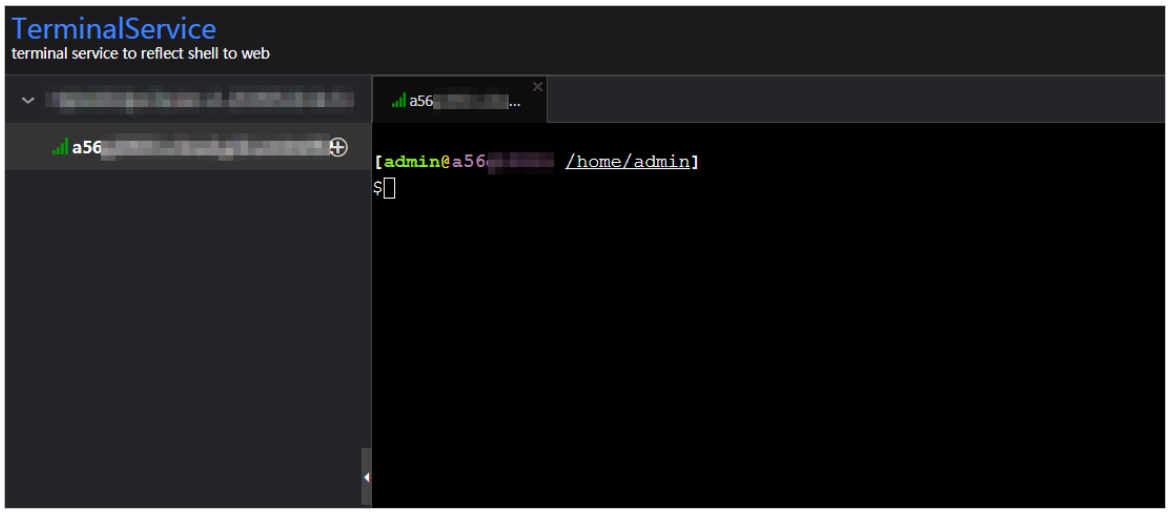

Servers

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 displays all services deployed in the cluster and the respective number 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.

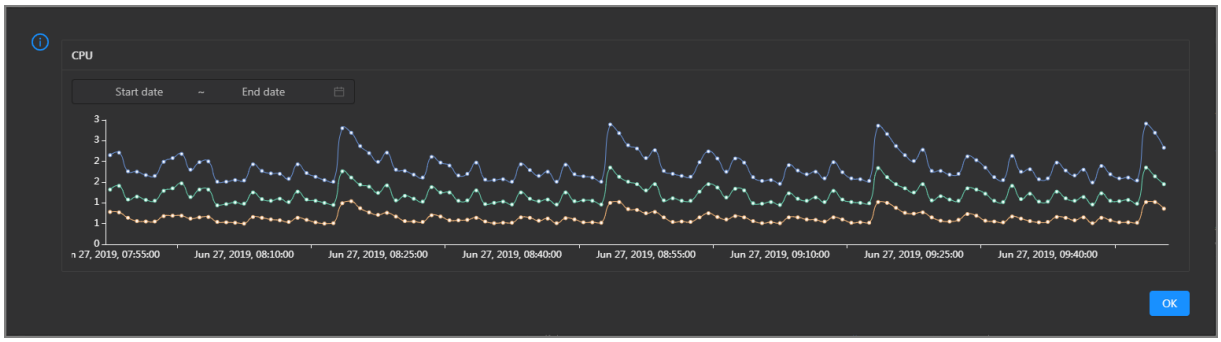

DISK

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

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

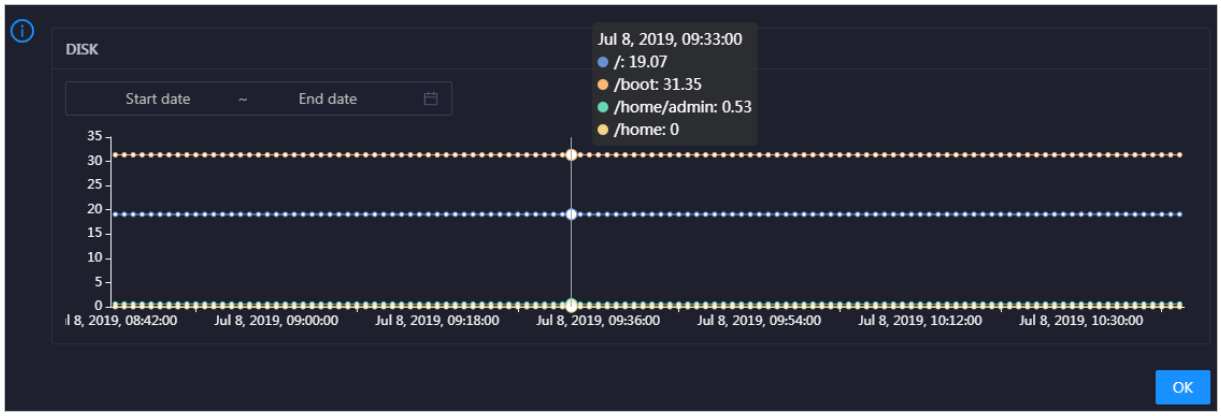

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

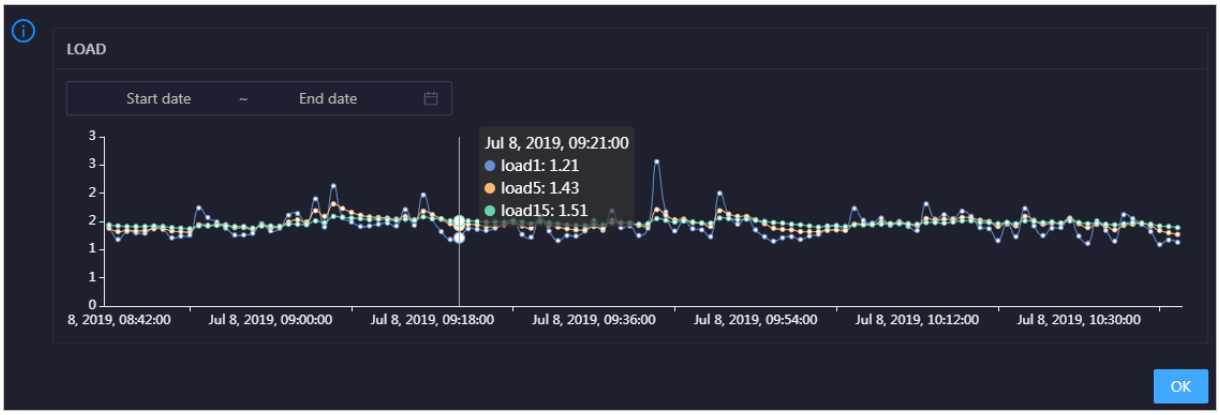

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

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.

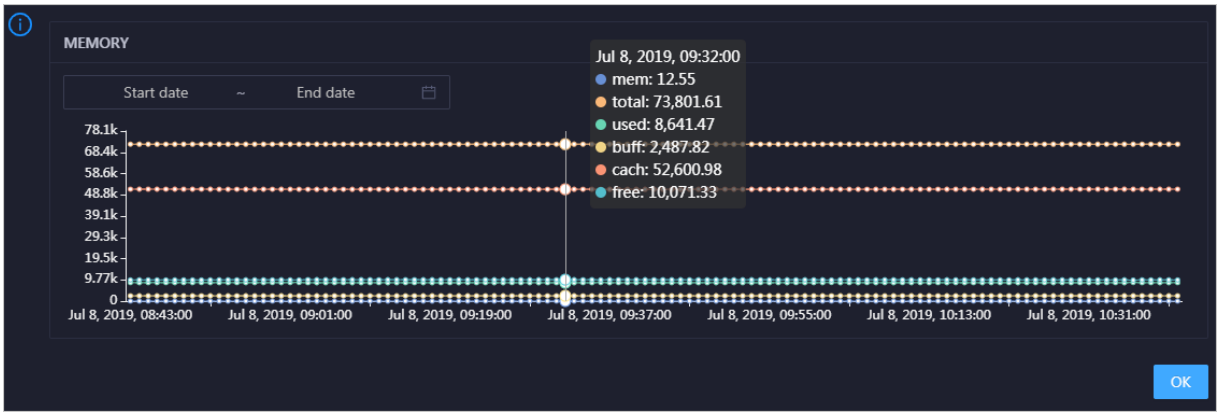

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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.

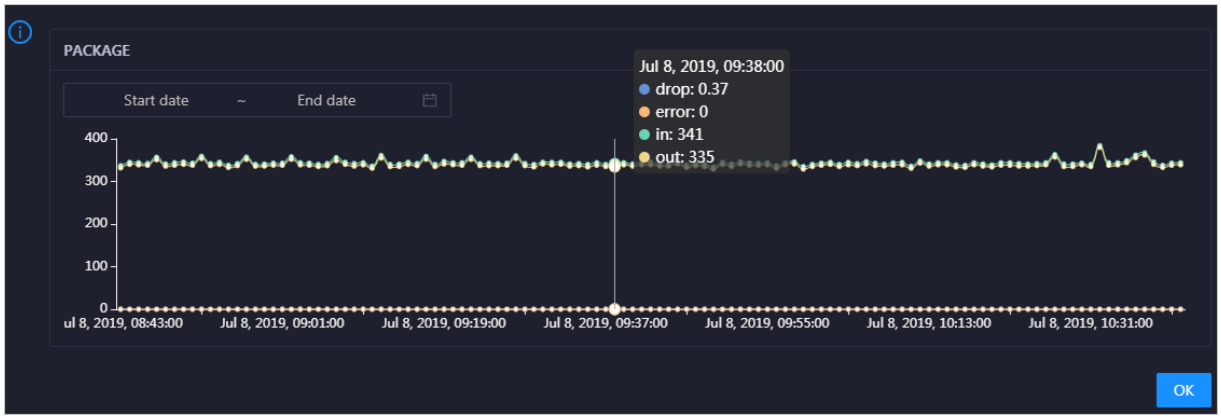

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

Health Check

This section shows the number of checkers for the cluster and the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

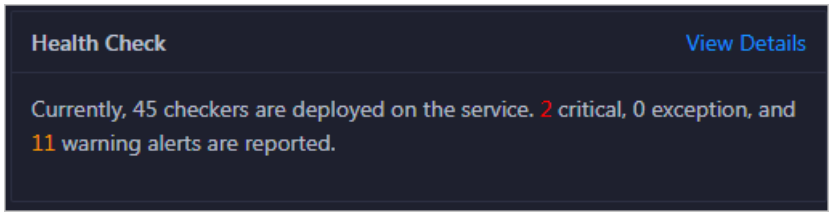

Click **View Details** to go to the Health Status tab. On this tab, you can view health check details. For more information, see Cluster health.

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.

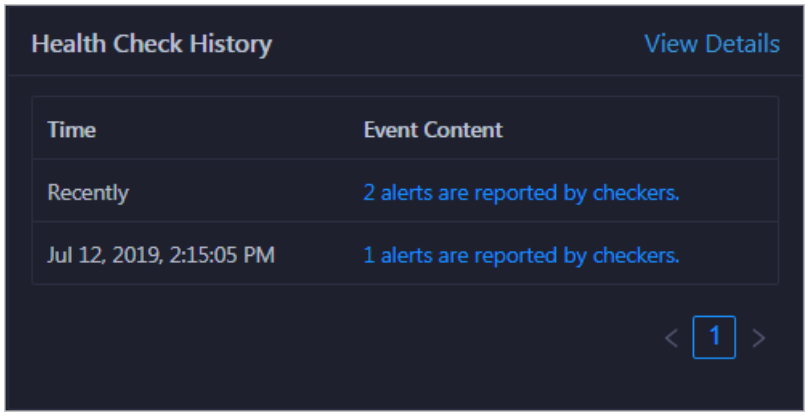

Click **View Details** to go to the Health Status tab. On this tab, you can view health check details. For more information, see Cluster health.

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

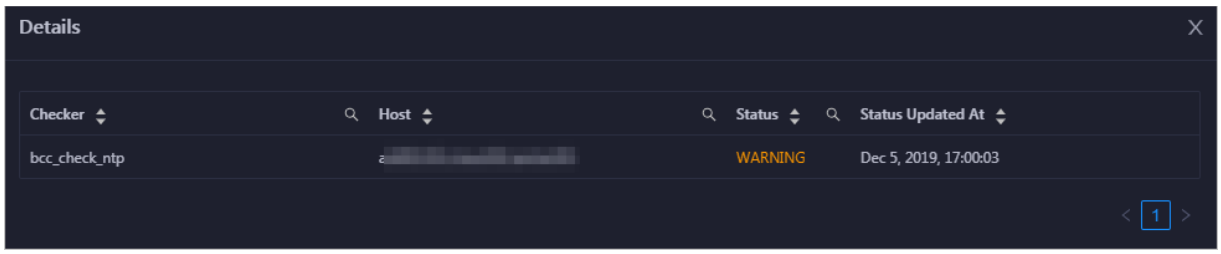

6.2.1.5.4.4. Servers

The Servers tab shows information about hosts. The information includes the hostname, IP address, role, type, CPU utilization, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

In the left-side navigation pane of the **Clusters** tab, click a cluster. Then, click the **Servers** tab. The **Servers** tab for the selected cluster appears.

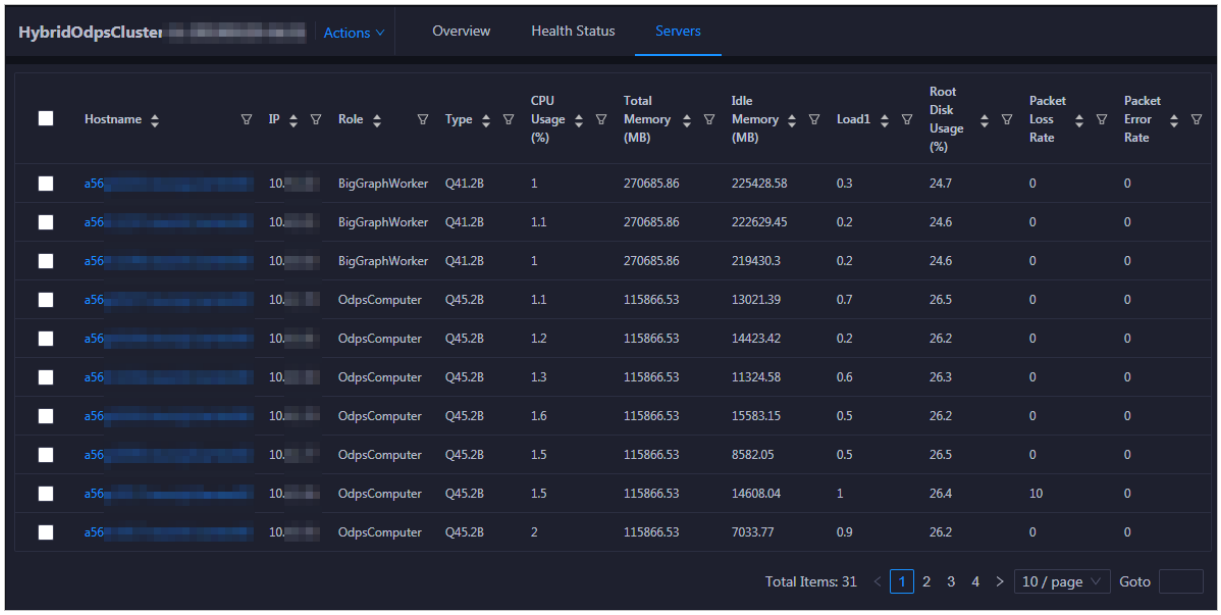

To view more information about a host, click the name of the host. The Hosts tab appears.

6.2.1.5.4.5. Scale in and scale out a MaxCompute cluster

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

Description

In Apsara Stack, scaling out a cluster involves complex operations. You must configure a new physical host on Deployment Planner and Apsara Infrastructure Management Framework 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 out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. If you want to scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

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 `rttrl` command to query and record SInstance hosts. For more information about how to log on to a host, see Log on to a host.

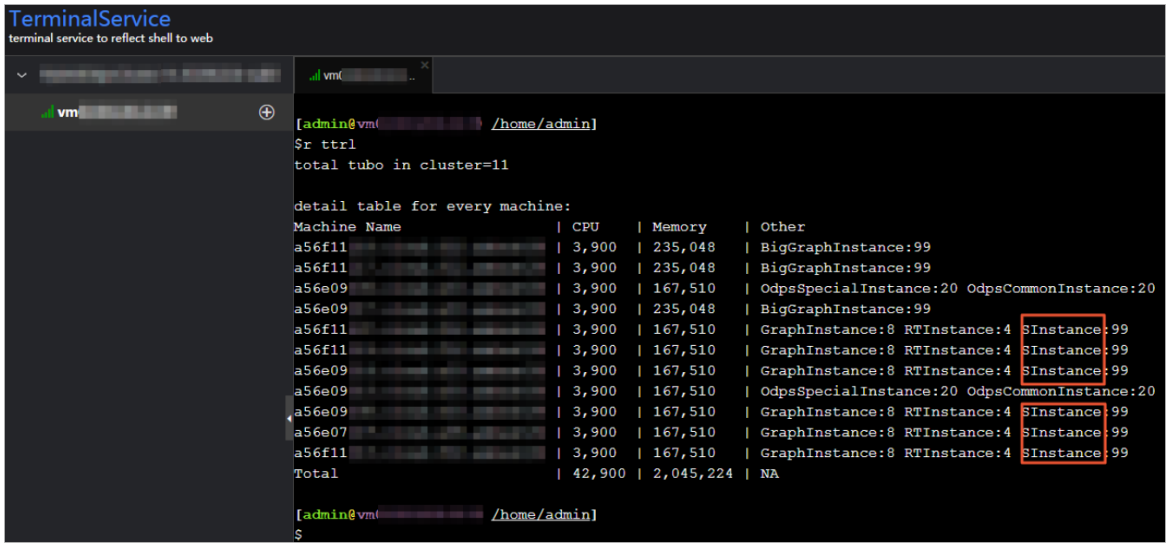

2. In the left-side navigation pane of the **Clusters** tab, click a cluster. Then, click the **Servers** tab. On the tab that appears, select an SInstance host and use it as the template host.

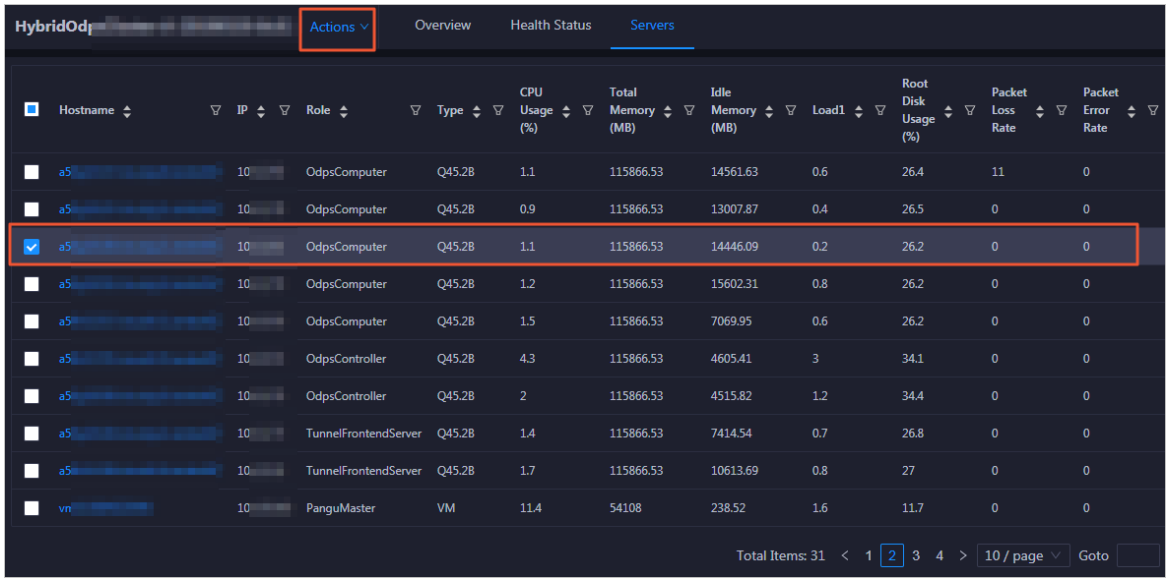

3. In the upper-right corner, choose **Actions > Scale out Cluster**. In the **Scale out Cluster** panel, configure the parameters.

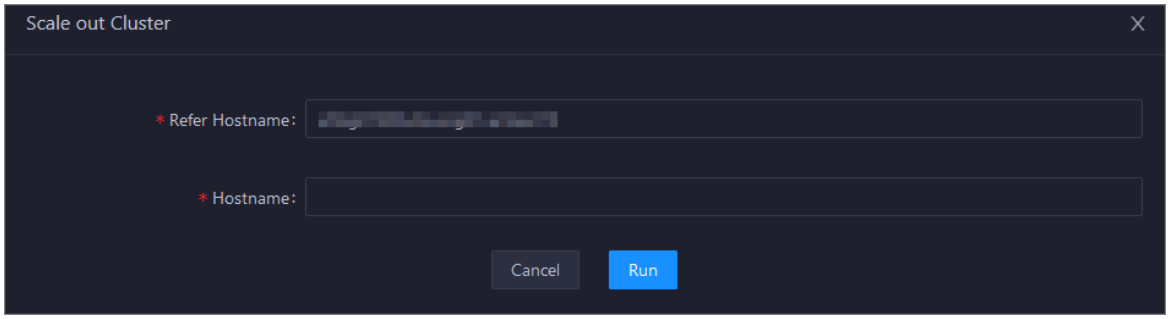

Parameters:

- 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**. A message appears, indicating that the request has been submitted.
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.

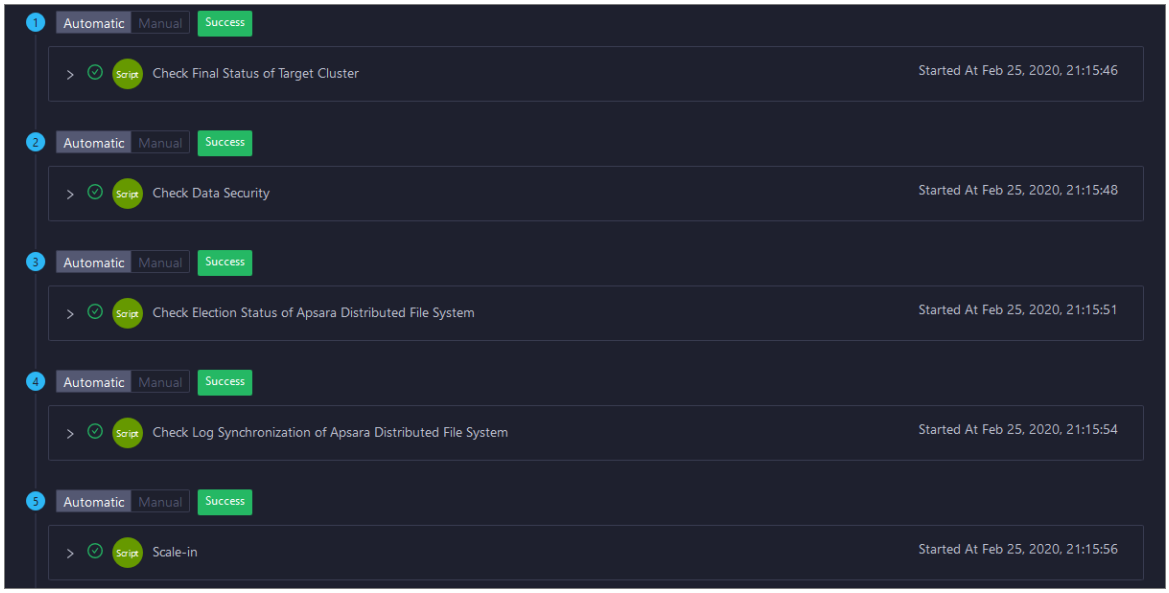

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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 `rttrl` command to query and record SInstance hosts. For more information about how to log on to a host, see Log on to a host.

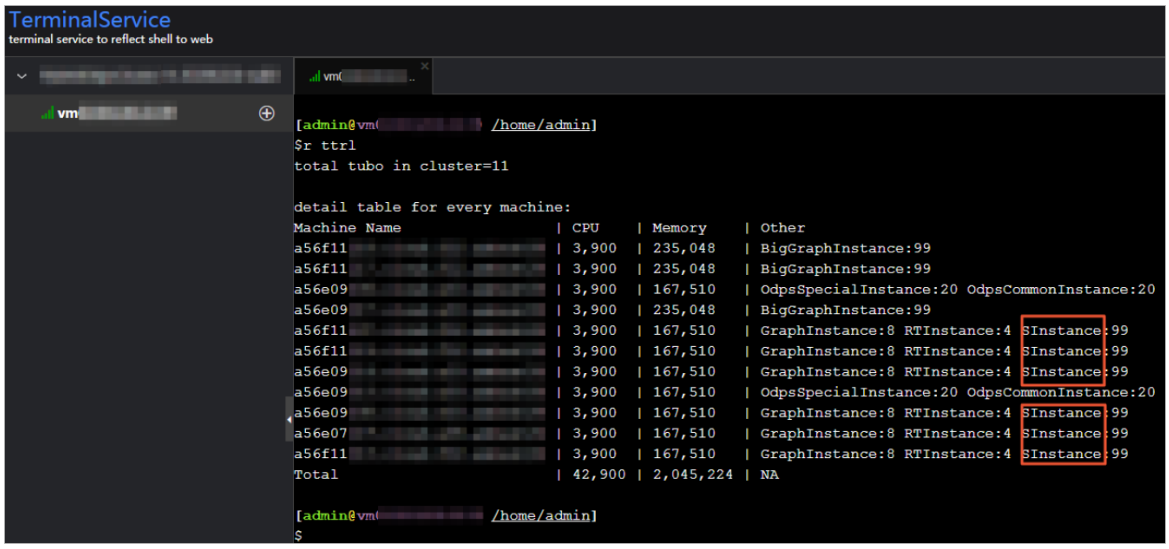

2. In the left-side navigation pane of the **Clusters** tab, click a cluster. Then, click the **Servers** tab. On the tab that appears, select one or more SInstance hosts that you want to remove.

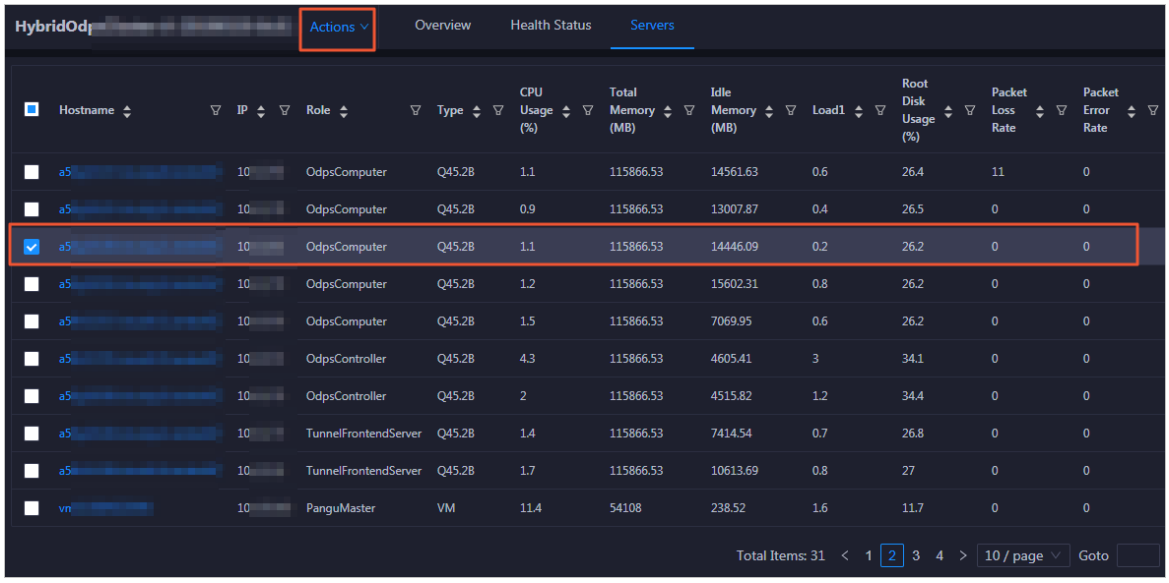

3. In the upper-right corner, choose **Actions > Scale in Cluster**. In the **Scale in Cluster** panel, configure the parameters.

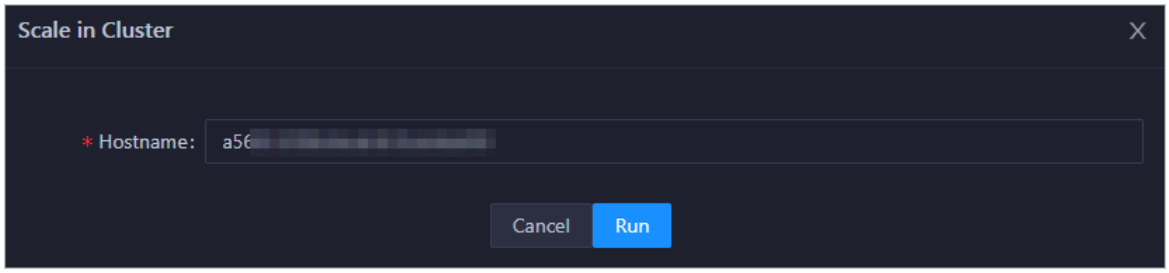

Parameters:

- 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**. A message appears, indicating that the request has been submitted.
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.

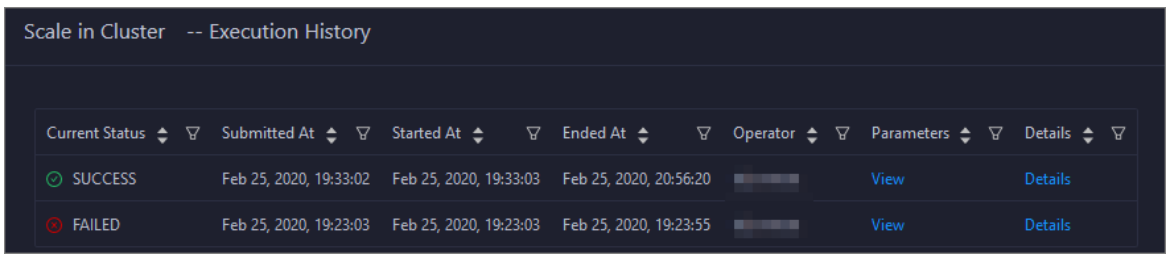

7. If the status is FAILED, click **Details** in the Details column to identify the cause of the failure.

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.

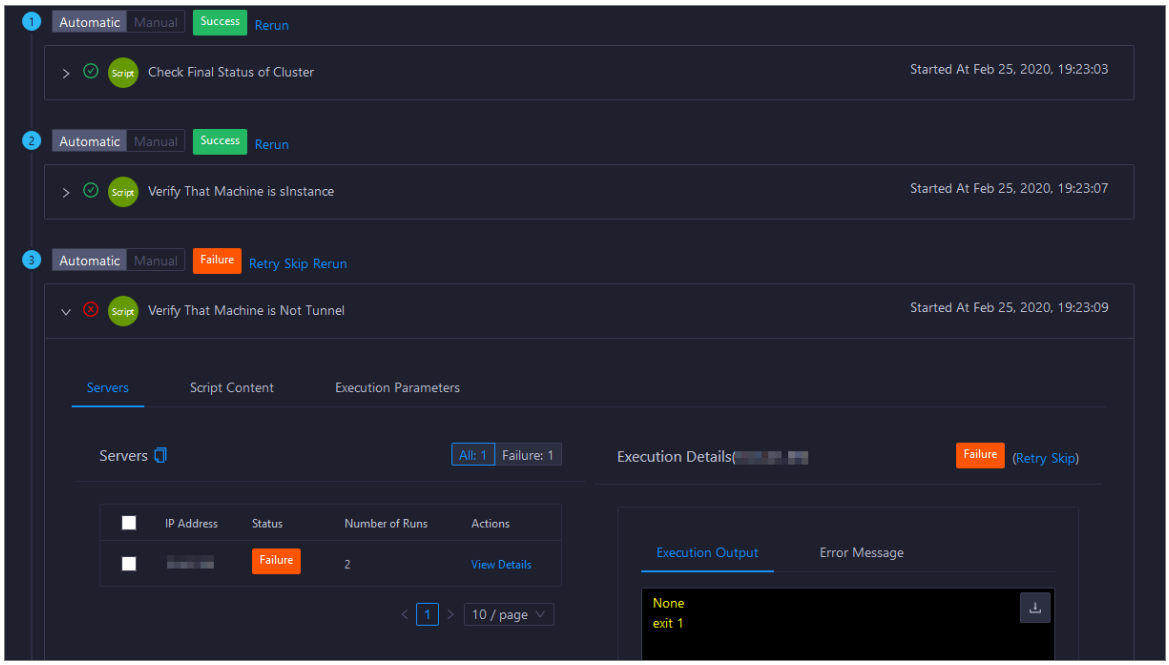

You can view information about parameter settings, host details, scripts, and runtime parameters to identify the cause of the failure.

6.2.1.5.4.6. Restore environment settings and enable auto repair

Apsara Big Data 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, choose **Actions > Restore Environment Settings**. In the **Restore Environment Settings** panel, set the Hosts parameter.

   > **Note** You can enter the names of multiple hosts and must separate the names with commas (,).

2. Click **Run**. A message appears, indicating that the request has been submitted.
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, choose **Actions > Enable Auto Repair**. In the **Enable Auto Repair** panel, set the Cluster parameter and select Enable for Auto Repair.

   Parameters:

   - Cluster: the name of the cluster for which you want to enable the auto repair feature.
   - Auto Repair: If you require the feature, select **Enable**. Otherwise, select **Disable**.
2. Click **Run**. A message appears, indicating that the request has been submitted.
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.

6.2.1.5.5. Host O&M

6.2.1.5.5.1. 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 Big Data 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.

6.2.1.5.5.2. Host overview

The Overview tab for a host shows brief information about the host in a MaxCompute cluster. On this tab, you can view server information, service role status, 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, and packet transmission for the host.

Go to the Overview tab

In the left-side navigation pane of the **Hosts** tab, click a host. Then, click the **Overview** tab. The **Overview** tab for the host appears.

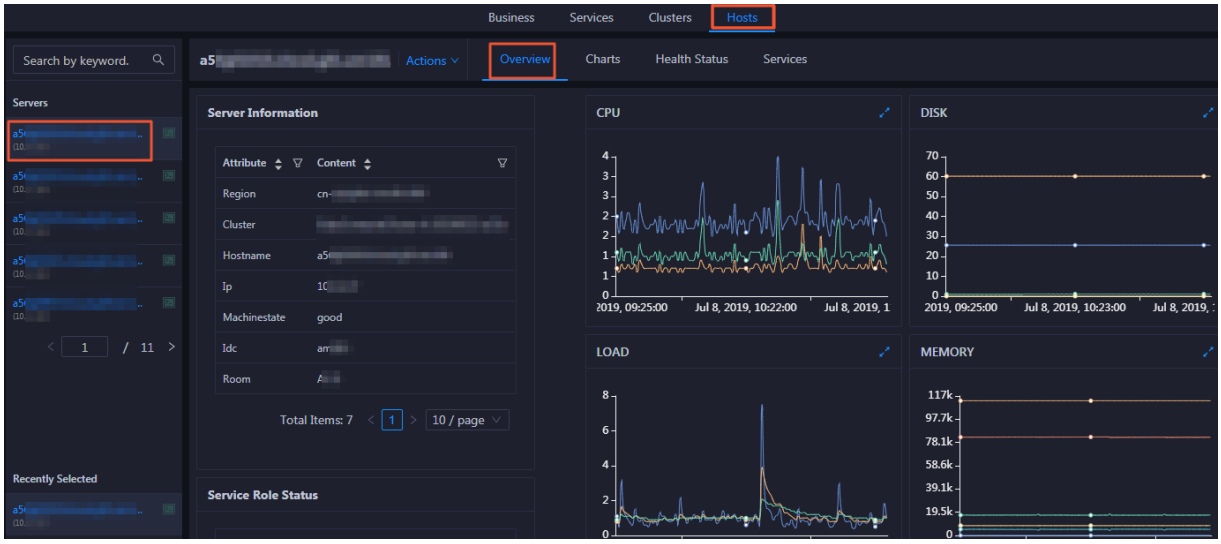

On the **Overview** tab, you can view server information, service role status, health check result, and health check history of the host. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, and packet transmission for the host.

Server Information

The Server Information section shows information about the host. Server information includes the region, cluster, name, IP address, data center, and server room.

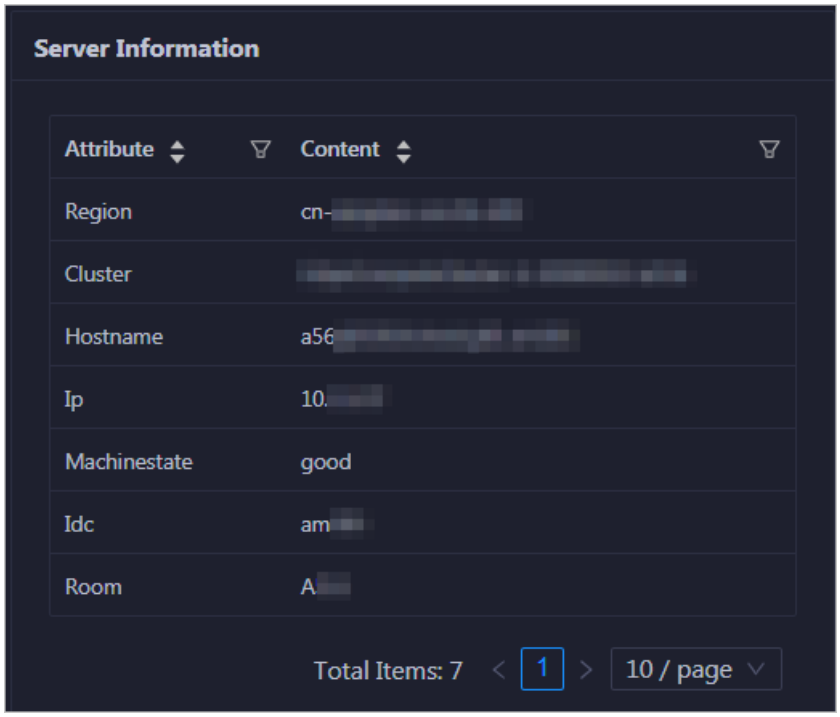

Service Role Status

The Service Role Status section shows information about the services deployed on the host, including the roles, status, and number of services.

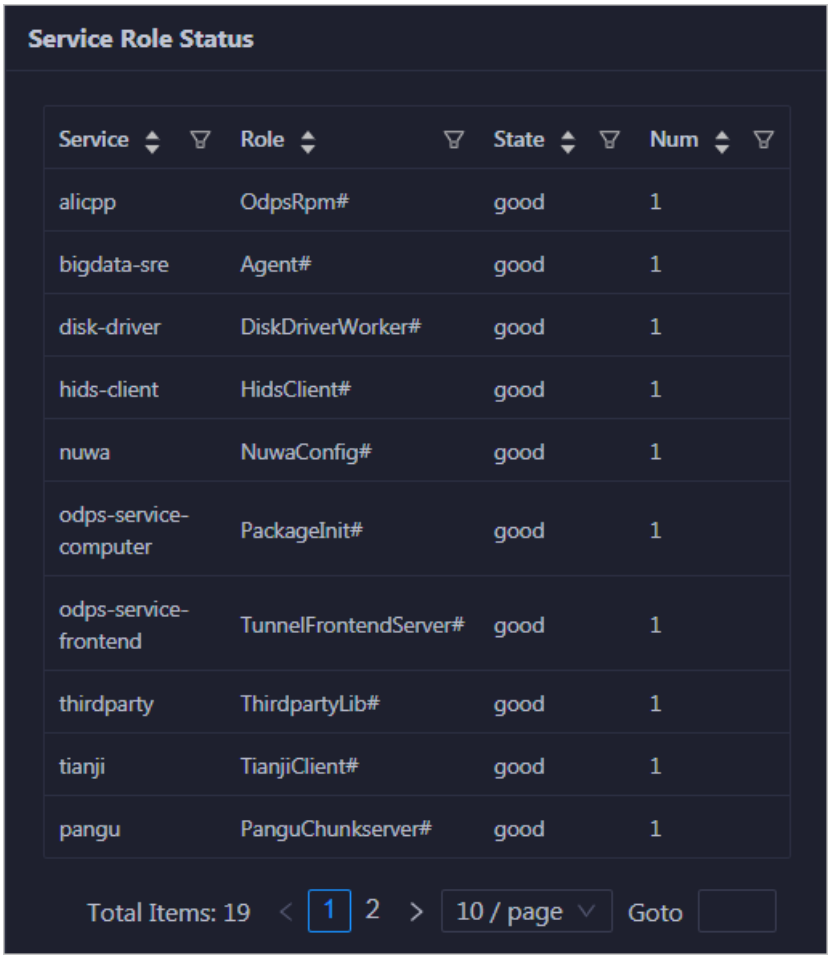

CPU

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

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

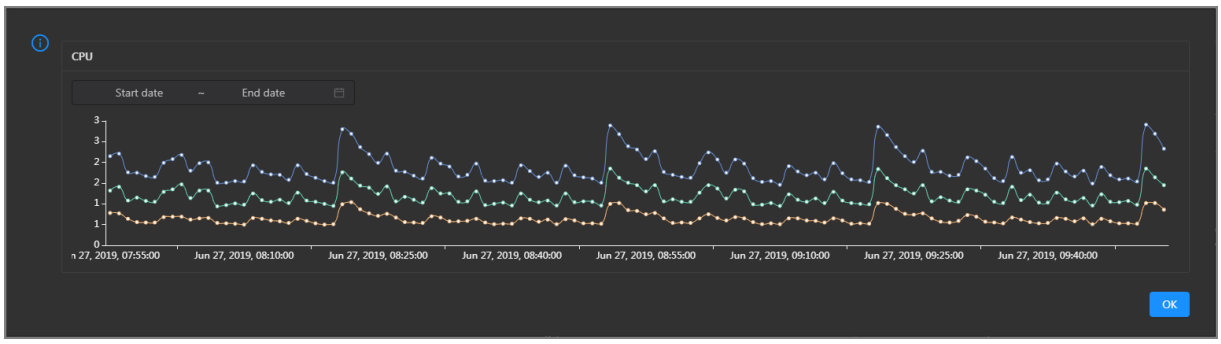

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

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

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

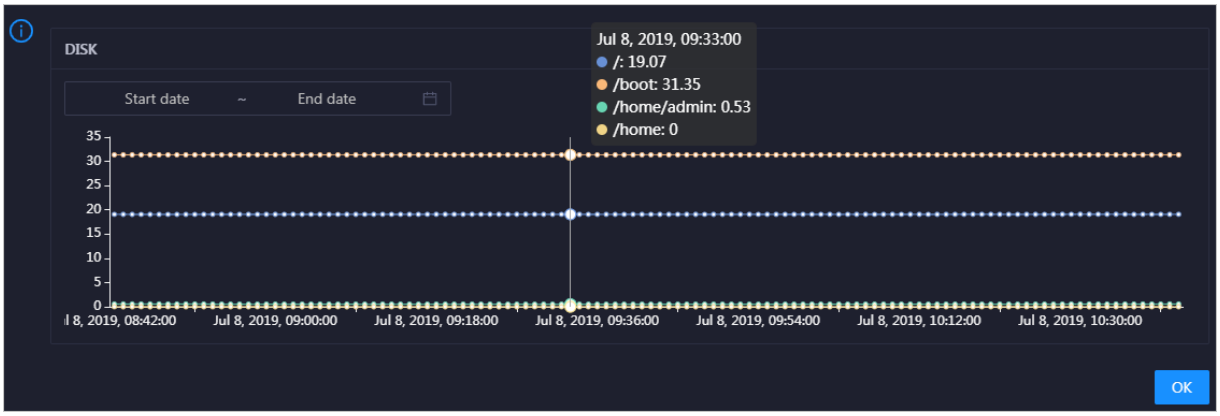

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

LOAD

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

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

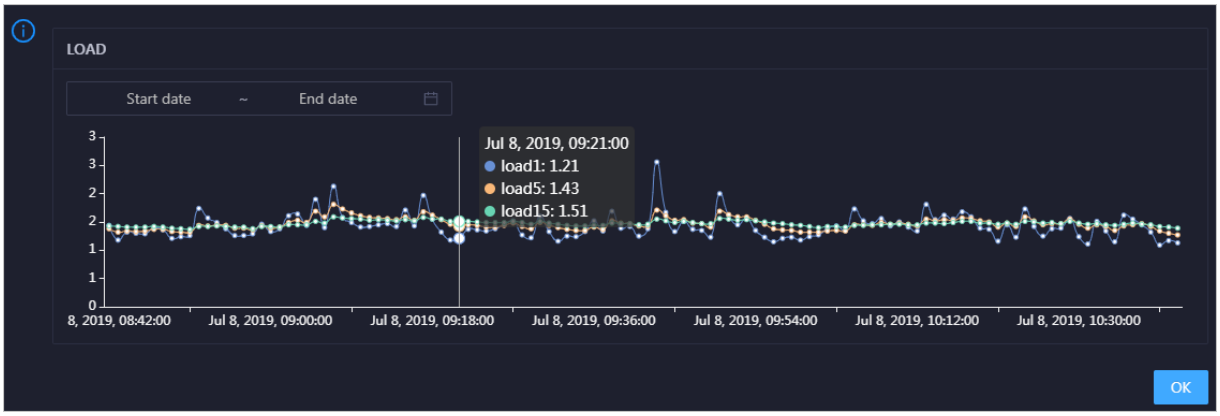

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

MEMORY

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

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

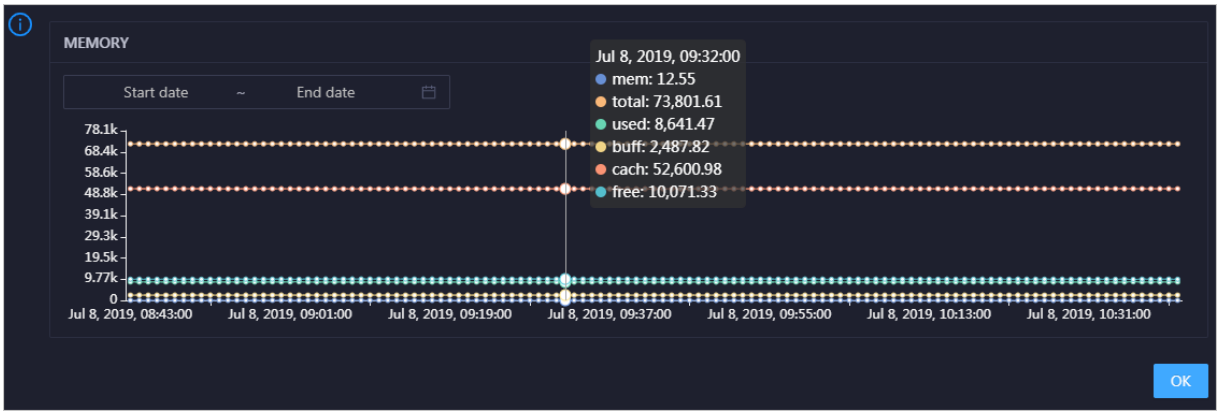

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

PACKAGE

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

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

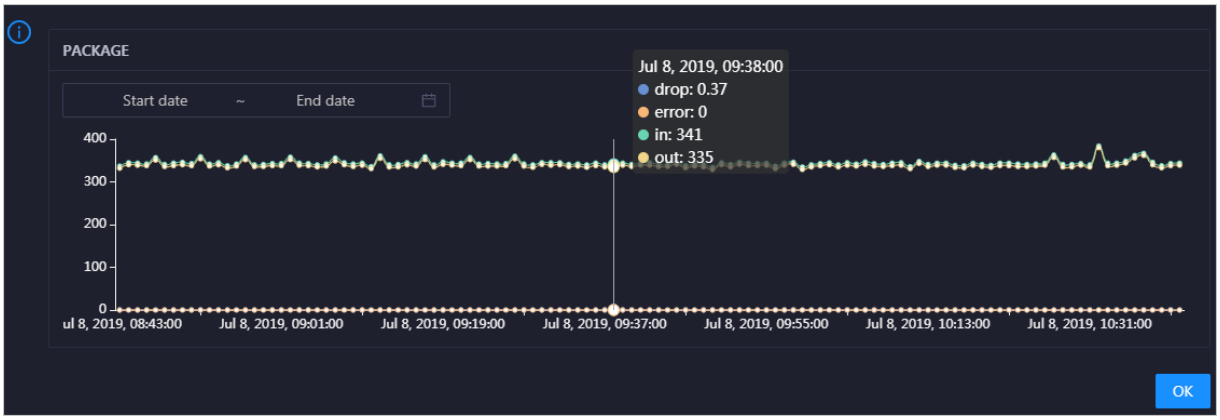

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

The Health Check section shows the number of checkers deployed for the host and the numbers of CRITICAL, WARNING, and EXCEPTION alerts.

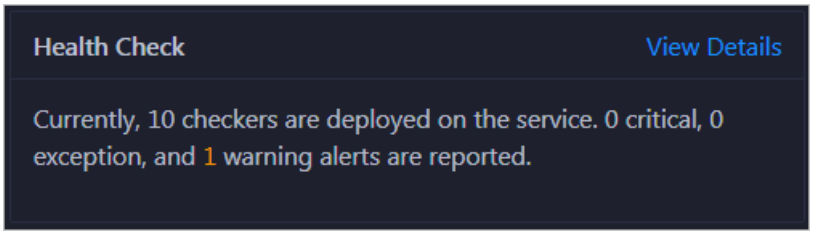

Click **View Details** to go to the Health Status tab. On this tab, you can view the health check details.

Health Check History

The Health Check History section shows the records of the health checks performed on the host.

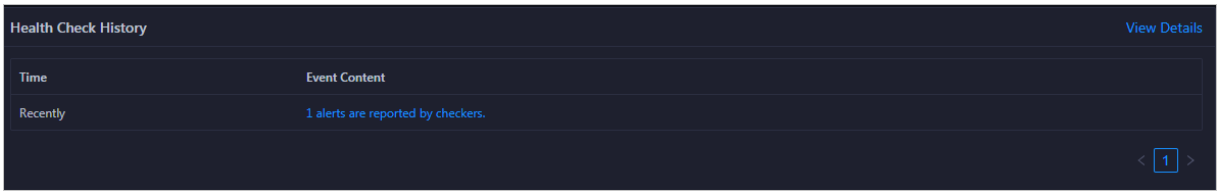

Click **View Details** to go to the Health Status tab. On this tab, you can view the health check details.

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

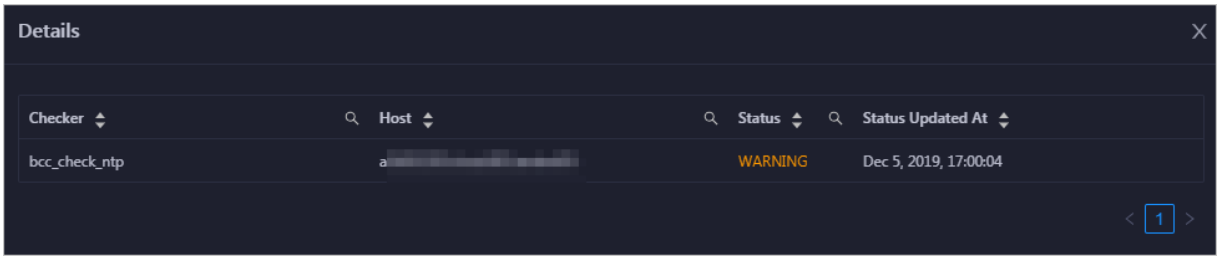

6.2.1.5.5.3. Host charts

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

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

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

6.2.1.5.5.4. Host health

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

Entry

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

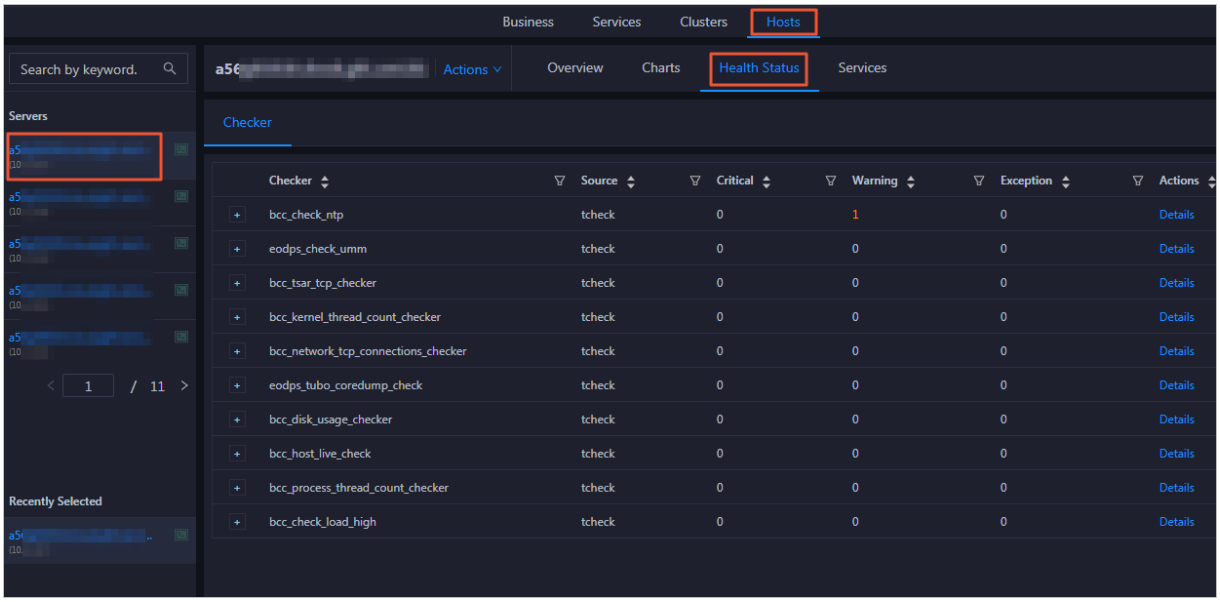

On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

View checker details

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

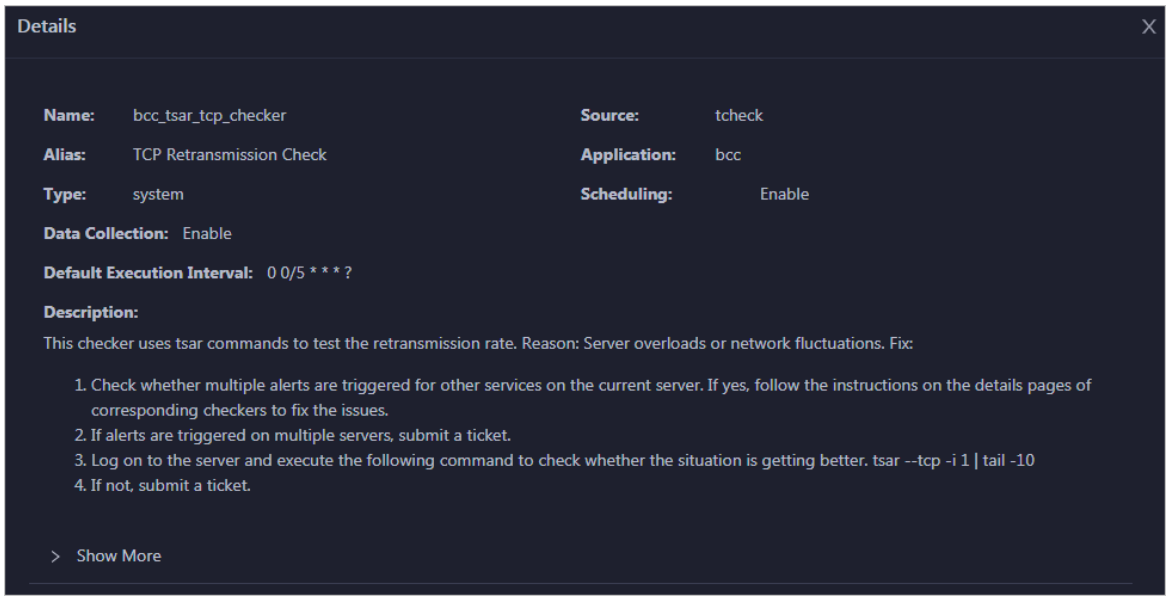

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

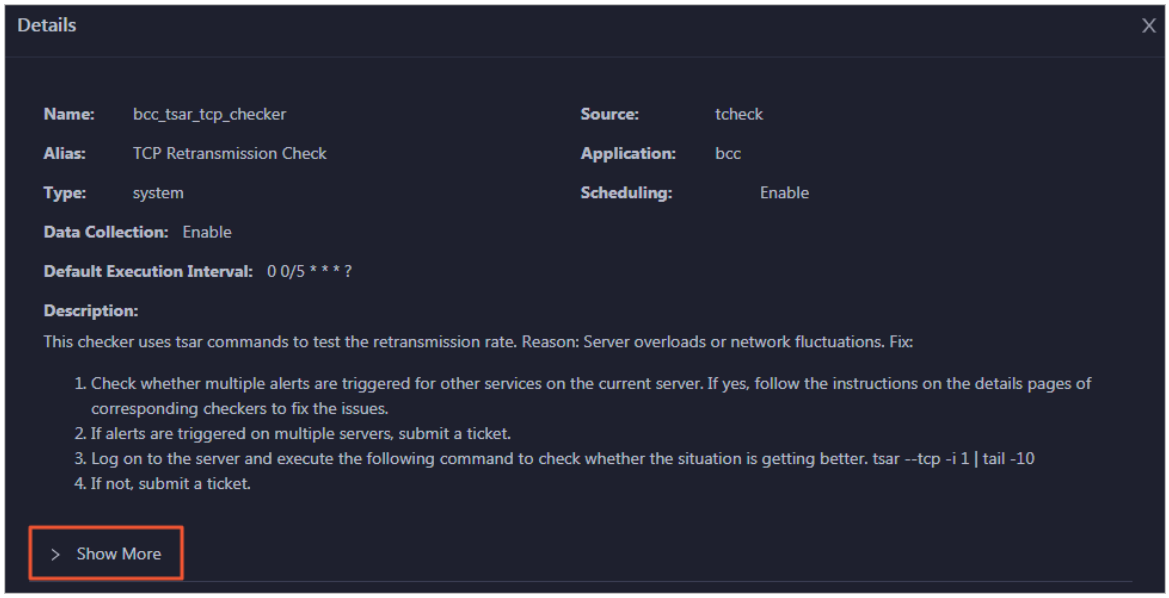

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

View alert causes

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

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

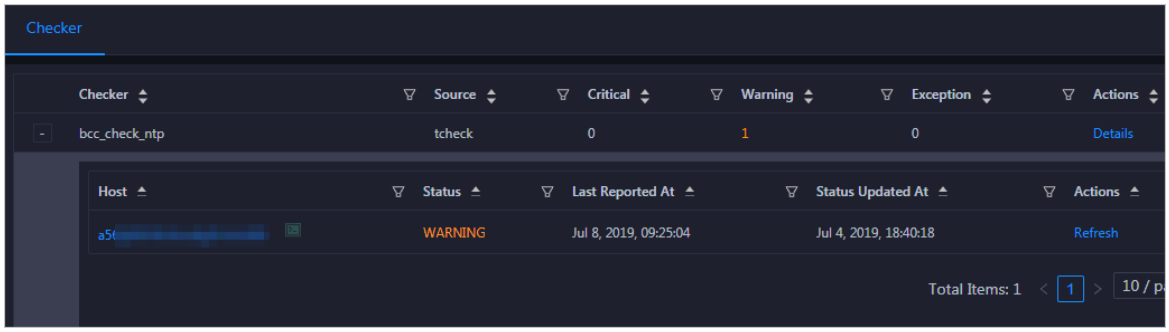

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

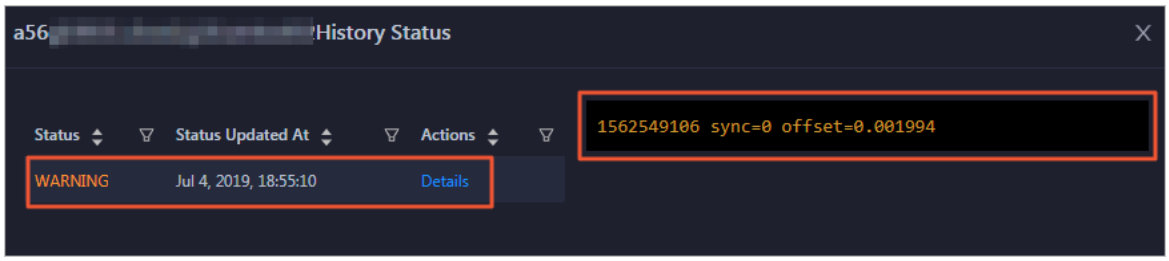

Clear alerts

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

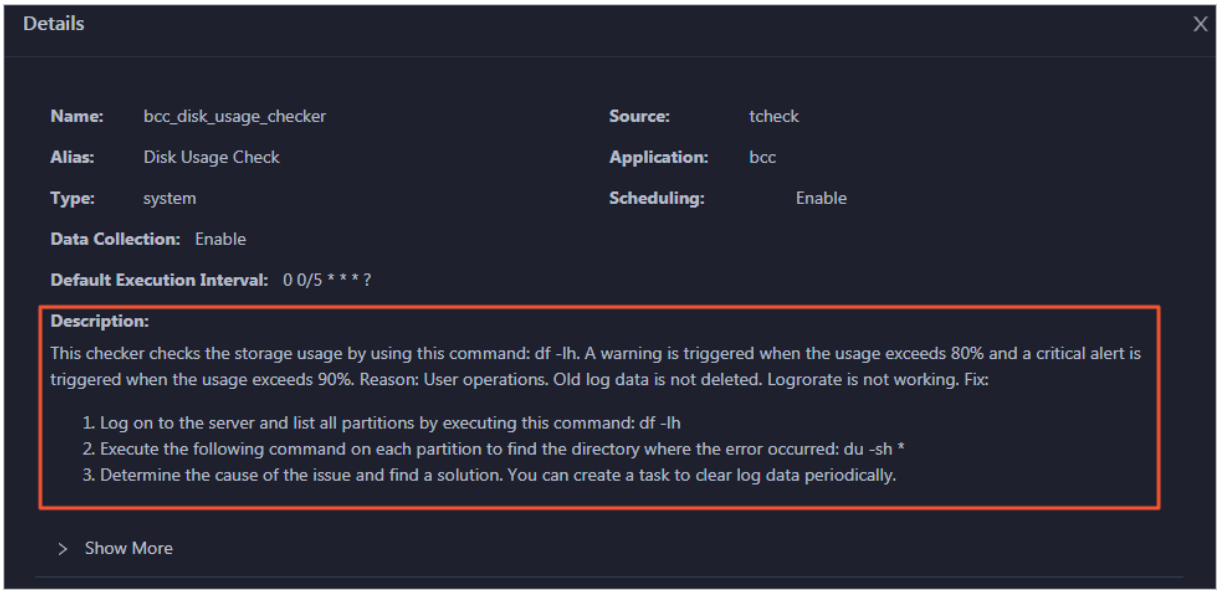

Log on to a host

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

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

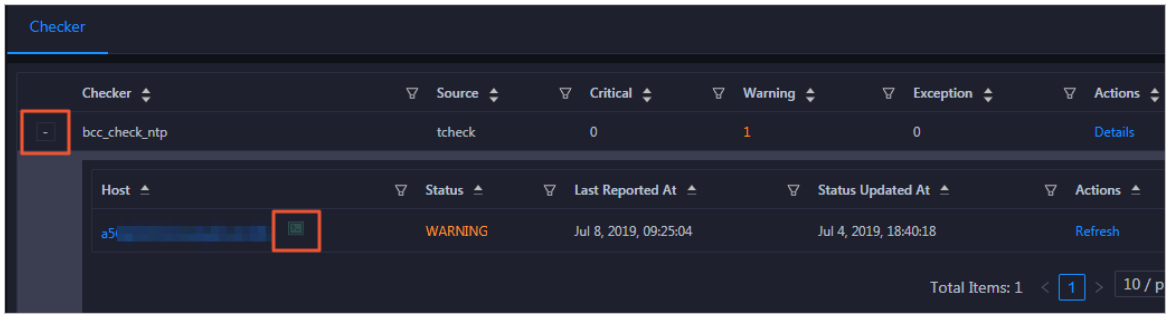

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

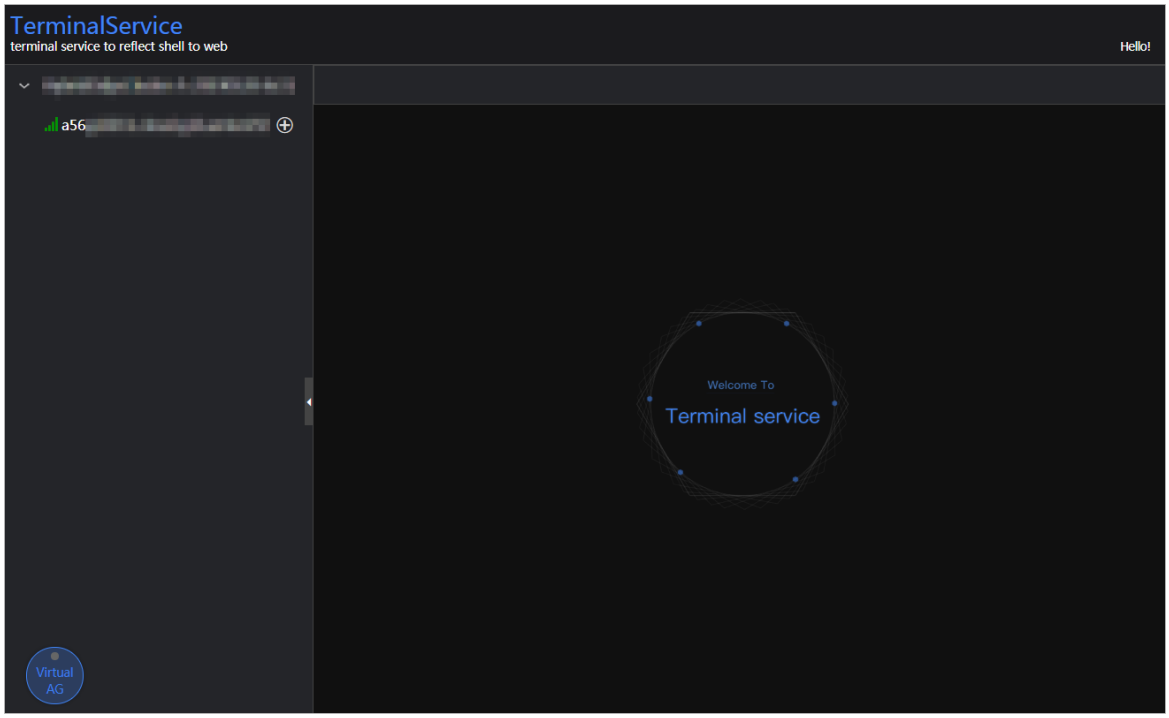

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

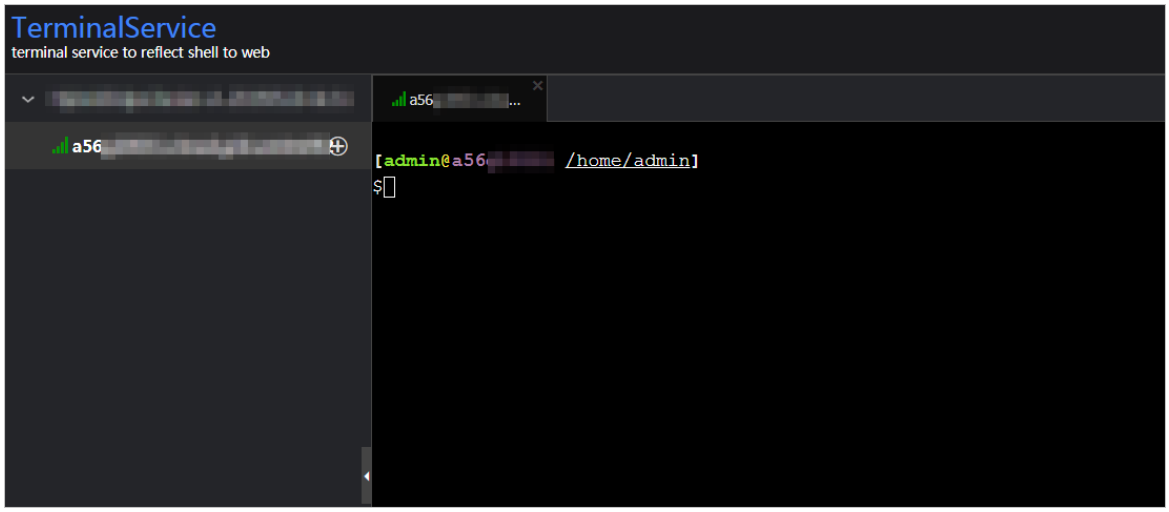

Run a checker again

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

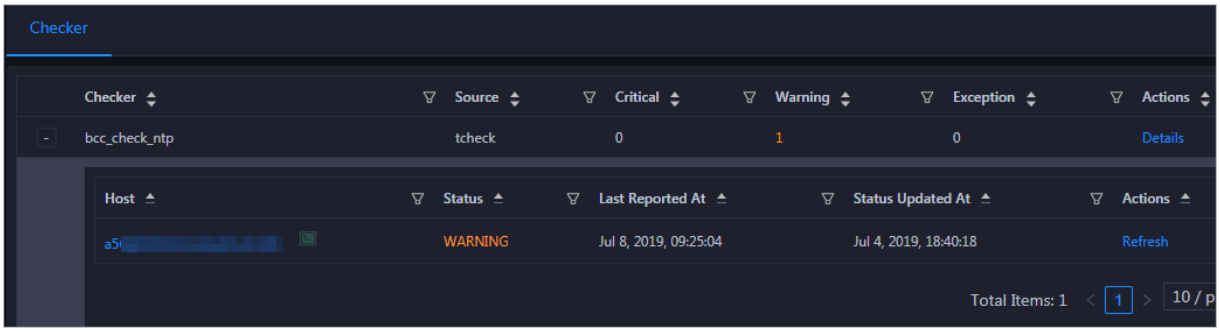

6.2.1.5.5.5. Host services

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

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

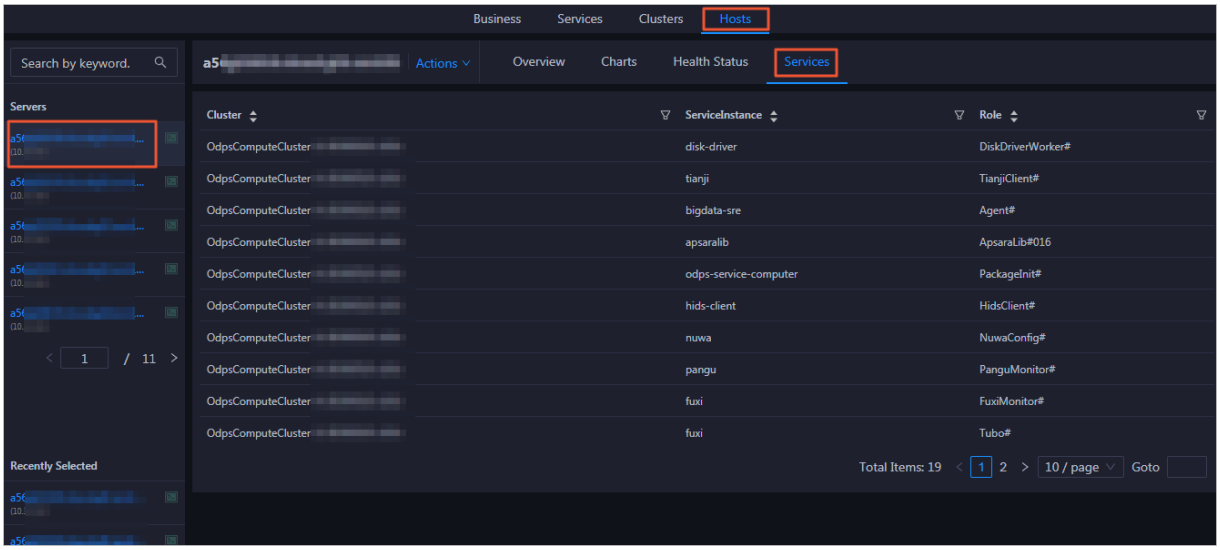

On the **Services** page, you can view the cluster, service instances, and service instance roles of the host.

6.3. DataWorks

6.3.1. Operations and Maintenance Guide

6.3.1.1. Basic concepts and structure

6.3.1.1.1. 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.

6.3.1.1.2. 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.

6.3.1.1.3. 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.

6.3.1.1.4. DataWorks architecture in Apsara Stack V3

This topic describes the framework and services of DataWorks.

DataWorks framework

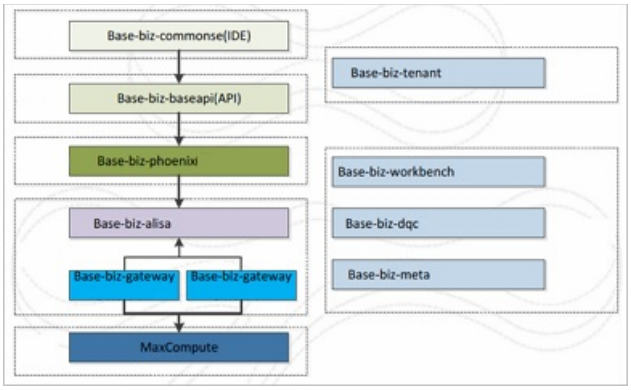

Services shown in the preceding figure play an important role for node scheduling and running. You can perform all O&M operations for DataWorks of Apsara Stack V3 in Apsara Infrastructure Management Framework. The following figure shows the services in DataWorks.

DataWorks services

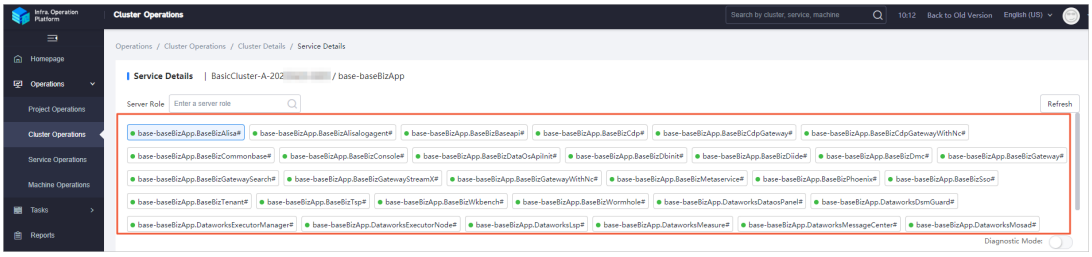

All services in DataWorks are deployed in Docker containers. You can log on to a host and run the docker ps command to view the containers in which the services are deployed.

Service architecture shows the architecture of each service except base-biz-gateway.

Service architecture

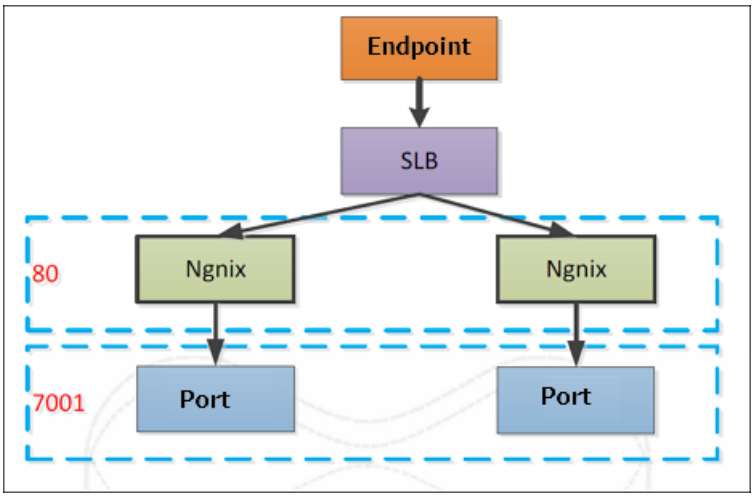

6.3.1.1.5. Service directories

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.

6.3.1.2. O&M by using Apsara Big Data Manager

6.3.1.2.1. Log on to the ABM console

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

Prerequisites

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

  The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id*.ops.console.*intranet-domain-id*.
- A browser is available. We recommend that you use Google Chrome.

Procedure

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

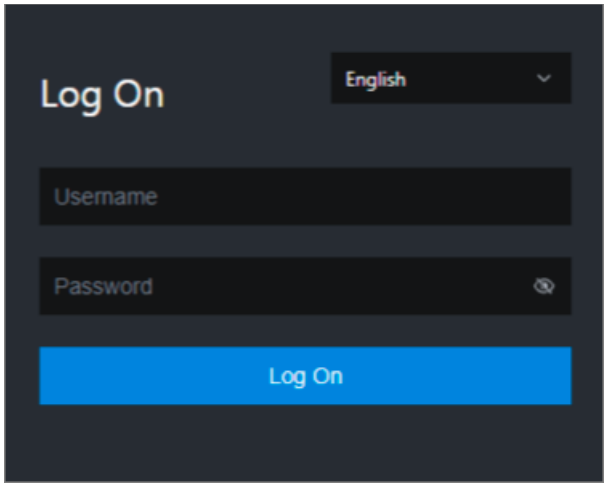

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.
4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Big Data Services** section, choose **General-Purpose O&M > Apsara Big Data Manager**.

6.3.1.2.2. 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
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.

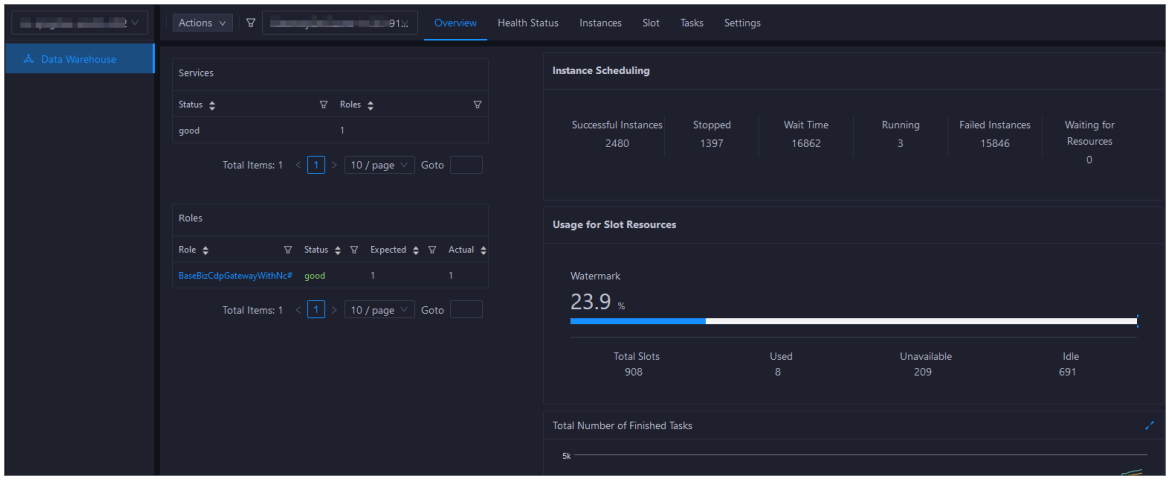

The **O&M** page includes three modules: **Services**, **Clusters**, and **Hosts**.

6.3.1.2.3. Service O&M

6.3.1.2.3.1. Data Warehouse

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.

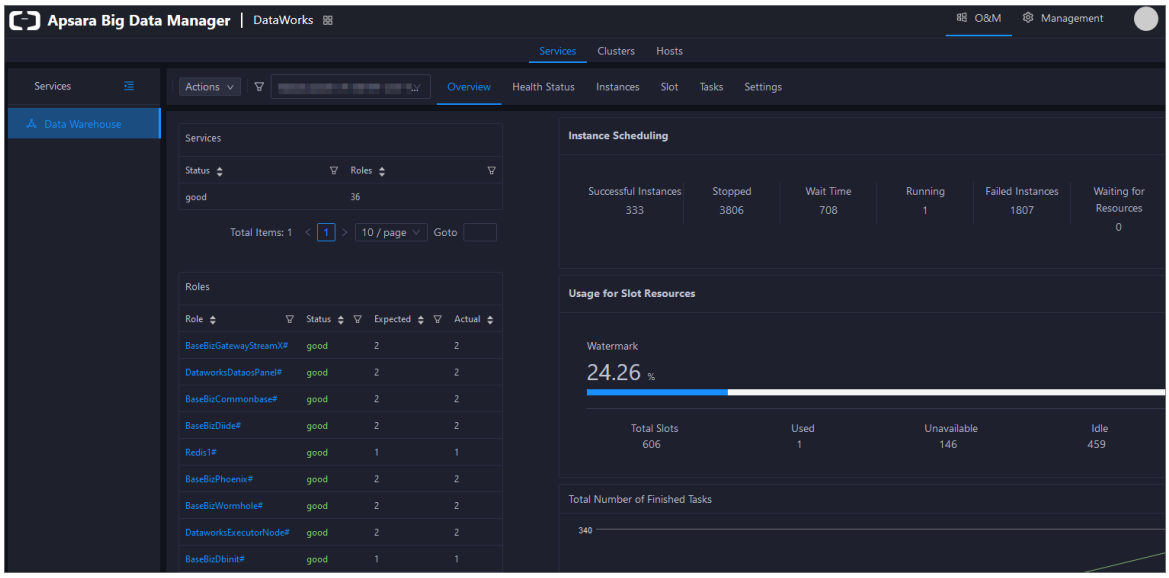

Services

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

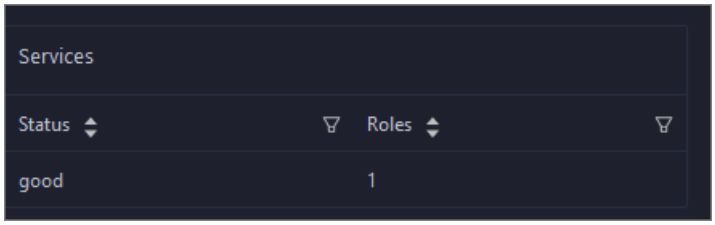

Roles

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.

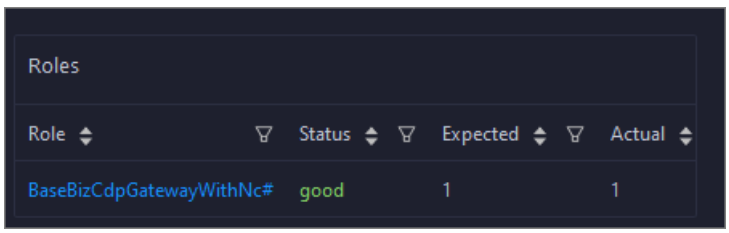

Instance Scheduling

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.

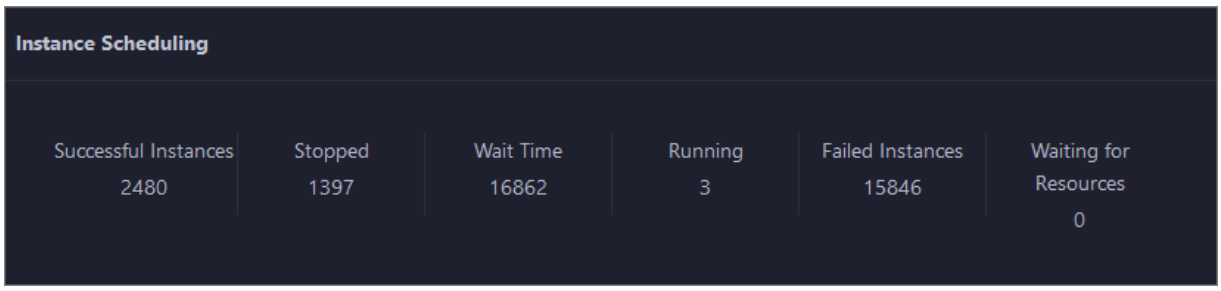

Usage for Slot Resources

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.

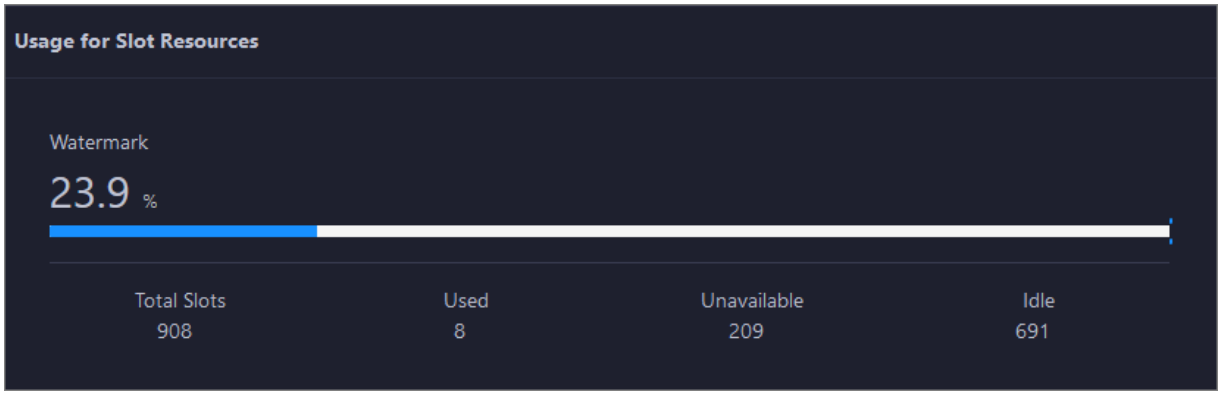

Note Slots are resources that can be used by DataWorks for instance scheduling.

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.

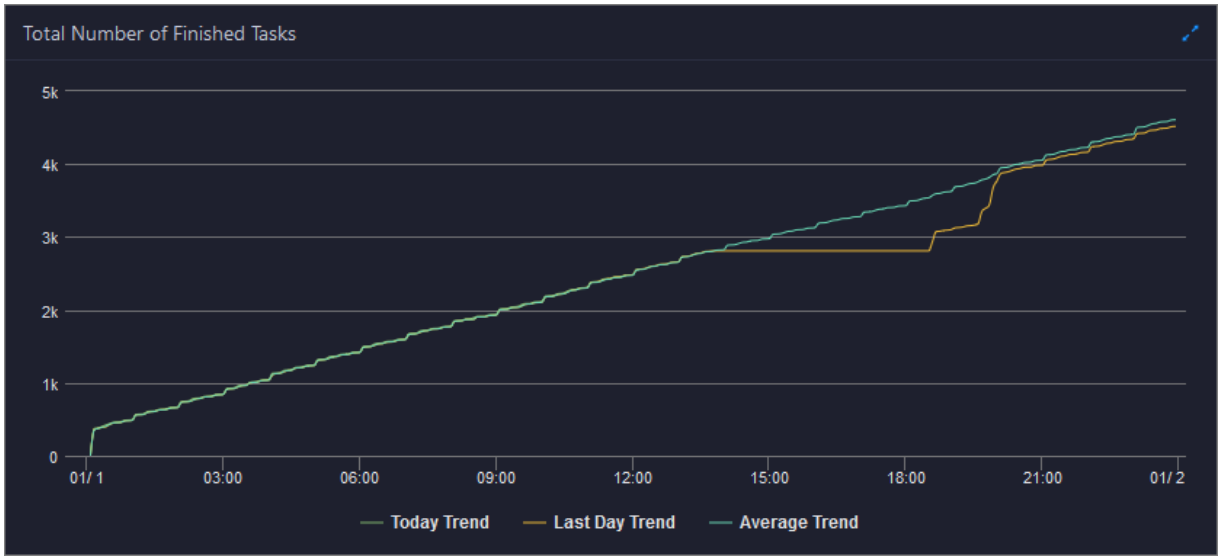

Service health

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.

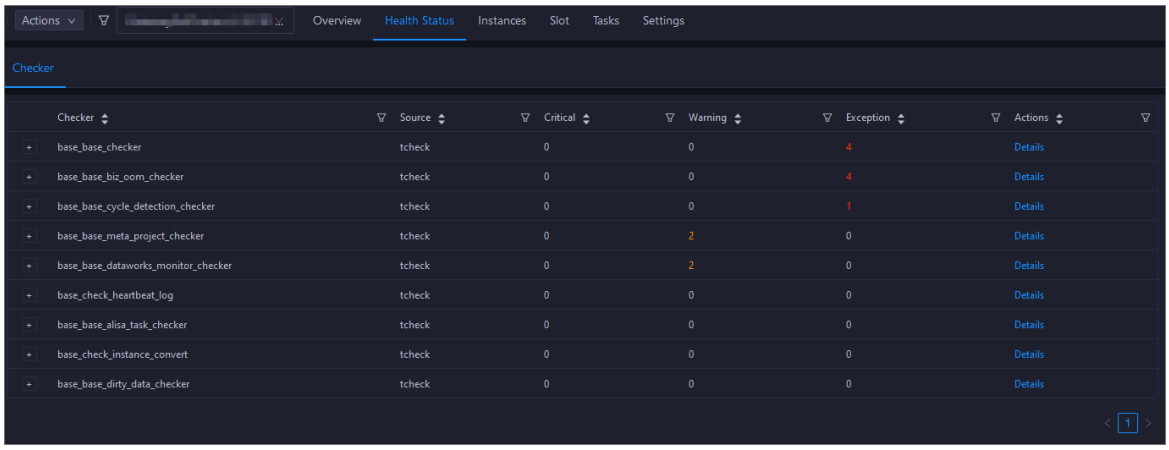

The **Health Status** page displays all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical**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

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.

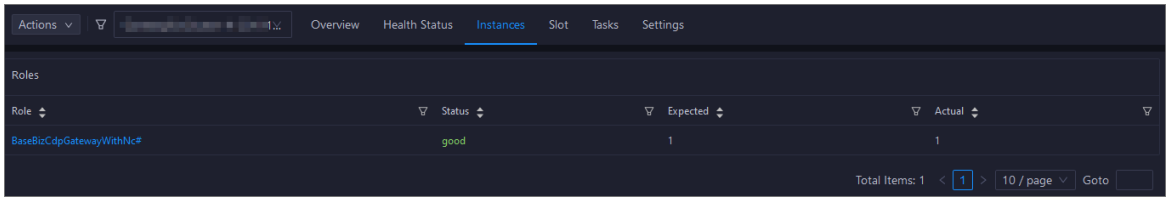

The **Instances** page displays information about all DataWorks service roles, including the status and the expected and actual numbers of hosts in the desired state. The statuses include **good**, **bad**, and **upgrading**.

Supported operations

You can filter or sort service roles based on a column to facilitate information retrieval on the **Instances** page.

Service slots

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

Terms

A data migration unit (DMU) represents the minimum operating capability required by a Data Integration task. This capability indicates the data synchronization processing capability in the case of 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 select **DataWorks**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Overview** page appears.
4. Select a cluster from the drop-down list and click the **Slot** tab.

Cluster slots

Click the **Cluster** tab on the **Slot** tab.

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

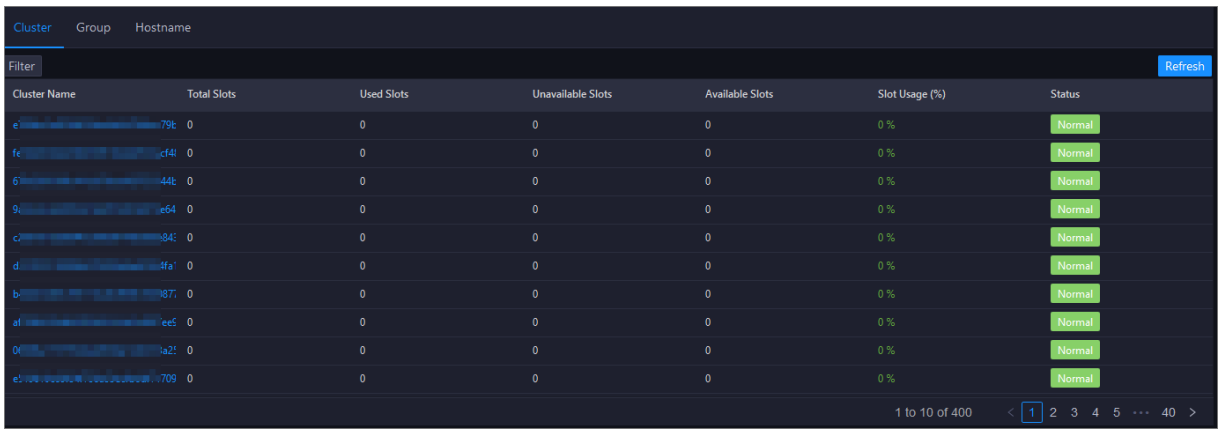

To view more information about the slots of a specified cluster, click the name of the cluster in the **Cluster Name** column.

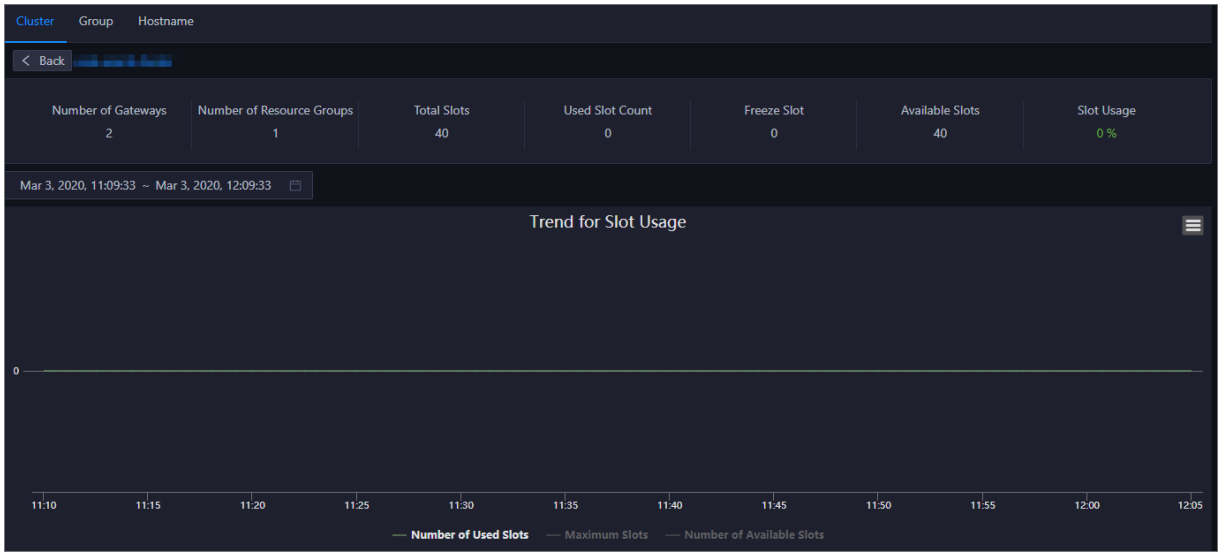

In the upper part of the page that appears, you can view the numbers of gateways, resource groups, slots, used slots, frozen slots, and available slots, and the slot usage of the cluster. You can also view the trend chart of slot usage over time in the lower part of the page. The trend chart displays trend lines for 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 related trend line in the chart. A highlighted metric name indicates that the related trend line is displayed, whereas a dimmed metric name indicates that the related trend line is not displayed.

Resource group slots

Click the **Group** tab on the **Slot** tab.

The **Group** tab displays the slot overview of all DataWorks resource groups, including the maximum number of slots, the numbers of used slots and available slots, and slot usage. The tab 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 in the **Resource Group ID** column.

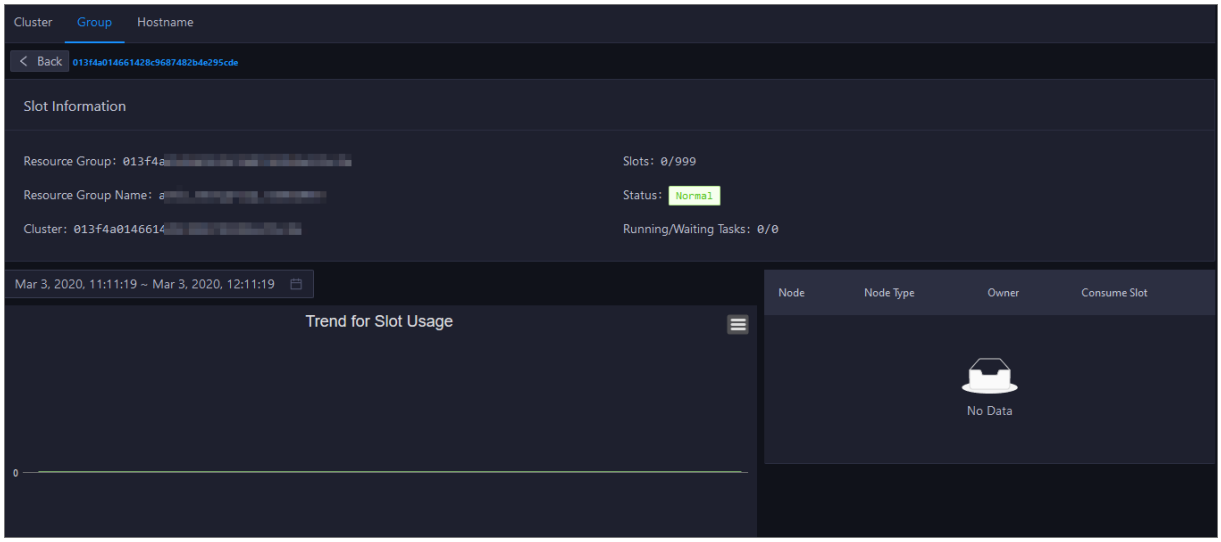

In the upper part of the page that appears, you can view the current slot information of the resource group, such as the number of used slots and the maximum number of slots. You can also view the trend chart of slot usage over time, the nodes that occupy the slots, and the owners in the lower part. The trend chart displays trend lines for 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 related trend line in the chart. A highlighted metric name indicates that the related trend line is displayed, whereas a dimmed metric name indicates that the related trend line is not displayed.

Change the number of slots for a resource group

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

1. On the **Group** tab, find the resource group for which you want to change the number of slots, and click **Change Maximum Slots** in the Actions column.
2. In the dialog box that appears, change the value of **Maximum Slots**.
3. Click **Run**. A message appears, indicating that the action is submitted.

Migrate a resource group

If the slots in a cluster that is associated with a resource group are insufficient and slots cannot be added for the cluster, you can associate the resource group with another cluster.

1. On the **Group** tab, find the resource group that you want to manage. Then, move the pointer over Change Maximum Slots in the Actions column and click **Bind Resource Group**.
2. In the dialog box that appears, change the value of **Target Cluster**.
3. Click **Run**. A message appears, indicating that the action is submitted.

Host slots

Click the **Hostname** tab on the **Slot** tab.

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

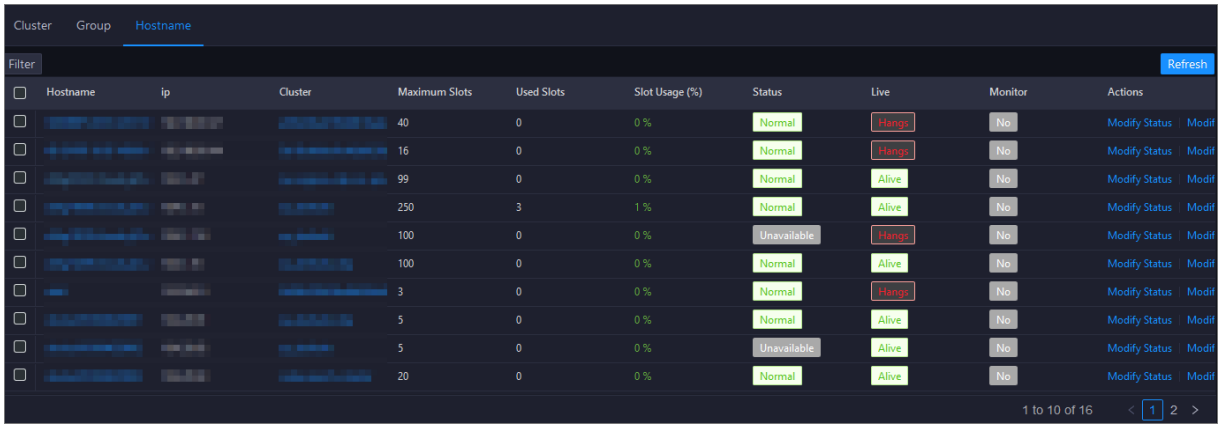

To view more information about the slots of a specified host, click the name of the host in the **Hostname** column.

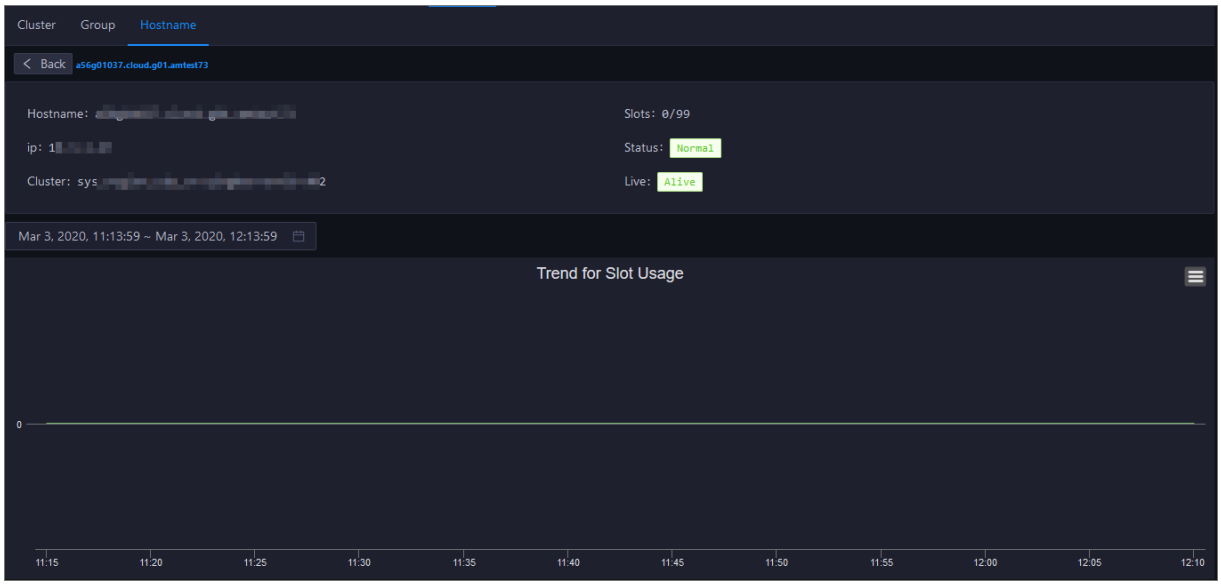

In the upper part of the host details page, you can view the current slot information of the host, such as the number of used slots and the maximum number of slots. You can also view the trend chart of slot usage over time in the lower part of the page. The trend chart displays trend lines for the number 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 related trend line in the chart. A highlighted metric name indicates that the related trend line is displayed, whereas a dimmed metric name indicates that the related trend line is not displayed.

Modify the host status

A host can be in the normal, unavailable, or suspended state. You can modify the host status based on your business requirements.

1. On the **Hostname** tab, find the host whose status you want to modify and click **Change Status** in the Actions column.
2. In the dialog box that appears, set **Status**.
3. Click **Run**. A message appears, indicating that the action is submitted.

Change the number of slots for a host

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

1. On the **Hostname** tab, find the host that you want to manage and click **Change Maximum Slots** in the Actions column.
2. In the dialog box that appears, change the value of **Maximum Slots**.
3. Click **Run**. A message appears, indicating that the action is submitted.

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.

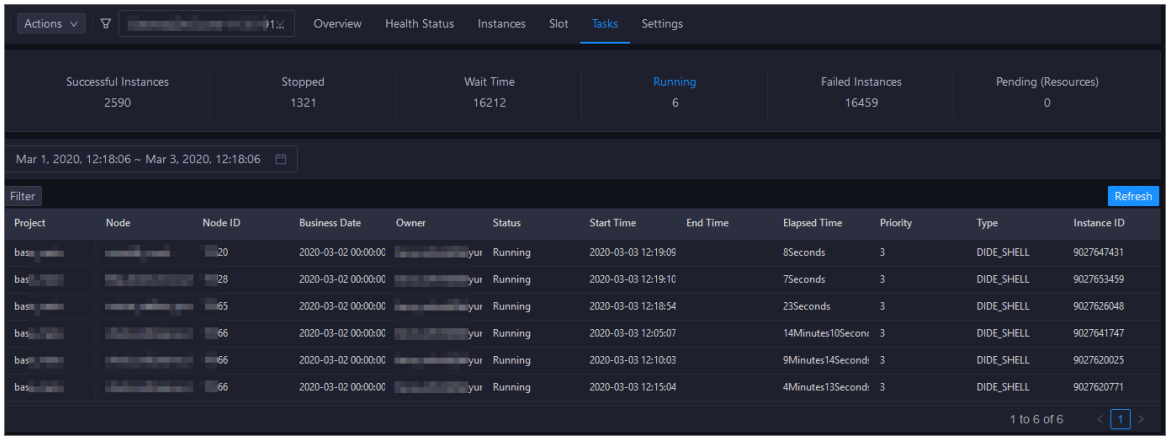

The **Tasks** page displays the node information of the current cluster, including the project name, node name, node ID, data timestamp, owner, running status, start time, end time, running duration, priority, type, and instance ID.

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.

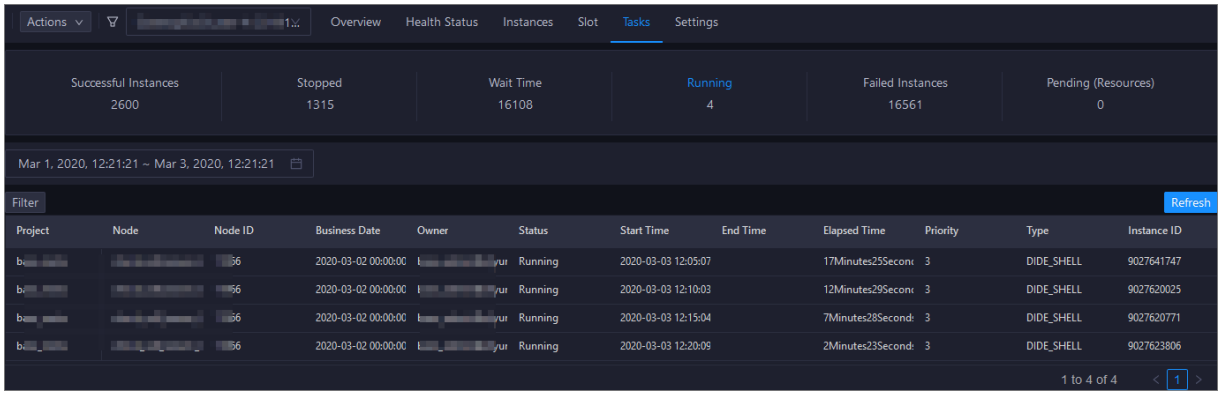

Filter nodes by time

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.

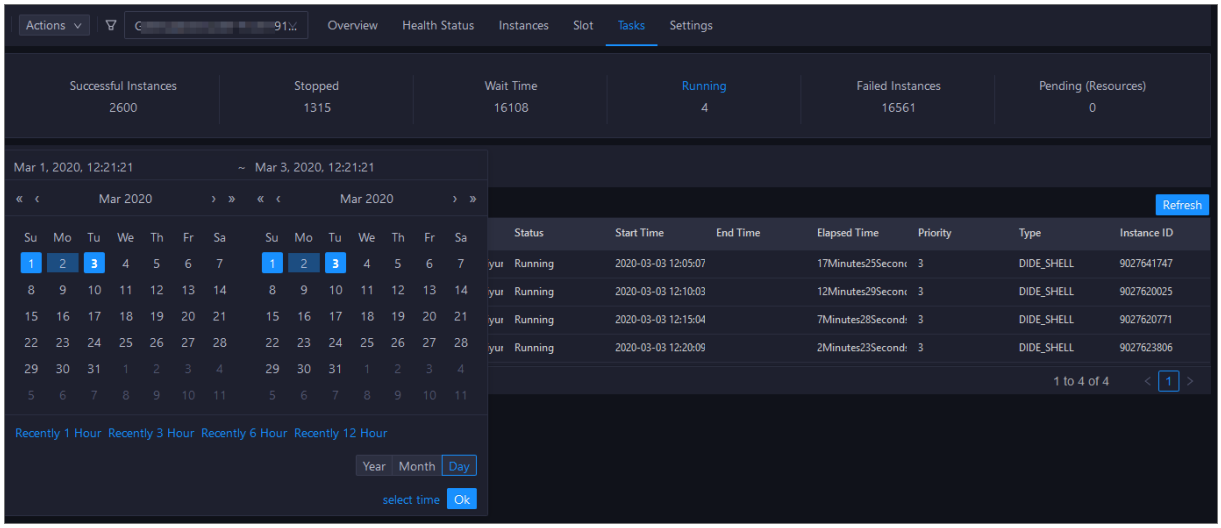

Other operations

You can filter nodes, sort nodes based on a column, and customize columns on the Tasks page.

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.

6.3.1.2.3.2. Data Integration

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**

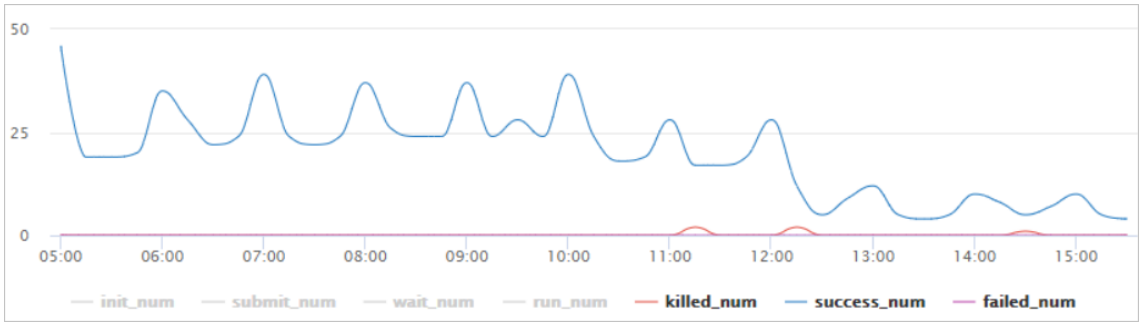

   - **Third-party Dependencies - Response Time (milliseconds)**

- Third-party Dependencies - Total Requests

- Third-party Dependencies - Request Error Rate

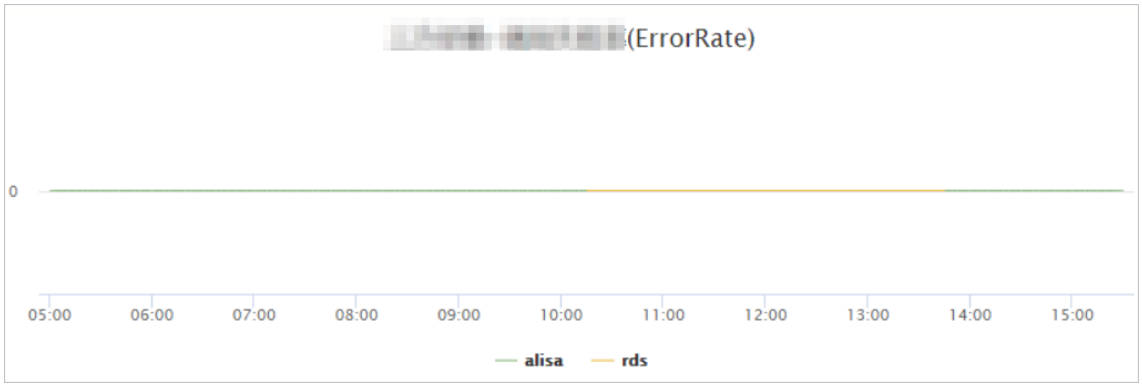

- Third-party Dependencies - Failed Requests

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**.

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**.

6.3.1.2.3.3. Cluster scaling

Apsara Big Data 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 information

In Apsara Stack, scaling out a cluster involves complex operations. You must configure new physical hosts on Deployment Planner so that they can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework is 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.

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 server roles of physical hosts in DataWorks include **BaseBizCdpGatewayWithNc#** and **BaseBizGatewayWithNc#**. DataWorks cluster scaling supports only the two server roles.

Prerequisites

- Scale-out
  - The physical hosts that you want to add to your DataWorks cluster are in the default cluster of Apsara Infrastructure Management Framework.
  - The server role of the template host is **BaseBizCdpGatewayWithNc#** or **BaseBizGatewayWithNc#**.

- Scale-in

  The server role of the template host is **BaseBizCdpGatewayWithNc#** or **BaseBizGatewayWithNc#**.

> **Note** You can go to the **DataWorks** page. Then, click **O&M** in the upper-right corner, and click the **Services** tab. In the left-side navigation pane, click **Data Warehouse**. On the page that appears, click the **Instances** tab. In the server role list, find the server role **BaseBizCdpGatewayWithNc#** or **BaseBizGatewayWithNc#**, and click the server role name to go to the Apsara Infrastructure Management Framework console and view the hosts with the server 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 must specify an existing host as the template host. During the scale-out, the configurations of the template host are copied to the hosts 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 select **DataWorks**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Overview** tab appears.
4. Select a cluster from the drop-down list and click the **Slot** tab.
5. On the **Slot** tab, click the **Hostname** tab. Then, select a physical host whose server role is **BaseBizCdpGatewayWithNc#** or **BaseBizGatewayWithNc#** in the host list as the template host.

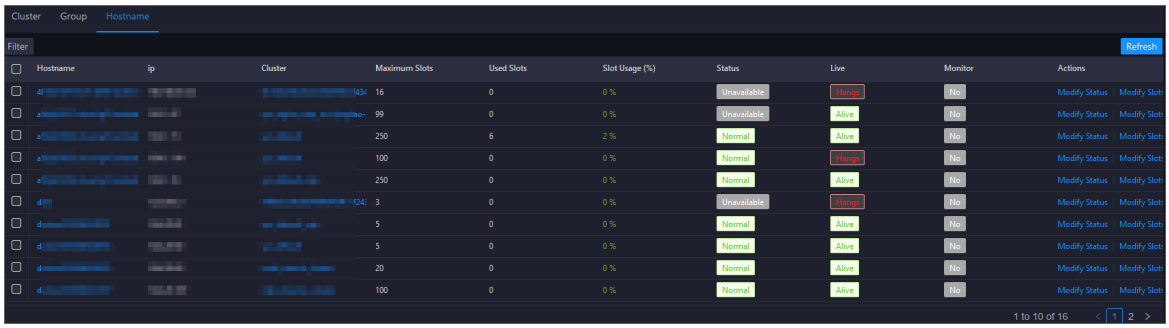

6. Move the pointer over **Actions** in the upper-right corner and click **Scale-up for Normal Hosts**. In the **Scale-up for Normal Hosts** panel, configure the parameters

   Parameters:

   - **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 DataWorks cluster. Enter the name of an available host in the default cluster for scale-out. If you want to add multiple hosts, enter multiple hostnames and separate the hostnames with commas (,).
7. Click **Run**. A message appears, indicating that the action is submitted.
8. View the scale-out status.

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

The scale-out may require a long period. 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 **RUNNING** is displayed in the Current Status column, you can click **Details** in the Details column to view the steps and progress of the scale-out.

If **FAILED** is displayed in the Current Status column, click **Details** in the Details column to locate the failure cause. For more information, see Locate the cause of a scaling failure.

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 select **DataWorks**.
3. On the page that appears, click **O&M** in the upper-right corner. The **Overview** page appears.
4. Select a cluster from the drop-down list and click the **Slot** tab.
5. On the **Slot** tab, click the **Hostname** tab. Then, select a physical host whose server role is **BaseBizCdpGatewayWithNc#** or **BaseBizGatewayWithNc#** in the host list as the template host.

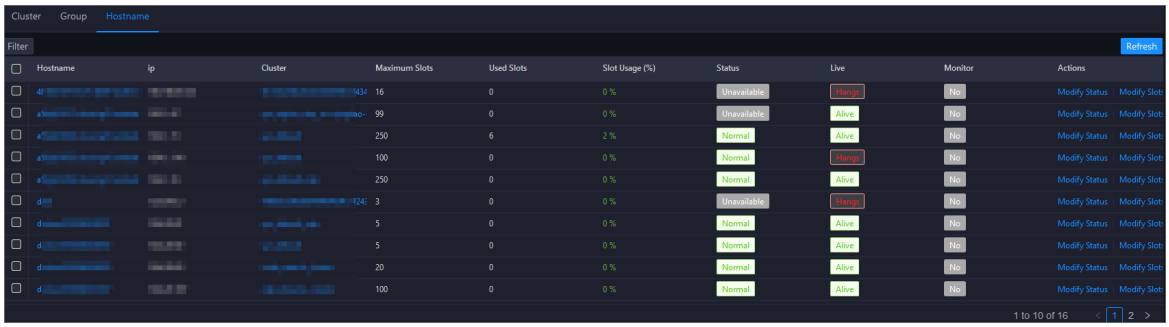

6. Move the pointer over **Actions** in the upper-right corner and click **Scale-down for Normal Hosts**. In the **Scale-up for Normal Hosts** panel, configure the parameters.

   Parameters:

   - **Hostname**: the name of the host that you want to remove from the DataWorks cluster. By default, the name of the selected host is used.
   - **Biz Name**: the server role of the host that you want to remove from the DataWorks cluster. Select the actual server 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 is submitted.
8. View the scale-in status.

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

   The scale-in may require a long period. 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 **RUNNING** is displayed in the Current Status column, you can click **Details** in the Details column to view the steps and progress of the scale-in.

If **FAILED** is displayed in the Current Status column, click **Details** in the Details column to locate the failure cause. For more information, see Locate the cause of a scaling failure.

Locate the cause of a scaling failure

The method for locating the cause of a scale-out failure and that of a scale-in failure are similar. This section describes how to locate the cause of a scale-out failure.

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** page appears.
4. Move the pointer over **Actions** in the upper-right corner and click **Execution History** next to **Scale-up for Normal Hosts** to view the scale-out history.
5. In the panel that appears, click **Details** in the Details column of a failed execution to locate the failure cause.

   You can locate the failure cause based on the following information: parameter settings, host details, scripts, and execution parameters.

6.3.1.2.4. Cluster O&M

6.3.1.2.4.1. 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.

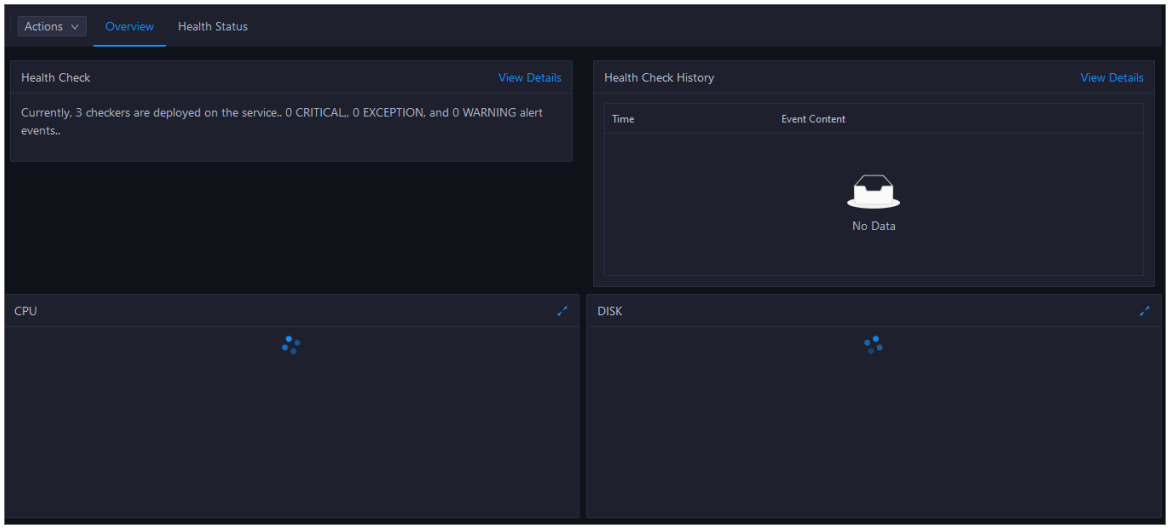

Health Check

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

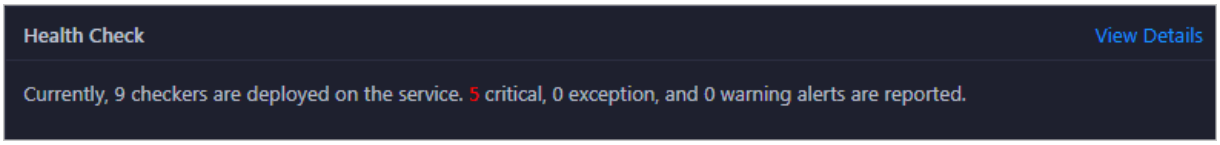

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Cluster health.

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.

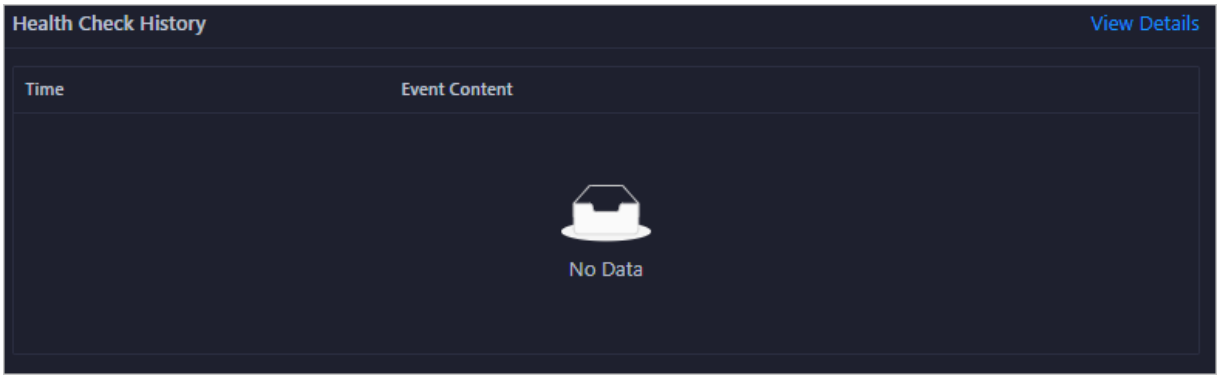

You can click the event content of a check to view the exception items.

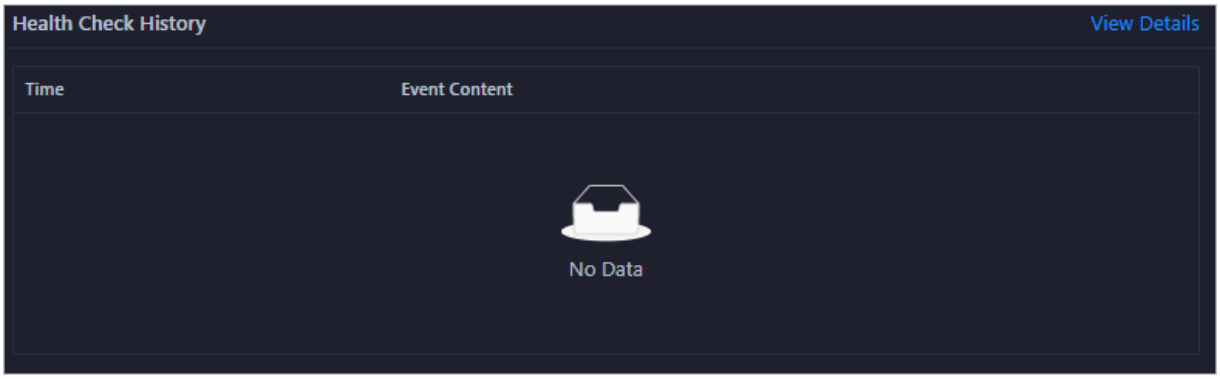

CPU

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.

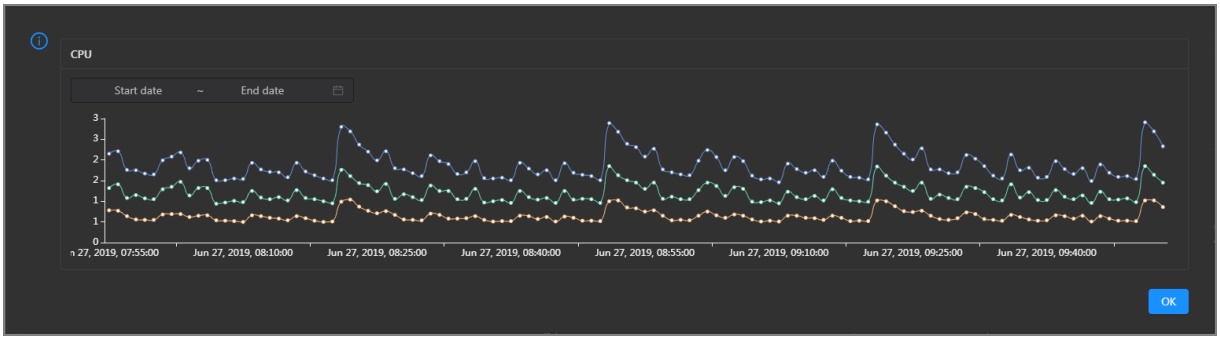

DISK

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

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

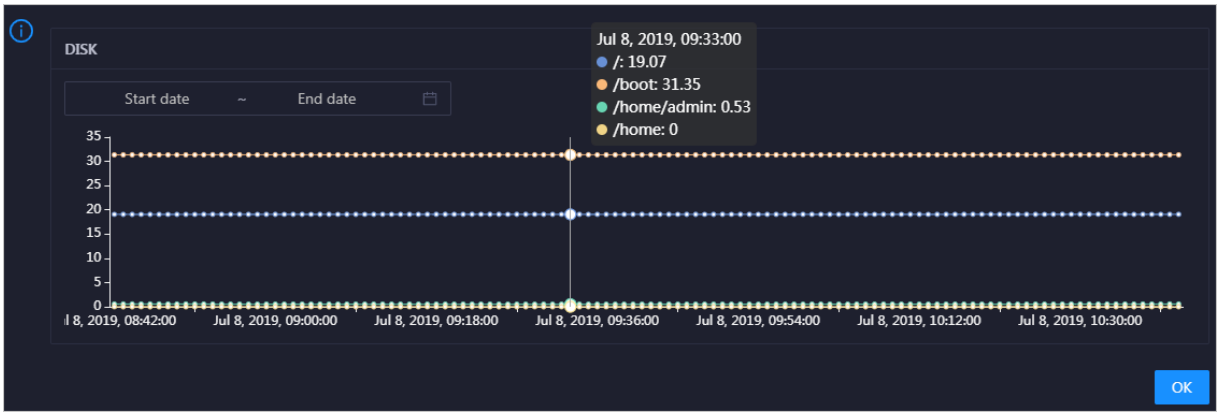

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

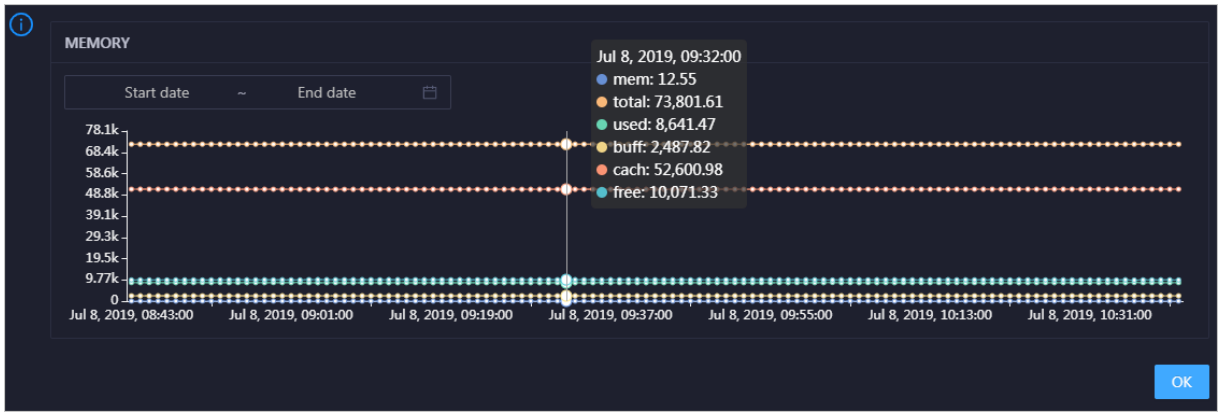

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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.

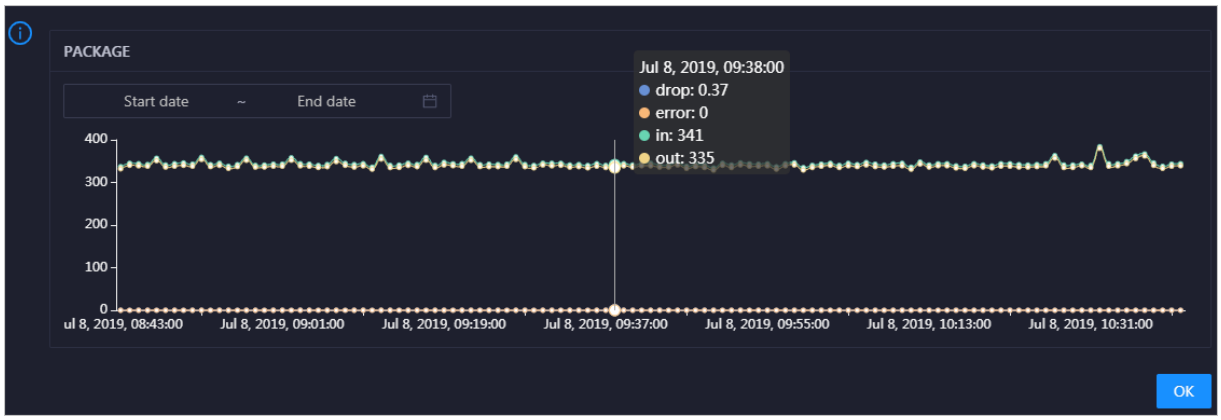

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

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.

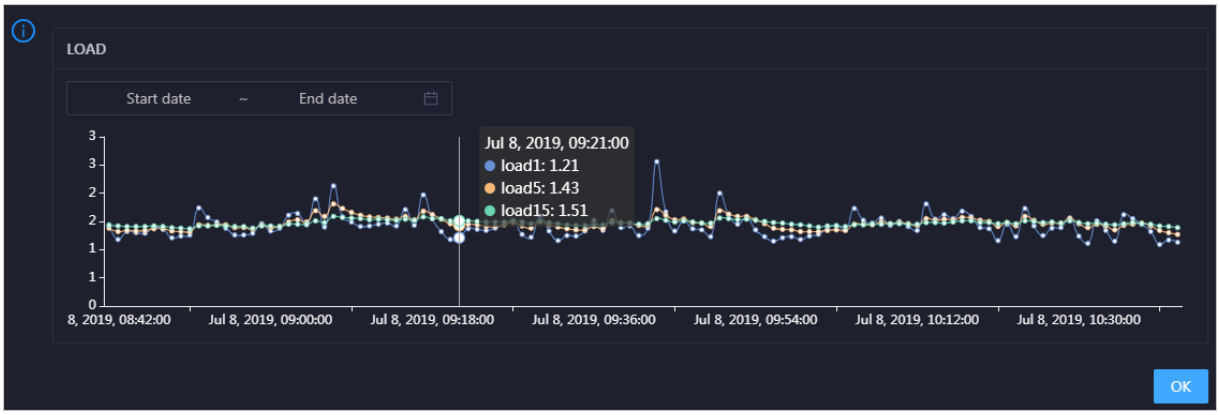

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

6.3.1.2.4.2. 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.

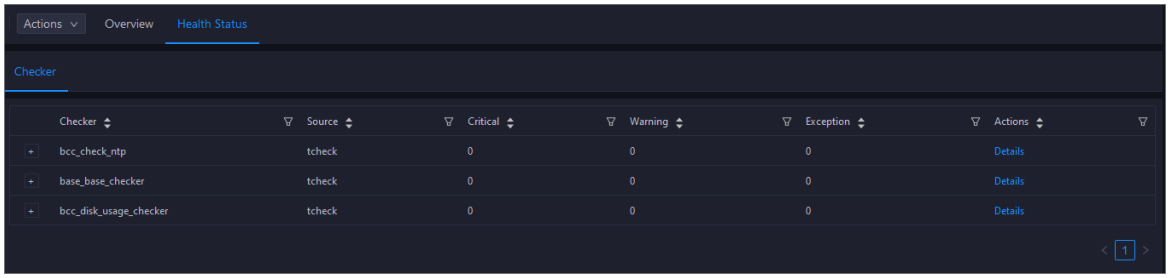

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are represented in different colors. You must handle the alerts in a timely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

View checker details

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

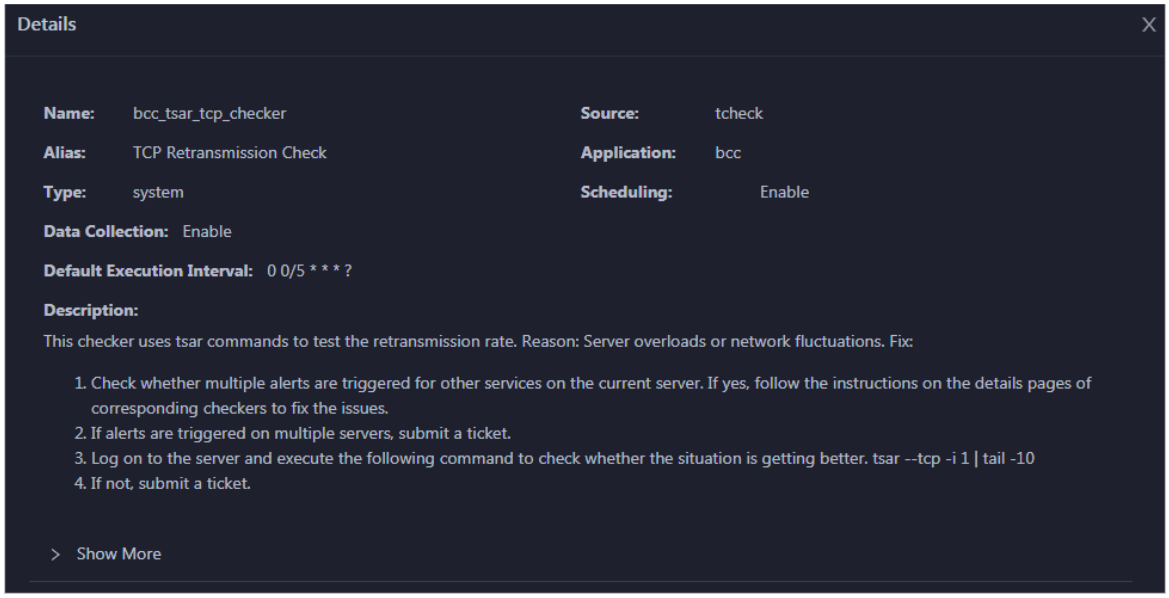

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

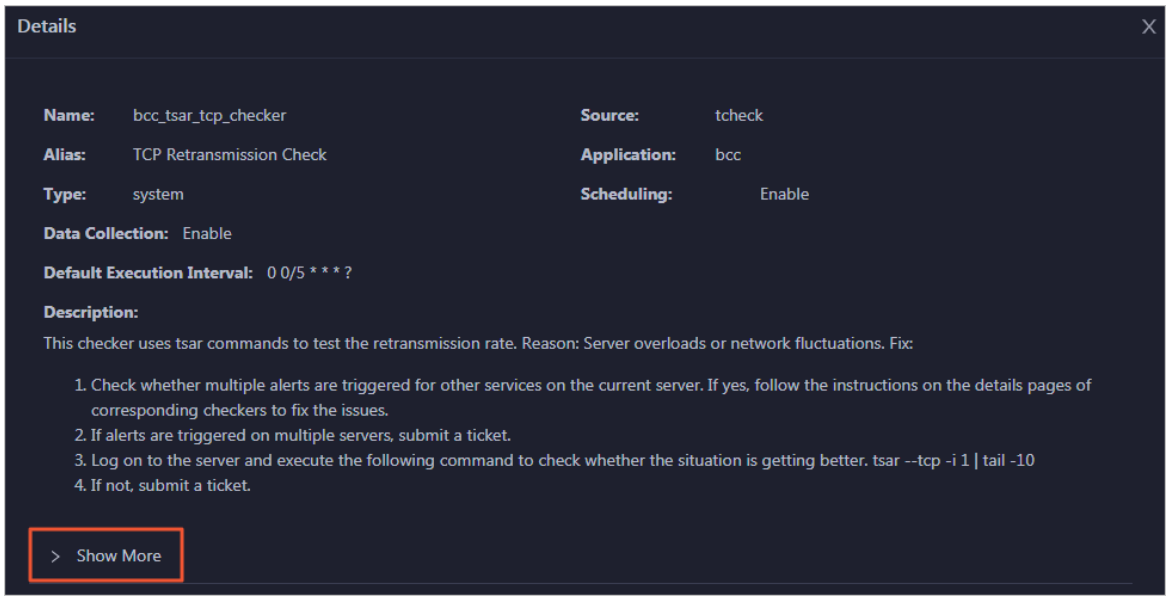

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

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

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

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

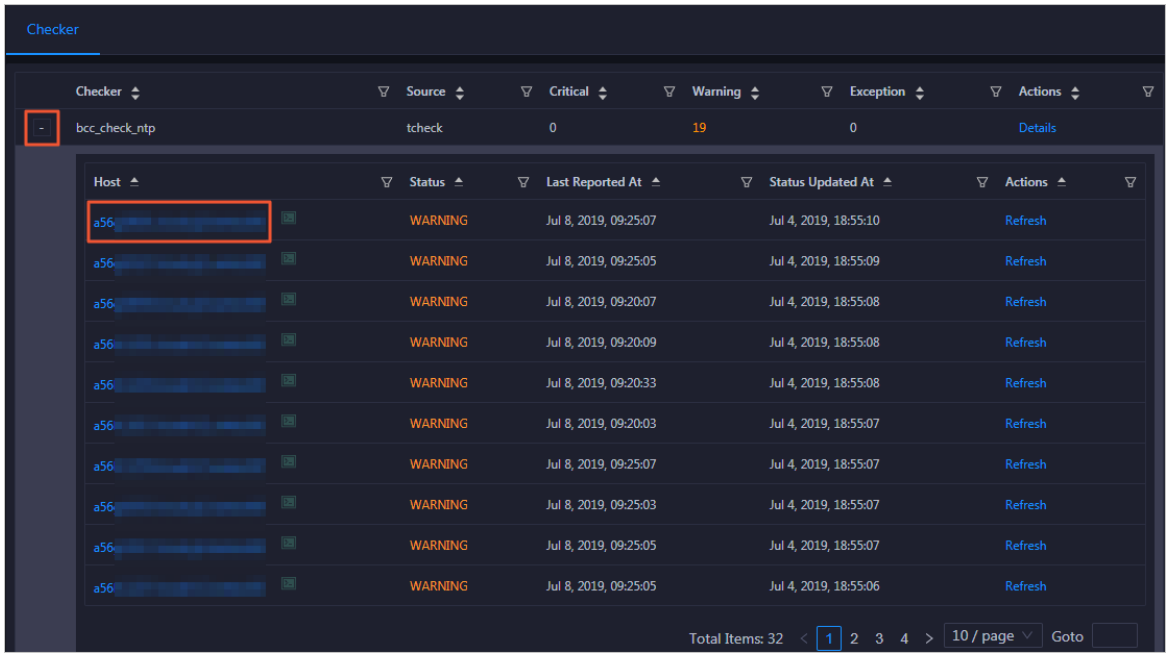

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

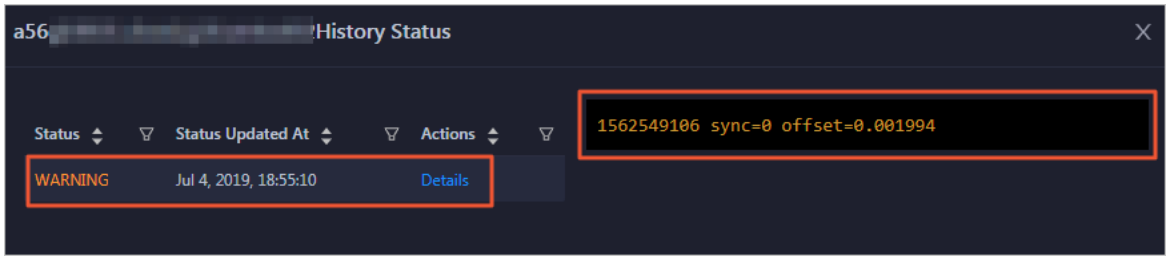

Clear alerts

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

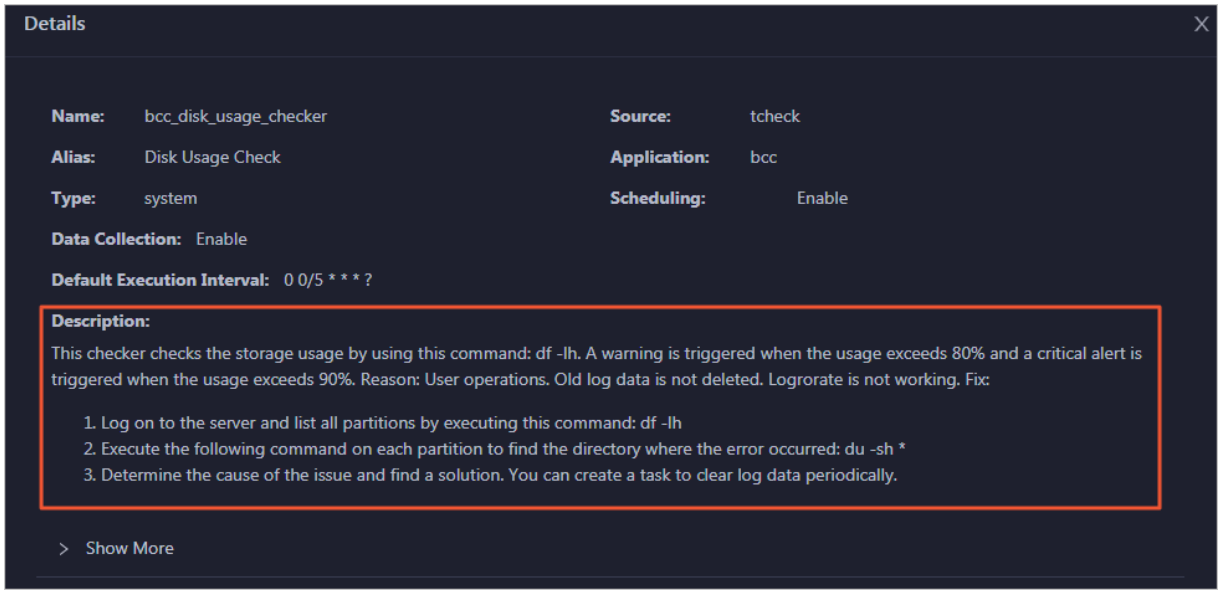

Log on to a host

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

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

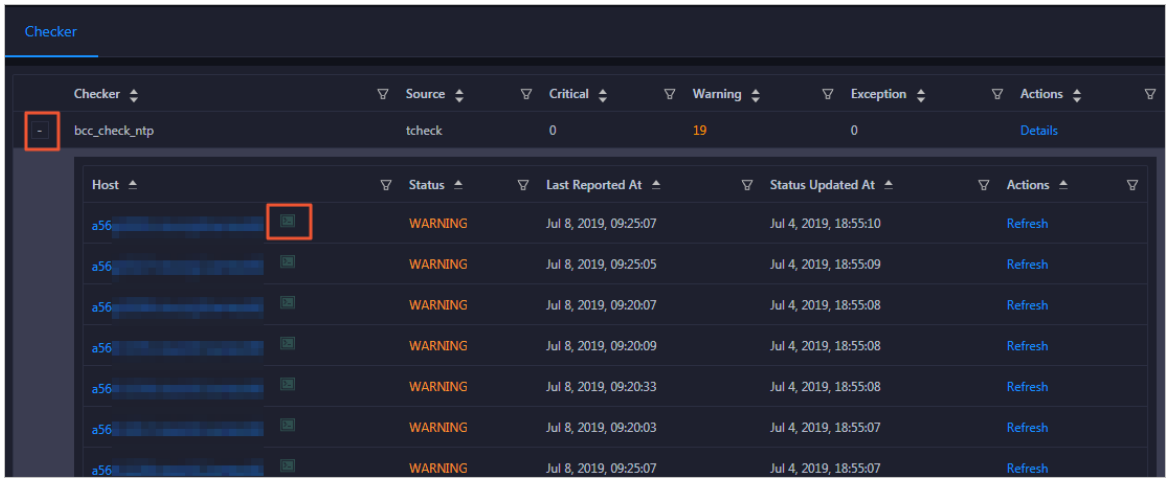

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

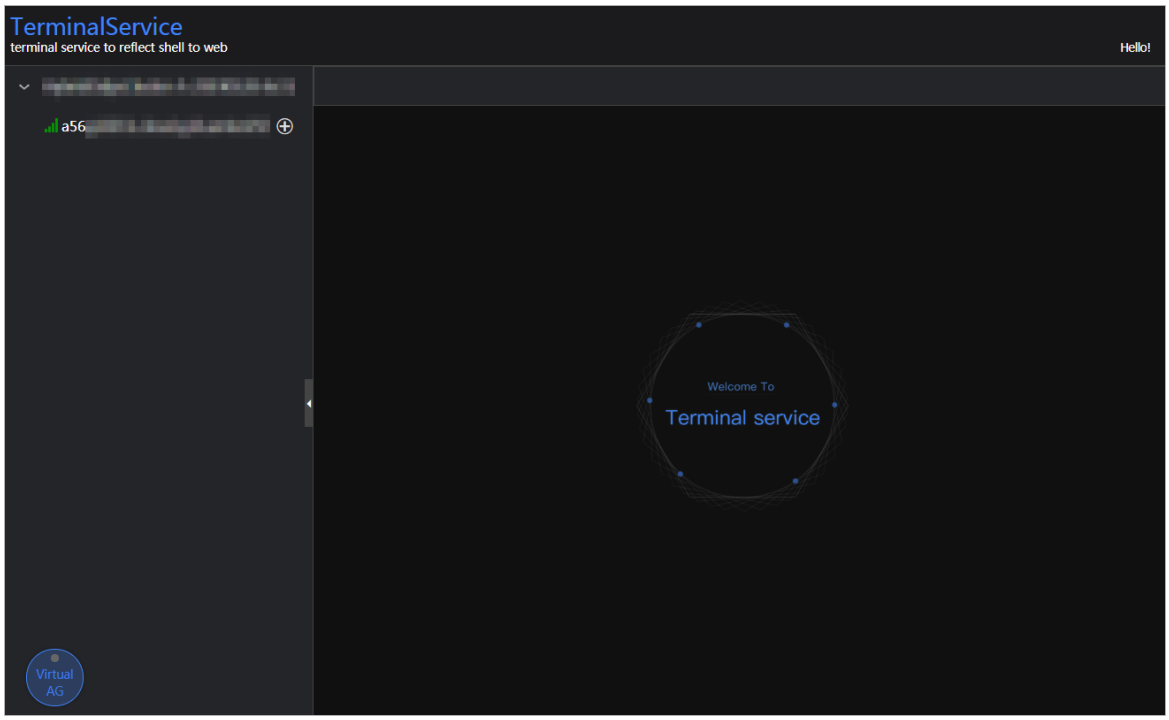

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

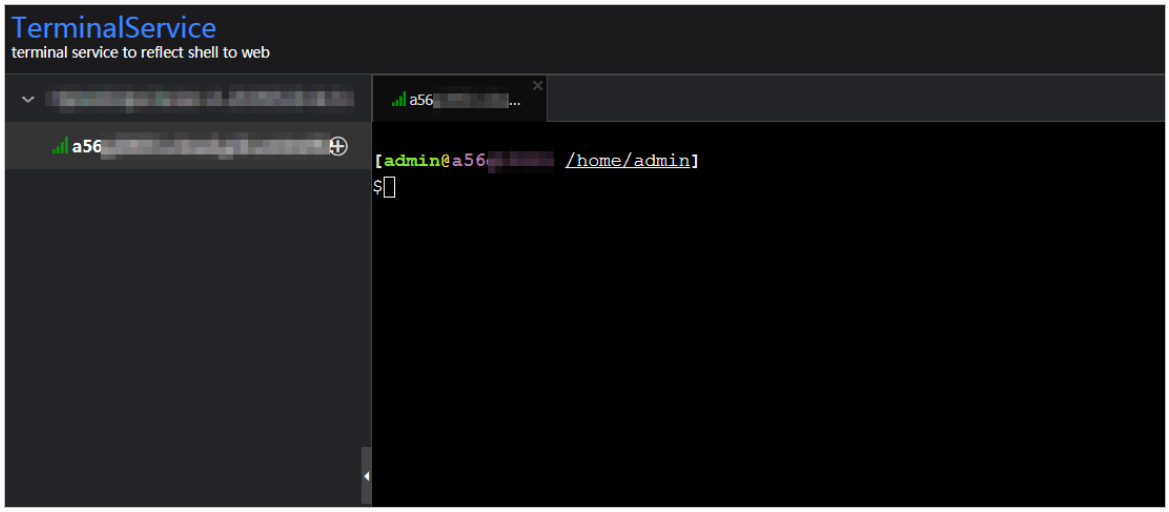

Run a checker again

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

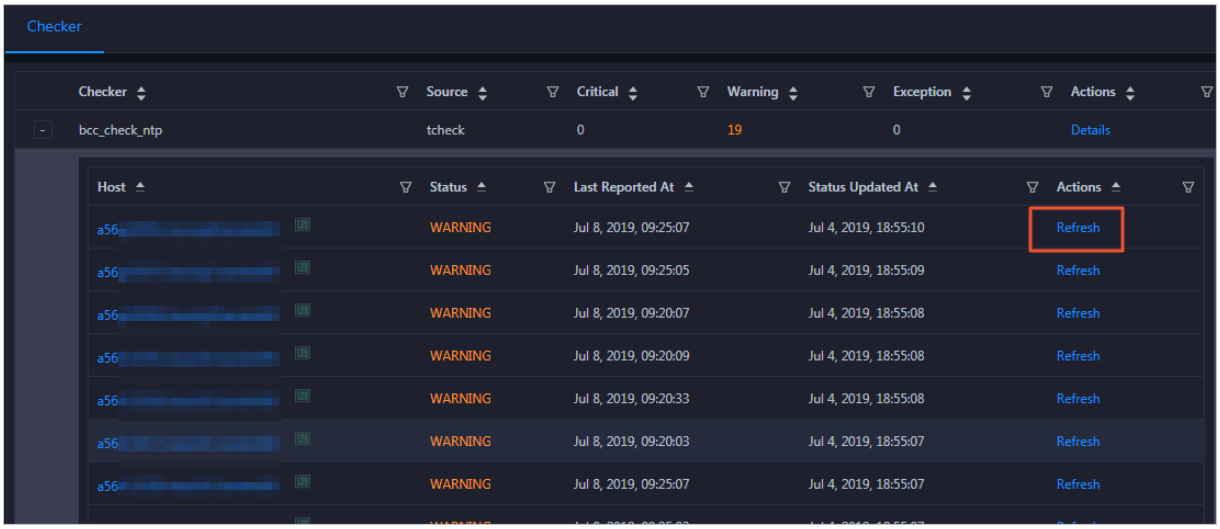

6.3.1.2.5. Host O&M

6.3.1.2.5.1. 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.

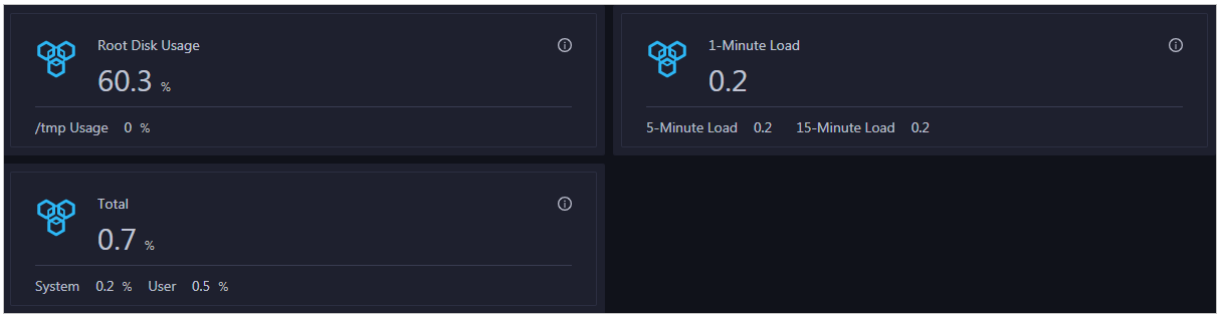

CPU

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

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

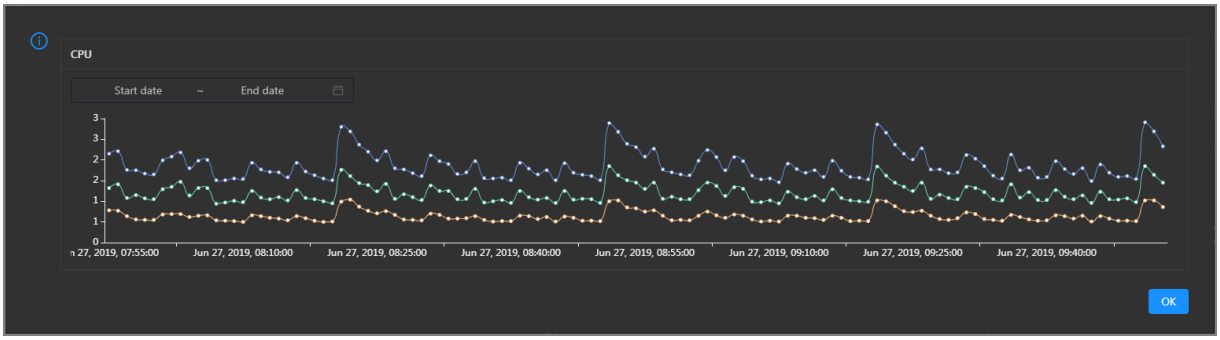

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

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

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

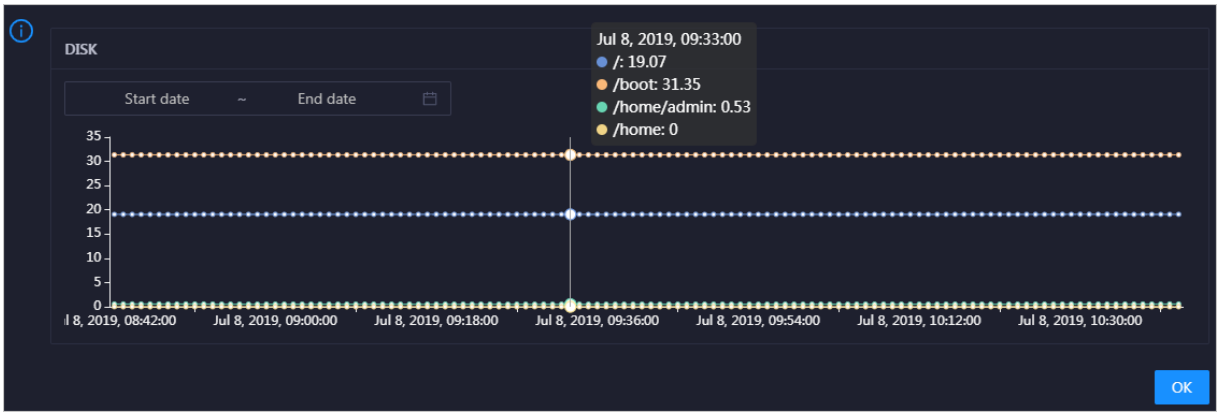

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

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

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

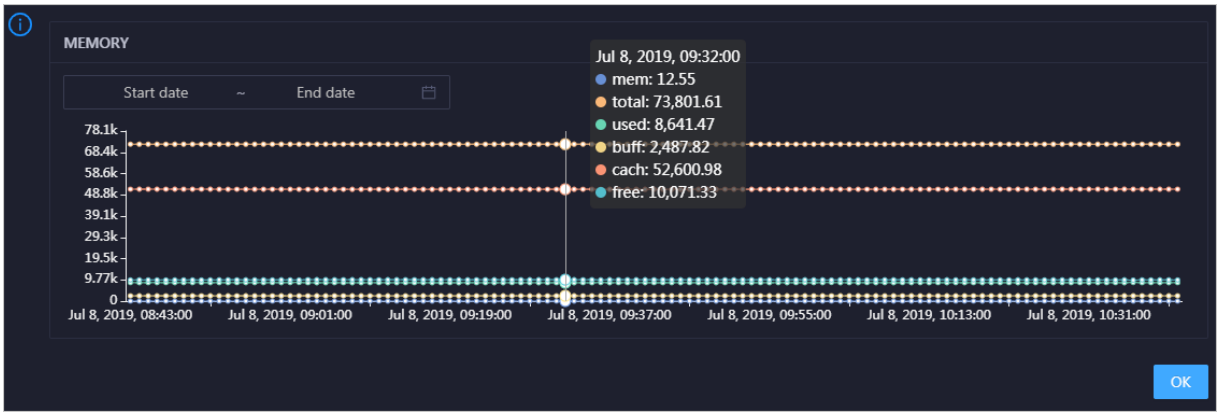

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

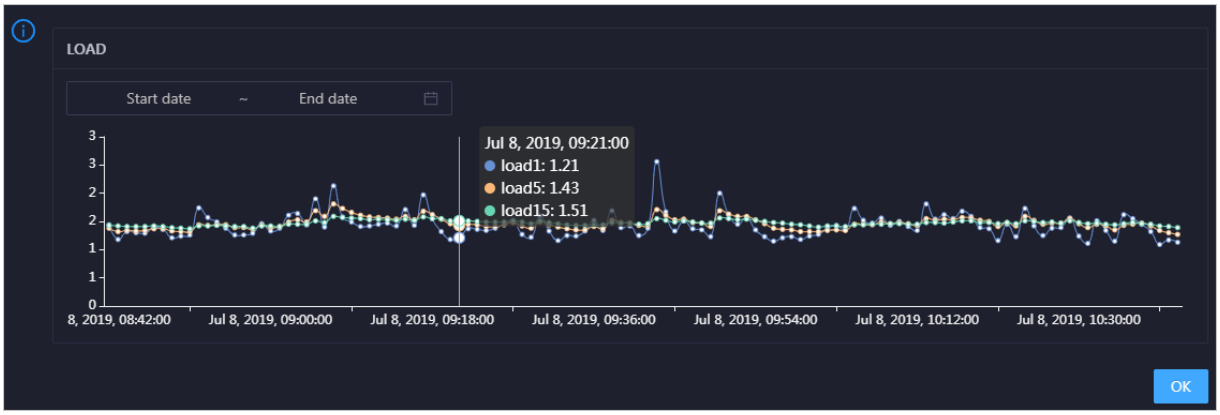

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

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

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

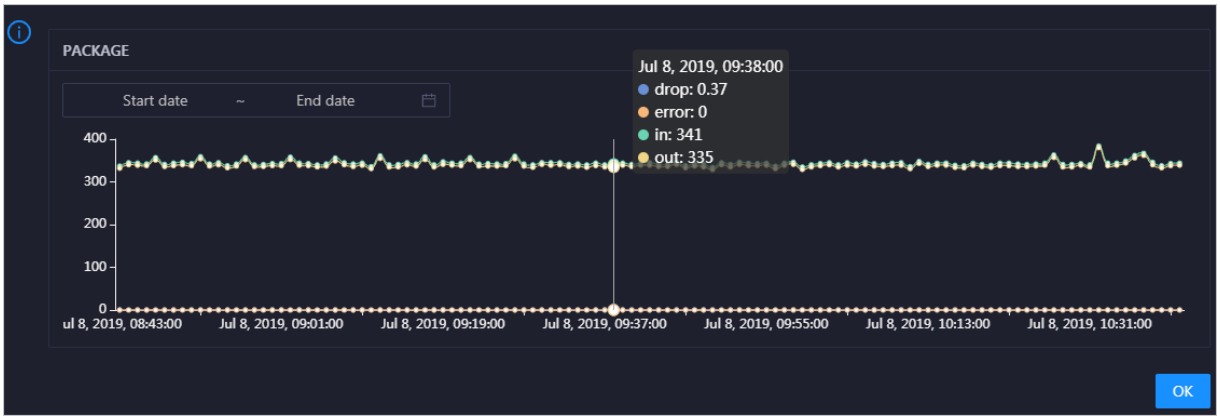

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

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

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

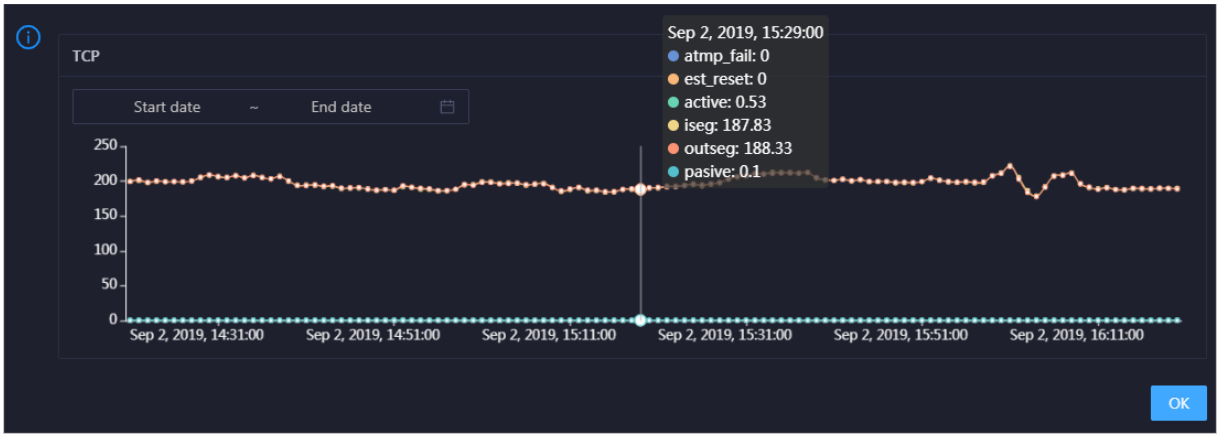

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

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

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

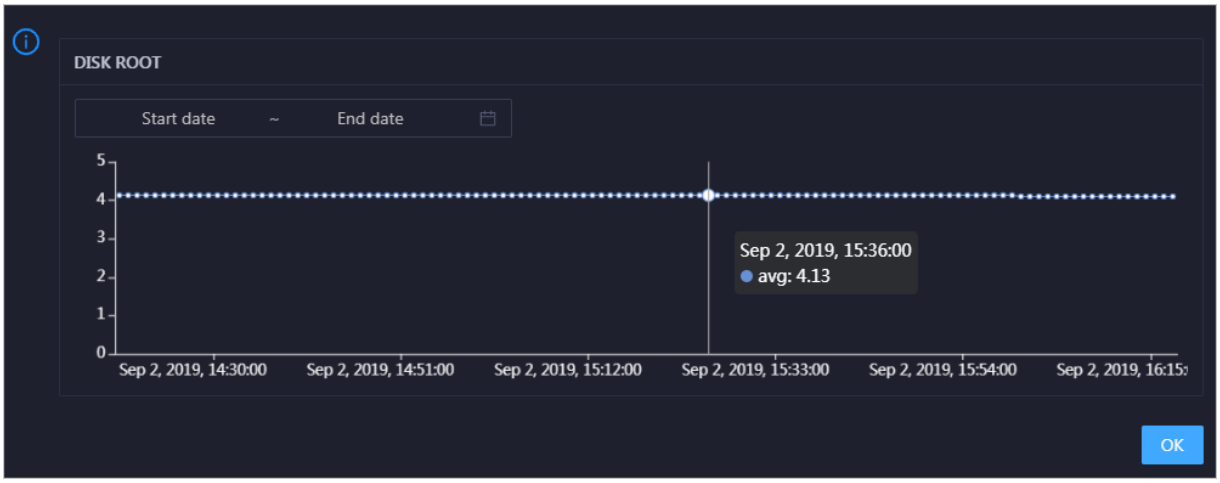

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

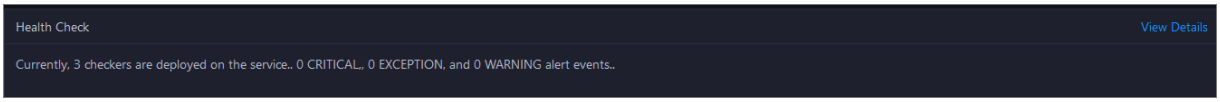

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

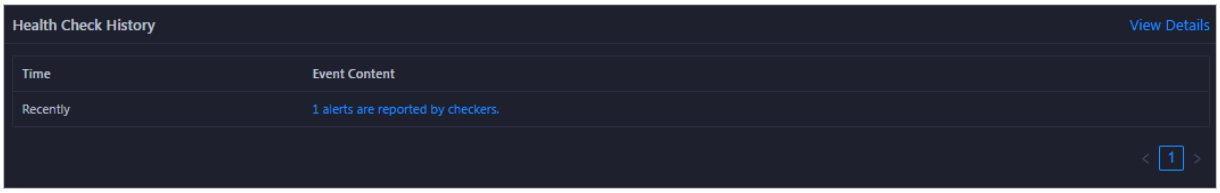

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

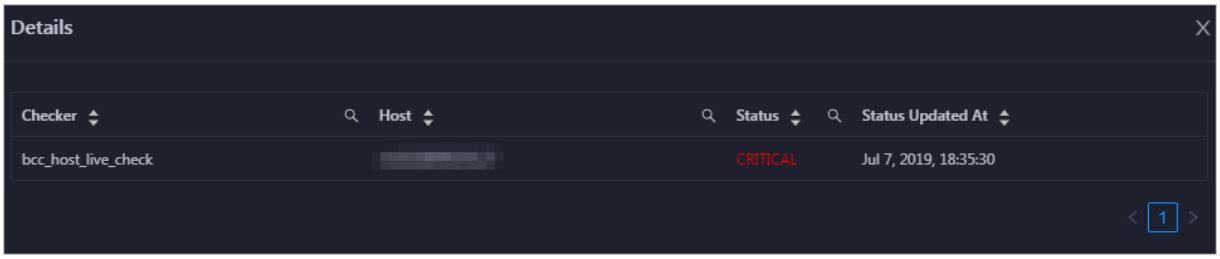

6.3.1.2.5.2. Host health

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.

   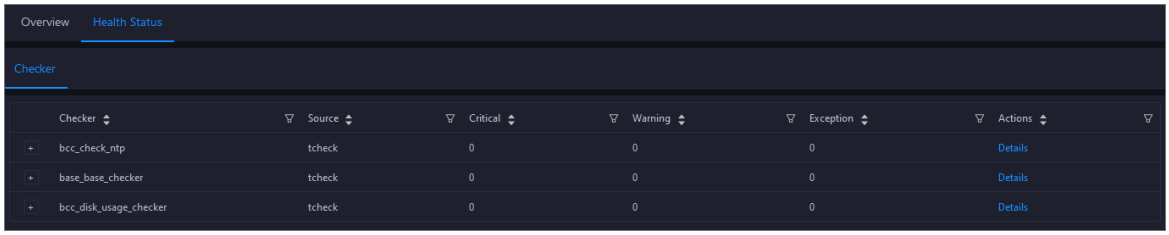

   On the **Health Status** page, you can view all checkers and the check results for the host. The check results are divided into **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results,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

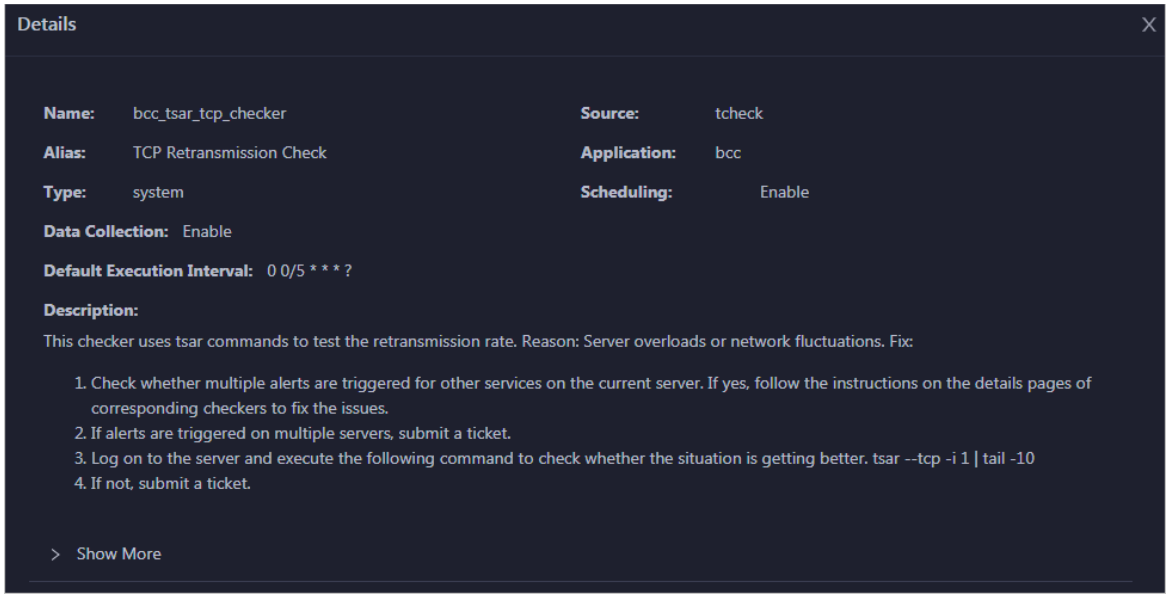

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

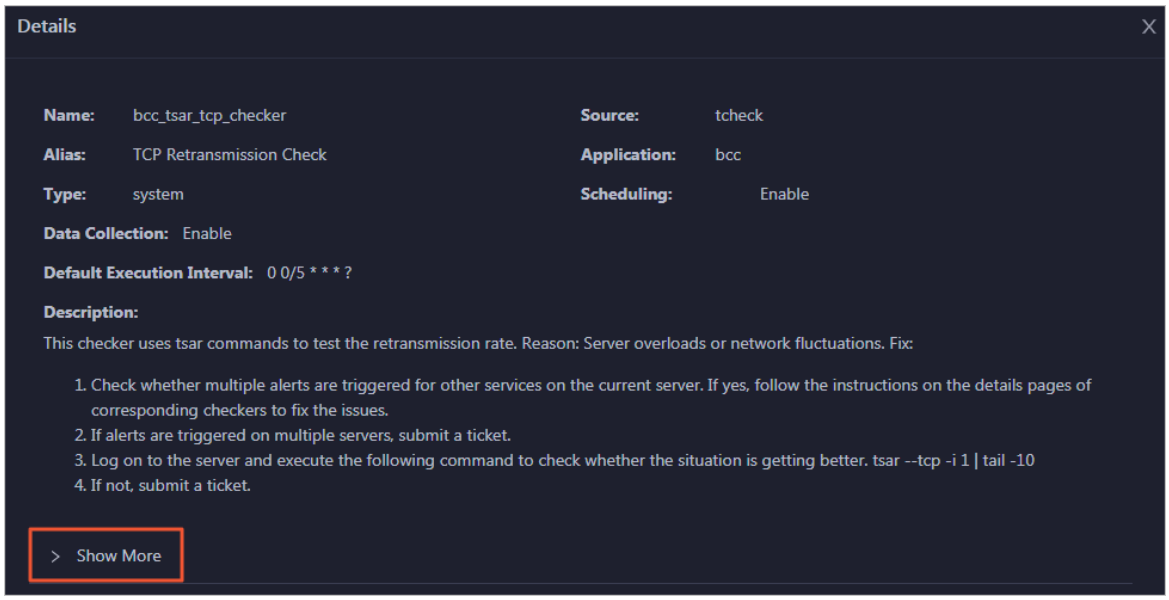

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

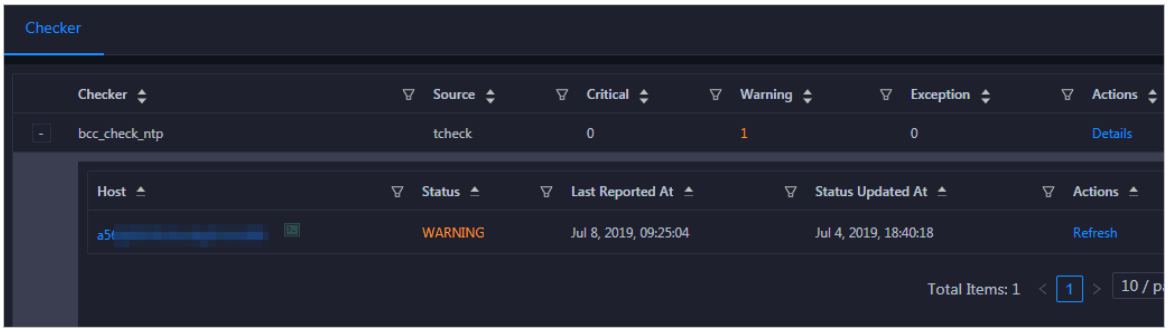

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

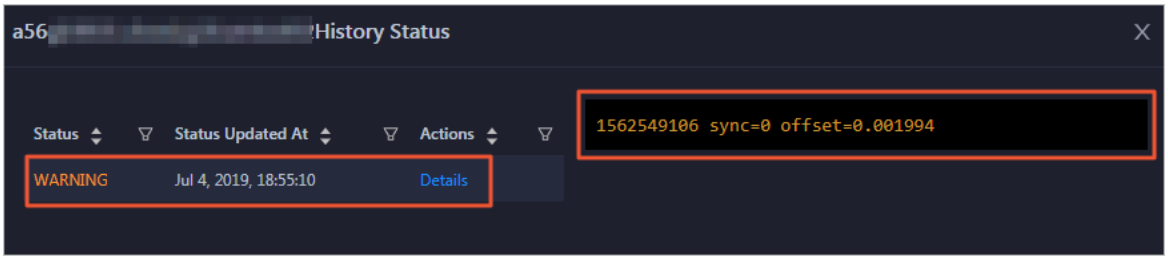

Clear alerts

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.

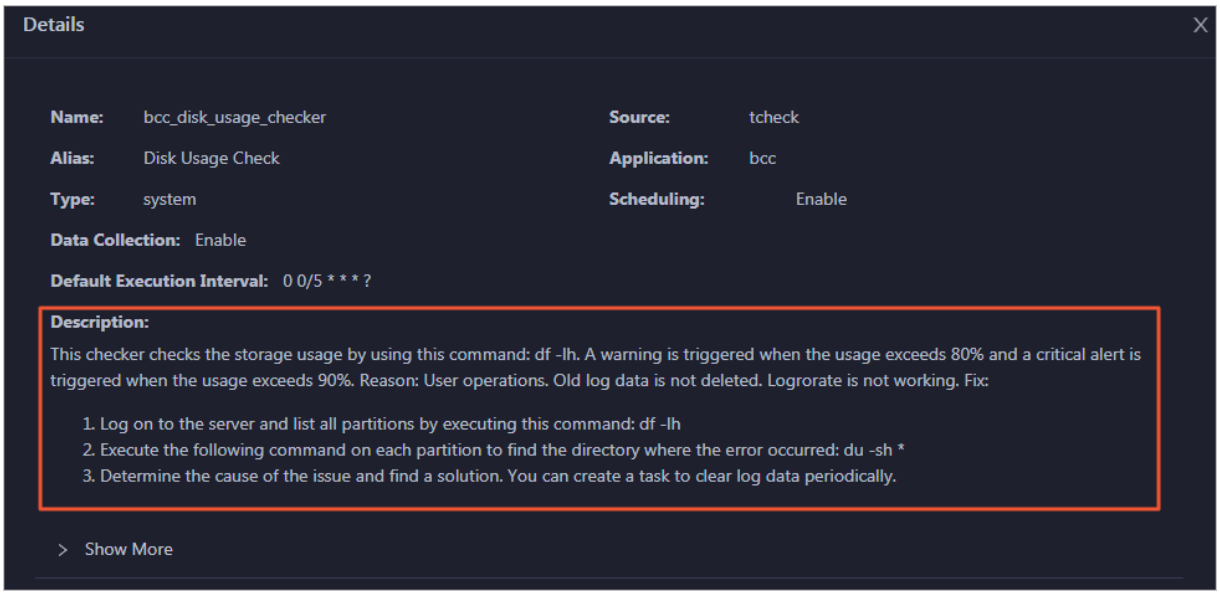

Log on to a host

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.

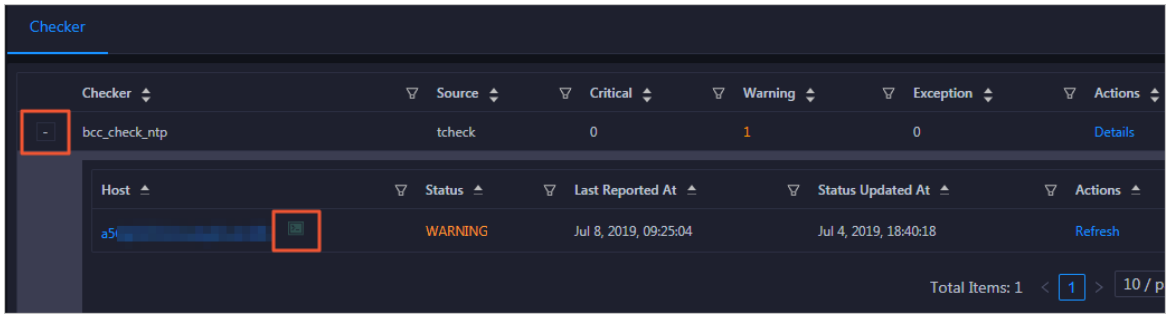

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

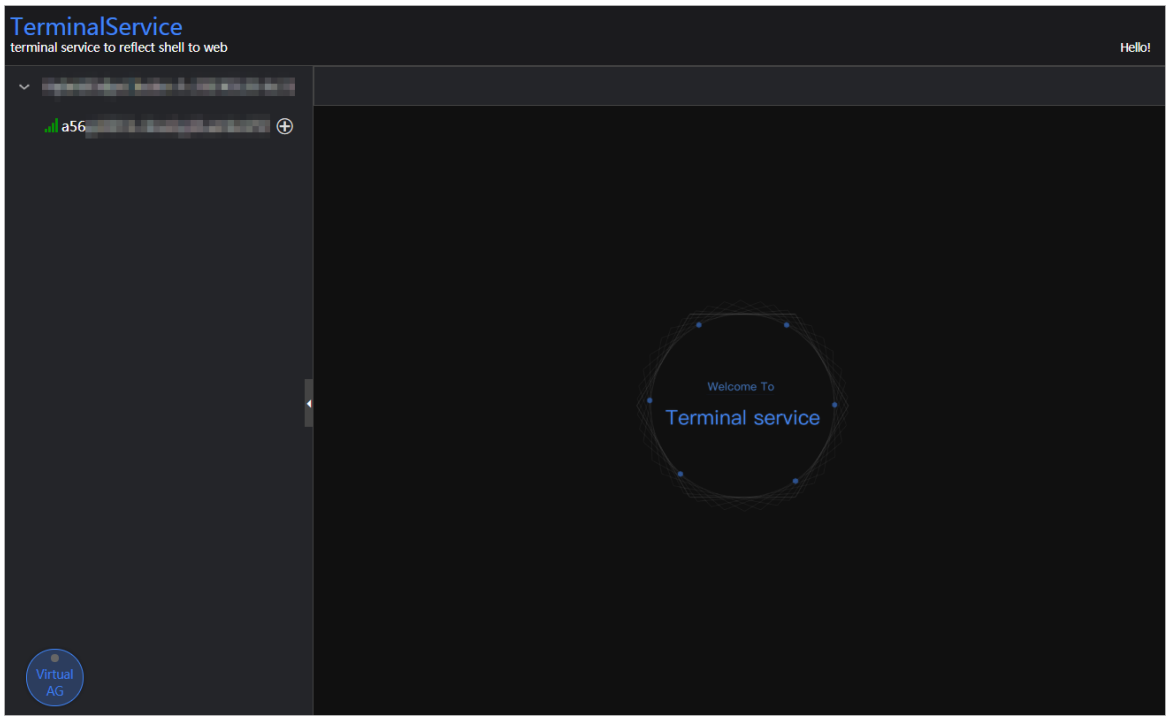

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

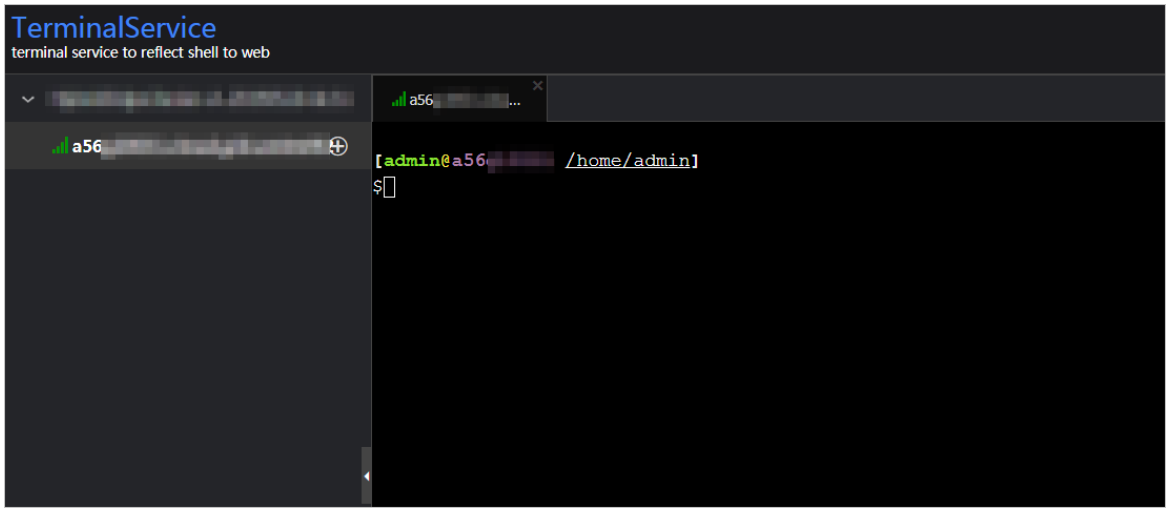

Run a checker again

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.

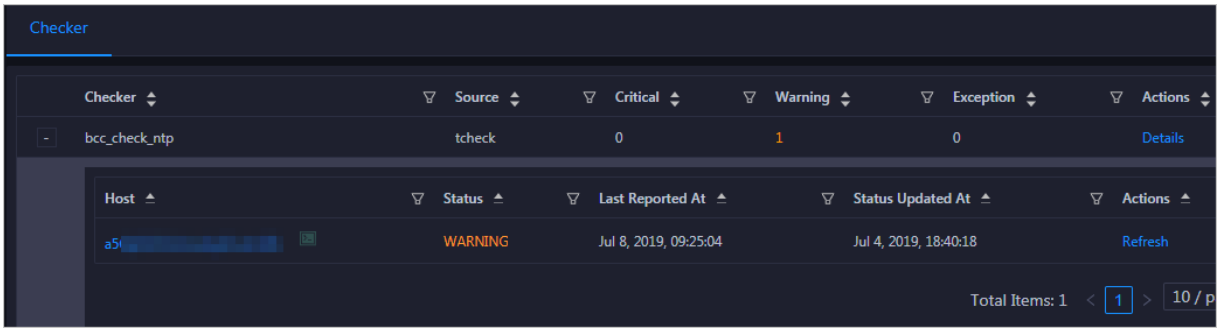

6.3.1.3. Common administration tools and commands

6.3.1.3.1. 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.

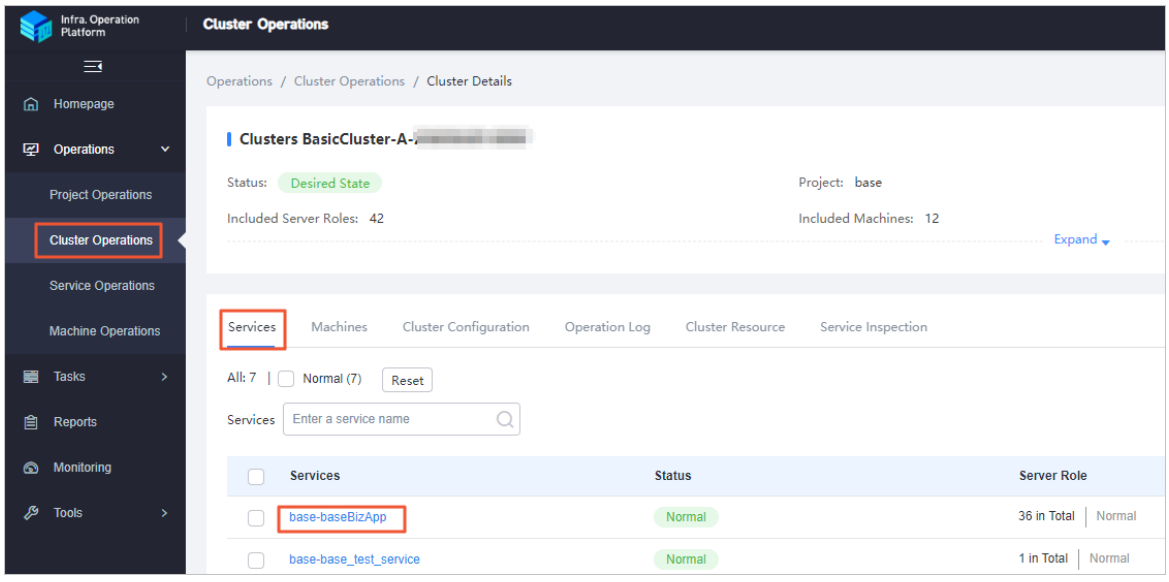

5. Click the **base-baseBizApp** service. On the **Machine Details** page, view information about the host where the service resides.

6.3.1.3.2. 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.

6.3.1.3.3. 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`.

6.3.1.3.4. 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**.

6.3.1.3.5. 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**.

6.3.1.3.6. 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**.

6.3.1.3.7. 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`

6.3.1.3.8. 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;
```

6.3.1.4. Process daily administration operations

6.3.1.4.1. Daily check

6.3.1.4.1.1. 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.

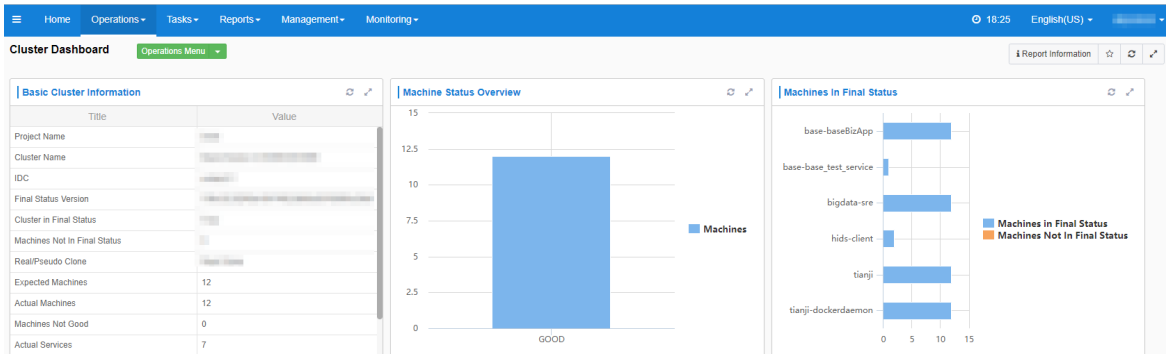

If only blue is shown in the Machine Status Overview section, all servers in the current cluster are running properly. If yellow is shown in the Machine Status Overview section, errors occur on servers.

6.3.1.4.1.2. 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**.

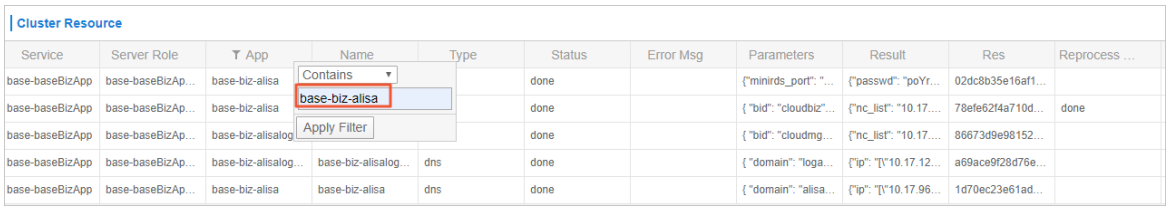

7. Filter services whose type is **db** in the same way.
8. Find the target service, right-click the information in the **Result** column, and then select **Show More** to view the endpoint, username, and password for logging on to the database.

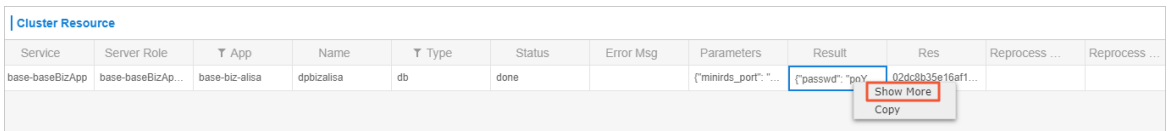

9. Connect to the database and run the following MySQL command to query the node information:

```
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.

6.3.1.4.1.3. 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.

6.3.1.4.2. 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`.

6.3.1.4.3. 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**.

  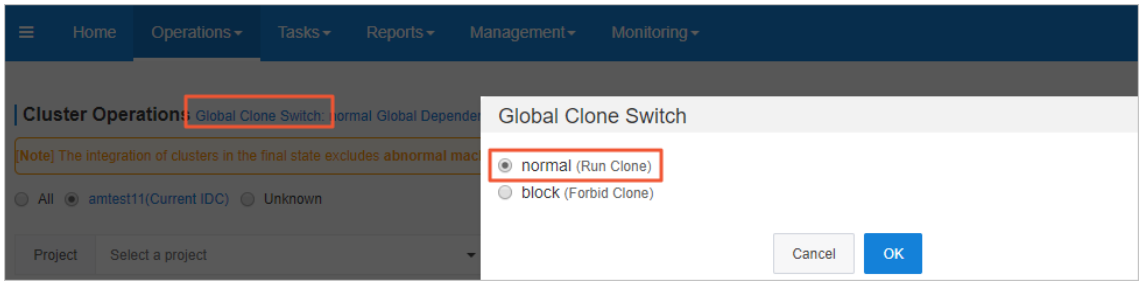

  v. Click **OK**.

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.

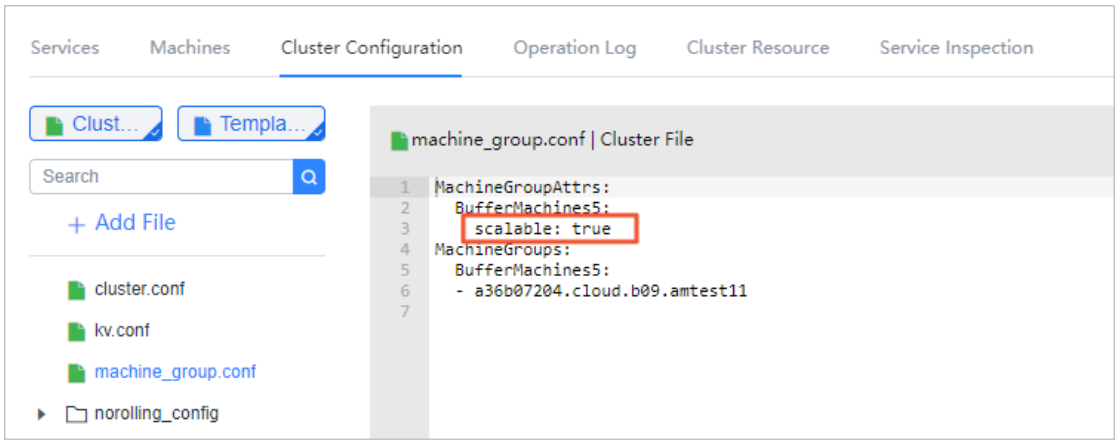

iv. Run the following commands on the OPS1 server to scale in the buffer cluster:

```
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**.

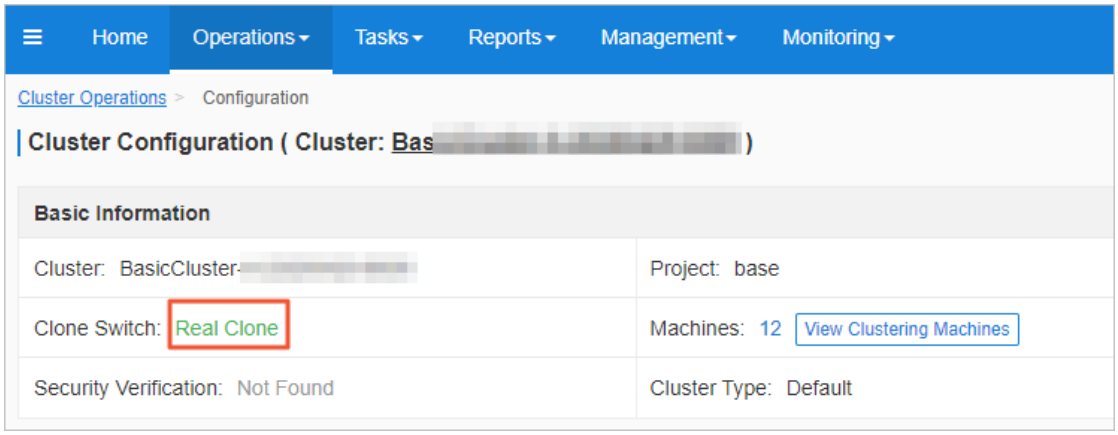

iv. Run the following commands to perform scaling. A rolling task is triggered after you run these commands.

```
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.

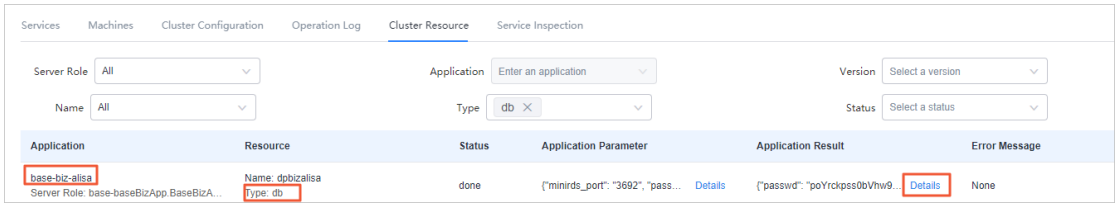

Connect to the database by using a MySQL command, and run the `select * from alisa_node;` command. The information of all servers that run the base-biz-gateway service appears.

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.

6.3.1.4.4. 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/servervtl stop` command.
4. Run the `ps -ef|grep java` command to check whether any process is running on the server. If any process is running, run the `kill -9 [process ID]` command to terminate the process.
5. Delete the program directories from the server.

   Clean up the disks of the server. Skip this step if you want to clone the server.

   ```
   #rm -rf /home/admin/*
   ```

   ```
   #rm -rf /opt/taobao/tbdpapp/
   ```

Move servers from the base-biz-gateway cluster to the default cluster in Apsara Infrastructure Management Framework

1. Log on to the ops1 server and run the following commands to remove a server from the base-biz-gateway cluster:

   ```
   cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
   ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.conf clt2.conf (After you run this command, if a message appears indicating that a symbolic link already exists, proceed with the next command.)
   ./tianji_ops_tool.py contract_nc -c [clusterName] -l [machineList] --config tianji_clt.conf -s [SRGname]
   ```

   The parameters are described as follows:

   - -c: Required. Set this parameter to the name of the base cluster to be scaled in. To obtain the cluster name, choose Operations > Cluster Operations in the top navigation bar and select base from the Project drop-down list.
   - -l: Required. Set this parameter to the hostname of the server to be removed. Separate multiple hostnames with commas (,).
   - -s: Required. Set this parameter to the SRG name of the server to be removed. Find the machine_group.conf file among the configuration files of the base cluster. In this file, find the SRG of the server to be removed.
   - -config: Required. Set this parameter to tianji_clt.conf.
2. After you run the preceding command, check whether the scale-in operation succeeds in Apsara Infrastructure Management Framework.

   Go to the Cluster Operation and Maintenance Center of the base cluster.
3. On the **Cluster Operation and Maintenance Center** page, check the number of servers that are being removed.
4. Click the number next to Machine: in: to identify the status of the servers that are being removed.

   If the scale-in operation succeeds, the number of servers that are being removed decreases to zero. Otherwise, check the server status on this page.

   You can follow the preceding steps to scale in a node cluster by moving servers to the default cluster in Apsara Infrastructure Management Framework. The following section describes how to remove servers from Apsara Infrastructure Management Framework.

Remove servers from Apsara Infrastructure Management Framework

1. In the top navigation bar, choose **Operations > Machine Operations**.

2. On the Machine Operations page that appears, click **Machine Online/Offline** in the upper-right corner.
3. In the Machine Online/Offline dialog box that appears, click Remove Machine.
4. On the Remove Machine tab, search for the server to be removed by hostname in the left-side Enter Machine List section. You can only remove servers in the default cluster.
5. Confirm the information of the server and click **Clear Machines** to remove it.

Verify the server removal result

1. Check whether the server is moved to the default cluster in Apsara Infrastructure Management Framework.

   In the top navigation bar, choose **Operations > Machine Operations**. On the Machine Operations page that appears, search for the target server by hostname and check whether it is in the default cluster.

2. Check whether the server is removed from the default cluster.

   In the top navigation bar, choose **Operations > Machine Operations**. On the Machine Operations page that appears, search for the server by hostname. If you cannot find the server in the search results, the server is removed.

3. To check whether the server is removed from the default cluster, run the following command on the ops1 server to call the GetMachineInfo operation:

   ```
   curl http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=$hostname
   ```

6.3.1.4.5. 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.

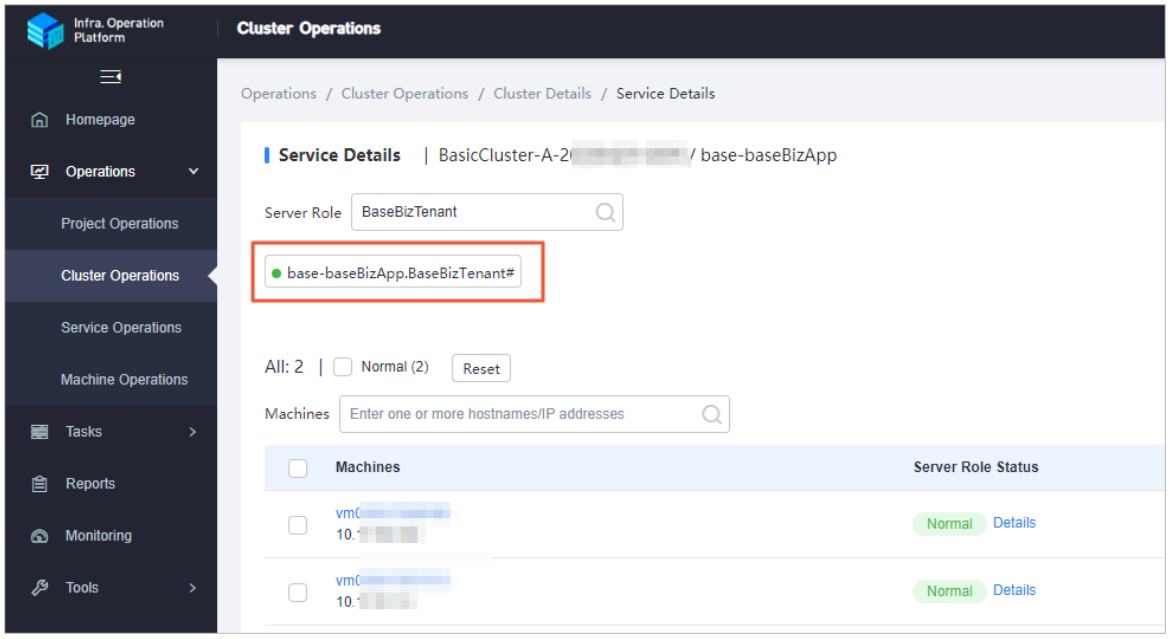

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep tenant` and press Enter to view the ID of the server.

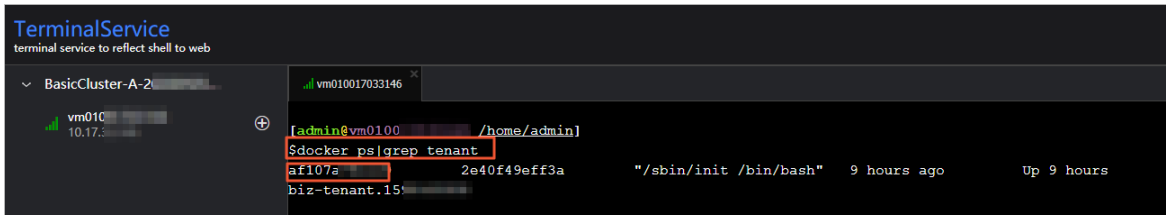

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **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-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.

6.3.1.4.6. 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
   ```

6.3.1.4.7. 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.

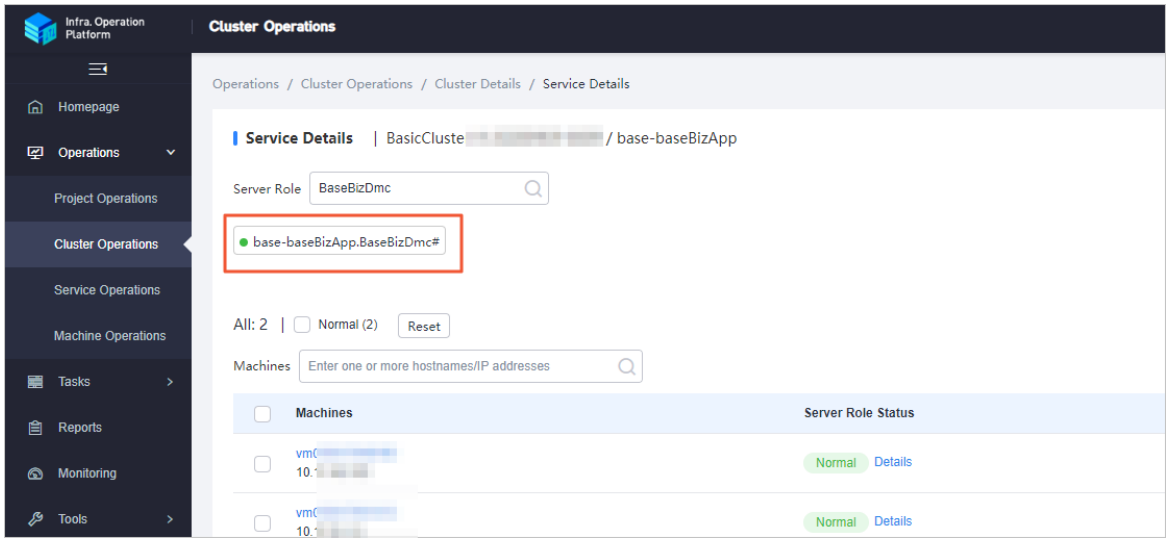

2. Click the name of the target server. The Machine Details page appears.
3. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
4. Enter `docker ps|grep dmc` and press Enter to view the ID of the server.

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.

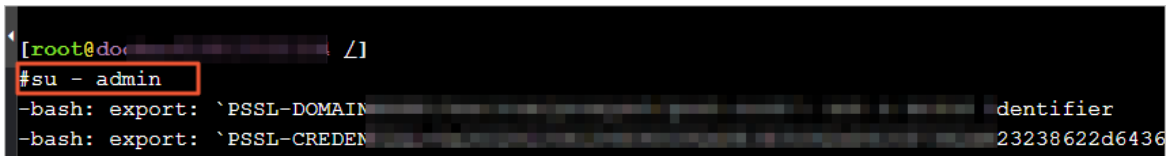

7. Enter `/home/admin/base-biz-dmc/bin/jbossctl restart` and press Enter to restart the base-biz-dmc service.

When the **OK** and **NGINX start Done** information appears, the base-biz-dmc service is restarted.

6.3.1.4.8. 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.

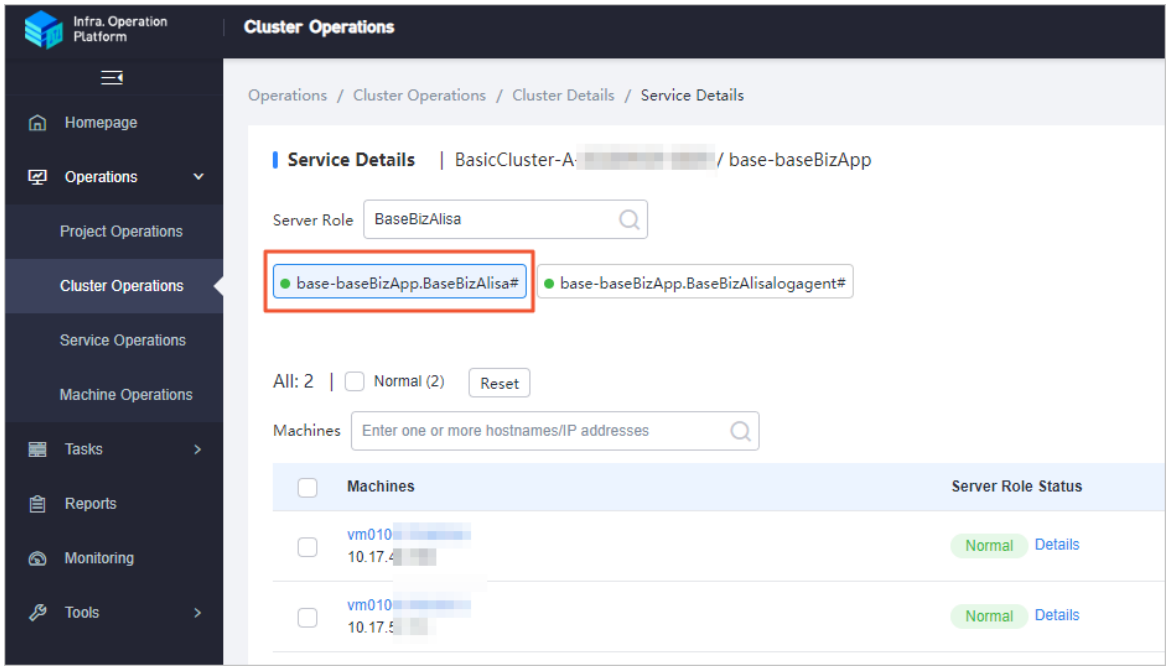

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep alisa` and press Enter to view the ID of the server.

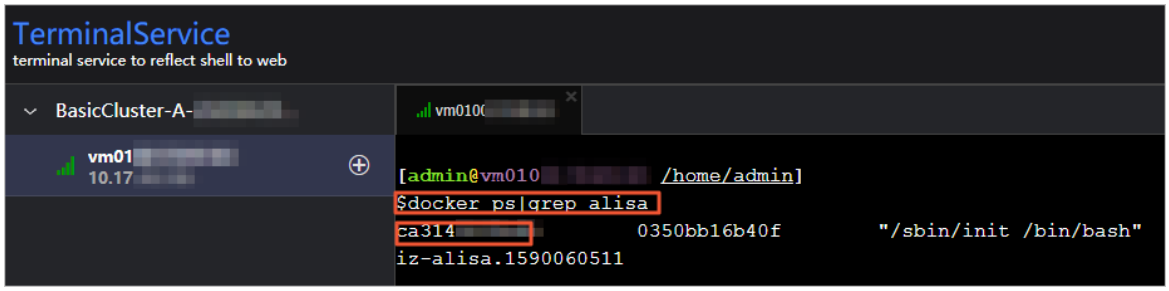

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

   > **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.

6.3.1.4.9. 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.

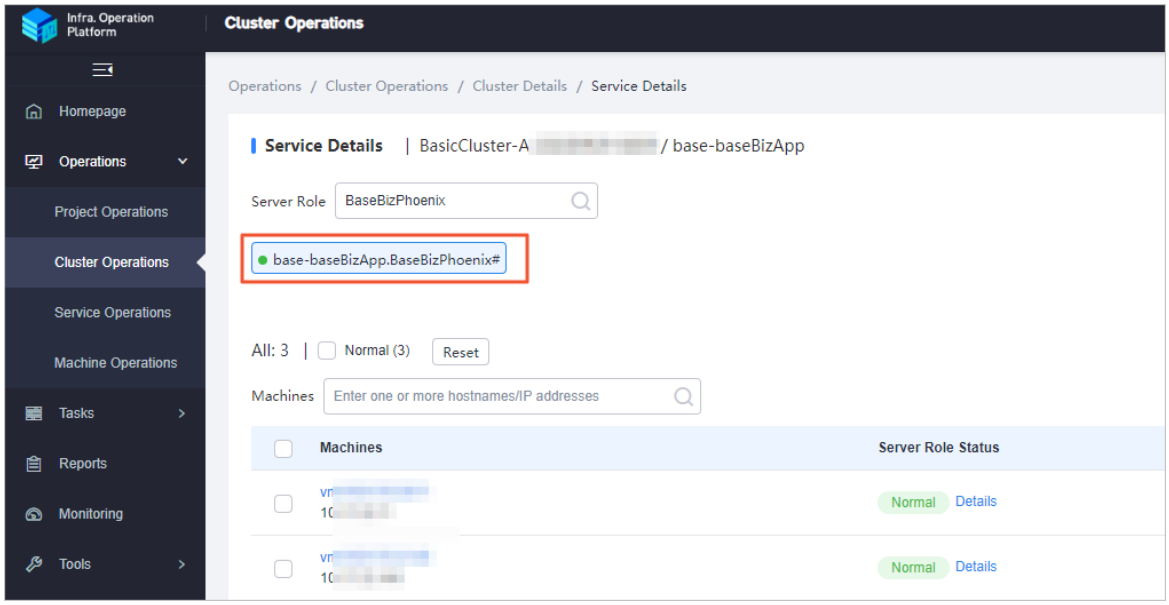

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The configuration tab appears on the right-side of the page.
5. Enter `docker ps|grep phoenix` and press Enter to view the ID of the server.

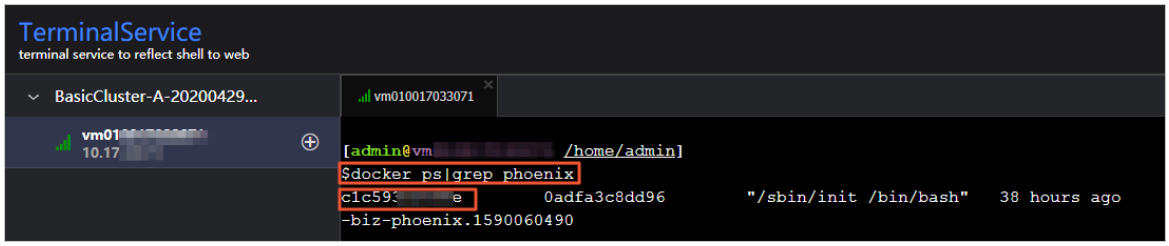

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **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.

6.3.1.4.10. 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.

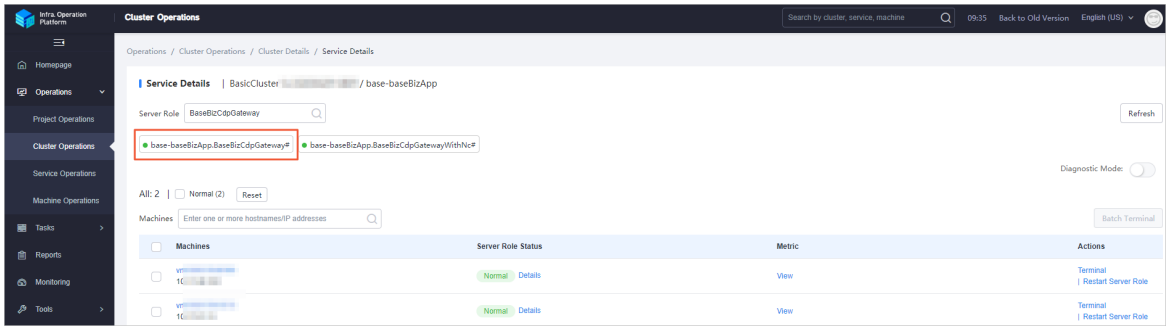

2. Click the name of the target server. The Machine Details page appears.
3. Click **Terminal** in the upper-right corner of the page.
4. In the left-side navigation pane of the **TerminalService** page, click the server name. The

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.

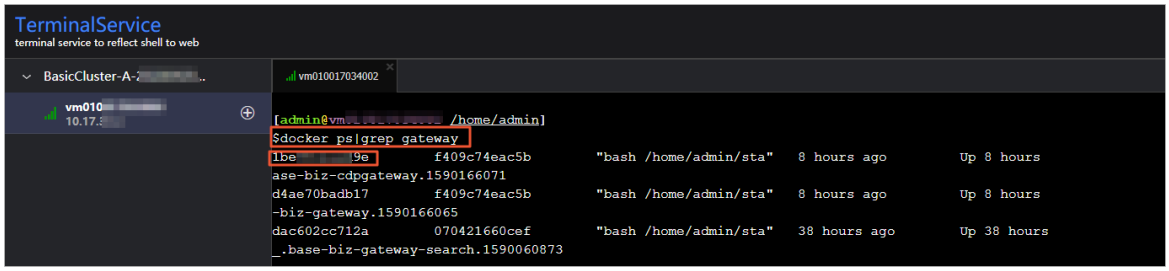

6. Enter `docker exec -it yourID bash` and press Enter to enter the corresponding container.

> **Note** The *yourID* parameter specifies the container ID.

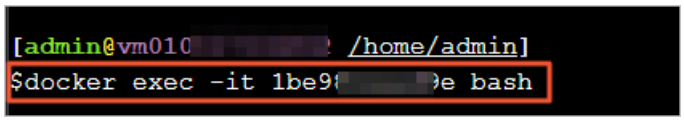

7. Enter `su - admin` and press Enter to switch to the admin user.
8. Enter `/home/admin/alisatasknode/target/alisatasknode/bin/serverctl restart` and press Enter to restart the base-biz-gateway service.

6.3.1.4.11. 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.

6.3.1.4.12. 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.

6.3.1.4.13. Configure a resource group in a multi-region environment

In a multi-region environment, nodes in all regions are run on the resource group named sys_default in the central region by default. This is the cause of resource preemption and node backlog. To resolve the issues, you must isolate resources and enable level-1 scheduling nodes in your region to run on your own resource groups.

Prerequisites

- Idle Elastic Compute Service (ECS) instance resources are available in a region rather than the central region. You can view the idle resources in the Apsara Infrastructure Management Framework console.
- Resource groups are created in DataWorks and resources are isolated.
- A custom resource group is specified as the default resource group in your workspace.
- Historical nodes that are run on the original default resource group are migrated to the new resource group. This is a high-risk operation.

Create a resource group

1. Go to the Schedule Resources page in DataWorks.
   i. Log on to the DataWorks console.
   ii. Click the icon in the upper-left corner and choose **All Products > Project Management**.
   iii. In the left-side navigation pane, click **Schedule Resources**.
2. On the **Schedule Resources** page, click **Add scheduling resources** in the upper-right corner.
3. In the **Add scheduling resources** dialog box, set the **Resource Name** and **Belonging workspace** parameters.

   In this example, the resource name is region_test. The resource group belongs to a tenant that is a level-1 department. All workspaces that are created by the tenant can use this resource group, whereas other tenants need to create their own resource groups.

   After a resource group is created, a cluster that corresponds to the resource group is also created. In the background database, you can view the resource group in the alisa_group table and view the cluster in the alisa_cluster table.
4. Click **Confirm**.

   After you create the resource group, you can log on to the dpbizalisa background database to view detailed information about the resource group and cluster.

```
select group_name,group_display_name,cluster_name,max_slot where group_display_name='xxxx';
select cluster_name,username,password from alisa_cluster where cluster_name ='xxxx';
// max_slot: the number of slots that are assigned to the current resource group. The number is determi
ned based on actual resources and requirements. You must set the number of slots to a proper value.
// group_display_name: the name of the custom resource group.
// cluster_name: the name of the cluster that corresponds to the custom resource group.
```

Migrate ECS instances to a new cluster

1. To migrate an ECS instance, run the following command:

```
update alisa_node set cluster_name='xxxxxx' where node_name='${hostname}';
```

2. To set the number of slots in the resource group to a proper value, run the following command:

```
update alisa_group set max_slot = xxxx where group_display_name='region_test';
```

3. Restart the base-biz-alisa service. For more information, see Restart the base-biz-alisa service.

   Note Both base-biz-alisa services need to be restarted. After one base-biz-alisa service is restarted, you can restart the other.

Specify the custom resource group as the default resource group in your workspace

1. Log on to the dwphoenix background database to view data.

```
select resource_group_id,name,identifier,is_default,project_env from phoenix_resgroup where name
= 'region_test';
// resource_group_id: the ID of the resource group in the phoenix database.
// name: the name of the resource group.
// identifier: the name of the resource group in the alisa database.
// project_env: the production environment or the development environment.
select * from phoenix_app_resgroup where app_id=xxxx;
// app_id: the ID of the project. The app_id parameter is equivalent to the project_id parameter that is d
escribed in other topics.
```

2. Specify a new resource group as the default resource group in your workspace.

```
create table phoenix_app_resgroup_20200108_bak as select * from phoenix_app_resgroup; // Back up t
he table.
update phoenix_app_resgroup set is_default=1 where is_default=0 and app_id=xxxx;
update phoenix_app_resgroup set is_default=0 where app_id=xxxx and resource_group_id=1;
select * from phoenix_app_resgroup where app_id=xxxx; // View data again after the change.
```

> **Note** After the change, one workspace has only two data entries whose is_default value is 1. If the preceding result is not displayed, the change is invalid. You must check the operations you have performed.
>
> The resource group belongs to a tenant so that you must change the default resource group for all workspaces that are created by the tenant.

Change the resource group for multiple nodes at a time

You can change the resource group for multiple nodes at a time on the Operation Center page. If a large number of nodes exist, you must perform the change operation multiple times on the page until the resource group is changed for all the nodes. If you do not want to change the resource group in that way, you can change the resource group for all the nodes at a time in the scheduled database. However, this is a high-risk operation. To change the resource group in the database, you must perform the operations based on your business requirements. You must write the code based on the data in the current environment by following the template. In this example, the resource group for nodes in the datam4 workspace is changed to the resource group named region_test in the following way:

1. Log on to the dwphoenix database that corresponds to the base-biz-phoenix service.
2. Query the ID of the resource group named sys_default. The obtained value is used as the default ID.

```
select resource_group_id from phoenix_resgroup where name='Default Group';
```

```
mysql> select resource_group_id from phoenix_resgroup where name='Default Group';
+-------------------+
| resource_group_id |
+-------------------+
|                 1 |
|                 1 |
+-------------------+
2 rows in set (0.00 sec)
```

3. Query the ID of the resource group named region_test. The obtained value is used as the region_test_id value.

```
select resource_group_id from phoenix_resgroup where name='region_test';
```

```
mysql> select resource_group_id from phoenix_resgroup where name='region_test';
+-------------------+
| resource_group_id |
+-------------------+
|            110001 |
|            110001 |
+-------------------+
2 rows in set (0.00 sec)
```

4. Query the ID of the required workspace. The obtained value is used as the app_id value.

```
select app_id from phoenix_app_config where name="datam4";
```

5. Change the resource group for historical nodes in the workspace that you want to change to the new resource group named region_test.

```
create table phoenix_node_def_20200108_bak as select * from phoenix_node_def; // Back up the table.
update phoenix_node_def set resgroup_id={region_test_id} where resource_group_id={Default ID} and app_id={app_id};
```

This way, the resource group for historical nodes in the datam4 workspace is changed to the resource group named region_test. After the resource group is changed, test whether all the historical nodes can be run as expected.

6.3.1.4.14. Configure a resource group for Data Integration in a multi-region environment

This topic describes how to configure a resource group for Data Integration in a multi-region environment.

Procedure

1. On the Data Integration page in DataWorks, create a custom resource group.

   This custom resource group must be real and can be used. In this example, you can specify the Elastic Compute Service (ECS) instance on which the region_group custom resource group is hosted as the default ECS instance to run nodes.

   i. Log on to the DataWorks console.

   ii. Click the ☰ icon in the upper-left corner and choose **All Products > Data Aggregation > Data Integration**.

   iii. In the left-side navigation pane, click **Custom Resource Group**.

   iv. Click **Add Resource Group** in the upper-right corner.

2. In the **Add Resource Group** wizard, perform the following steps:

   i. In the **Create Resource Group** step, set the **Resource Group Name** parameter.

   > **Note** The name can contain letters, digits, and underscores (_), and must start with a letter.

ii. Click **Next**.

iii. In the **Add Server** step, set the parameters as required.

iv. Click **Next**.

v. Perform the steps that are described in the **Install Agent** step.

> **Note** If an error occurs when you run the `install.sh` command or you need to run it again, run the `rm -rf install.sh` command in the same directory as the `install.sh` command to delete the generated file. Then, run the `install.sh` command again.
>
> The commands to run during the installation and initialization process differ for each user. Run relevant commands based on the instructions on the initialization interface.

vi. Click **Next**.

vii. In the **Test Connection** step, click **Refresh** and check the status of the ECS instance.

viii. Click **Complete**.

Resource groups that you have created on the Data Integration page can be used only in workspaces to which they belong. If you create a resource group in Workspace A, this resource group cannot be used in Workspace B. If you need to use a resource group in Workspace B, perform the preceding steps to create one in Workspace B.

3. View the newly created resource group in the alisa database.

```
select * from alisa_group where group_display_name='region_group'\G;
```

```
mysql> select * from alisa_group where group_display_name='region_group'\G;
*************************** 1. row ***************************
        group_name: 4be29b9408f24bd4be1566131f97afb4
group_display_name: region_group
      cluster_name: 4be29b9408f24bd4be1566131f97afb4
          max_slot: 999
       active_type: 1
       create_time: 2019-11-07 17:01:41
  last_modify_time: 2019-11-07 17:01:41
        group_type: NULL
1 row in set (0.00 sec)

ERROR:
No query specified
```

4. Create an alisa cluster.

   To deploy services in a new region, you must first create an alisa cluster. You can find the dpbizalisa database and execute the following SQL statements in this database:

```
insert ignore into alisa_cluster(
cluster_name,
cluster_display_name,
username,
password,
create_time,
last_modify_time
)
values(
'sys_region_cdp_${regionId}',
'sys_region_cdp_${regionId}',
'private_cluster',
'tyn3n71c2oyhah447m6fnfuq',
now(),
now()
);
```

Replace ${regionId} in the code with an actual value. You can log on to the Apsara Infrastructure Management Framework console and go to the **Cluster Configuration** tab to view the cluster file.

5. Migrate the ECS instance that you have bound to the custom resource group to this cluster.

```
update alisa_node set cluster_name='sys_region_cdp_${regionId}' where node_name='ECS instance na
me';
```

> **Note** You must check the name of the ECS instance in the current region.

6. Mount the cluster that is created in this region to the resource group.

```
update alisa_group set cluster_name='sys_region_cdp_${regionId}' where group_name='4be29b9408f2
4bd4be1566131f97afb4';
```

In the code, group_name indicates the resource group name that is specified in Step 2.

7. Log on to the Apsara Infrastructure Management Framework console and restart the BaseBizAlisa service.
8. On the configuration tab of a sync node, select the name of the required custom resource group, and save and commit the node.

The preceding steps describe how to create a resource group in a region for the first time. If you need to create resource groups for multiple workspaces, repeat steps 1, 2, 3, 6, and 7 to create resource groups as required.

To perform the rollback operation, run the following command:

```
update alisa_group set cluster_name='4be2***' where group_name='4be2***';
```

> **Note** group_name and cluster_name indicate the name and cluster name of the custom resource group.

6.3.1.5. Common issues and solutions

6.3.1.5.1. 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.

6.3.1.5.2. 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';
```

6.3.1.5.3. 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.

6.3.1.5.4. 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.

6.3.1.5.5. 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.

6.3.1.5.6. 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 {} \;
   ```

6.3.1.5.7. 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.

6.4. Realtime Compute

6.4.1. Operations and Maintenance Guide

6.4.1.1. Job status

6.4.1.1.1. Overview

RealtimeCompute 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.

6.4.1.1.2. 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.

6.4.1.1.3. 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.

6.4.1.1.4. 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.

Name	Description
Input RPS	Indicates the number of data records that are read from the source table per second.
Output RPS	Indicates the number of data records that are written into result tables per second.
Input BPS	Indicates the data transmission rate per second, which is measured in bytes per second.
CPU usage	Indicates the CPU usage of the job.
Start time	Indicates the start time of the job.
Running duration	Indicates the duration during which the job has been running.

6.4.1.1.5. Running topology

A running topology shows the execution of the computing logic of Realtime Compute. Each component corresponds to a task. Each data stream starts from one or more sources and ends in one or more result tables. The data flow resembles an arbitrary directed acyclic graph (DAG). To enable efficient distributed execution, Realtime Compute chains operator subtasks together into tasks if possible. Each task is executed by one thread.

You can chain operators together into tasks for the following benefits:

- Reduces thread-to-thread handovers.
- Reduces message serialization and deserialization.
- Reduces data handovers in the buffer zone.
- Increases the overall throughput while decreasing the delay.

An operator describes the computing logic, and a task is a collection of operators.

View mode

The computing logic is visualized in the view mode for an intuitive display, as shown in the View mode figure.

View mode

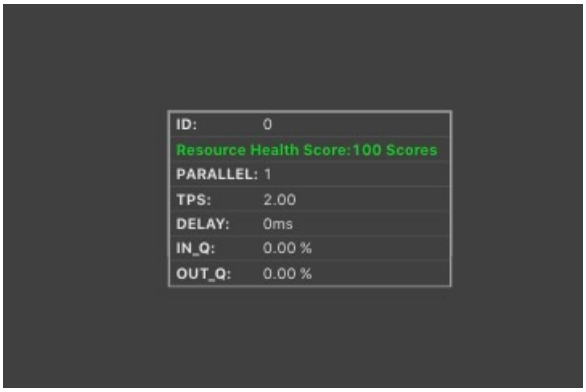

You can view the detailed information about a task by moving the pointer over the task. The Task parameters table describes the task parameters.

Task parameters

Parameter	Description
ID	The ID of the task in the running topology.
PARALLEL	The number of requests that are concurrently processed.
CPU	The CPU usage of a parallel instance for the task.
MEM	The memory usage of a parallel instance for the task.
TPS	The amount of data that is read from the upstream. Unit: blocks per second.
LATENCY	The compute time consumed at the task node.
DELAY	The processing delay that occurs at 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 the Task details figure.

Task details

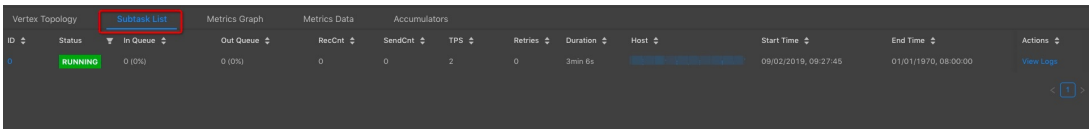

The **Curve Charts** tab provides curve charts that show the metrics of each task, as shown in the Curve charts for task metrics figure.

Curve charts for task metrics

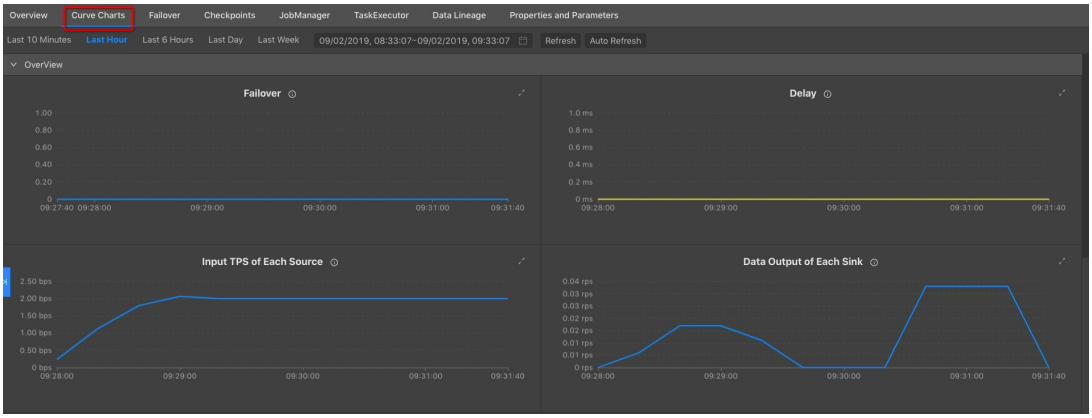

List mode

In addition to the view mode, Realtime Compute allows you to view each task in the list mode, as shown in the List mode figure.

List mode

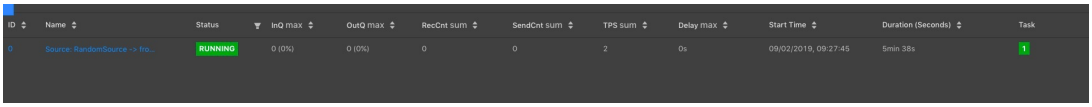

The Task parameters table describes the task parameters.

Task parameters

Parameter	Description
ID	The ID of the task in the running topology.
Name	The name of the task. The name shows the task details.
Status	The status of the task.
InQ max	The percentage of input queues for the task node.
OutQ max	The percentage of output queues for the task node.
RecCnt 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 upstream per second.
Delay max	The processing delay that occurs at the task node.
Task	The status of parallel instances for the task node.
Start Time	The start time of the task node.
Duration (Seconds)	The running duration of the task node.

6.4.1.2. Curve charts

6.4.1.2.1. 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

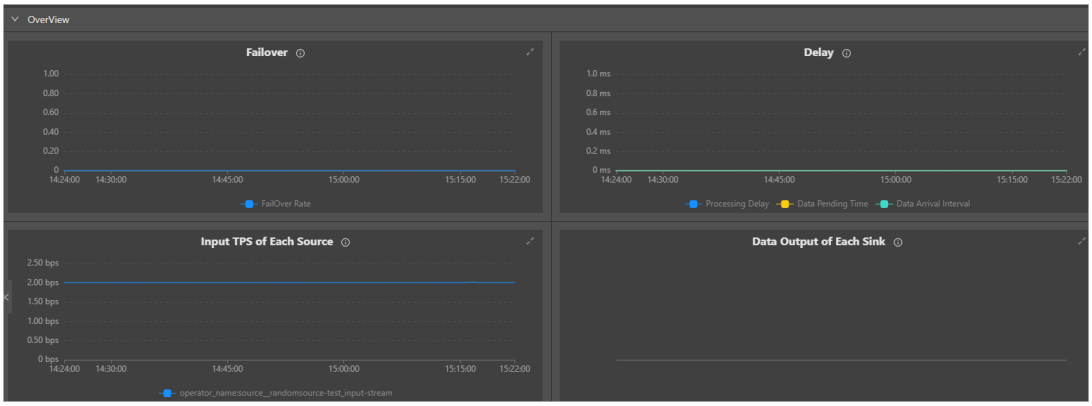

> **Note**
>
> - The metrics shown in this figure are displayed only when the job is in the running status.
> - The metrics are asynchronously collected in the background, which results in delays. The metrics can be collected and displayed only after a job has been running for more than 1 minute.

6.4.1.2.2. Overview

Failover

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 timestamp carried by the streaming data in the source table and the time when Realtime Compute processes the streaming data. If no field in the source table indicates the streaming data timestamp, the delay is calculated based on the system timestamp assigned by source data stores to the data. Examples of the source data stores include DataHub and LogHub. The processing delay shows the timeliness of end-to-end data processing. For example, if the current processing time is 05:00 and the timestamp of the stored data is 01:00, the processing delay is four hours. In this example, the data to be processed was stored at 01:00, which is four hours earlier than the current processing time. The data processing progress is based on the processing delay. For example, if data fails to flow into the DataHub source data store because of faults, the processing delay increases, which causes the processing progress to be delayed. Processing delay shows the processing delay.

Processing delay

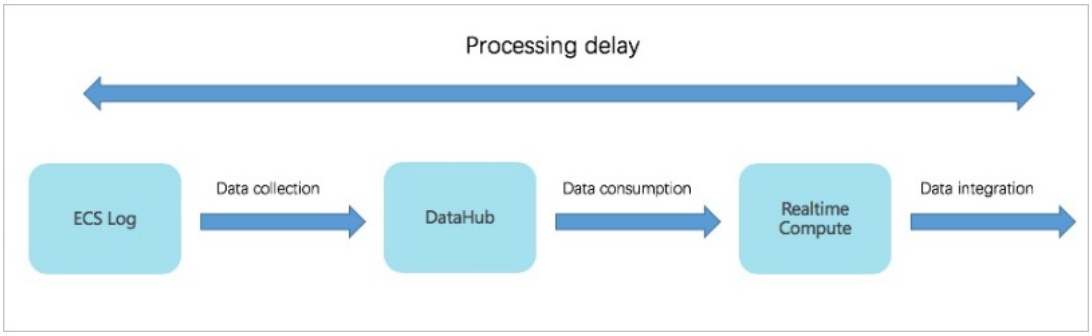

The processing delay can be categorized into the following three types:

- Shortest delay: indicates the shortest processing delay that a shard in each data source experiences.

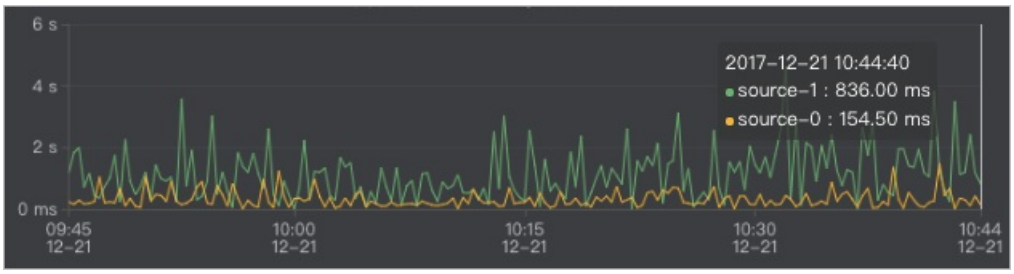

- Longest delay: indicates the longest processing delay that a shard in each data source experiences.

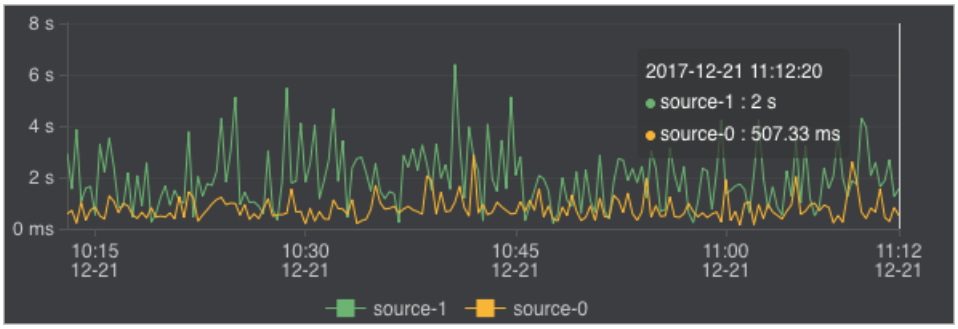

- Average delay: indicates the average processing delay of shards in each data source.

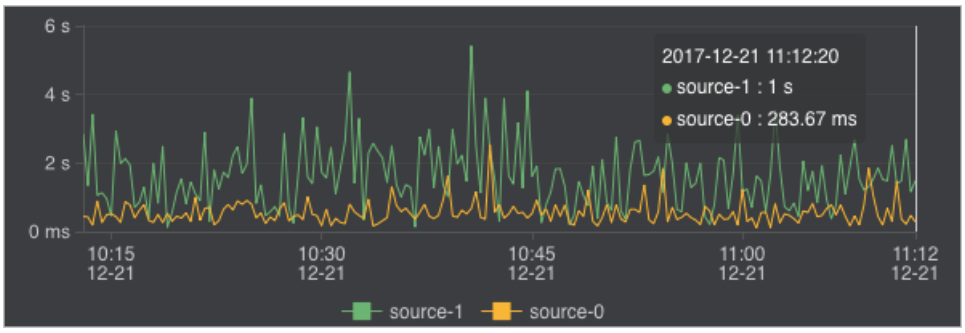

Input TPS of each source

Realtime Compute collects statistics about streaming data inputs of each job to visualize input transactions per second (TPS). The input TPS describes the amount of data read from the source table, which is measured in blocks per second. Unlike 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, X log groups are read per second and Y log records are parsed based on the X log groups. In this example, the input TPS is X, and the output RPS is Y.

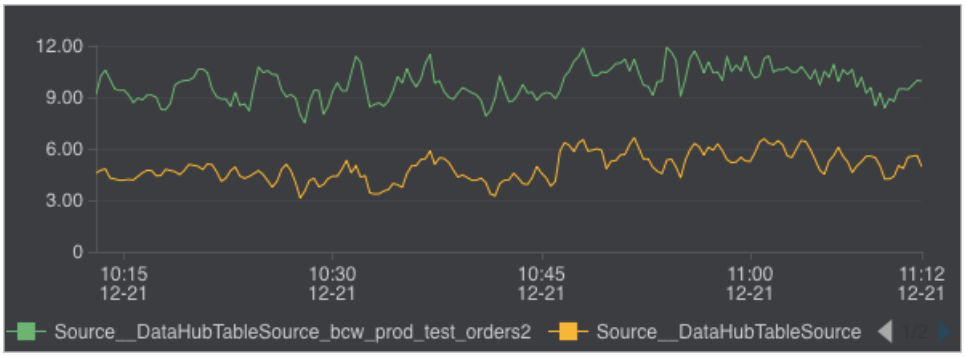

Data outputs of each sink

Realtime Compute collects statistics about data outputs of each job to visualize the output RPS.

> **Note** You can view the data outputs for all the target data stores, including data stores for streaming and non-streaming data.

If you find that no data outputs are detected during job administration, you must check whether data inputs exist in the upstream. You must also check whether data outputs exist in the downstream.

Data outputs of each sink

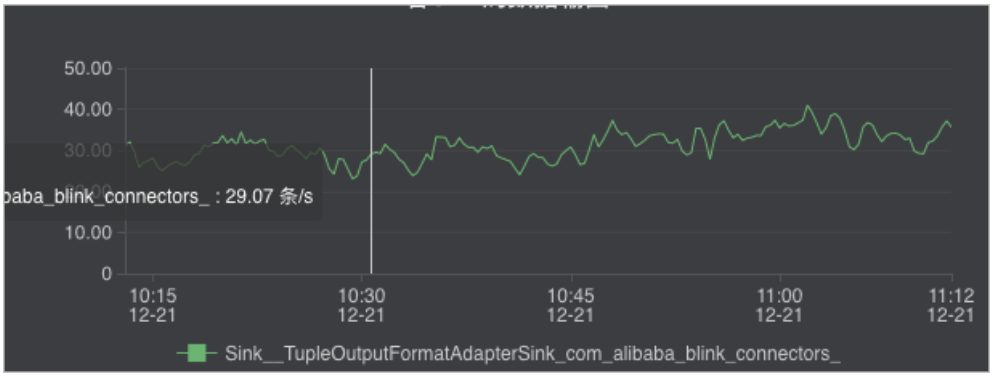

Input RPS of each source

Realtime Compute collects statistics about streaming data inputs of each job to visualize the input data records per second. If you find that no data outputs are detected during job administration, you must check whether data inputs from the source exist.

Input RPS of each source

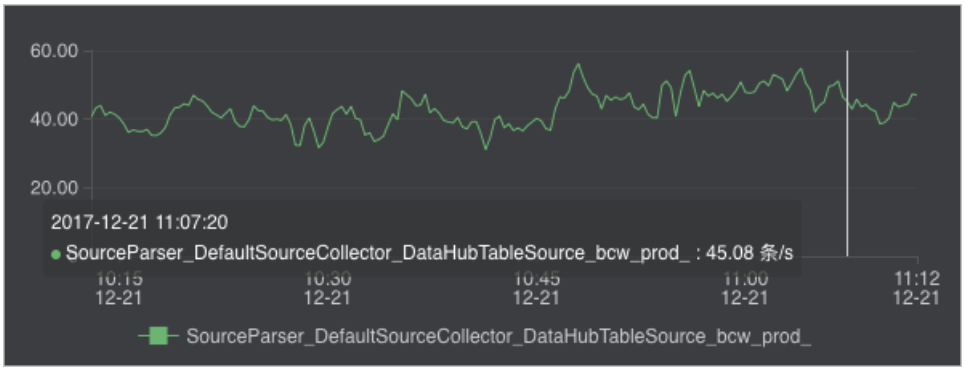

Input BPS of each source

Realtime Compute collects statistics about streaming data inputs of each job to visualize the input data bytes per second (BPS). The input BPS indicates the amount of data that is read from the source table per second.

Input BPS of each source

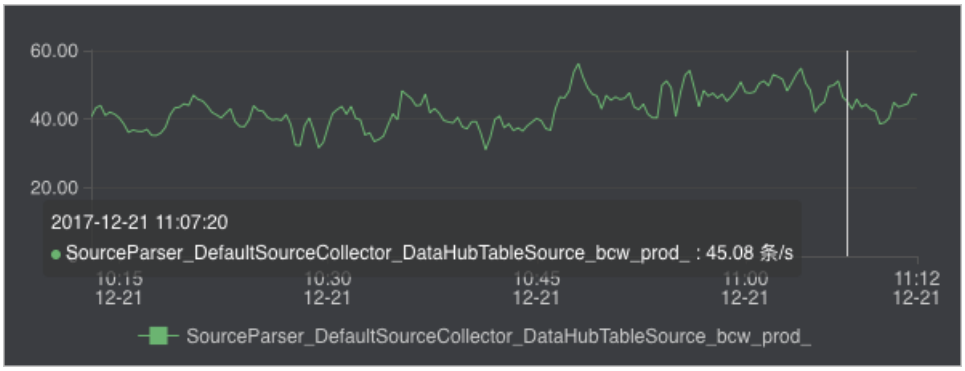

CPU usage

The CPU usage describes the CPU resources that are 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 requested.
- The CPU usage of the current job at a specified time. You can view the CPU usage from the curve chart.

Memory usage

The memory usage describes the memory resources that are consumed by a Realtime Compute job. Realtime Compute provides the following two metrics to reflect the memory usage:

- The memory size that you have requested.
- The memory usage of the current job at a specified time. You can view the memory usage from the curve chart.

Dirty data from each source

Realtime Compute allows you to view the statistics of dirty data from each source in a curve chart.

Dirty data from each source

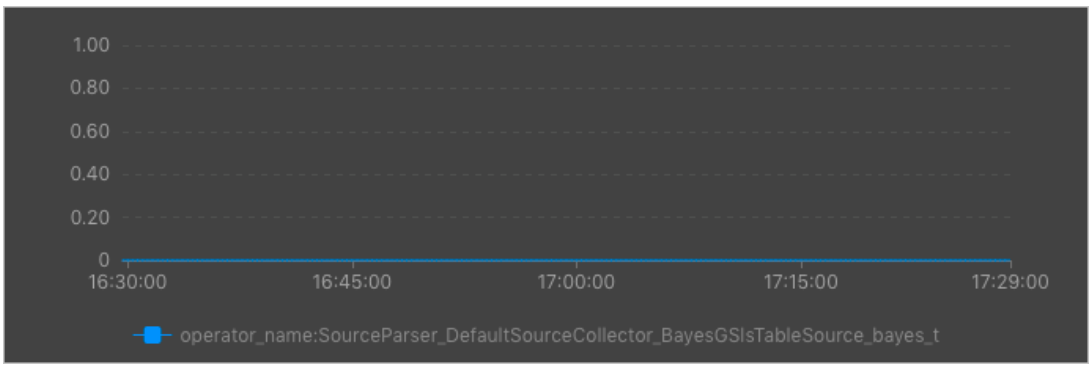

6.4.1.2.3. Advanced view

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 creating consistent snapshots of distributed data streams 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 do not overtake the source streaming data records, and the records flow strictly in line. A barrier separates the records in a data stream into two sets of records.

- One set of records goes into the current snapshot.
- The other set of records goes into the next snapshot.

Each barrier carries the ID of a snapshot that contains the records pushed before the barrier. Barriers do not interrupt the flow of the stream, and therefore are very lightweight. Multiple barriers from different snapshots can be found in the stream at the same time, which means that multiple snapshots may be concurrently created.

Barriers

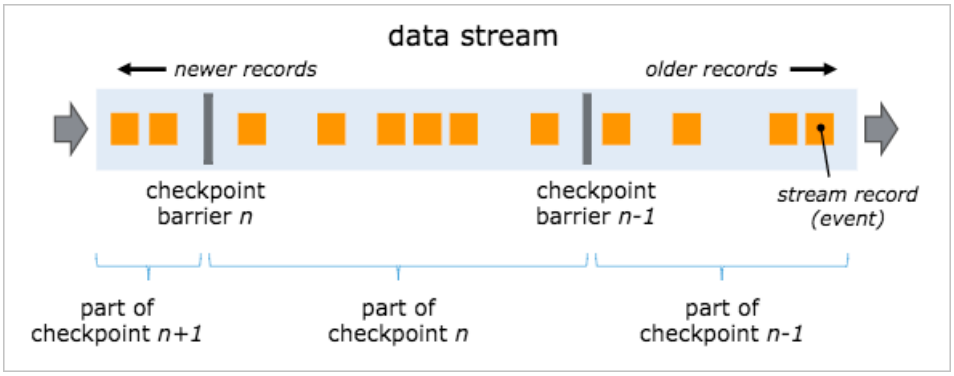

Barriers are inserted into data streams at data sources. The point where the barriers for snapshot n are inserted is the position in the source stream, up to which the snapshot covers the data. This point is indicated by Sn. The barriers then flow to the downstream. When an intermediate operator has received a barrier for snapshot n from all of its input streams, the operator emits a barrier for snapshot n into all of its outgoing streams. When a sink operator has received barrier n from all of its input streams, the sink operator acknowledges that snapshot n to the checkpoint coordinator. A sink operator is the end of a streaming directed acyclic graph (DAG). After all sink operators have acknowledged a snapshot, the snapshot is considered completed.

Barrier

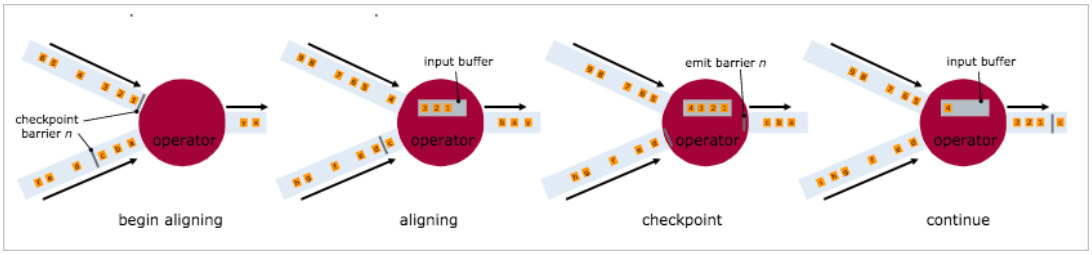

Checkpoint parameters

- **Checkpoint Duration**

  This parameter indicates the time spent on saving the state for each checkpoint. The duration is measured in milliseconds.

- **Checkpoint Size**

  This parameter indicates the state size of a checkpoint, which is measured in MiB.

- **Checkpoint Alignment Time**

This parameter indicates the time that is spent on receiving and acknowledging barrier n from all input streams. When a sink operator has received barrier n from all of its input streams, it acknowledges the snapshot n to the checkpoint coordinator. After all sink operators have acknowledged the snapshot n to the checkpoint coordinator, this snapshot is considered completed. The time consumed by the acknowledgement is included in the checkpoint alignment time.

- **Checkpoint Count**
- **Get**

  This parameter indicates the longest time that a subtask spends on performing a GET operation on the RocksDB database within a specified period.

- **Put**

  This parameter indicates the longest time that a subtask spends on performing a PUT operation on the RocksDB database within a specified period.

- **Seek**

  This parameter indicates the longest time that a subtask spends on performing a SEEK operation on the RocksDB database within a specified period.

- **State Size**

  This parameter indicates the state size of a job. If the size increases excessively fast, we recommend that you check and resolve potential issues.

- **CMS GC Time**

  This parameter indicates the garbage collection (GC) time 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.

6.4.1.2.4. Processing delay

Top 15 source subtasks with the longest processing delay

This metric describes the processing delays of each parallelism of a source.

6.4.1.2.5. 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.

6.4.1.2.6. 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.

6.4.1.2.7. 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.

6.4.1.2.8. 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.

6.4.1.2.9. 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.

6.4.1.3. 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.

6.4.1.4. 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.

Parameter	Description
Durations(ms)	The time that is spent on creating the checkpoint.

6.4.1.5. 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.

6.4.1.6. 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.

6.4.1.7. 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.

6.4.1.8. 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.

6.4.1.9. 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 a data sink, such as a deadlock in ApsaraDB RDS.
- Performance issues of user-defined extensions (UDXs) such as user-defined functions (UDFs), user-defined aggregate functions (UDAFs), and user-defined table-valued functions (UDTFs) in your Realtime Compute job.

Description

- 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.

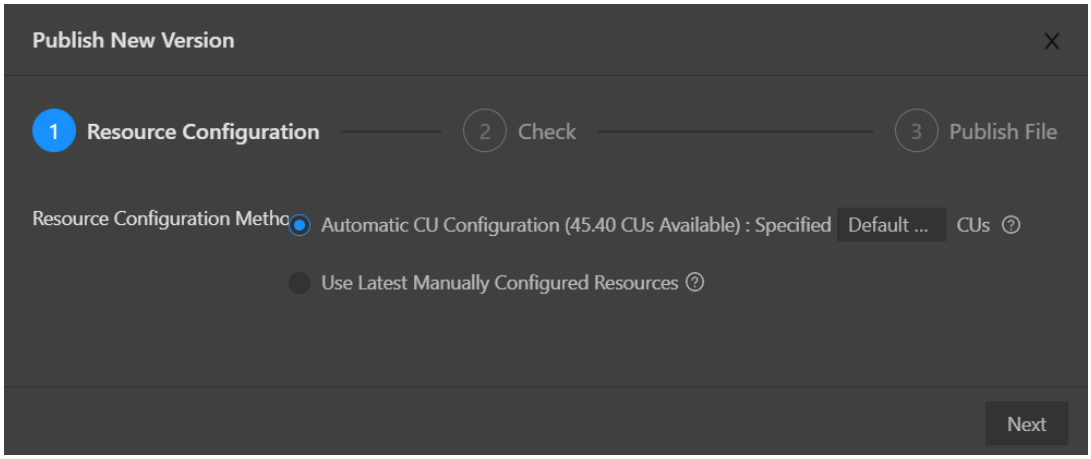

b. Specify **Resource Configuration Method**.

- **Automatic CU Configuration**: 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 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.

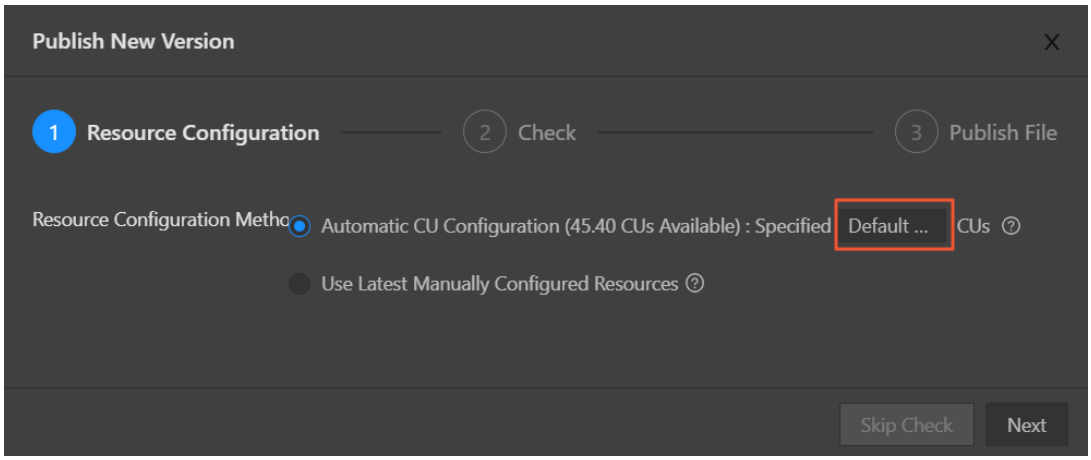

b. On the Administration page, find the job and click **Start** in the Actions column to start the job.

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 for Resource Configuration Method.

iii. Use the automatic CU configuration to start a job.

a. Resource performance optimization

If you select Automatic CU Configuration 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.

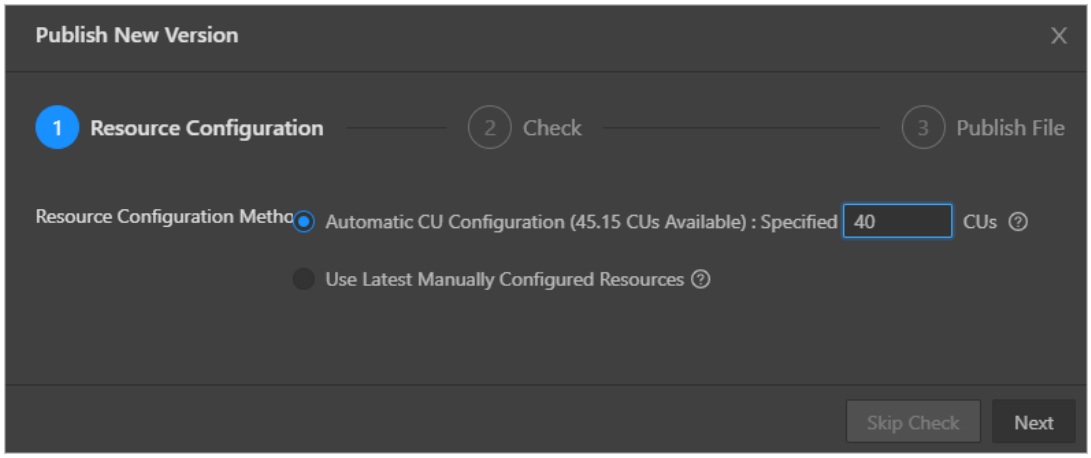

b. View the result of optimization. The following figure shows an example.

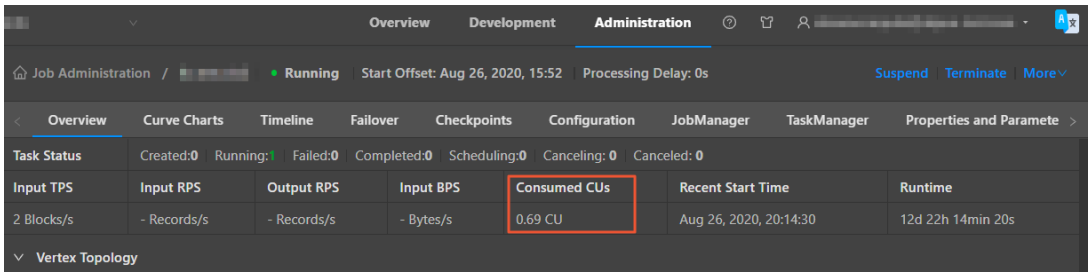

> **Note** Do not select **Use Latest Manually Configured Resources** for a new job. Otherwise, an error is returned.

- Existing jobs

- The following figure shows the optimization process of automatic configuration.

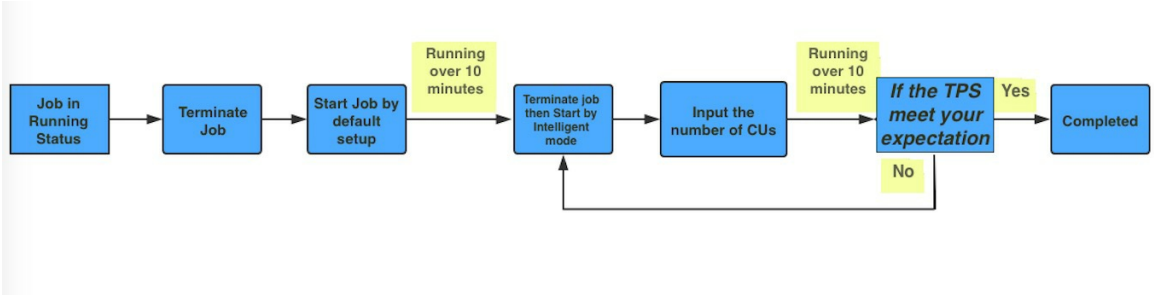

> **Note**
>
> - Before you use automatic configuration for a job that is in the running state, check whether stateful operations are involved. This is because the saved state data of a job may be cleared during the optimization process of automatic configuration.
> - If you make changes to a job, for example, modifying SQL statements or changing the Realtime Compute version, automatic configuration may fail. These changes may lead to topology changes, which results in some 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.

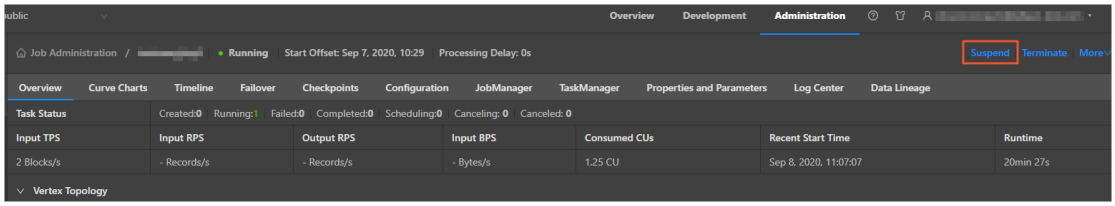

  b. Repeat the steps performed for new jobs and resume the job with the latest configuration.

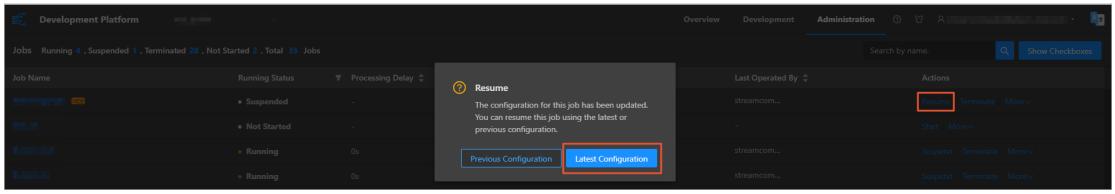

FAQ

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.

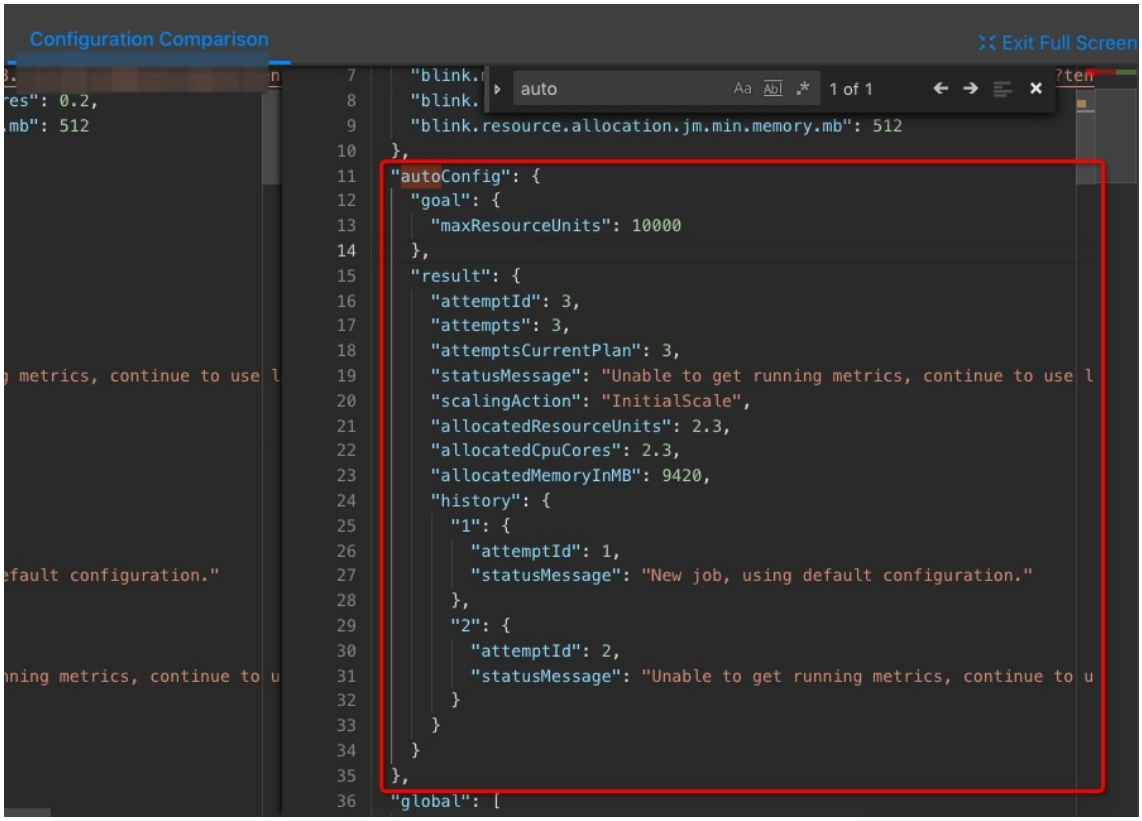

  ii. View the file on the Job Administration page, as shown in the following figure.

- 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 appears, automatic configuration runs properly. If some messages appear, 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 job to a backup node, perform the following steps to mitigate the negative impact of service interruption: Roll back the job to an earlier version and specify the start offset to re-read the data during off-peak hours. To roll back the 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(AbstractStreamOpe
rator.java:687)
   at org.apache.flink.streaming.api.operators.AbstractStreamOperator.initializeState(AbstractStreamOper
ator.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 o
bject.
   at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSecondaryState.deserializeStateEntry(
AbstractRocksDBRawSecondaryState.java:167)
   at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRestoreOperation.restoreRawStateDat
a(RocksDBIncrementalRestoreOperation.java:425)
   at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRestoreOperation.restore(RocksDBIncr
ementalRestoreOperation.java:119)
   at com.alibaba.blink.contrib.streaming.state.RocksDBKeyedStateBackend.restore(RocksDBKeyedStateBa
ckend.java:216)
   at org.apache.flink.streaming.api.operators.AbstractStreamOperator.createKeyedStateBackend(Abstract
StreamOperator.java:986)
   at org.apache.flink.streaming.api.operators.AbstractStreamOperator.initKeyedState(AbstractStreamOpe
rator.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
```

6.4.1.10. Improve performance by manual configuration

6.4.1.10.1. 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.

6.4.1.10.2. 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 RealtimeCompute 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.

6.4.1.10.3. 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
```

6.4.1.10.4. 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 configure batch size parameters to specify the number of data records that are read from a source table or written to 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.
AnalyticDB for MySQL result table	batchSize	The number of data records that are written at a time.	Optional. Default value: 1000.
ApsaraDB 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 parameters in the table to the WITH parameter list in DDL statements for the related data storage. For example, add `batchReadSize=' 500'` to the WITH parameter list in DDL statements for the DataHub source table.

6.4.1.10.5. 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.

6.4.1.10.6. 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.

6.4.1.11. O&M of Apsara Big Data Manager

6.4.1.11.1. What is Apsara Big Data Manager?

Apsara Big Data 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 in the ABM console, see the following topics.

6.4.1.11.2. Log on to the ABM console

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

Prerequisites

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

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

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

Procedure

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

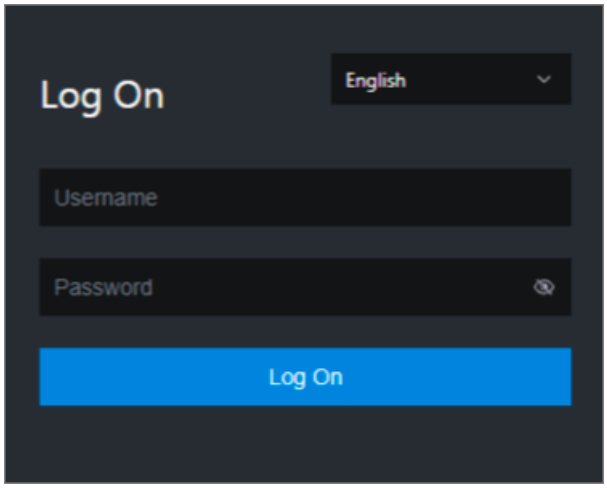

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.
4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Big Data Services** section, choose **General-Purpose O&M > Apsara Big Data Manager**.

6.4.1.11.3. O&M overview of Realtime Compute for Apache Flink

This topic describes the O&M features of Realtime Compute for Apache Flink supported by Apsara Bigdata Manager (ABM). It also shows how to access the O&M page of Realtime Compute for Apache Flink.

Modules

O&M of Realtime Compute for Apache Flink includes business O&M, service O&M, cluster O&M, and host O&M. The following table describes these modules.

Module	Feature	Description
Business	Item	Displays information about all projects in Realtime Compute for Apache Flink.
	Job	Displays information about all jobs in Realtime Compute for Apache Flink and supports job diagnosis and analysis.
	Queues	Displays information about all queues in Realtime Compute for Apache Flink.
Services	Blink	Displays the overview of the Blink service in Realtime Compute for Apache Flink.
	Yarn	Displays the overview and health status of the YARN service in Realtime Compute for Apache Flink.
	HDFS	Displays the overview and health status of the HDFS service in Realtime Compute for Apache Flink.
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 utilization, disk usage, memory usage, load, and packet transmission for the cluster.
	Health Status	Displays all check items of a cluster, including the check item details, check results for the hosts in the cluster, and methods to handle alerts. In addition, you can log on to a host and perform manual checks on the host.
	Hosts	Displays the information about hosts in a cluster, including the hostname, IP address, role, type, CPU utilization, memory usage, root disk usage, packet loss rate, and packet error rate.
Hosts	Overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU utilization, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.
	Health Status	Displays the check items of the selected host, including the check item details, check results for the host, and methods to handle alerts. 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 of the ABM console, and then click **RealtimeCompute**.
3. On the RealtimeCompute page, click **O&M** in the upper-right corner. The **Business** page appears by

default.

The **O&M** page includes four modules, **Business**, **Services**, **Clusters**, and **Hosts**.

6.4.1.11.4. Business O&M

6.4.1.11.4.1. Projects

This topic describes how to view information about the projects in Realtime Compute and how to go to the Queue Analysis page from the Projects page.

Projects

On the **Business** page, click **Projects** in the left-side navigation pane. The **Projects** page for Realtime Compute appears.

The **Projects** page displays information about the projects in Realtime Compute, including the name, BRS, queue, used compute units (CUs), total CUs, CU usage percentage, and number of jobs.

Go to the Queue Analysis page

The Queue Analysis page displays the status and resource usage of a queue, and information about jobs running in the queue, so that you can quickly know the running status of the queue.

On the **Projects** page, click a queue in the Queue column. The **Queue Analysis** page of the queue appears. For more information about the Queue Analysis page and operations that you can perform on this page, see Queues.

6.4.1.11.4.2. Jobs

This topic describes how to view information about the jobs in Realtime Compute and how to go to the Job Analysis page, Queue Analysis page, and Realtime Compute console from the Jobs page.

Jobs

On the **Business** page, click **Jobs** in the left-side navigation pane. The **Jobs** page for Realtime Compute appears.

The **Jobs** page displays information about the jobs in Realtime Compute, including the job names, users who created the jobs, projects to which the jobs belong, queues where the jobs are running, transactions per second (TPS) in the inbound direction, job latency, requested compute units (CUs), job statuses, and start time.

Go to the Realtime Compute console

On the **Jobs** page, click the content in the Failover column of a job to go to the Realtime Compute console.

Go to the Job Analysis page

The job analysis feature allows you to diagnose jobs to quickly troubleshoot job failures.

On the **Jobs** page, click a job in the Name column. The **Job Analysis** page of the job appears. For more information about the Job Analysis page and operations that you can perform on this page, see Job analysis.

Go to the Queue Analysis page

The Queue Analysis page displays the status and resource usage of a queue, and information about jobs running in the queue, so that you can quickly know the running status of the queue.

On the **Jobs** page, click a queue in the Queue column. The **Queue Analysis** page of the queue appears. For more information about the Queue Analysis page and operations that you can perform on this page, see Queues.

6.4.1.11.4.3. Queues

Apsara Big Data Manager (ABM) allows you to view the information about the queues in Realtime Compute, including the queue names, queue statuses, minimum numbers of CPU cores and minimum memory capacity guaranteed for the queues, maximum numbers of CPU cores and maximum memory capacity available for the queues, and numbers of jobs running in the queues.

Queues

On the **Business** page, click **Queues** in the left-side navigation pane. The **Queues** page for Realtime Compute appears.

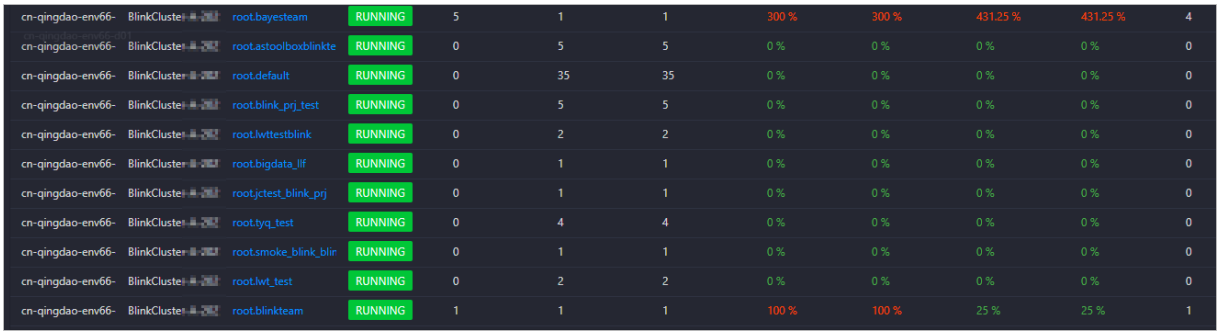

cn-qingdao-env66-	BlinkCluste	root.bayesteam	RUNNING	5	1	1	300 %	300 %	431.25 %	431.25 %	4
cn-qingdao-env66-	BlinkCluste	root.astoolboxblinkte	RUNNING	0	5	5	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.default	RUNNING	0	35	35	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.blink_prj_test	RUNNING	0	5	5	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.lwttestblink	RUNNING	0	2	2	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.bigdata_llf	RUNNING	0	1	1	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.jctest_blink_prj	RUNNING	0	1	1	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.tyq_test	RUNNING	0	4	4	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.smoke_blink_blir	RUNNING	0	1	1	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.lwt_test	RUNNING	0	2	2	0 %	0 %	0 %	0 %	0
cn-qingdao-env66-	BlinkCluste	root.blinkteam	RUNNING	1	1	1	100 %	100 %	25 %	25 %	1

The **Queues** page displays information about the queues in Realtime Compute, including the clusters to which the queues belong, queue names, queue statuses, requested compute units (CUs), minimum CUs guaranteed, maximum CUs available, and numbers of jobs running in the queues.

Go to the Queue Analysis page

The Queue Analysis page displays the status and resource usage of a queue, and information about jobs running in the queue, so that you can learn the running status of the queue.

On the **Queues** page, click a queue in the Queue column. The **Queue Analysis** page of the queue appears. For more information about the Queue Analysis page and operations that you can perform on this page, see Queues.

6.4.1.11.5. Service O&M

6.4.1.11.5.1. Blink

Apsara Big Data 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.

6.4.1.11.5.2. Yarn

Apsara Big Data 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 on the top of the Services page. The **Health Status** page for the YARN service appears.

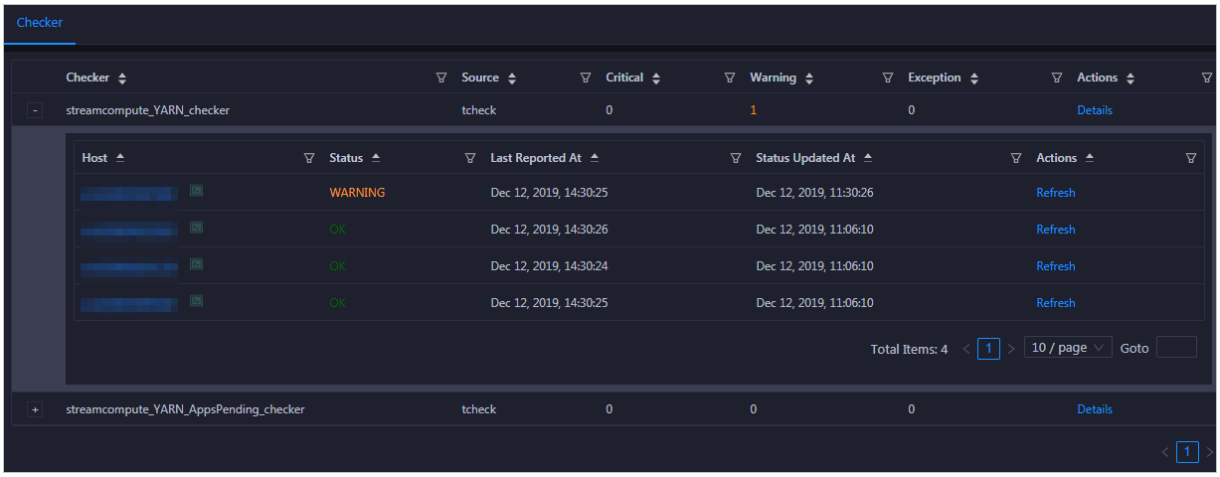

On the **Health Status** page, you can view all checkers of the YARN service and the check results for all hosts. The following alerts may be reported on a host: **Critical**, **Warning**, and **Exception**. The alerts are repesented in different colors. Pay attention to the check results, especially the **Critical** and **Warning** results, and handle them in a timely manner.

The operations you can perform on the **Health Status** page for the YARN service are the same as those on the Health Status page for Realtime Compute clusters. For more information, see Cluster health.

6.4.1.11.5.3. HDFS

Apsara Big Data Manager (ABM) allows you to view the overview and health status of the Hadoop Distributed File System (HDFS) service in Realtime Compute.

Overview

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

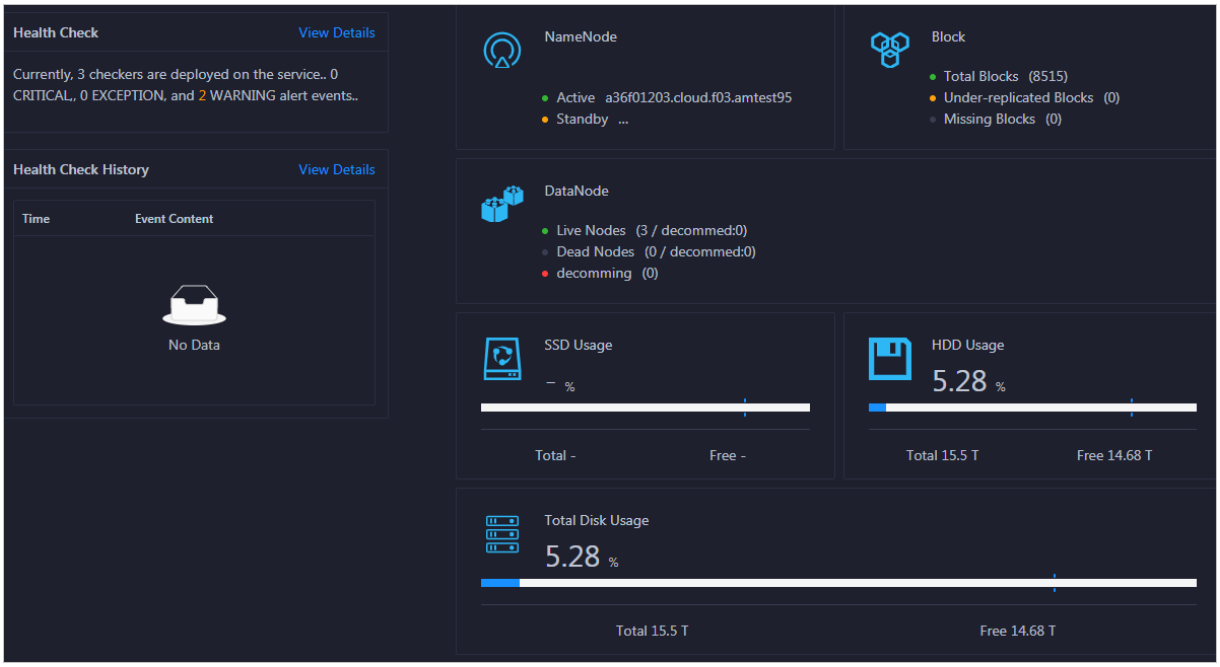

The **Overview** page displays the health check result, health check history, the information of NameNode, blocks, and DataNode, solid-state disk (SSD) usage, hard disk drive (HDD) usage, and total disk usage.

Click **View Details** in the **Health Check** or **Health Check History** section. The **Health Status** page for the HDFS service appears. On this page, you can view more details about the health check.

Health status

On the **Services** page, click **HDFS** in the left-side navigation pane. Click the **Health Status** tab on the top of the Services page. The **Health Status** page for the HDFS service appears.

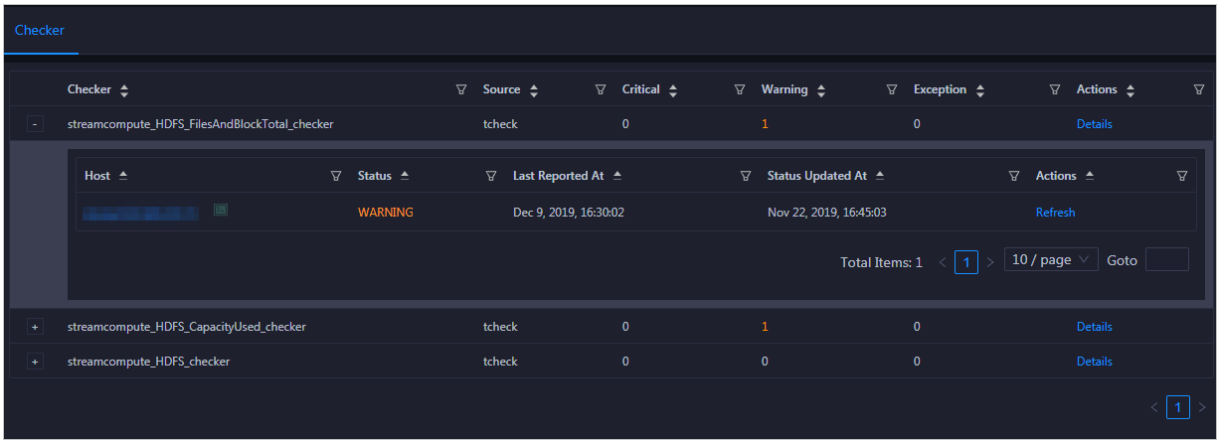

On the **Health Status** page, you can view all checkers of the HDFS service and the check results for all hosts in the cluster. The following check results can be returned: **Critical**, **Warning**, and **Exception**. They are displayed in different colors. Pay attention to the check results, especially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

The operations you can perform on the **Health Status** page for the HDFS service are the same as those on the Health Status page for Realtime Compute clusters. For more information, see Cluster health.

6.4.1.11.6. Cluster O&M

6.4.1.11.6.1. 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, and packet transmission for the cluster.

Entry

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

Hosts

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

Services

This section displays all services deployed in the cluster and the respective number of available and unavailable services.

CPU

This chart 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.

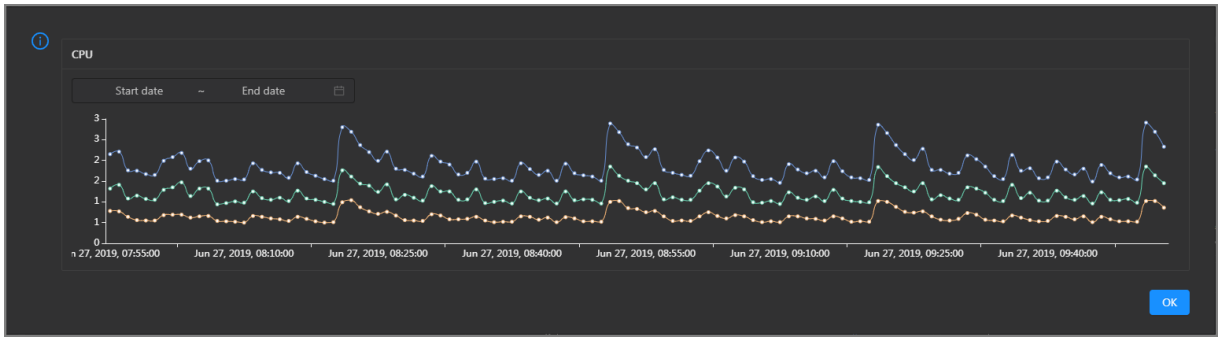

DISK

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

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

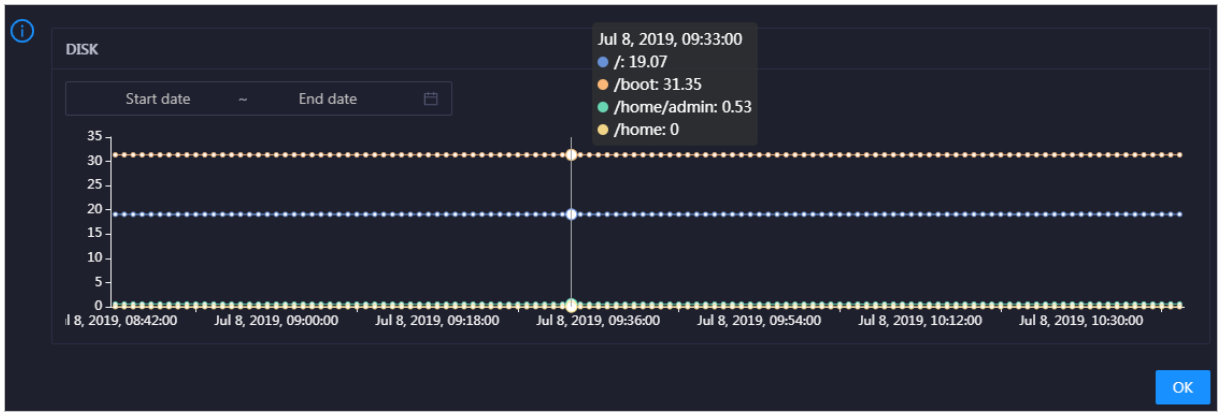

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

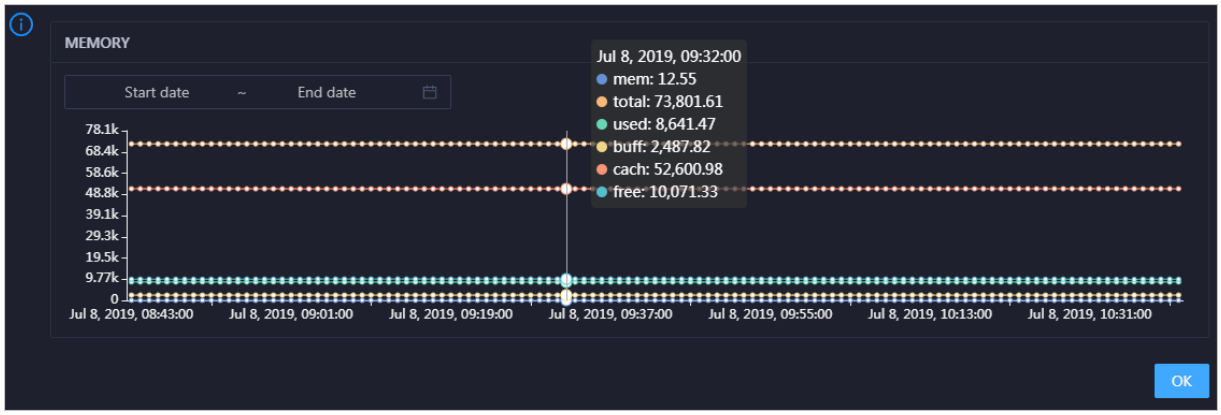

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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.

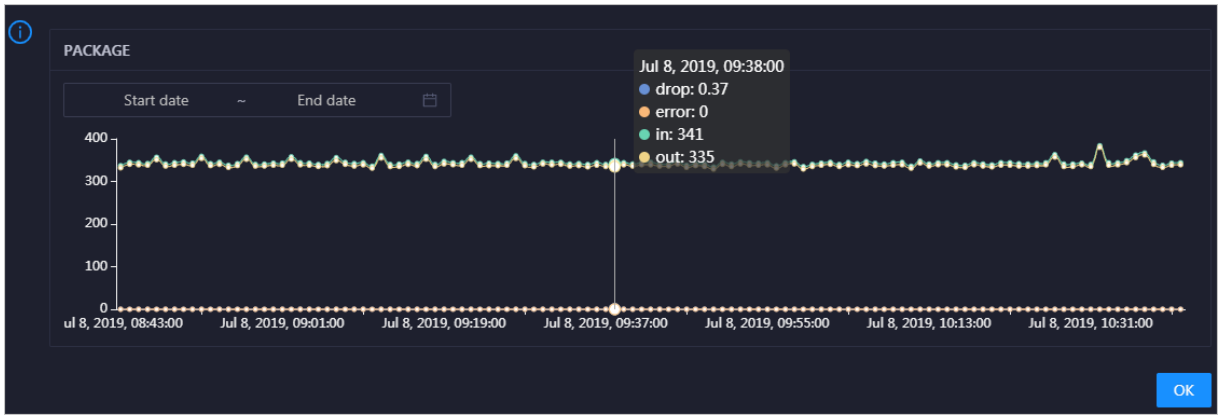

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

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.

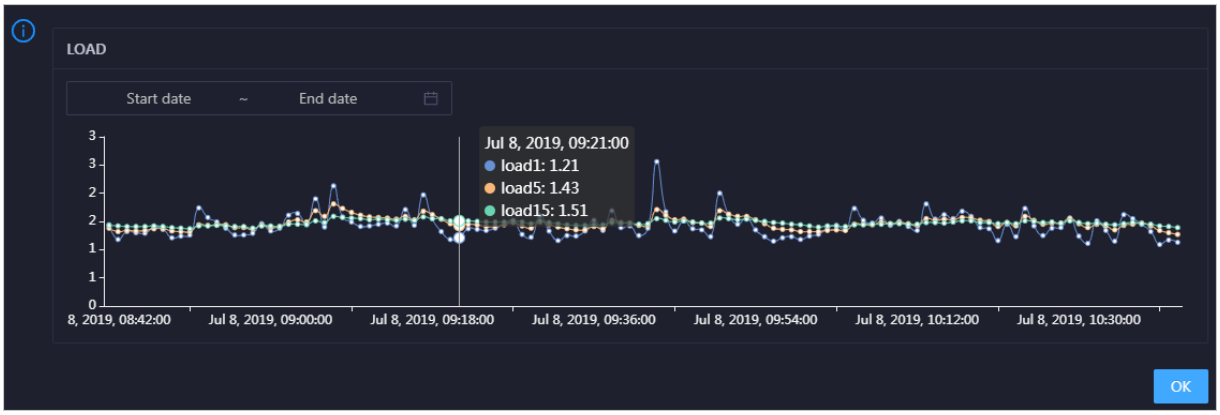

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

Health Check

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

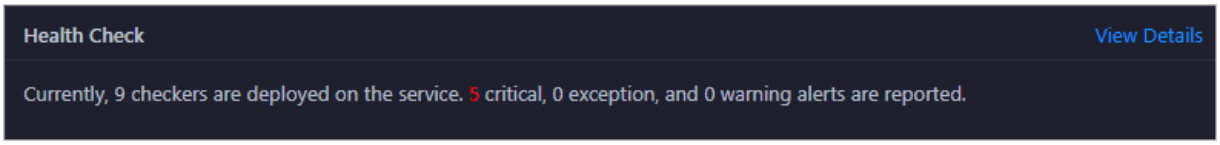

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Cluster health.

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.

6.4.1.11.6.2. 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.

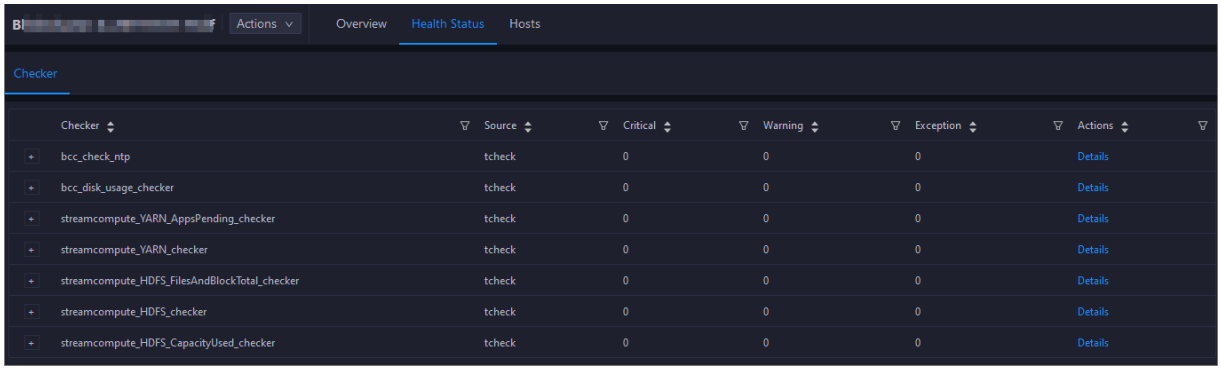

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are represented in different colors. You must handle the alerts in a timely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

View checker details

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

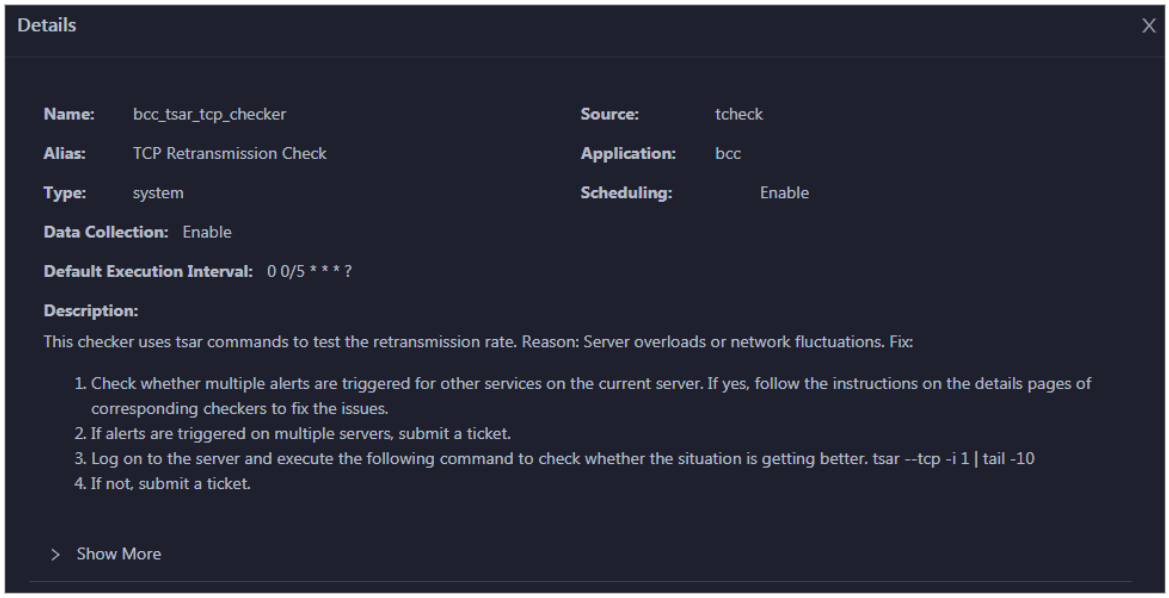

   The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

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

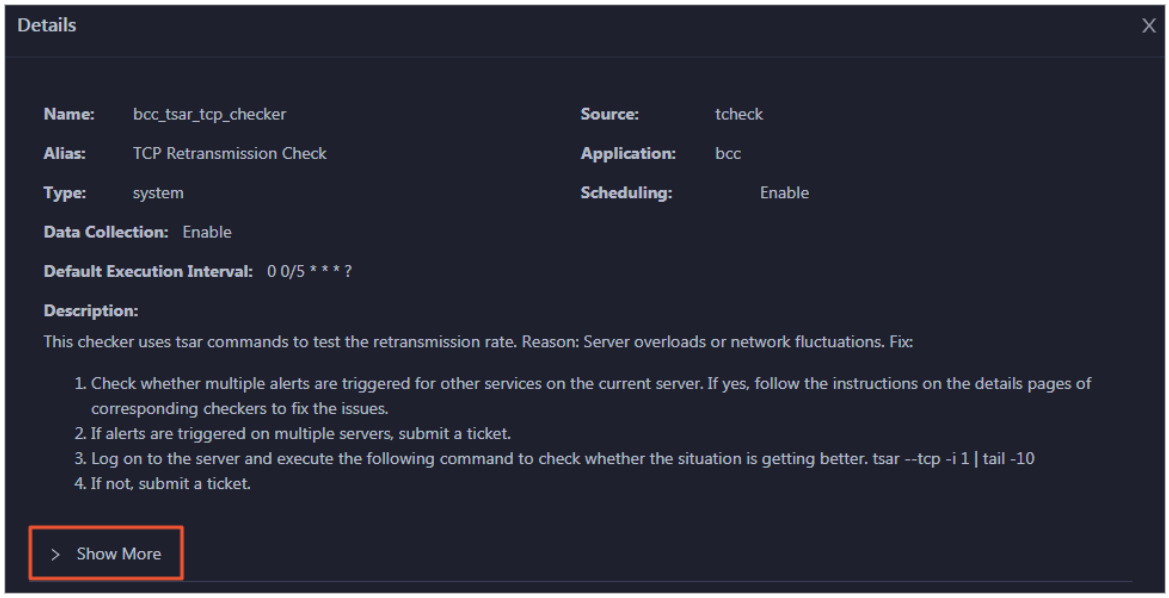

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

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

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

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

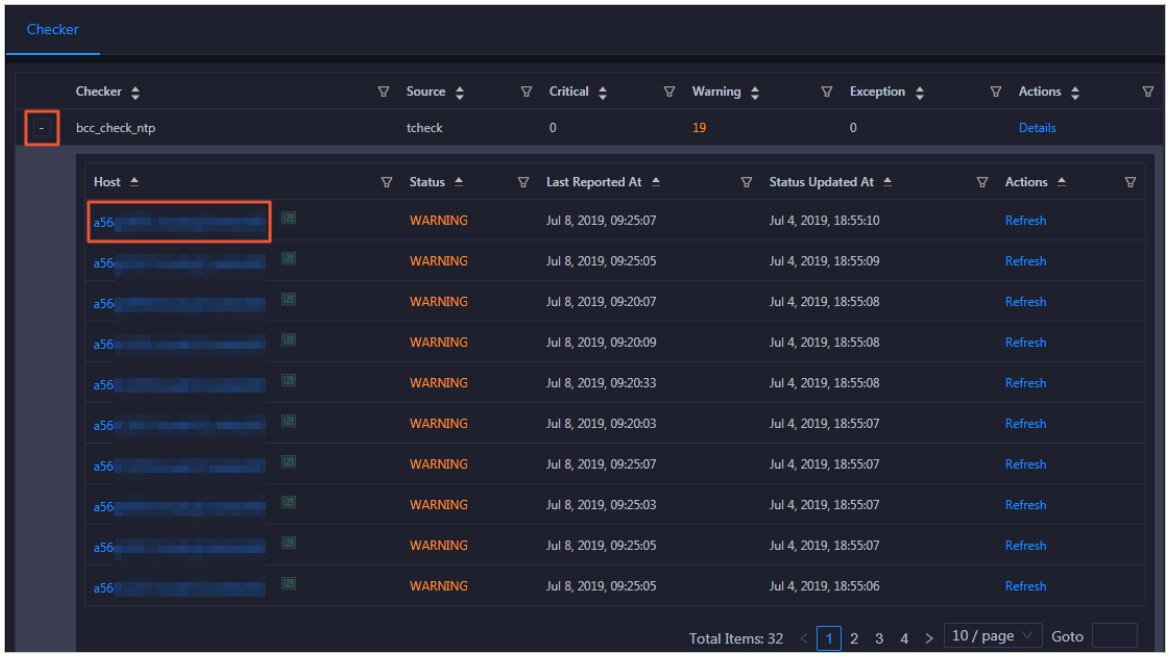

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

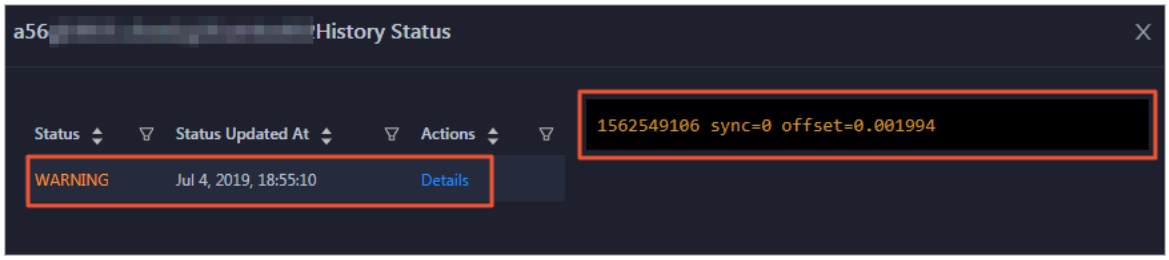

Clear alerts

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

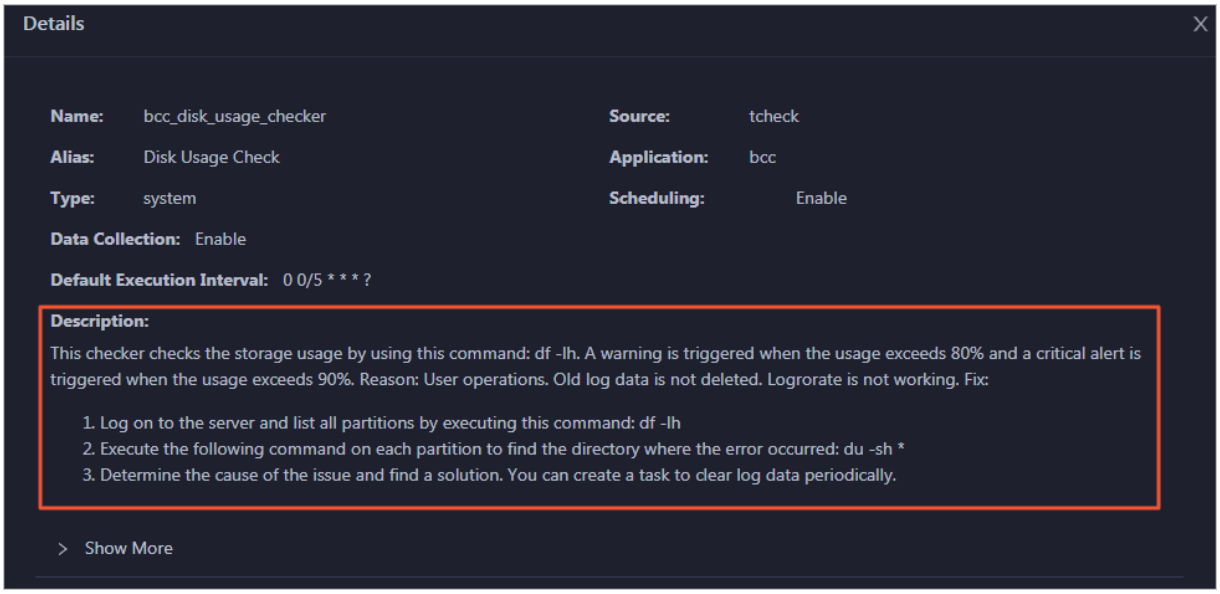

Log on to a host

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

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

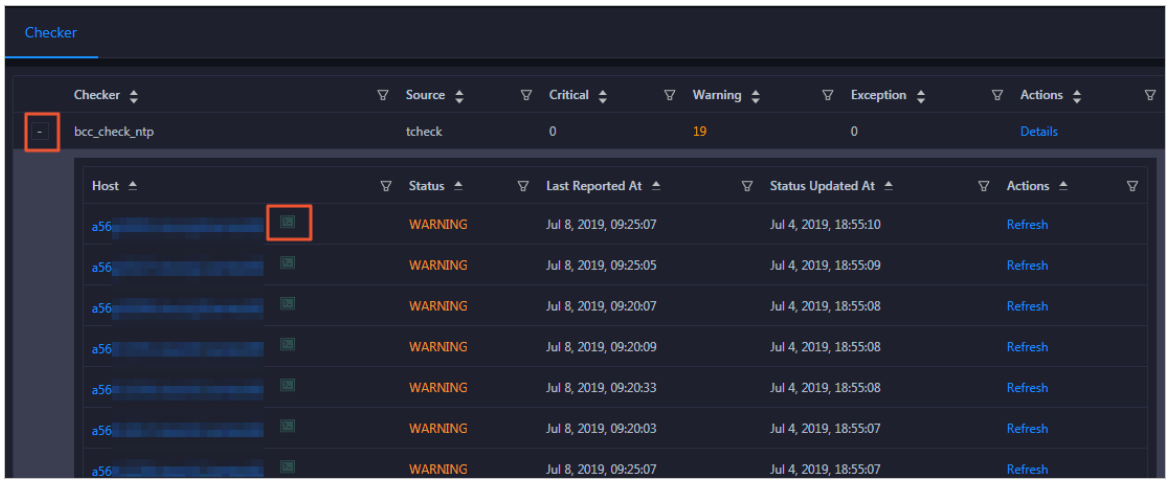

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

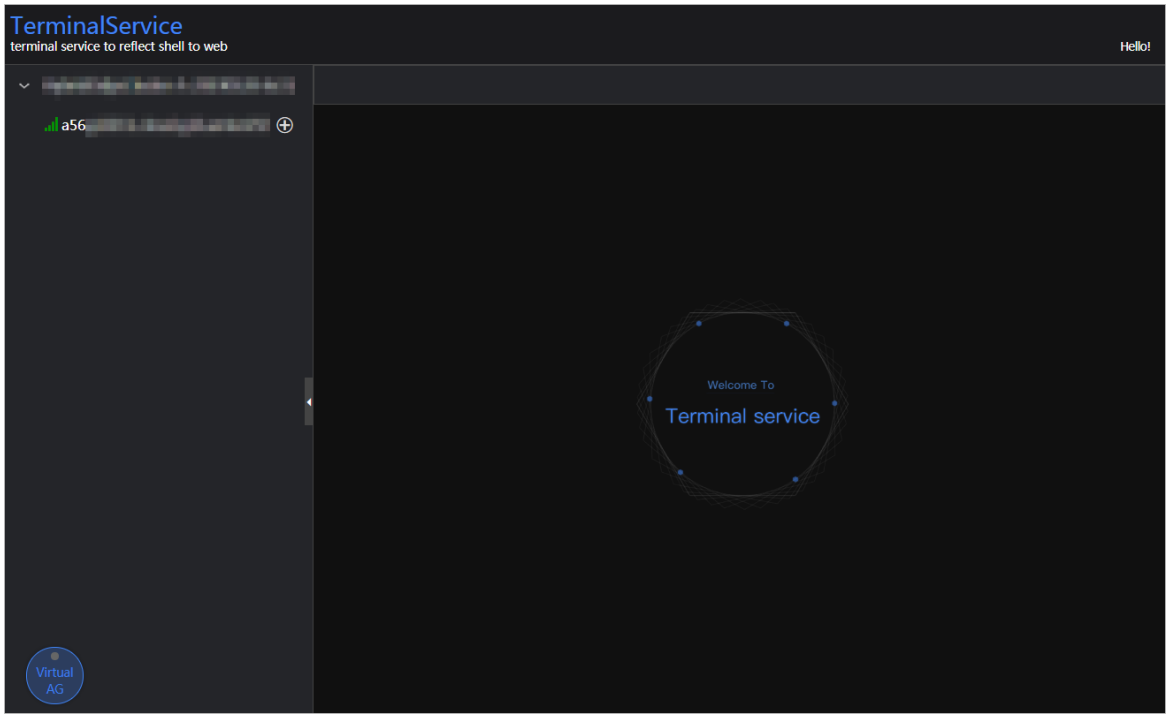

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

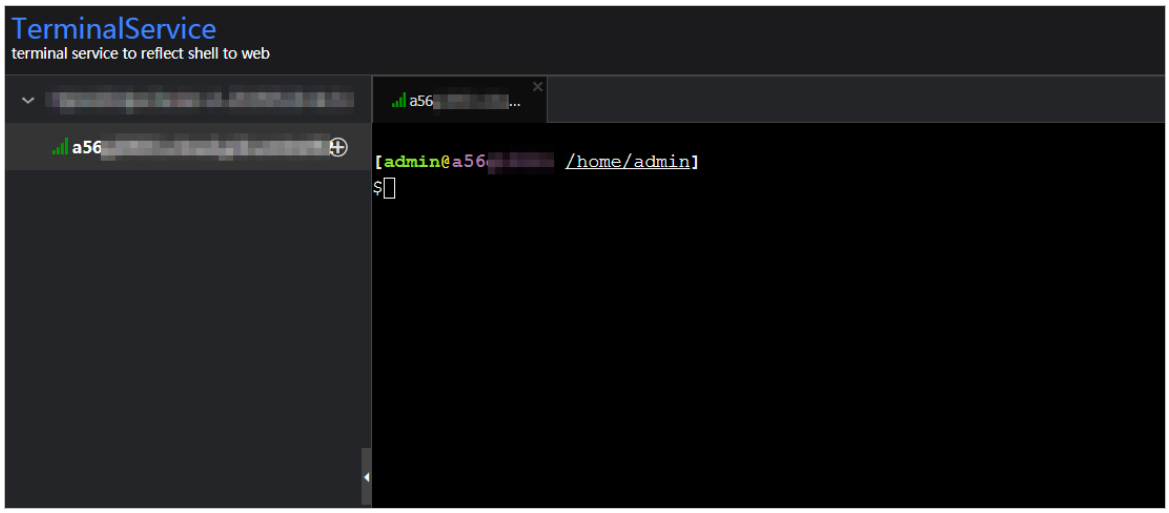

Run a checker again

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

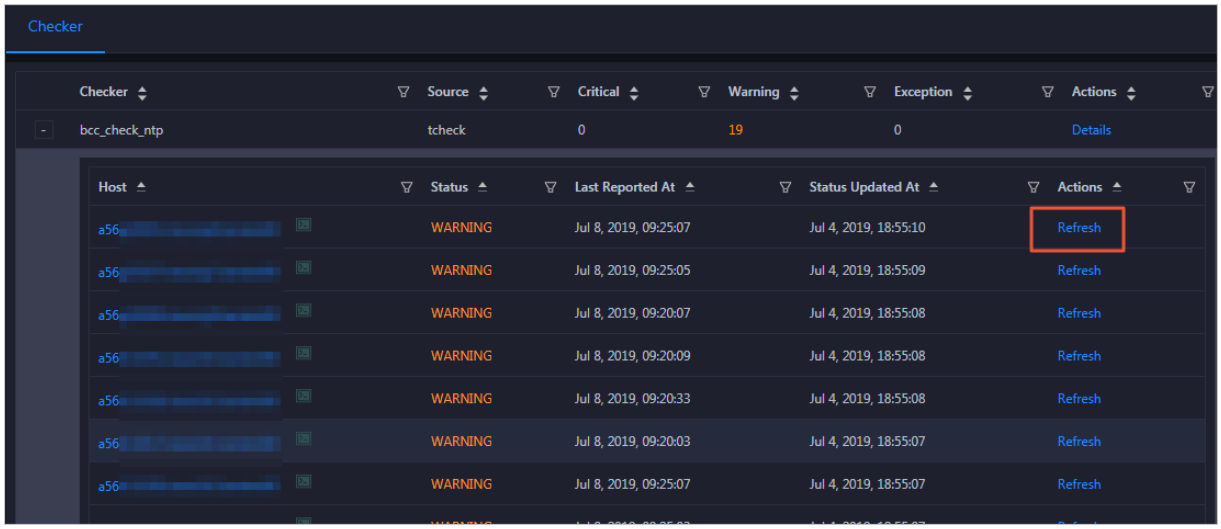

6.4.1.11.6.3. Hosts

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.

6.4.1.11.6.4. Cluster scale-out

Apsara Big Data Manager (ABM) allows you to scale out a Realtime Compute cluster by adding physical hosts. Cluster scale-out refers to the process of adding physical hosts in the default cluster of Apsara Infrastructure Management Framework to a Realtime Compute cluster. You can perform the scale-out operation only for **worker** nodes in a Realtime Compute cluster.

Prerequisites

- Your ABM account is granted the required permissions to perform O&M operations on Realtime Compute.
- Hosts whose service type is **blink** are deployed in the default cluster of Apsara Infrastructure Management Framework.

Background 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 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 the icon in the upper-left corner, and 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 cluster

You can add multiple hosts to a Realtime Compute 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 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 the icon in the upper-left corner, and 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** pane, configure the required parameters.

   You must configure the following parameters in this step:

   - **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 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 is submitted.
7. View the scale-out status.

   Move the pointer over **Actions** in the upper-left corner, and 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 succeeds, 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.

6.4.1.11.6.5. Cluster scale-in

Apsara Big Data 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. Scale-in operations can be performed only on the **worker** nodes in a Realtime Compute cluster.

Prerequisites

- Your ABM account is granted the required permissions to perform O&M operations on Realtime Compute.
- More than three **worker** nodes are deployed in the current cluster. A Realtime Compute cluster creates three replicas for data by default. At least three **worker** nodes are required. Make sure that the cluster has at least three worker nodes after scale-in.
- Resources of the cluster, including the disk, CPU, and memory, are checked and still sufficient if the cluster is scaled in. For more information about how to check CPU and memory usage, see Yarn. You can run the **df** command to check disk usage.

> **Notice** Scale-in triggers a job failover on hosts. If the cluster resources are insufficient after scale-in, the failover fails. This leads to negative effects on your business.

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

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

Procedure

(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 select one or more hosts whose role is **Worker**.
2. On the Clusters page, choose **Actions > Scale in Cluster**. The **Scale in Cluster** pane appears.

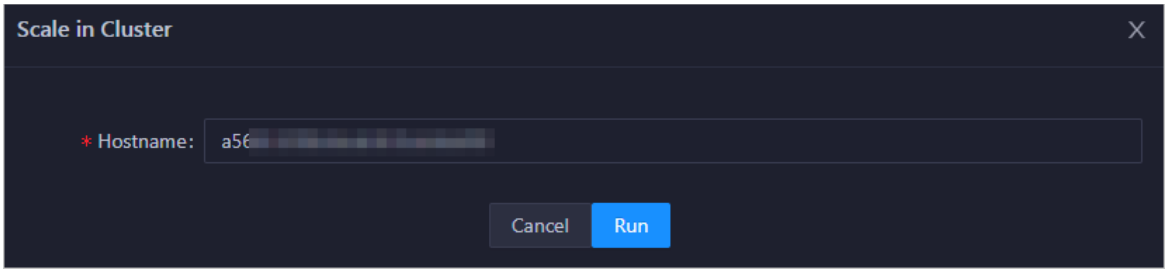

Hostname: the name of the host to be removed from the cluster. The name of the selected host is used by default.

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

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

   It may take some time for the cluster to be scaled in. In the Current Status column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution succeeds, and **FAILED** indicates that the execution fails.
5. (Optional)View the scale-in progress.

   If the status is **RUNNING**, click **Details** in the Details column to view the steps and progress of the scale-in operation.
6. Locate the cause of a scale-in failure.

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

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

6.4.1.11.7. Host O&M

6.4.1.11.7.1. 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 for the host.

Entry

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

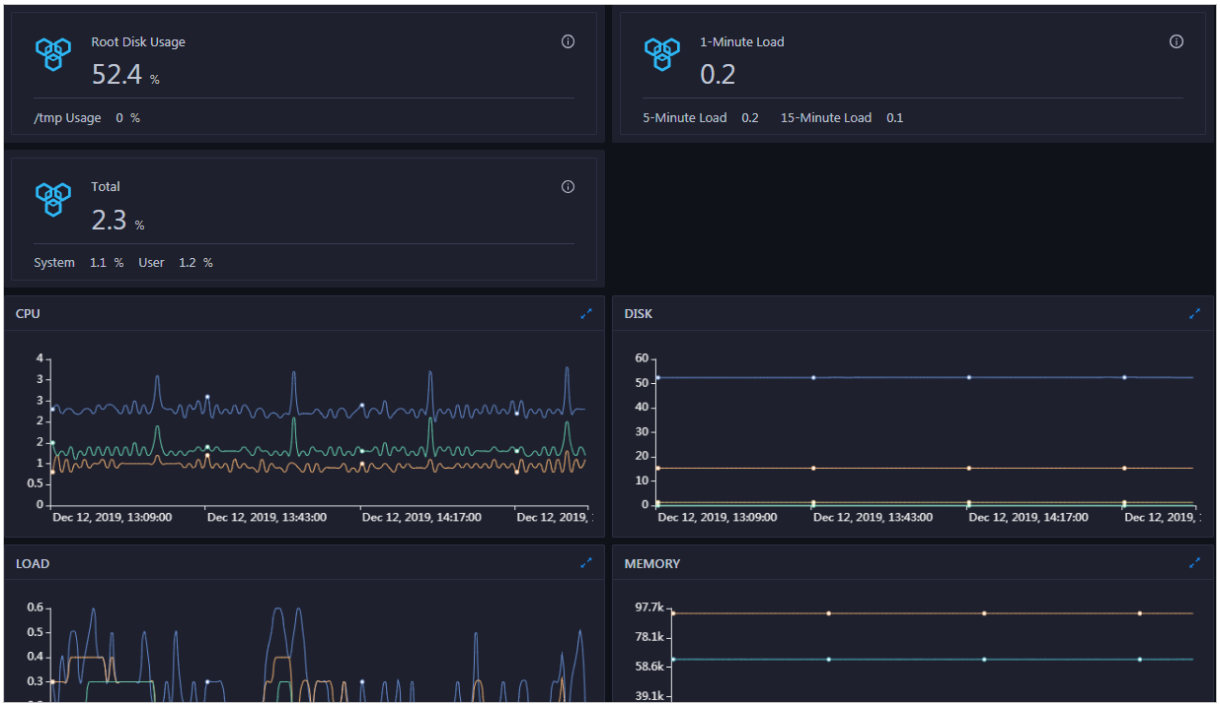

On the **Overview** page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.

Root Disk Usage, Total, and 1-Minute Load

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

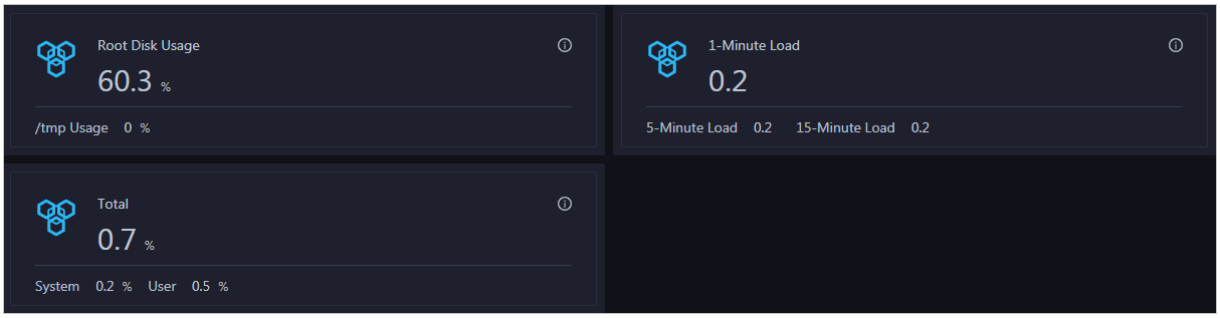

CPU

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

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

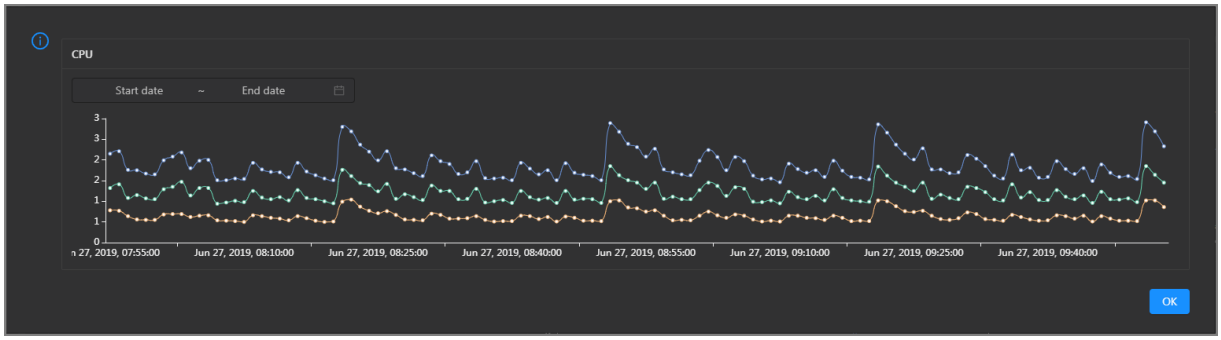

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

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

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

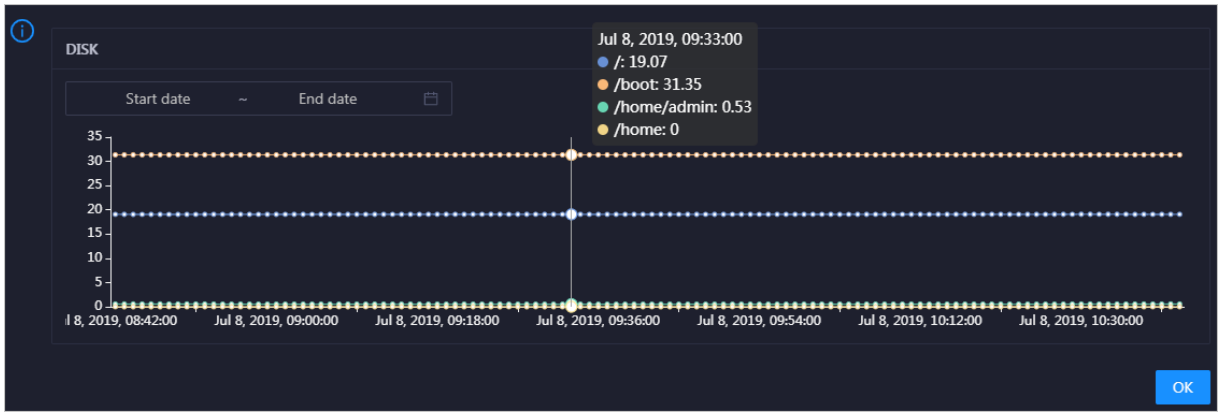

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

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

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

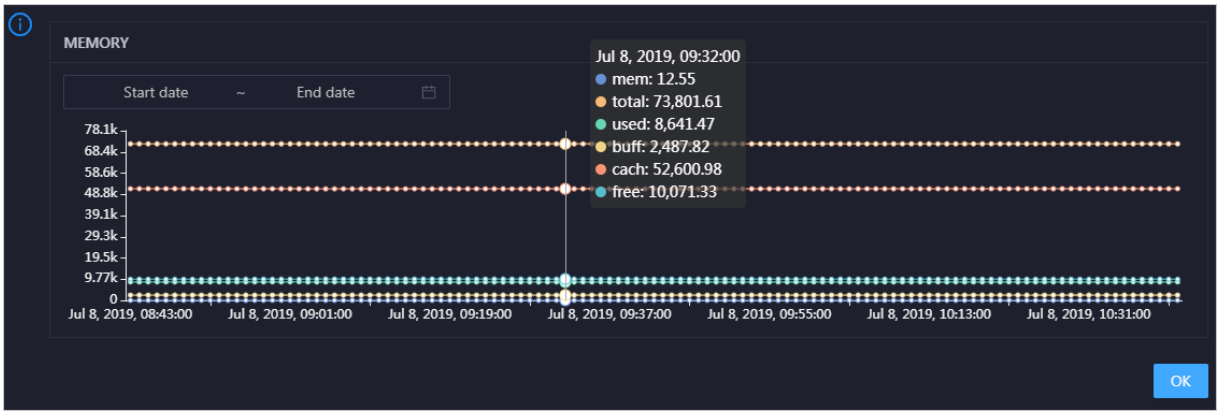

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

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

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

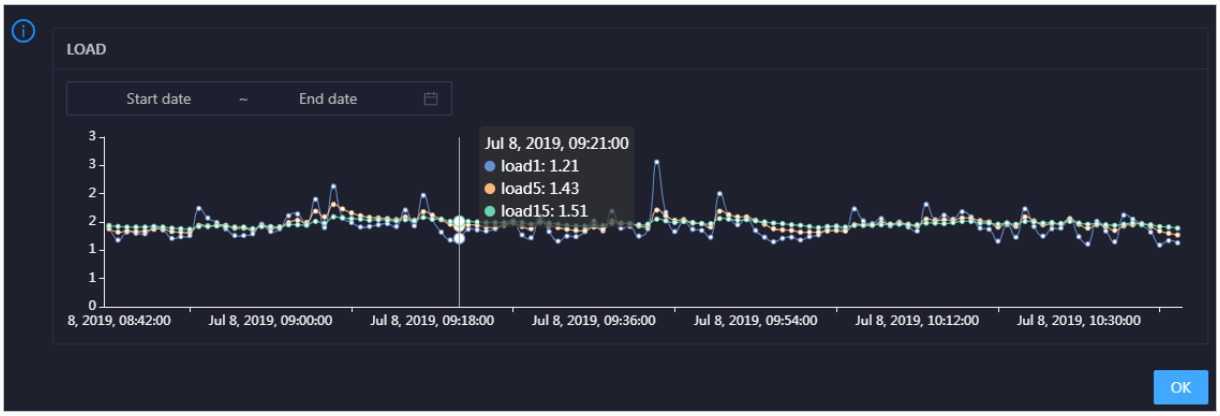

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

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

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

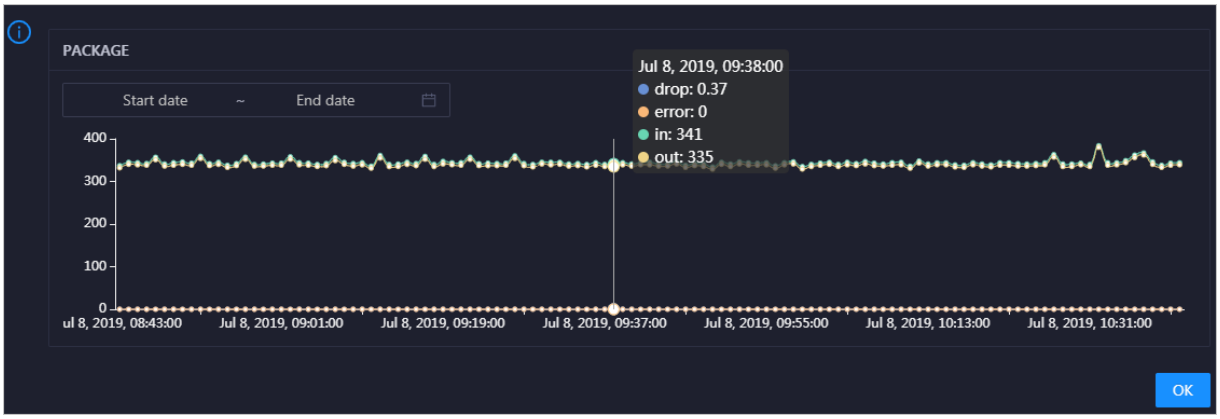

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

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

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

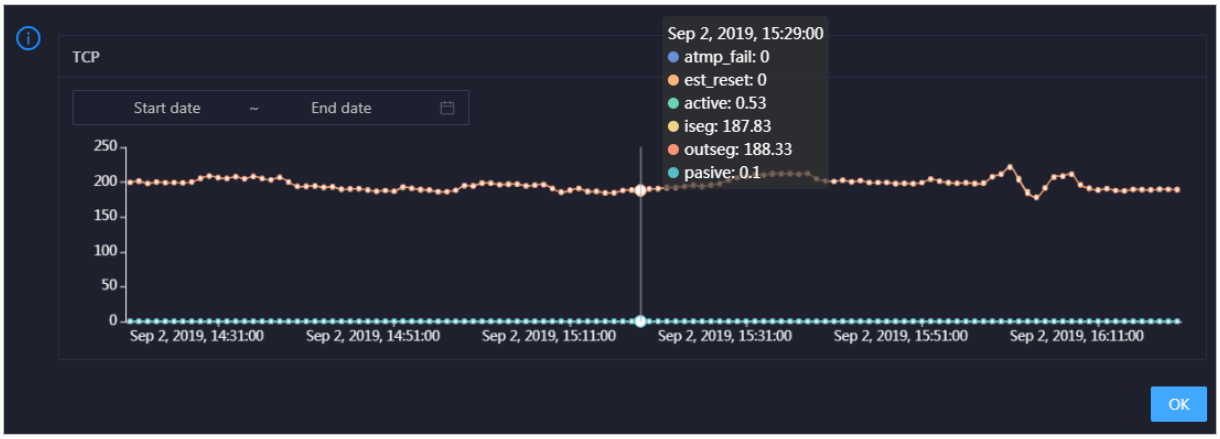

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

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

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

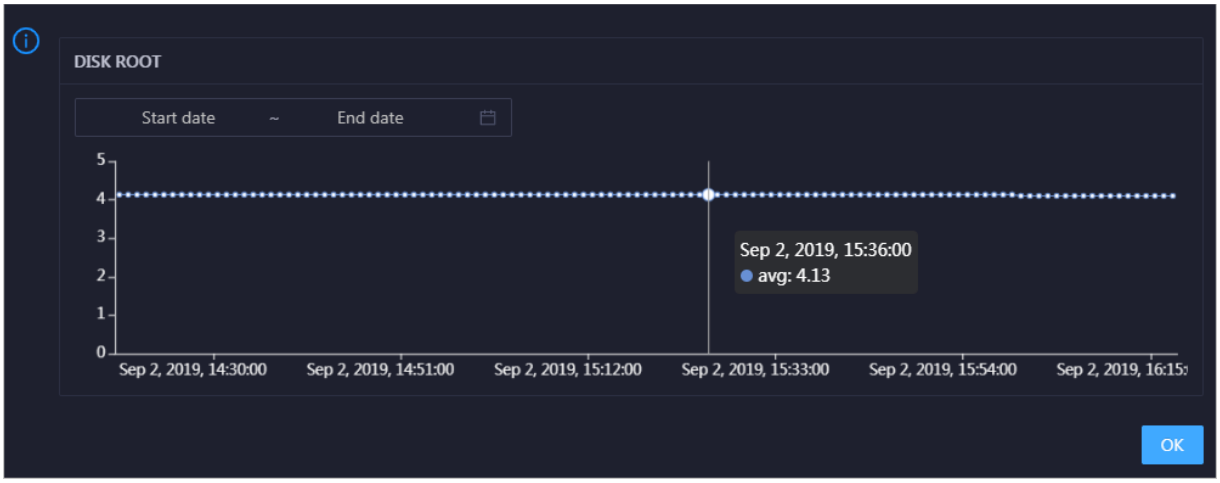

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

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

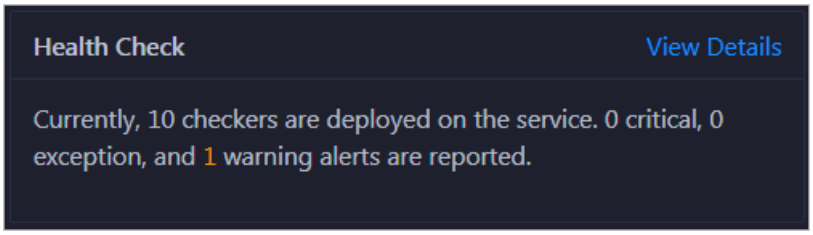

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

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

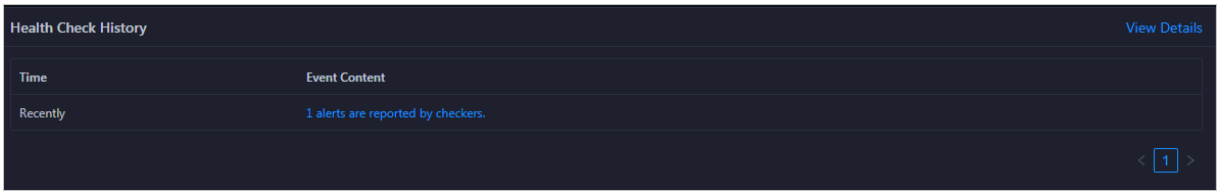

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

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

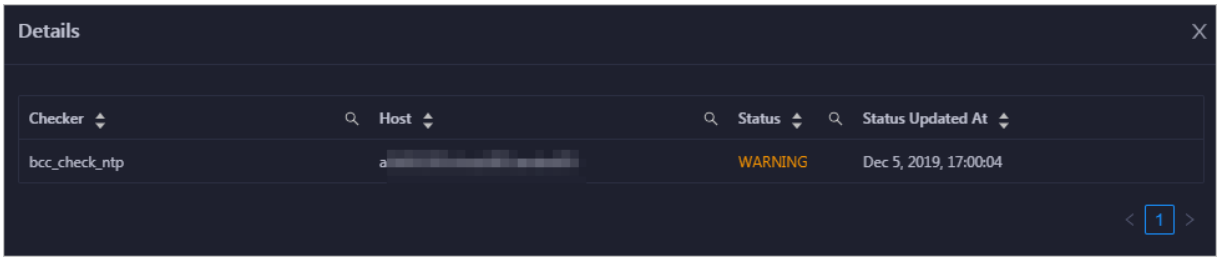

6.4.1.11.7.2. Host health

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.

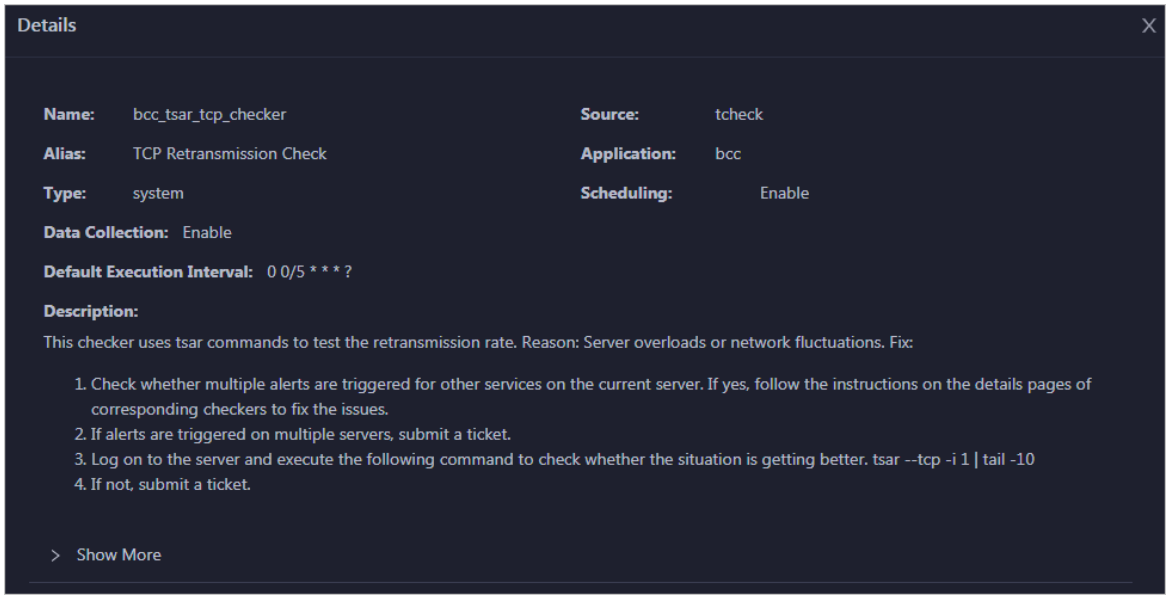

   The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

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

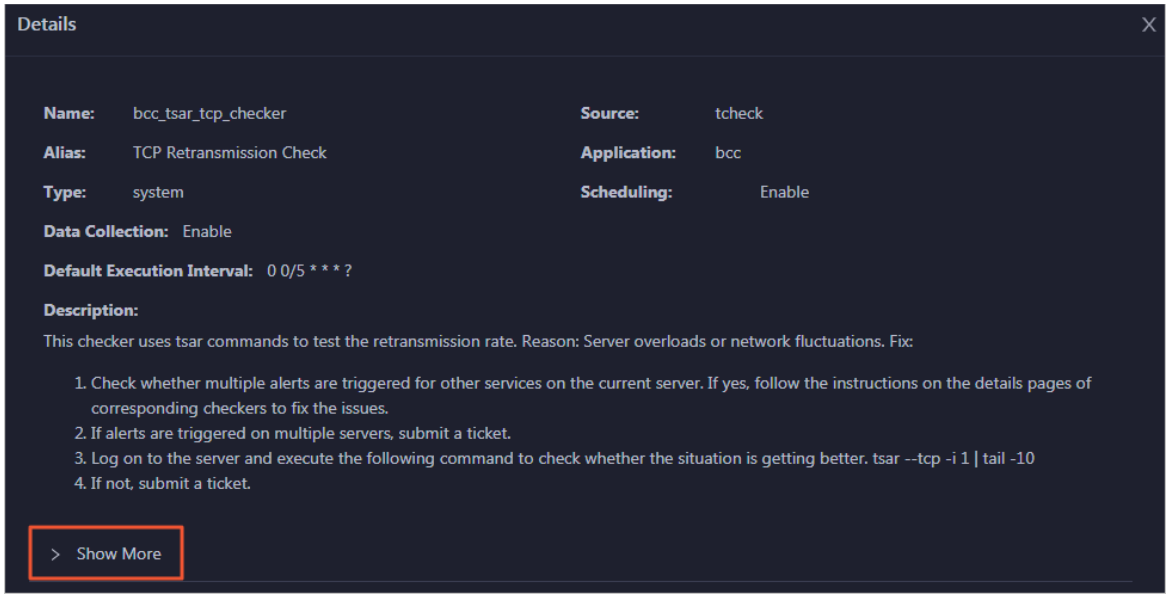

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

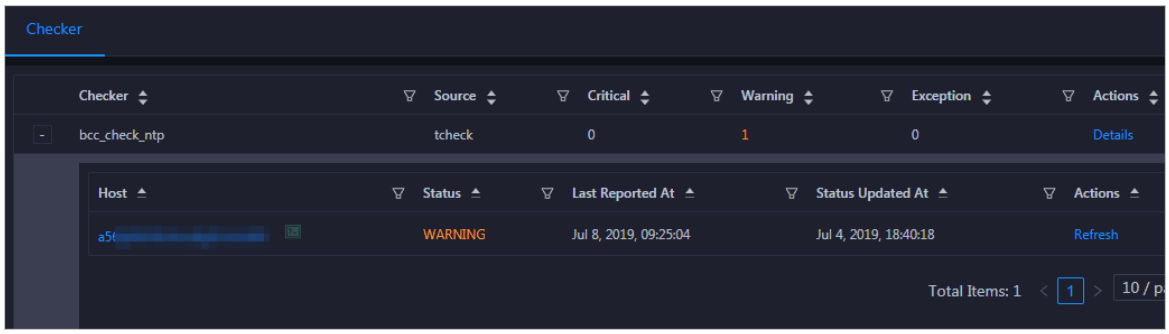

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

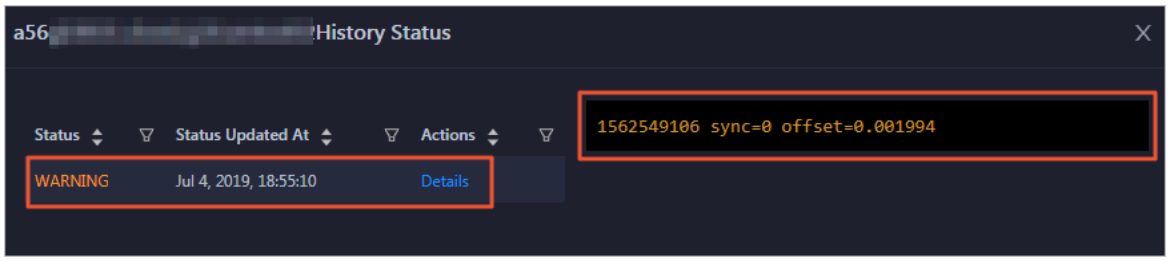

Clear alerts

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.

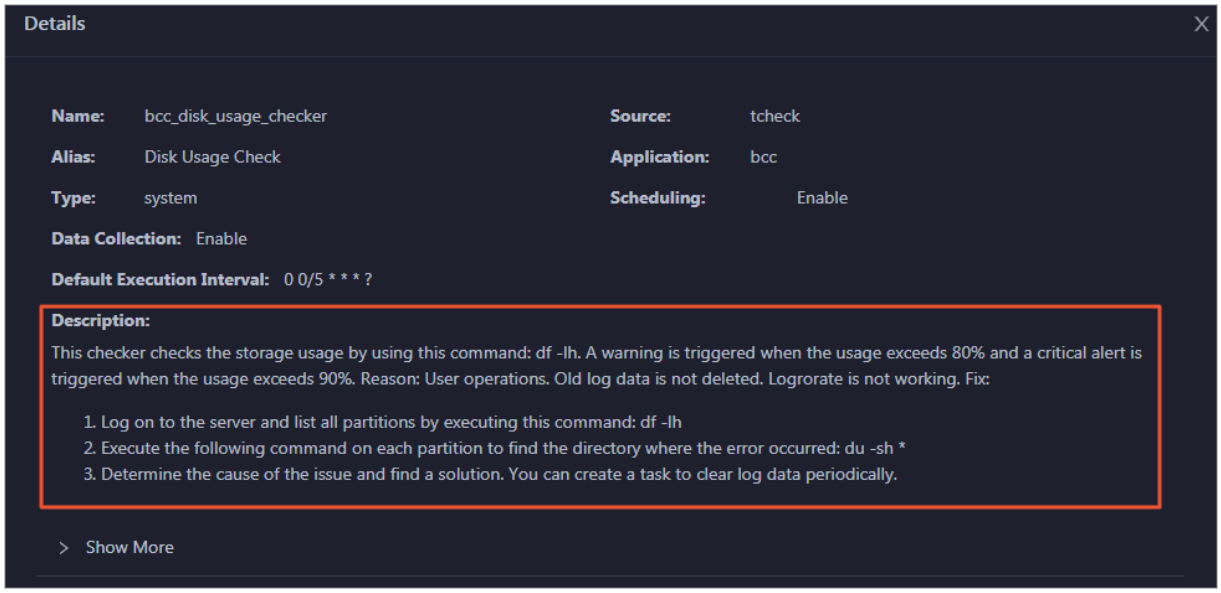

Log on to a host

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.

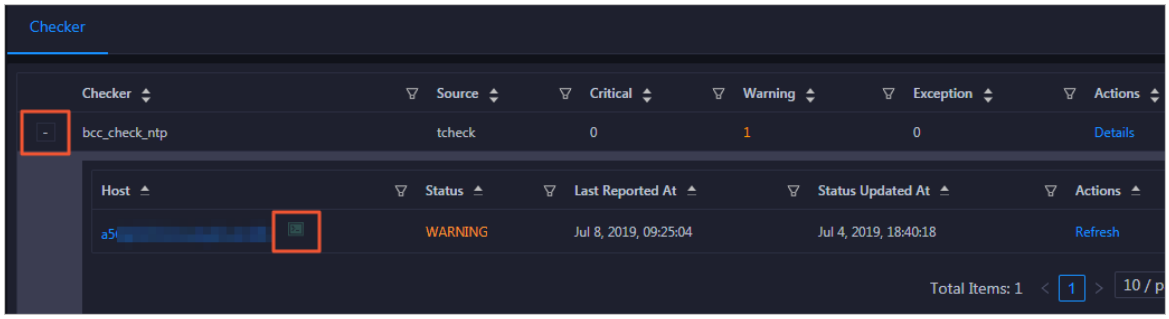

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

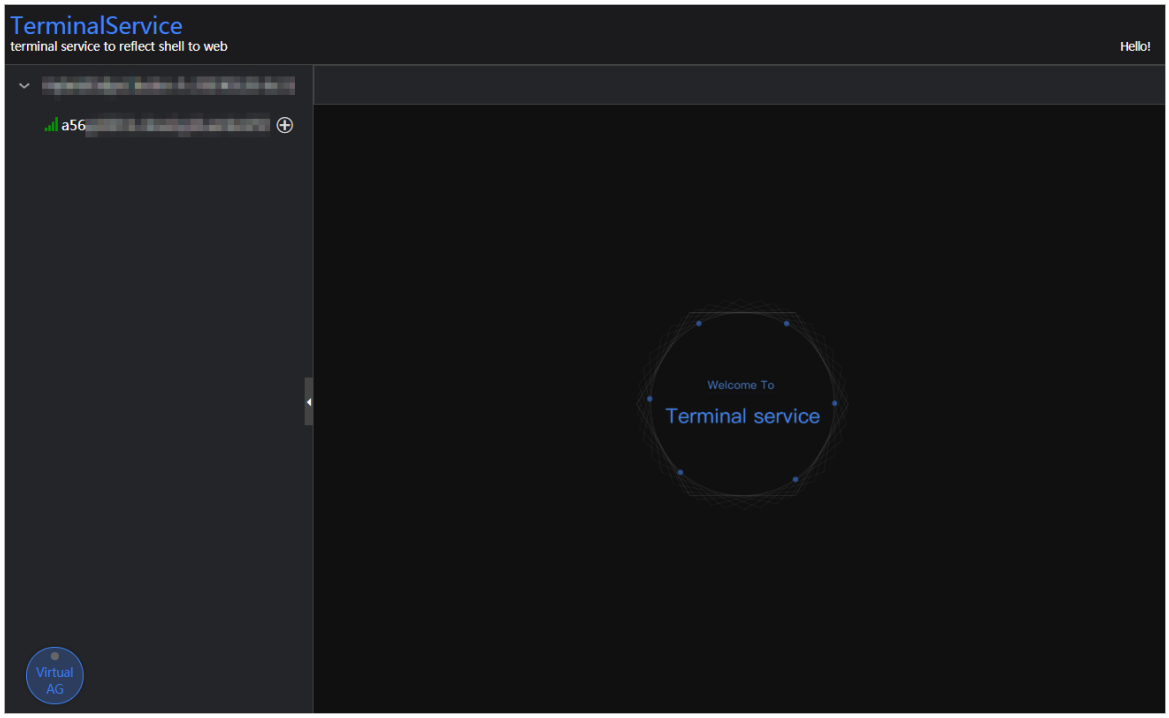

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

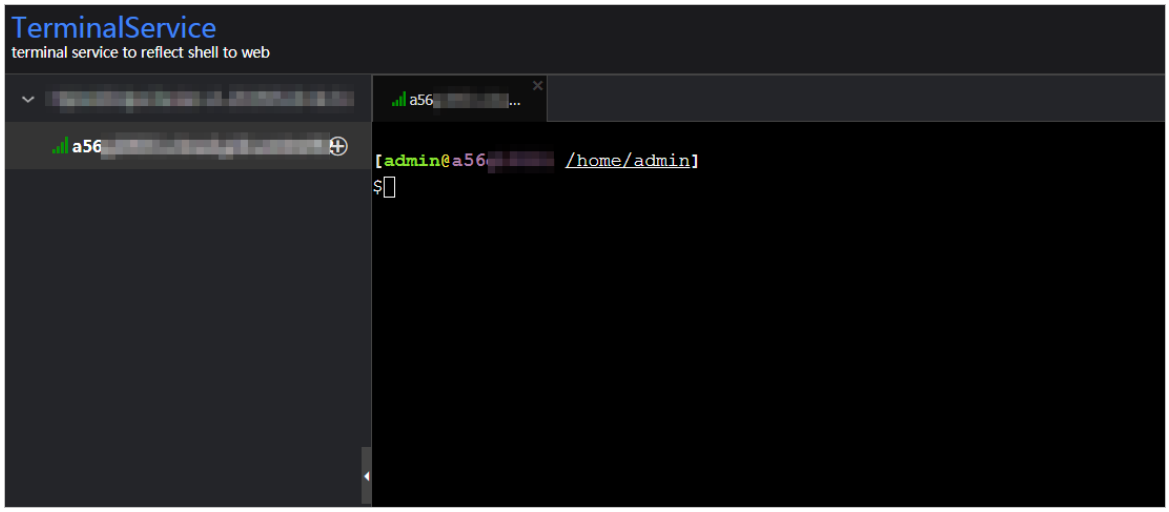

Run a checker again

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.

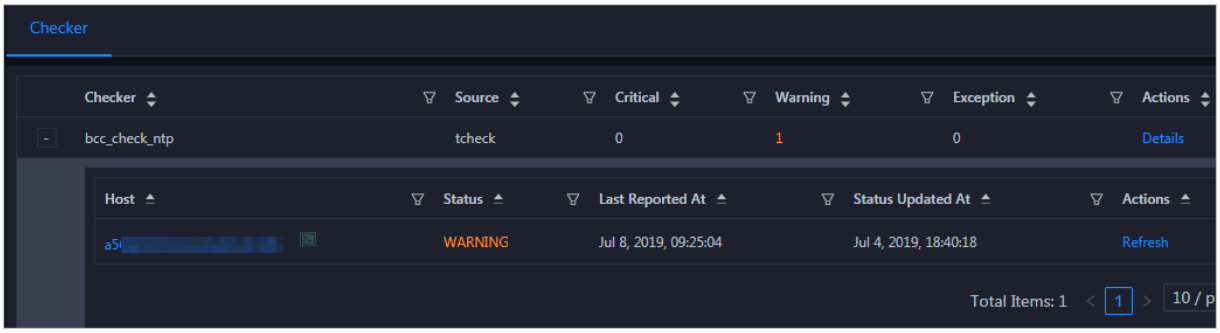

6.4.1.11.7.3. Host charts

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.

6.4.1.11.7.4. 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.

6.4.1.11.8. Job and queue analysis

6.4.1.11.8.1. 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.

6.4.1.11.8.2. 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.

6.5. Apsara Big Data Manager (ABM)

6.5.1. Operations and Maintenance Guide

6.5.1.1. Routine maintenance

6.5.1.1.1. 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.

6.5.1.1.2. 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.

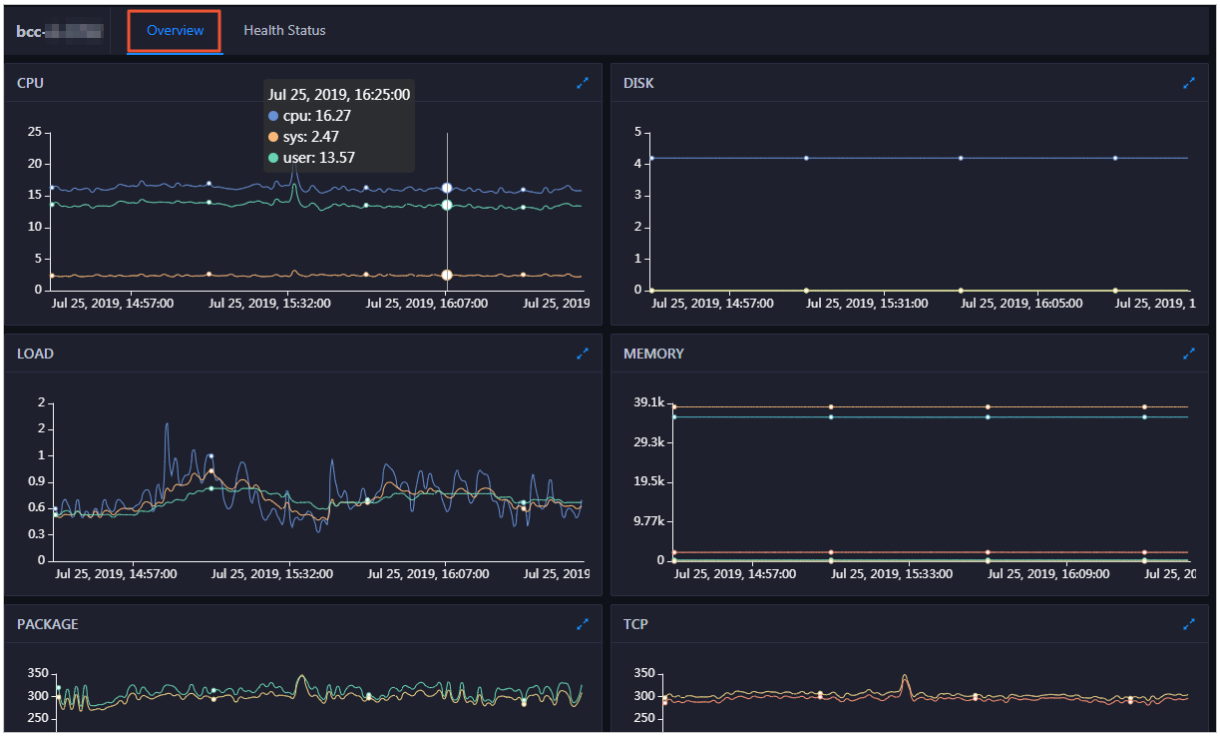

The **Overview** tab displays tendency charts for cluster metrics, including the CPU, memory, disk, load, package, TCP, and disk root directory usage. You need to regularly view and record these metrics to evaluate system downtime risks in the future.

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.

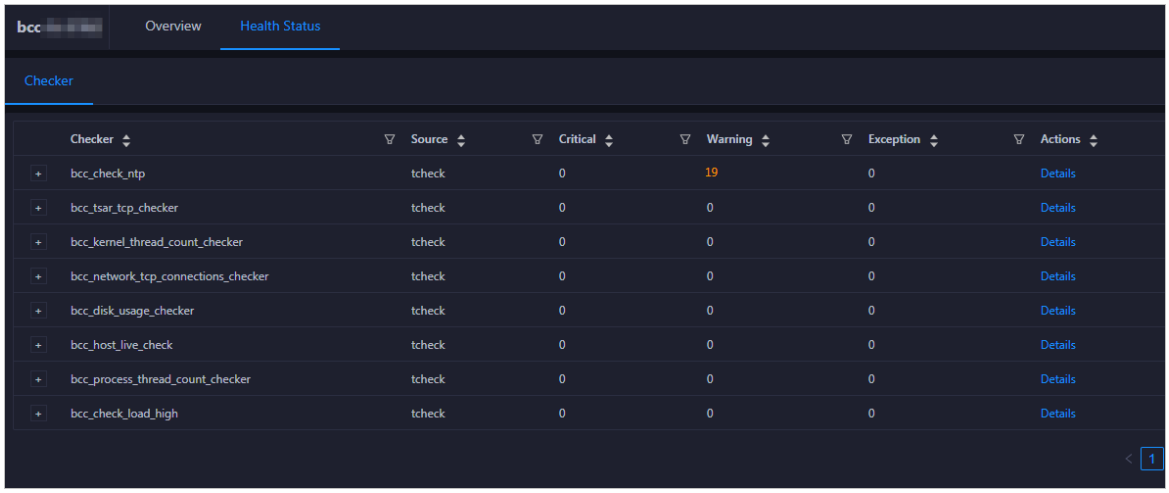

The **Health Status** tab displays all check items and the alerts that were generated during the check.

2. Click the **Fold** icon for a check item with alerts. All hosts on which the check item was performed appear.

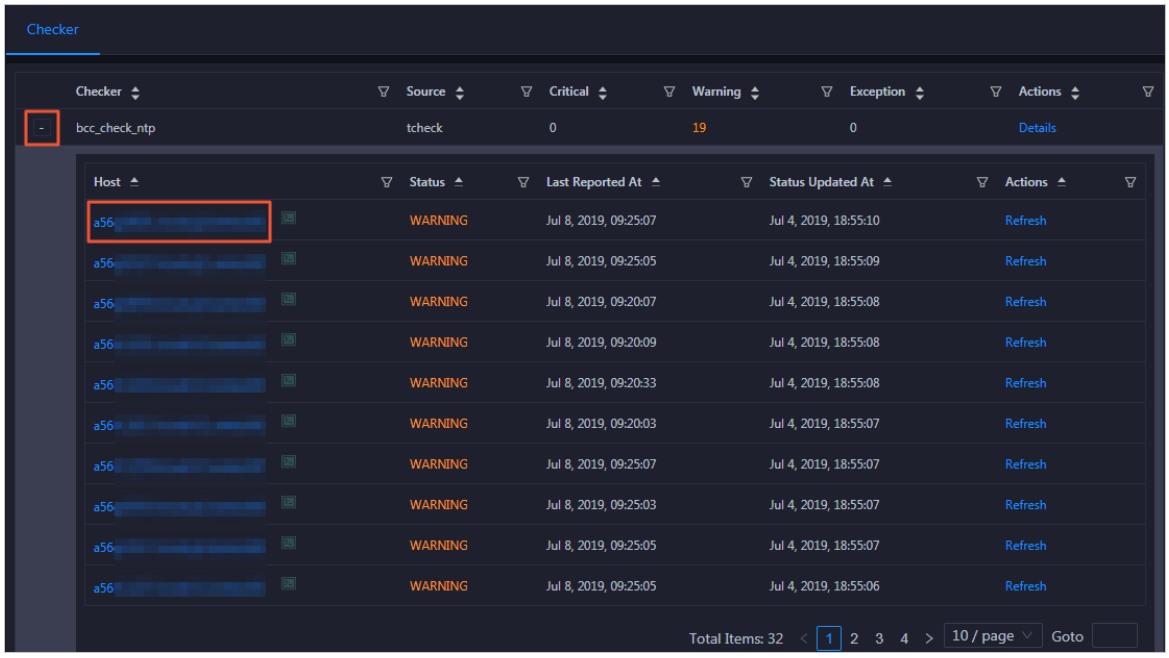

3. Click a host. In the dialog box that appears, click **Details** for an alert. The alert cause appears on the right.

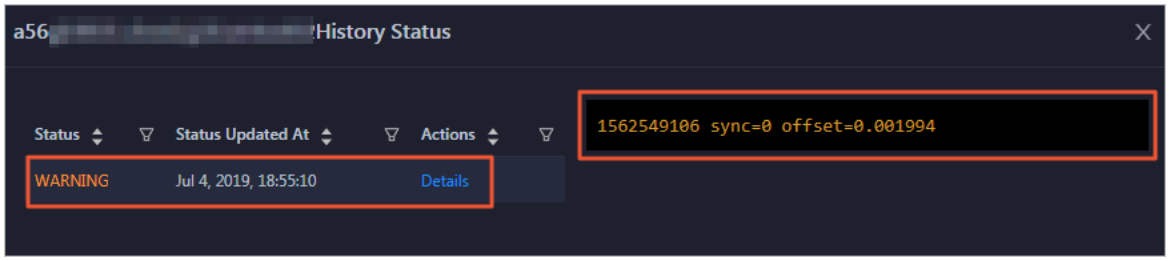

4. Click **Details** for a check item with an alert and view the fix method for the alert in the dialog box that appears.

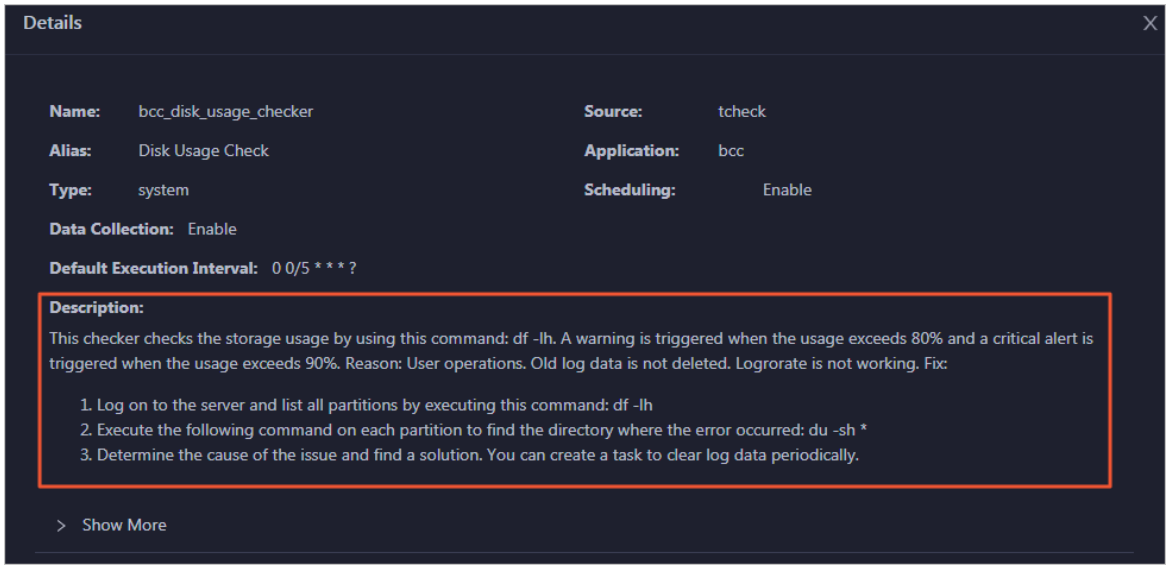

5. Handle the alert based on the fix method.

   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.

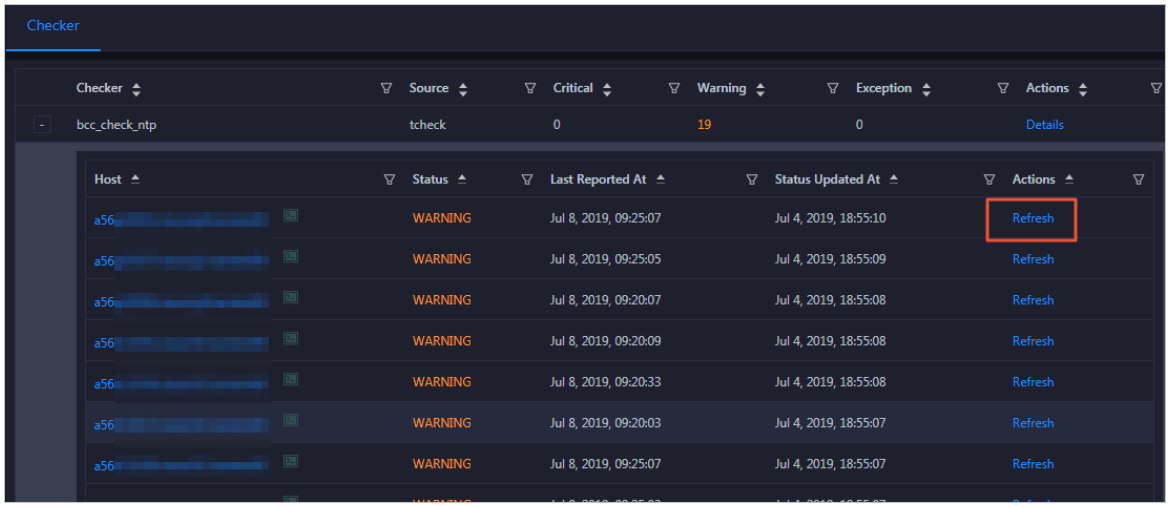

Log on to a host

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.

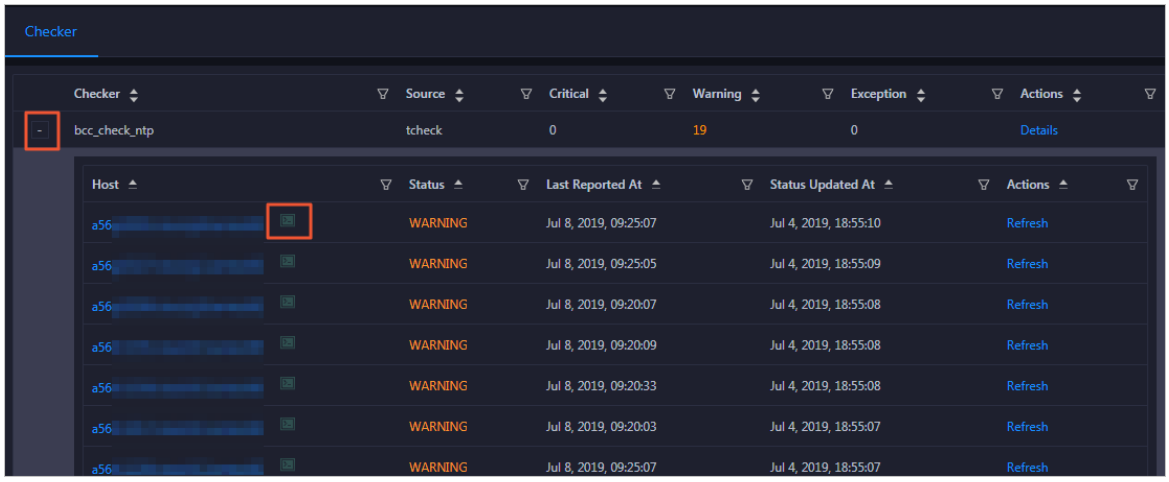

2. Click the **Logon** icon for a host. The **TerminalService** page appears.

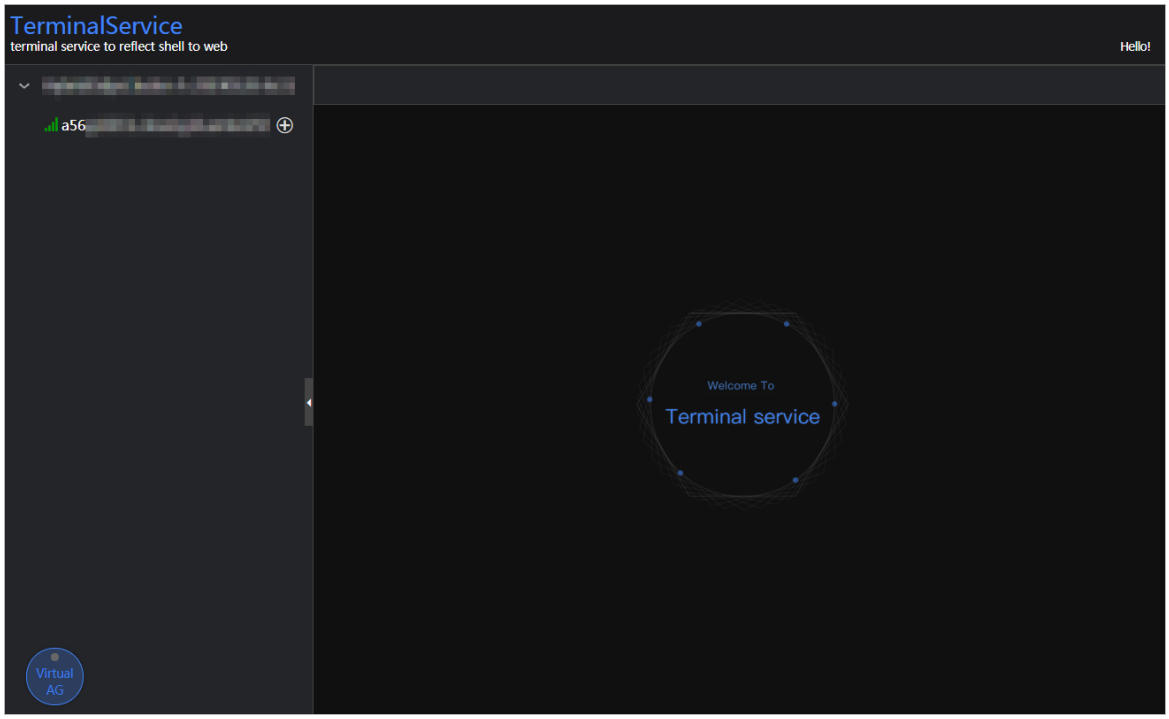

3. On the **TerminalService** page, select the host on the left to log on to it.

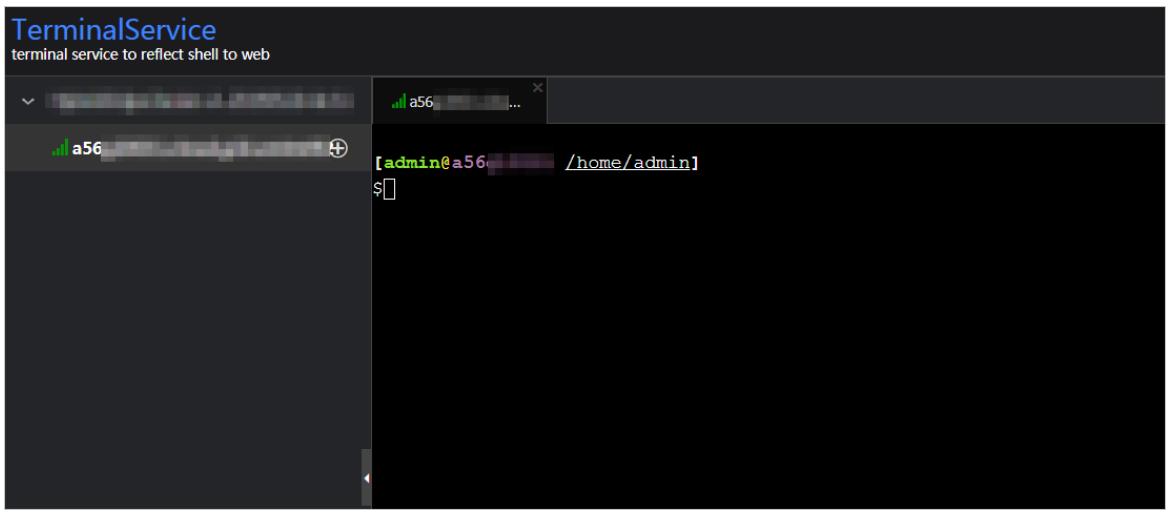

6.5.1.1.3. Troubleshooting

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.

6.5.1.2. Backup and restore

Back up data

ABM uses a high-availability database. You do not need to manually back up data. To obtain full backup data, contact technical support.

Restore data

You do not need to restore data for ABM.

6.6. Machine Learning Platform for AI

6.6.1. Operations and Maintenance Guide

6.6.1.1. Query server and application information

6.6.1.1.1. Apsara Stack Machine Learning Platform for AI

6.6.1.1.1.1. 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.

6.6.1.1.1.2. Log on to a server

Machine Learning Platform for AI is deployed based on Apsara Infrastructure Management Framework. Its application information and database information can be found by accessing the corresponding Apsara Infrastructure Management Framework address. This topic describes how to log on to a server.

Context

Each module is deployed on two servers with the same application package and configuration. You can log on to the back-end server through the server IP address and perform operations.

Procedure

1. Ensure that the network is connected and the IP address of the jump server has been obtained.
2. Log on to the jump server.
3. Switch to the root account.
4. All applications are deployed by using a Docker container. You can run the following command to view the current container:

   ```
   sudo docker ps
   ```

5. Run the following command to go to the container:

   ```
   sudo docker exec -ti container_id /bin/bash
   ```

   The application log is stored in the */home/admin/logs/${app}* path.

6.6.1.1.1.3. 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
   ```

6.6.1.1.1.4. Restart an application service

The application structures and directories of the PaiCap, PaiDmscloud, and PaiJcs modules are almost the same. You can restart an application service in either of the following ways:

- Log on to the container and run the following command to restart the service:

  ```
  sudo -u admin /home/admin/pai-dms/bin/appclt.sh restart
  ```

- Run the following command on the server to restart the container:

  ```
  sudo docker restart $container_id
  ```

Run the following command to check whether the service is restarted:

```
curl localhost/status.taobao
```

6.6.1.1.2. Online model service

6.6.1.1.2.1. Query online model service information

Check the online model service status

Online model services are deployed in the Kubernetes cluster. Log on to the master node in the Kubernetes cluster and run the following command to query the service deployment status:

```
kubectl get pod -n eas-system
```

If no errors occur, all pods in the STATUS column display *Running*.

If not, run the following command to perform troubleshooting:

```
kubectl describe pod ${pod_name} -n eas-system
```

View the online model service configurations

1. Log on to the homepage of Apsara Infrastructure Management Framework.
2. Click the **C** tab and search for **pai**. Hover over the dots next to the PAI cluster, and choose **Dashboard** from the shortcut menu.
3. Search for the *eas-sentinel* role and log on to the VM from the terminal.
4. Run the `docker ps |grep eas-sentinel` command to view the ID of the container for the sentinel.
5. Run the `docker logs ${sentinelcontainerid}` command to view the output log, which contains the

configuration information of the online model service.

6.6.1.1.2.2. Log on to the online model service container

Prerequisites

Ensure that the network is connected and the IP address of the jump server has been obtained.

Procedure

1. Log on to the jump server.
2. Switch to the root account.
3. All applications are deployed with a container. Run the following command to log on to the current pod:

   ```
   kubectl exec -ti ${pod_name} -n ${pod_namespace} – bash
   ```

6.6.1.1.2.3. Restart a pod

Procedure

1. Log on to the master node in the Kubernetes cluster.
2. Run the `kubectl get` command to find the corresponding *pod name*.
3. Run the following command to restart the pod:

   ```
   kubectl delete ${pod_name}
   ```

6.6.1.1.3. DSW service

6.6.1.1.3.1. View resources and application configurations

This topic describes how to view the resources in the control plane and data plane, and the application configurations in Data Science Workshop (DSW).

Context

DSW can be deployed only in Kubernetes clusters. To view DSW resources, you must log on to the master node of the Kubernetes cluster that you use to deploy DSW.

Procedure

1. Log on to the master node of the Kubernetes cluster.
   i. Log on to the Oracle parallel server (OPS) in the Apsara Stack environment. The OPS serves as a jump server.
   ii. Run the following command to navigate to the master node of the Kubernetes cluster from the OPS:

   ```
   ssh master
   ```

2. View the resources in the control plane.

   You can run the following command to query all resources in the control plane:

   ```
   kubectl get all -n pai-dsw-system
   ```

   The following code provides an example of the return result:

   ```
   NAME                  READY  STATUS   RESTARTS  AGE
   pod/pai-dsw-gateway-686cf7****-6v7c2  1/1    Running  1      69d   // The DSW gateway, which forward
   s traffic to the user pod.
   pod/pai-dsw-gpu-q****           1/1    Running  1      22d   // The Kubernetes plug-in that provides supp
   ort for GPUs.
   pod/pai-dsw-init-6fd784****-tjl9f     1/1    Running  0      16d   // Initializes data in DSW.
   pod/pai-dsw-notebook-679d6f****-lxnth  1/1    Running  0      2d16h  // The backend application for DS
   W management.
   pod/pai-dsw-redis-749bbc****-mczdn    1/1    Running  0      69d   // Provides cache services for DSW.
   pod/pai-dsw-vip-5f9d5f****-5rmvn      1/1    Running  0      10d   // Initializes video-based intelligent alg
   orithms.
   NAME             TYPE     CLUSTER-IP     EXTERNAL-IP  PORT(S)   AGE
   service/pai-dsw-gateway   ClusterIP  XX.XX.XX.XX  <none>     80/TCP    36d  // The cluster IP address of t
   he pod where the DSW gateway is deployed.
   service/pai-dsw-notebook  ClusterIP  XX.XX.XX.XX  <none>     80/TCP    36d  // The cluster IP address of
   the pod where the backend application for DSW management is deployed.
   service/pai-dsw-redis     ClusterIP  XX.XX.XX.XX  <none>     6379/TCP  36d  // The cluster IP address of th
   e pod where DSW cache services are deployed.
   NAME              DESIRED  CURRENT  READY  UP-TO-DATE  AVAILABLE  NODE SELECTOR  AGE
   daemonset.apps/pai-dsw-gpu  0      0      0    0       0       dsw=gpu      22d  // The DaemonSet of the K
   ubernetes plug-in that provides support for GPUs.
   NAME                 READY  UP-TO-DATE  AVAILABLE  AGE
   deployment.apps/pai-dsw-gateway   1/1    1       1       69d
   deployment.apps/pai-dsw-init      1/1    1       1       69d
   deployment.apps/pai-dsw-notebook  1/1    1       1       69d
   deployment.apps/pai-dsw-redis     1/1    1       1       69d
   deployment.apps/pai-dsw-vip       1/1    1       1       11d
   // The Deployments for the preceding pods. Application auto-recovery is supported and you can add m
   ore replicas.
   ```

3. View the resources in the data plane.

   You can run the following command to query all resources in the data plane:

   ```
   kubectl get all -n dsw-resource
   ```

   The following code provides an example of the return result:

   ```
   NAME                            READY  STATUS   RESTARTS  AGE
   pod/b2021031812484239c0d979871d1000055853-5c6c84****-wpfwf  1/1    Running  0      16d  // The DS
   W instance.
   NAME                       TYPE     CLUSTER-IP     EXTERNAL-IP  PORT(S)  AGE
   service/b2021031812484239c0d979871d100005****  ClusterIP  11.166.126.132  <none>     80/TCP   21d
   // The cluster IP address of your DSW instance.
   NAME                           READY  UP-TO-DATE  AVAILABLE  AGE
   deployment.apps/b2021031812484239c0d979871d100005 **** 1/1 1 1 21d // The Deployment for your D
   SW instance. Instance auto-recovery is supported.
   ```

4. View application configurations.

You can run the following command to query the application configurations of DSW:

```
kubectl get cm -n pai-dsw-system
```

The following code provides an example of the return result:

```
NAME                        DATA  AGE
pai-dsw-controller          34    22d  // The application configurations of DSW.
container-runtime-image-limits  1     70d
dsw-base                    1     21d
init.check                  1     70d
pai-dsw-base                1     69d
pai-dsw-cpu                 1     70d
pai-dsw-gateway             4     69d
pai-dsw-gpu                 1     22d
pai-dsw-init                1     69d
pai-dsw-notebook            1     69d
pai-dsw-redis               2     69d
pai-dsw-vip                 1     69d
// Other configmap objects, which are provided as extensions.
```

6.6.1.1.4. GPU cluster and task information

6.6.1.1.4.1. 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 `rttrl` command to view all GPU workers in the current GPU cluster.

   If the Other column displays `FUXI_GPU:200`, the worker has two GPUs. If the column displays `FUXI_GPU:800`, the worker has eight GPUs.

6.6.1.1.4.2. 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 `r al` command.

4. Run the `r cru` command to view the resources allocated to the current cluster, including CPU, memory, and FUXI_GPU resources.

5. **Notice** Use caution when performing this step.

   Run the `r jstop WorkItemName` command to stop a Fuxi task.

   `WorkItemName`: specifies the values in the first column displayed by the `r al` command.

6.6.1.2. Maintenance and troubleshooting

6.6.1.2.1. Machine Learning Platform for AI maintenance

6.6.1.2.1.1. 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

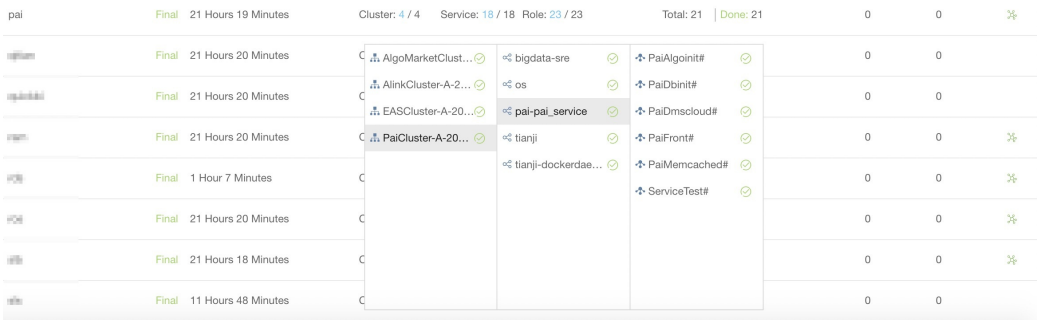

6. Run the `sudo docker restart e90f70353031` command to restart the ServiceTest service, as shown in the following figure.

   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

   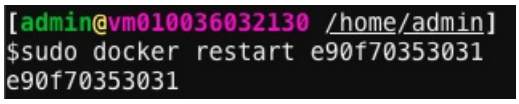

   - PASS: The algorithm is running properly.
   - SKIP or FAIL: The algorithm fails.

6.6.1.2.1.2. Common faults and solutions

Maintenance commands

nc, telent, curl, ping, mysql

`docker images` : shows all images on a server.

`docker ps` : shows the running images on a server.

`docker exec –ti containerID /bin/bash`

`docker log containerID` : shows the container log.

`curl http://localhost/status.taobao` : determines whether the SpringBoot service is started.

pai.xx.xx access failures

Procedure

1. Run the `ping pai.xx.xx` command to check whether the domain name has been translated to the corresponding VIP.

   If the domain name cannot be resolved properly, contact the on-site engineer to check the network configurations.

2. Run the `curl http://ip/status.taobao` command to check whether all service modules are running normally.

   If the status.taobao module fails the check, perform the following operations:

   i. Log on to the server to check whether the container is active.

   ii. Go to the container and run the following command to check whether the service process is active:

   `ps –lef | grep java`

   iii. View the */home/admin/{app}/logs/err_pai-dms.log* file to locate causes, such as dependent tenant service request timeout, dependent OCS timeout, and database connection exceptions.

   We recommend that you view the log after checking all items in the checklist to verify whether the malfunction was not caused by a component exception.

3. Verify whether ApsaraDB RDS is accessible.

   i. Run the following command to check whether the port is active:

   `nc –v –z $rds_host $port`

ii. Run the following command to check whether the database is accessible:

```
mysql -h$Host -P$Port -u$user -p$password
```

4. Verify whether the caching service is functioning properly.

   Run the following command to check whether port 11211 is active:

   ```
   nc -v -z $ocs_host 11211
   ```

   Search for ocs_host as follows:

   i. Search for the dmscloud instance, as shown in the following figure.

   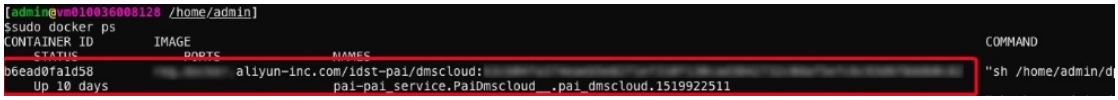

   ii. Run the `sudo docker inspect b6ead0fa1d58 | grep ocs` command to view the ocs_host information, as shown in the following figure.

   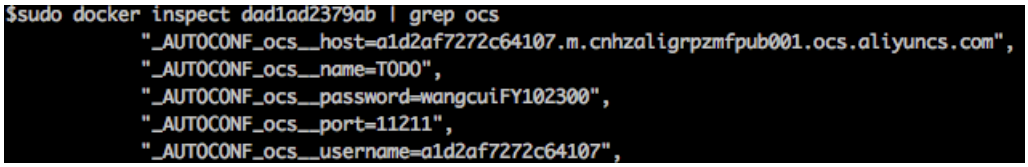

   `host` is a list of servers on which OCS (caching service) is deployed. `port` indicates the port number.

   Machine Learning Platform for AI in Apsara Stack typically uses the built-in memcached service as the dependent caching service. If port 11211 is inaccessible, log on to the server and run the following command to restart the memcached service:

   ```
   docker restart containerid
   ```

Experiment failures

We recommend that you run a Machine Learning Platform for AI experiment in Google Chrome version 66 or later. Google Chrome is the only supported browser.

- Components cannot be dragged and dropped.

  Clear cookies and caches, and then retry. Check the version of Chrome. If the problem persists, it is due to a service failure. Log on to the container to view the log.

- An error message is displayed while an algorithm is running.

  If an error message is displayed, the task has been submitted to MaxCompute. Check the parameters and source data against the user guide and algorithm descriptions to locate the error.

Other failures

If a problem persists after you have checked all items by referring to pai.xx.xx access failures, troubleshoot the underlying dependency services, including MaxCompute and DataWorks (tenants and metadata).

- MaxCompute: Make sure that MaxCompute can pass the pai_console test.
- DataWorks: Make sure the configured domain name is accessible, and verify the application log.

If no errors are found, restart the service.

6.6.1.2.2. Online model service maintenance (must be activated separately)

Node maintenance

Online model service nodes are Kubernetes nodes. You can run the `kubectl get node` command to view all nodes in a cluster. A healthy node is in the Ready state. When a node is not in the Ready state, the one of the following errors may have occurred:

- Node failures

  There are many reasons that may cause a node to fail. Typically, a node fails when the kernel crashes or the disk does not have sufficient space. If the node can be restarted properly, it rejoins the cluster after it is restarted. If the node cannot be restarted properly, contact the corresponding ECS support personnel.

- Docker daemon exceptions

  A Docker daemon exception rarely occurs. Docker daemon exceptions are typically caused by storage issues. Run the `systemctl restart docker` command to restart the Docker daemon.

Online model service maintenance

- A service cannot be created or deleted.
  - If Error 500 is returned while an operation is called, the configurations of the eas-ui component are incorrect. Contact Apsara Stack delivery engineers.
  - If a creation or deletion operation is called but no response is returned in a timely manner, the jobworker of the service does not work properly. Check whether the KVStore for Redis service in the cluster is normal. If not, restart the pod for KVStore for Redis.
- The system fails to read the monitoring data.

  Check whether the influxdb-0 pod under *eas-system* is created properly. If the pod is not in the running state, an influxdb out of memory error has occurred. You can expand the influxdb-0 memory.

Service maintenance

- Service creation failures.

  The request is sent but the service creation result displays **Failed**. A model error has caused a crash. The system then fails to create the model. Check whether the model code contains any null pointers or has any other problems.

- The system fails to obtain the monitoring data.

  Check whether the influxdb-0 of each service is normal. The service cannot be created because a persistent volume cannot be created. Check whether the Apsara Stack environment has sufficient disk space. If influxdb-0 runs properly but you cannot obtain the monitoring data, restart the influxdb-0 pod.

6.6.1.2.3. FAQ about DSW O&M

This topic answers frequently asked questions (FAQ) about the operations and maintenance (O&M) of DSW.

What can I do if I fail to start my DSW instance?

Run the `kubectl get all -n dsw-resource` command to query the instance resources in the data plane. The following code provides an example of the return result:

```
NAME                                READY  STATUS   RESTARTS  AGE
pod/b2021031812484239c0d979871d1000055853-5c6c84****-wpfwf  1/1   Running  0     16d
NAME                          TYPE     CLUSTER-IP    EXTERNAL-IP  PORT(S)  AGE
service/b2021031812484239c0d979871d100005****  ClusterIP  XX.XX.XX.XX  <none>    80/TCP  22d
NAME                                READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/b2021031812484239c0d979871d100005****  1/1   1      1      22d
```

If the pod is abnormal, run the following command to view the value of the `Events` parameter and the error cause in the return result:

```
kubectl describe -n dsw-resource po b2021031812484239c0d979871d1000055853-5c6c84****-wpfwf
...
Events:      <none>
```

The following figure shows an example of the return result.

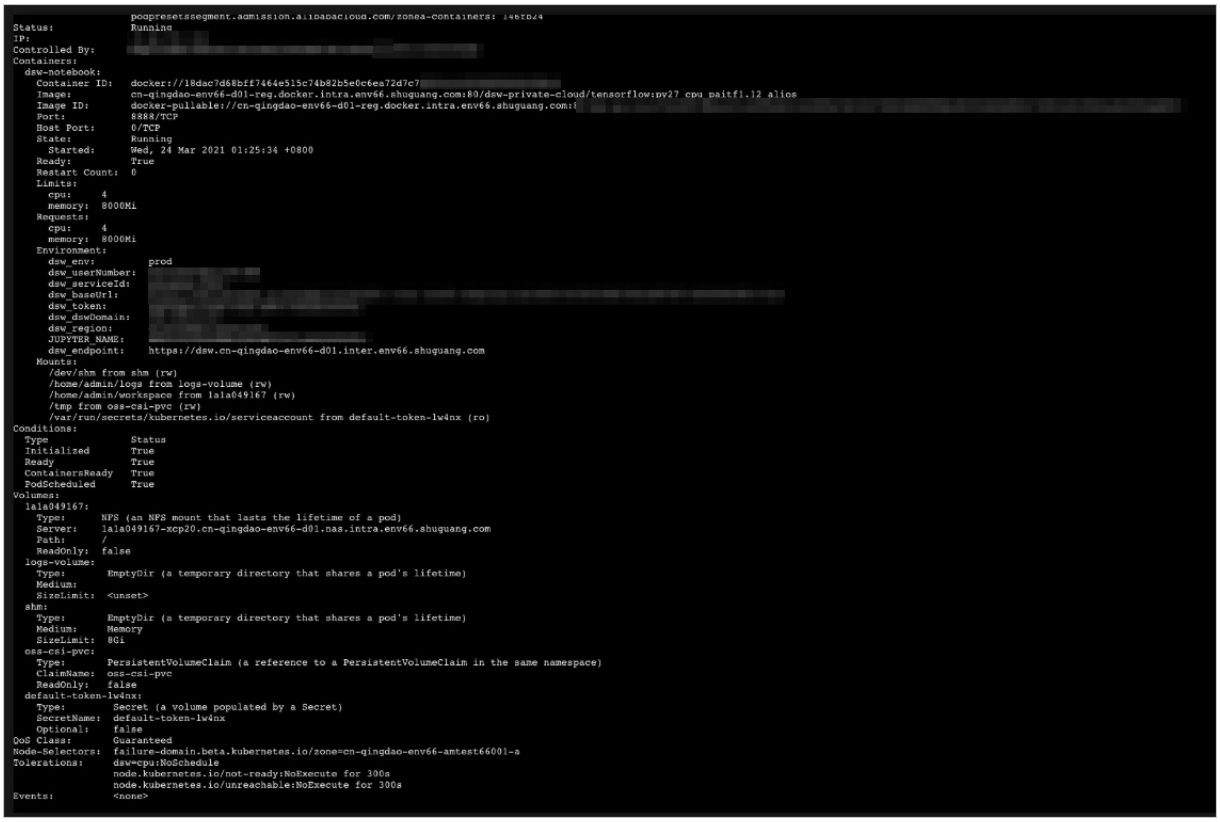

What can I do if the access to my DSW instance fails after the instance is started in the cluster?

1. Run the `kubectl get ingress -n dsw-resource` command in the Kubernetes cluster to query the Ingress. The following code provides an example of the return result:

```
NAME                  HOSTS                          ADDRESS     PORTS  AGE
b2021031812484239c0d979871d100005****  dsw-instance.cn-qingdao-env66-d01.intra.env66.shuguan
g.com  XX.XX.XX.XX    22d
```

2. Run the following command to check whether the Ingress works as required. If the Ingress does not work, contact Kubernetes engineers of Alibaba Cloud to troubleshoot the error.

   ```
   ping <Ingress Host>
   ```

   You must replace `Ingress Host` with the value of the `HOSTS` field that is obtained in the preceding step.

6.6.1.2.4. GPU cluster maintenance (deep learning must be activated separately)

Node maintenance

A deep learning node is a server where a GPU cluster runs.

1. Perform steps 1 through 3 in Query GPU cluster information and log on to ApsaraAG of the GPU cluster.
2. Run the `r ttrl` command to view all nodes that support deep learning tasks.

- Node failures

  There are many reasons that may cause a node to fail. Typically, a node fails when the kernel crashes or the disk does not have sufficient space. If the node can be restarted properly, it rejoins the cluster after it is restarted. If the node cannot be restarted properly, contact the corresponding service support team.

- Docker daemon exceptions

  A Docker daemon exception rarely occurs. Docker daemon exceptions are typically caused by storage issues. Run the `systemctl restart docker` command to restart the Docker daemon.

Service maintenance

Failure to allocate resources to a task

Perform the following steps for troubleshooting:

1. Perform steps 1 through 3 in Query GPU cluster information and log on to ApsaraAG of the GPU cluster.
2. Run the `r quota` command to view the quota information of the GPU cluster.
3. Run the `r cru` command to view the resources allocated to each task in the current cluster.
4. Run the `r al` command to view all tasks submitted to the cluster.
5. Run the `r wwl WorkItemName` command to view the status of a specific task.
   - If only **ChildMaster** is displayed, no resources are allocated to the worker.
   - If **worker name** is displayed but no **hostname** is displayed, service resuming is pending or has failed. Log on to the server of the ChildMaster and locate the error. You can also contact the service support team.
6. Run the `r ttrl` command to check the value of **FUXI_GPU** in the **Other** column. If the value is 200,

the worker has two GPUs. If the value is 800, the worker has eight GPUs.

7. Log on to a GPU worker in the worker list obtained in Step 3 over SSH. Run the `nvidia-smi` command to view the GPU status. If an exception occurs, contact the relevant service support personnel.

6.7. DataHub

6.7.1. Operations and Maintenance Guide

6.7.1.1. Concepts and architecture

6.7.1.1.1. 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

```
[admin@datahub-ext-ay03-st3-ag /home/admin]
$r al
WorkItemName                    | NuwaAddress
Datahub/ShipperServiceEXTAY03   | nuwa://datahub-ext-ay03-st3:10240/Datahub/ShipperServiceEXTAY03/ServiceMaster
Datahub/XStreamServiceEXTAY03   | nuwa://datahub-ext-ay03-st3:10240/Datahub/XStreamServiceEXTAY03/ServiceMaster
Datahub/CoordinatorServiceEXTAY03 | nuwa://datahub-ext-ay03-st3:10240/Datahub/CoordinatorServiceEXTAY03/ServiceMaster
```

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

```
$r wwl Datahub/XStreamServiceEXTAY03
total resource planned for the workitem:
[('CPU', 1600), ('Memory', 111616)]
detail:
worker name                     | process start time       | status  | tubo's address
ChildMaster                     | Fri Jan 19 10:48:12 2018 | Running | tcp:
XStreamBroker@b25f09396.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09397.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09399.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09402.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09407.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09416.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09424.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f09430.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f12348.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f12359.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f12363.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamBroker@b25f12373.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamMeter@b25f09397.cloud.st3    | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamMetric@b25f09397.cloud.st3   | Fri Jan 19 10:48:18 2018 | Running | tcp:
XStreamRecycler@b25f09397.cloud.st3 | Fri Jan 19 10:48:18 2018 | Running | tcp:
```

6.7.1.1.2. Architecture

6.7.1.1.2.1. Architecture

Architecture shows the architecture of DataHub.

Architecture

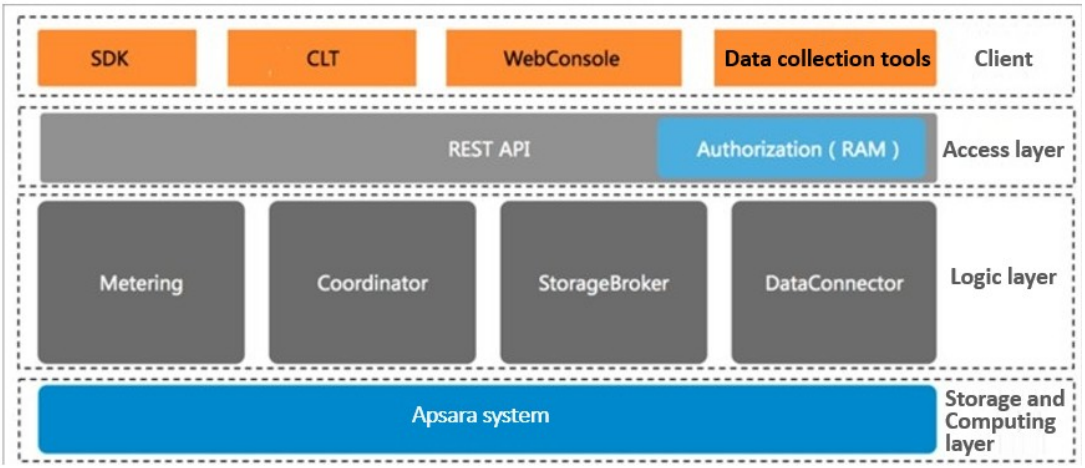

The architecture of DataHub consists of four layers: **clients**, **access layer**, **logic layer**, and **storage and scheduling layer**.

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.

6.7.1.1.2.2. Technical architecture

Technical architecture of DataHub shows the technical architecture of DataHub.

Technical architecture of DataHub

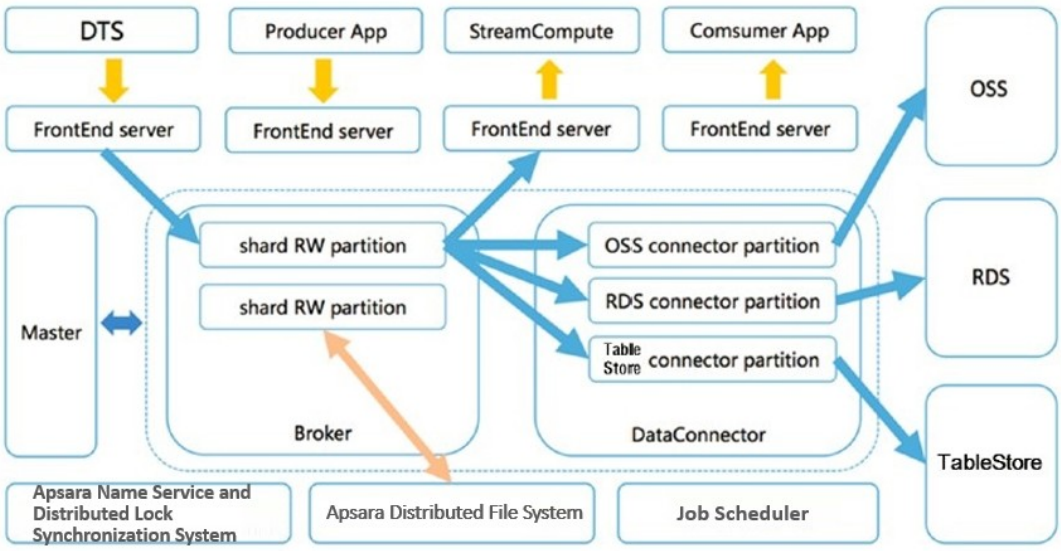

The figure shows the process from data ingestion to consumption.

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.

6.7.1.2. Commands and tools

6.7.1.2.1. 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#
```

6.7.1.2.2. 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
  ```

6.7.1.2.3. 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.

6.7.1.2.4. 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

```
[admin@docker010001040152 /home/admin]
$r wwl Datahub/XStreamServicex
total resource planned for the workitem:
[('CPU', 1100), ('Memory', 76800)]
detail:
worker name                          | process start time       | status  | tubo:o address
ChildMaster                          | Sat Feb 24 12:04:27 2018 | Running | tcp:/
XStreamBroker@rc8255132.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8255133.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8255134.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8255138.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8255140.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8255141.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8255142.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamBroker@rc8f73140.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamMetric@rc8255141.cloud.nu17   | Sat Feb 24 12:04:44 2018 | Running | tcp:/
XStreamRecycler@rc8255141.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp:/
```

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

```
[admin@rc8255138 /apsara/tubo/TempRoot/Datahub/XStreamServicex/tool]
$ ./xstream_tool -x x mo
{"XStreamBroker": {
        "LoadingPartitions": {},
        "OverloadTopic": {},
        "PartitionNumber": {
            "rc8255132.cloud.nu17": 11,
            "rc8255133.cloud.nu17": 12,
            "rc8255134.cloud.nu17": 12,
            "rc8255138.cloud.nu17": 11,
            "rc8255140.cloud.nu17": 11,
            "rc8255141.cloud.nu17": 10,
            "rc8255142.cloud.nu17": 10,
            "rc8f73140.cloud.nu17": 11},
        "StartingWorker": {},
        "UnloadingPartitions": {}}}
```

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

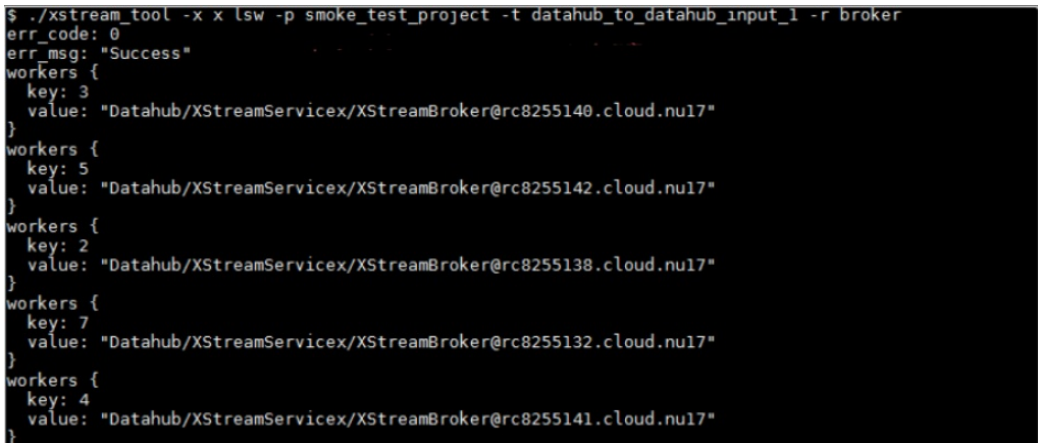

From the command output, you can find the anomalous shards.

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.

6.7.1.2.5. View performance statistics in the DataHub console

In the DataHub console, you can obtain performance statistics to facilitate O&M.

For information about how to log on to the DataHub console, see the Log on to the DataHub console topic in *User Guide*.

To view 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 project whose performance statistics you want to view and click **View** in the Actions column.
2. On the project details page, find the topic whose performance metrics you want to view and click **View** in the Actions column.
3. On the topic details page that appears, click the **Metric Statistics** tab to view the charts that display the performance statistics of this topic.

6.7.1.2.6. Apsara Big Data Manager

Apsara Big Data Manager provides O&M on big data services from the perspective of business, services, clusters, and hosts. You can upgrade big data services, customize alert configurations, and view the O&M history in the Apsara Big Data Manager console.

Apsara Big Data Manager allows onsite Apsara Stack engineers to manage big data services. For example, they can view resource usage, check and handle alerts, and modify configurations.

For information about how to log on to the Apsara Big Data Manager console and the O&M operations of DataHub, see the relevant topics in *DataHub O&M*.

6.7.1.3. Routine maintenance

6.7.1.3.1. 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.

6.7.1.3.2. 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

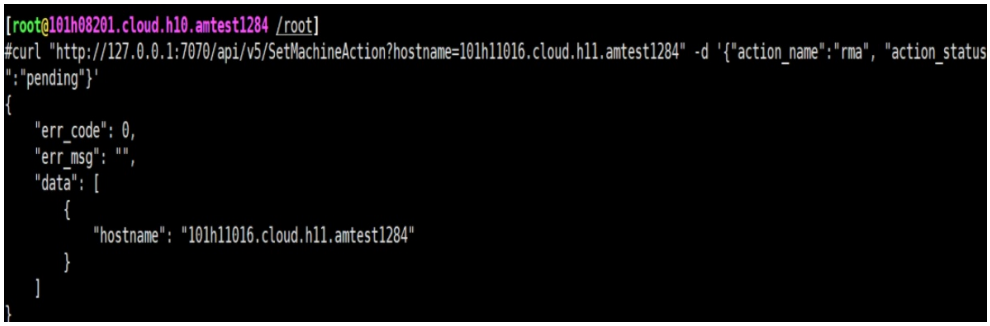

> **Note** Replace the IP address and hostname in the sample code with those of your anomalous chunkserver host.

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

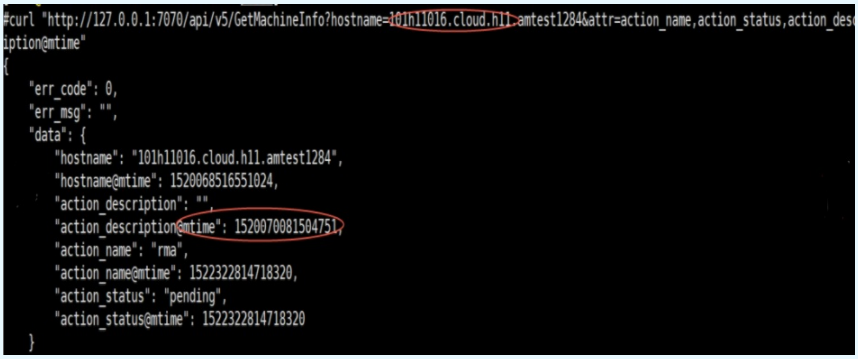

2. Wait for approval.
   i. Check the action status of the host.

   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.

6.7.1.3.3. 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"
```

6.7.1.3.4. 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.

6.7.1.4. DataHub O&M

6.7.1.4.1. Log on to the ABM console

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

Prerequisites

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

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

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

Procedure

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

   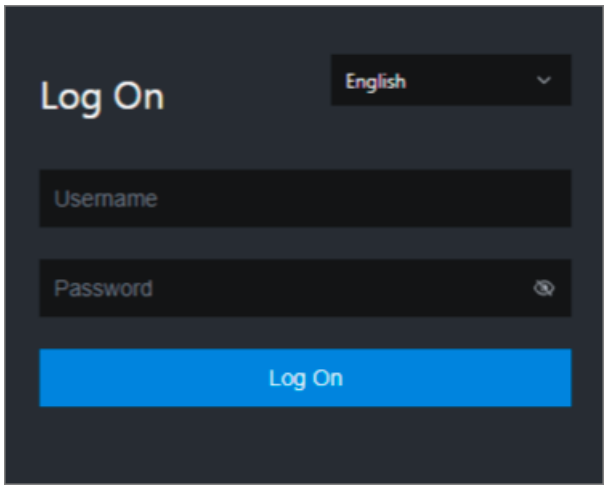

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.
4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**.
6. In the left-side navigation pane, choose **Product Management > Products**.
7. In the **Big Data Services** section, choose **General-Purpose O&M > Apsara Big Data Manager**.

6.7.1.4.2. Common operations

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

Search for a project quickly

You can quickly search for a project based on the project name.

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

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

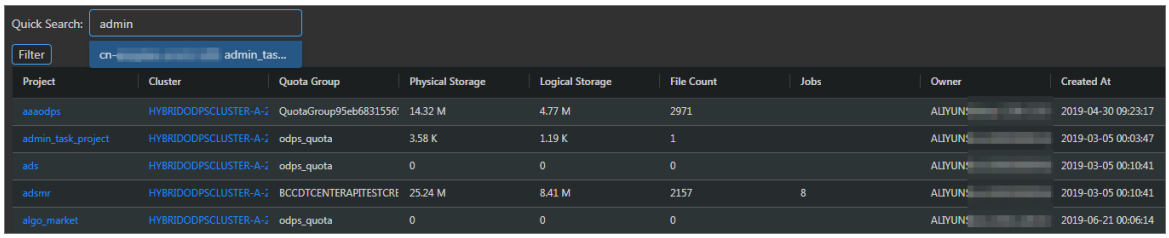

Example:

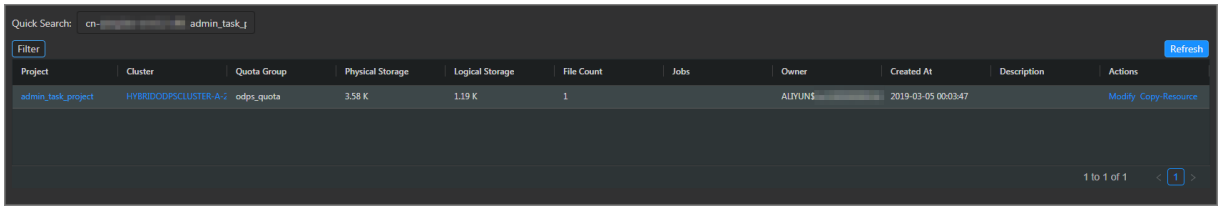

Filter projects

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

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

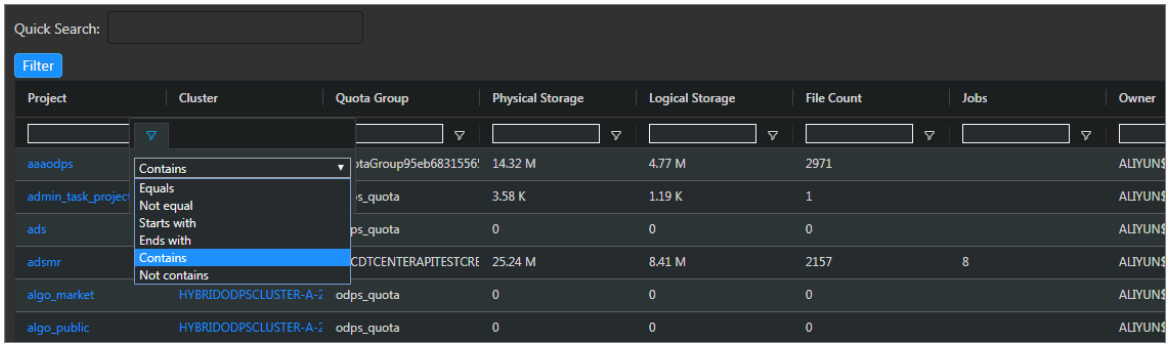

Optional filtering methods include:

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

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

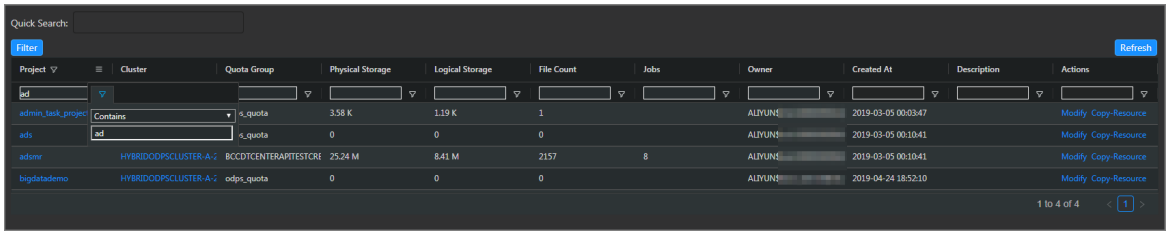

5. If the filtering result is not accurate, you can continue performing this operation on other columns.

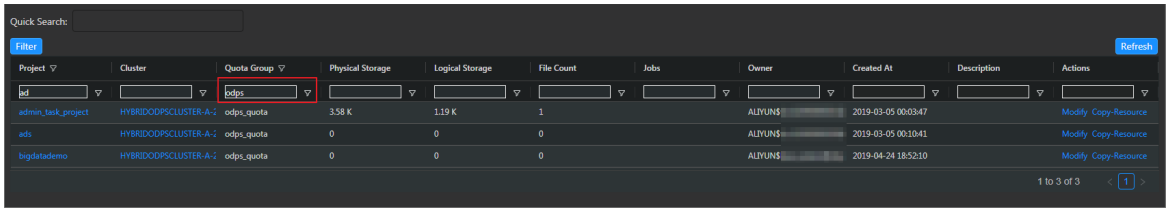

After you set the filter conditions for the projects, the **Filter** button is highlighted. If you need to cancel filtering, click the highlighted **Filter** button.

Search for items

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

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

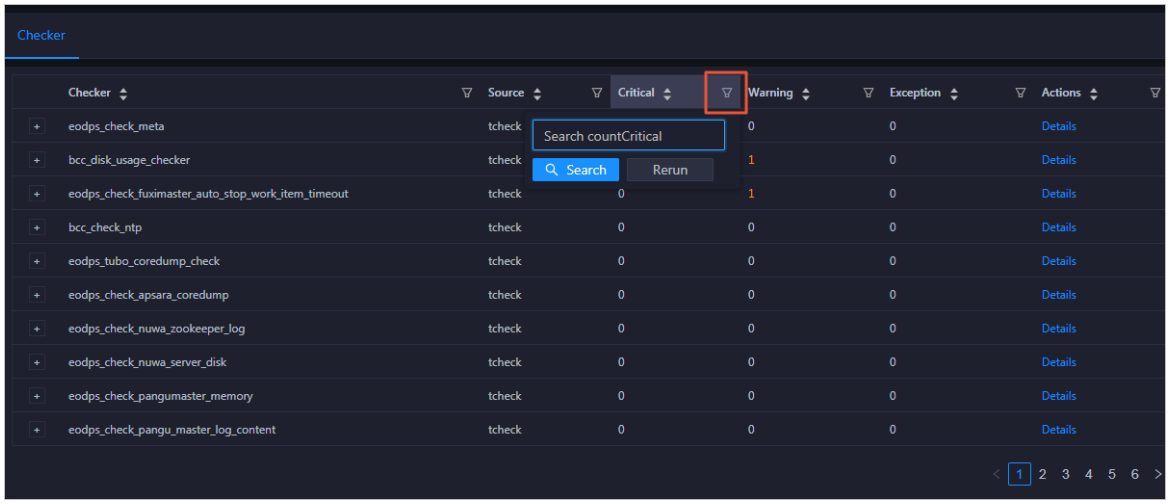

3. Click **Search**. The checkers that meet the requirements appear.
4. If the search result is not accurate, you can continue performing this operation on other columns.

Customize a column

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

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

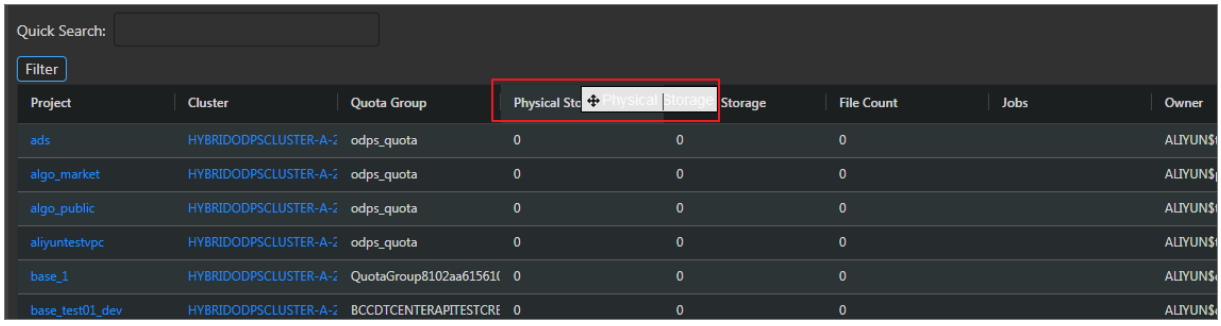

You can click ☰ in a column heading to customize the column.

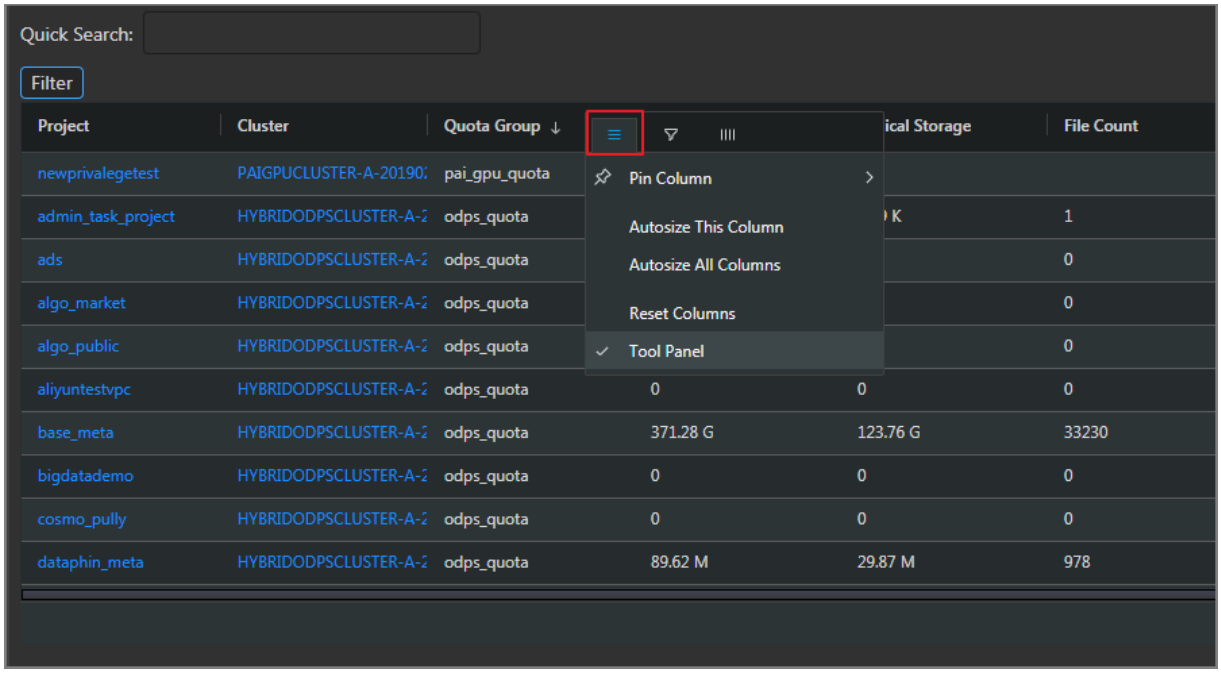

- **Pin Column**: allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.
- **Autosize This Column**: allows you to adjust the width of a column automatically.
- **Autosize All Columns**: allows you to adjust the width of all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

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

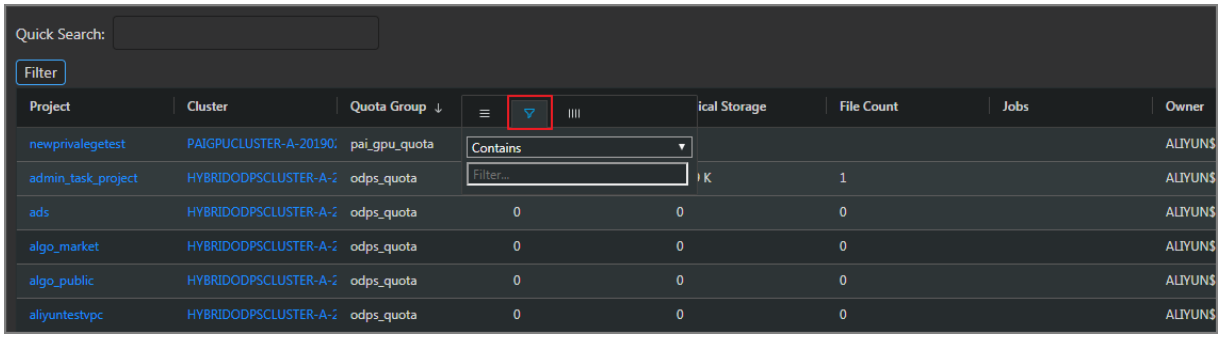

Click in a column heading and select the columns to be displayed.

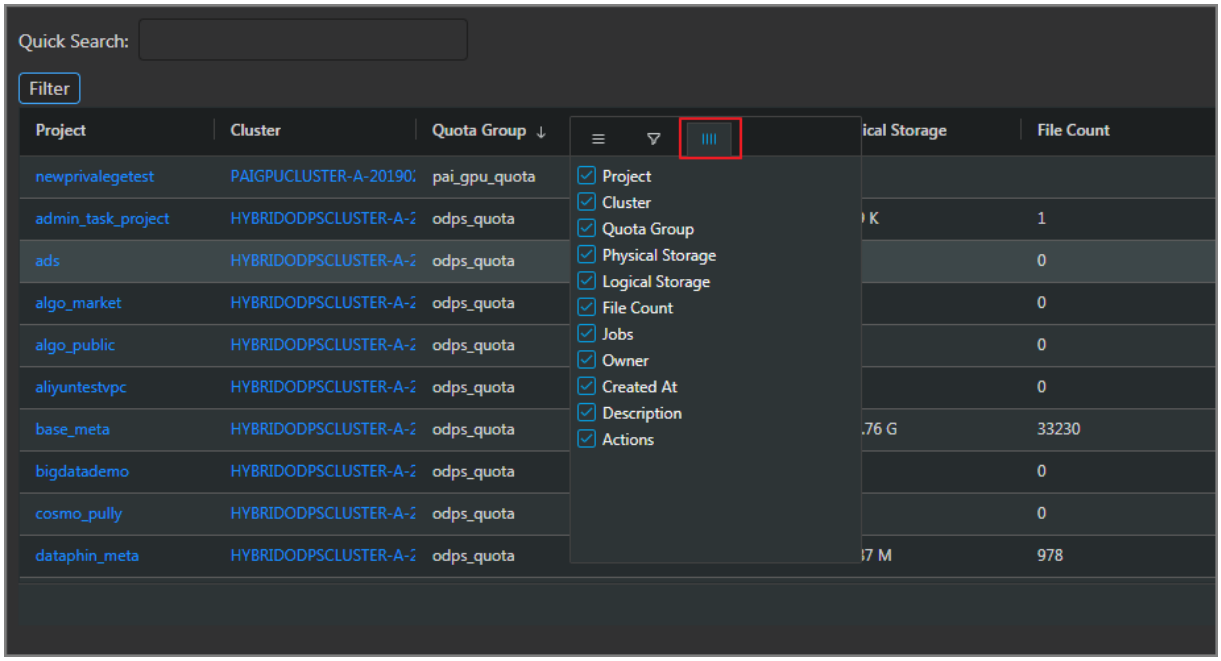

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

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

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

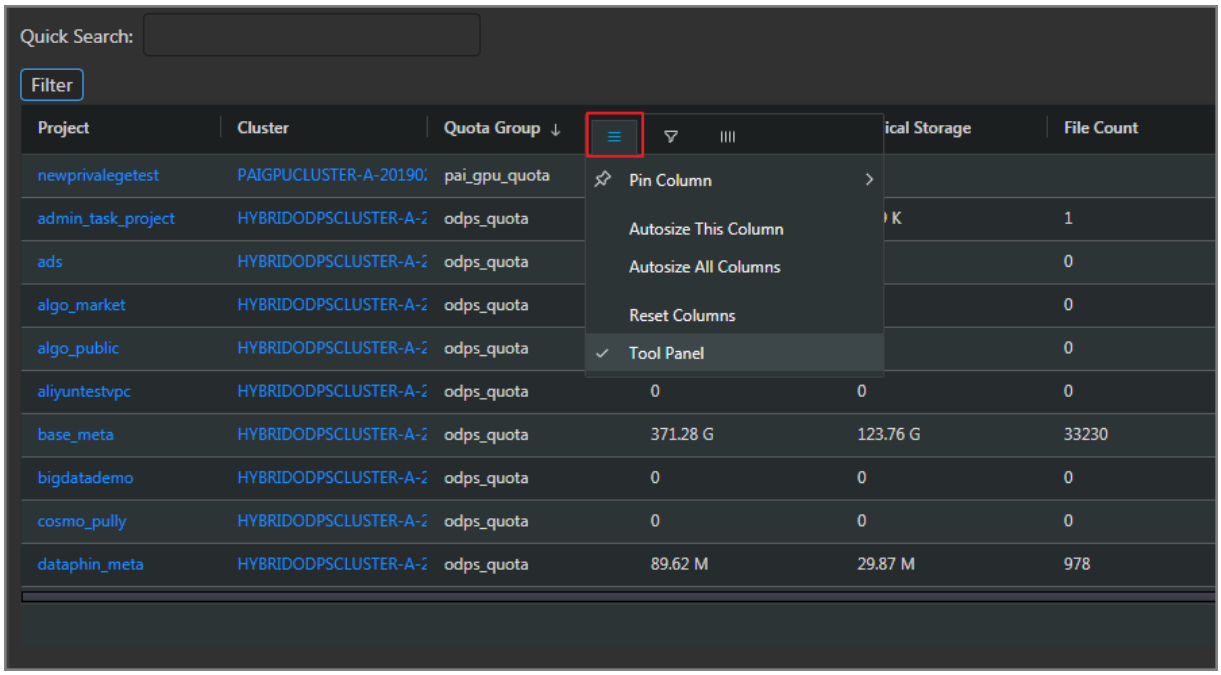

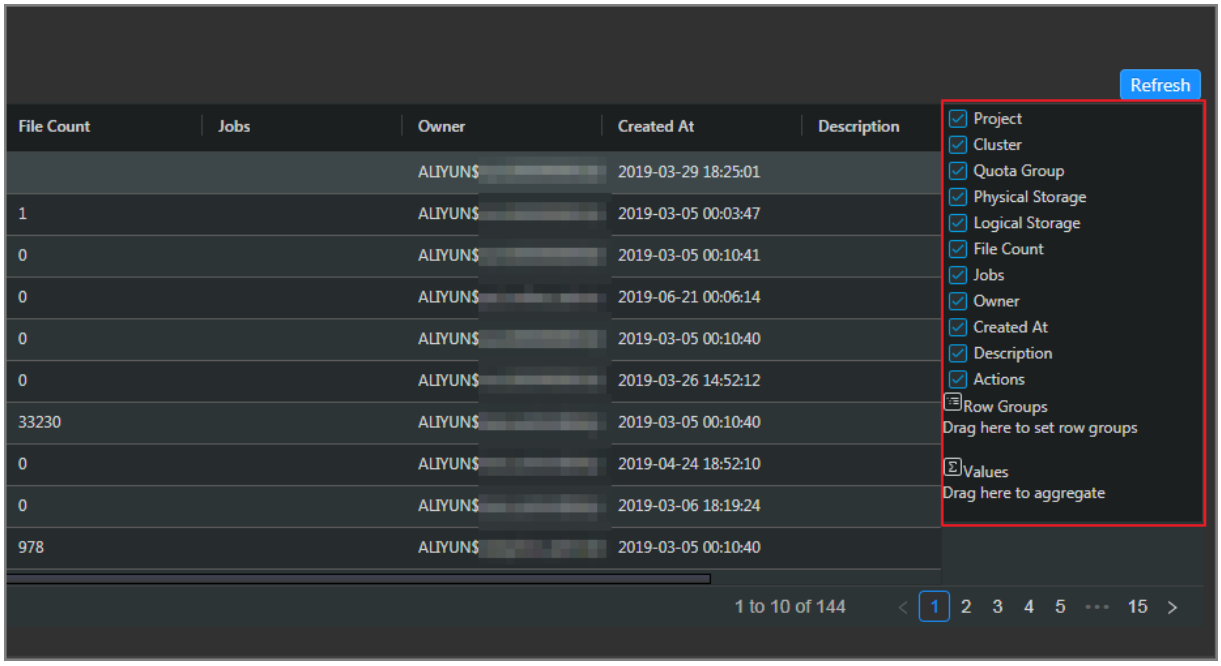

Sort projects based on a column

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

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

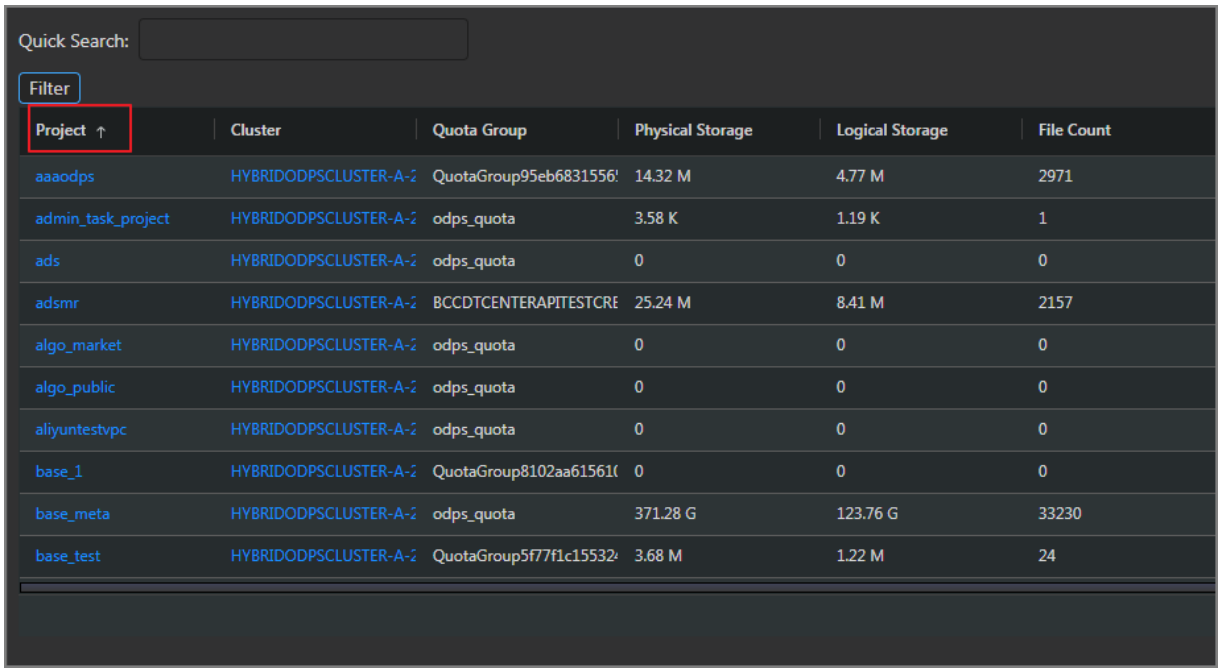

Sort items based on a column

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

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

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

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

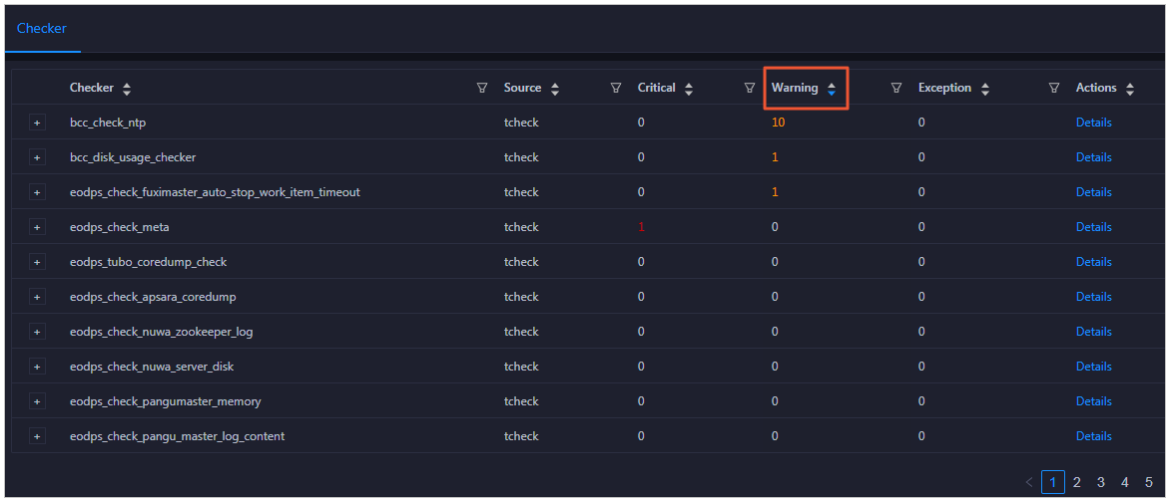

The highlighted up arrow indicates that the checkers are sorted in ascending order. The highlighted down arrow indicates that the checkers are sorted in descending order.

Trend chart 1

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

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

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

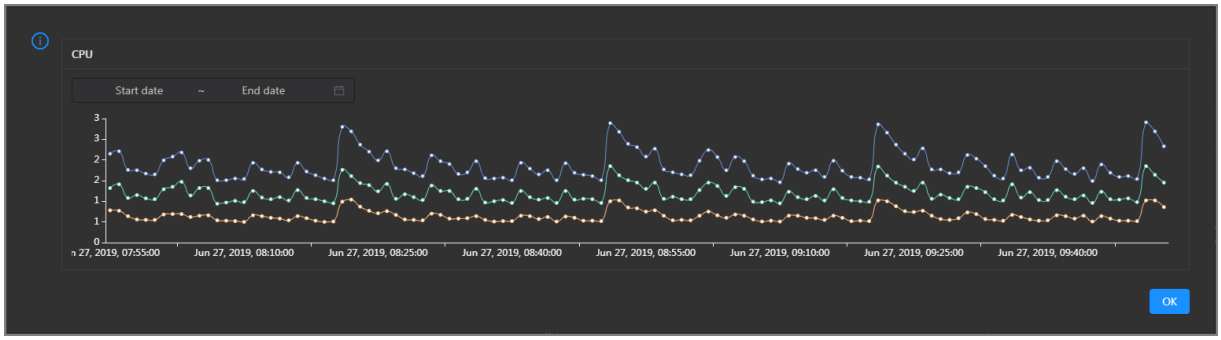

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

6.7.1.4.3. 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.
	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.

Module	Submodule or feature		Description
Service O&M		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.
		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.

Module	Submodule or feature	Description
	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.
	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.

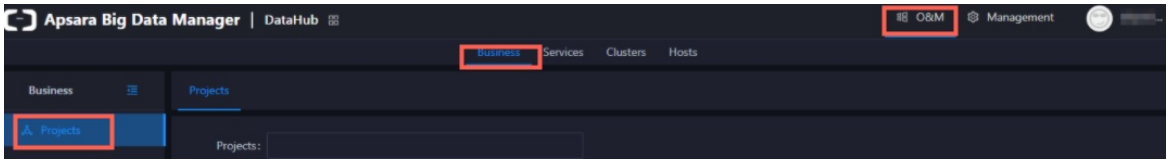

The **O&M** page includes four modules, namely, **Business**, **Services**, **Clusters**, and **Hosts**.

6.7.1.4.4. Business O&M

6.7.1.4.4.1. Business O&M

This topic describes how to go to the business O&M page for DataHub in the Apsara Big Data Manager console.

1. Log on to the Apsara Big Data Manager console.
2. Click the icon in the upper-left corner, and then click **DataHub**.
3. On the DataHub page, click **O&M** in the upper-right corner, and then click the **Business** tab. The **Projects** page appears.

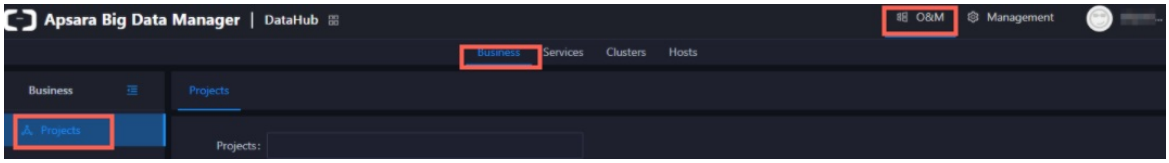

6.7.1.4.4.2. Projects

The Projects page displays the name, owner, number of topics, read traffic, write traffic, storage, and creation time of each project.

Go to the Projects page

On the **Business** tab, click **Projects** in the left-side navigation pane to view the information of all projects.

View project overview

On the **Projects** page, click the name of a project that you want to view. The **Overview** tab 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. On this tab, you can view the information of all topics in the project.

6.7.1.4.4.3. Topics

The topic list displays the name, project, number of shards, storage, read traffic, write traffic, and creation time of each topic.

View the topic list

On the **Business** tab, click **Topics** in the left-side navigation pane to view the topic list.

View the details of a topic

On the **List** 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, the current storage usage, the lifecycle, the type, and the description of the topic. You can also view more details about monitoring metrics, shards, subscriptions, DataConnectors, and schema.

The Metric, Shard, Subscriptions, DataConnector, and Schema tabs provide more details about the topic.

- Metric: On the **Metric** tab, you can view the throughput and latency of the topic in near real time.
- Shard: Shards are concurrent tunnels used for data transmission in the topic.

  On the Shard tab, you can view the ID, status, and active time of each shard.

- Subscriptions: The subscription feature allows you to save consumption offsets to the server and resume 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: Schemas define the data types of fields.

  On the Schema tab, you can view the data type and name of each field.

Disable or enable a topic

When you view the topic information, you may find that the DataHub instance has abnormal traffic or its cluster resources are nearly full. In this case, you need to temporarily disable data read/write for abnormal topics and low-priority topics. When abnormal Meta requests increase, you also need to temporarily disable requests related to abnormal topics to ensure cluster stability.

In this case, you can click **Disable** in the Actions column. In the dialog box that appears, specify the topic that you want to disable. After you disable a topic, all requests for the topic trigger errors, including read and write requests.

After the issue is solved, you can click **Enable** in the Actions column to enable the topic. After the topic is enabled, you can use the topic again.

You can view the information in the Status column to check whether a topic is disabled or enabled. The off status indicates that the topic is disabled, and the on status indicates that the topic is enabled.

6.7.1.4.4.4. Hotspot analysis

The Hotspot Analysis page displays the distribution of shards on the servers of a cluster.

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 servers 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.

qps indicates the number of machine-level requests per second in the cluster. If the qps of a server is high, the request quantity of this server may be higher than that of other servers. In this case, hot spots may exist.

Note

- If a cluster is running, at least four rows of data are displayed. This is because a cluster can contain a minimum of four servers.
- At least two qps values in the four rows are not N/A, that is, at least two of the four servers play Frontend roles. If the qps value of a server is N/A, the server plays the Chunckserver role. In this case, the server is not in a Frontend hybrid deployment or the server is not running.

6.7.1.4.4.5. Archiving latency

The Archiving Tasks page displays the latency during data archiving.

View the archiving tasks

On the **Business** tab, click **Archiving Latency** in the left-side navigation pane to view the archiving tasks.

View the archiving latency

On the **Archiving Tasks** page, the archiving latency of each topic is displayed in the min_done_time column.

For archived MaxCompute data (sink_type is sink_odps), if the time in the min_done_time column is more than 30 minutes later than the current time, check whether the task encounters exceptions.

For other archived data, if the time in the min_done_time column is more than 5 minutes later than the current time, check whether the task encounters exceptions.

You can click the **topic name** to go to the DataConnector tab on the **Topics** page. On this tab, you can view the detailed archive information and perform O&M.

6.7.1.4.5. Service O&M

6.7.1.4.5.1. 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.

6.7.1.4.5.2. Service O&M for Job Scheduler

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.

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.

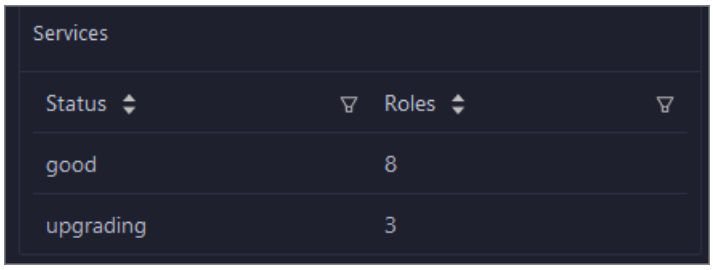

Roles

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.

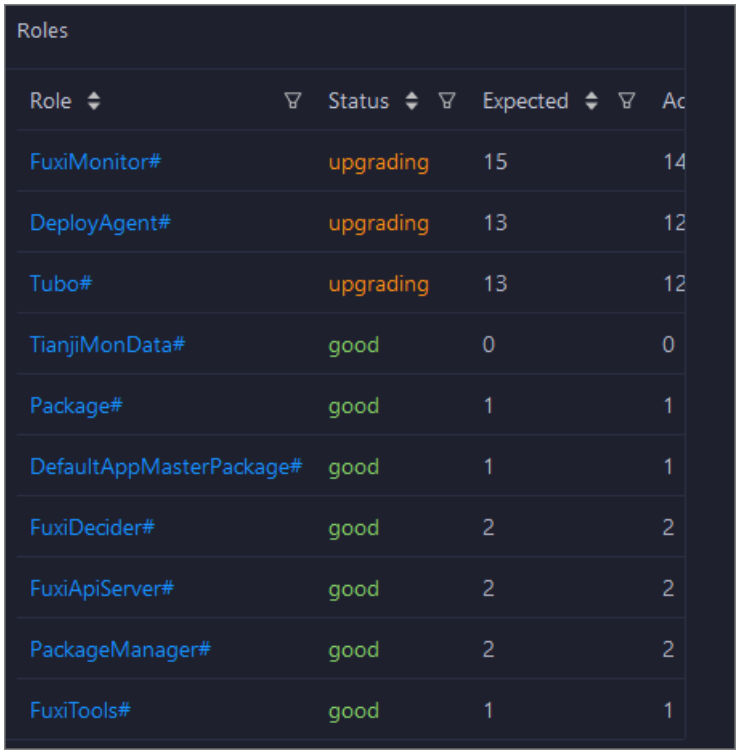

Click the name of a server role to go to the Apsara Infrastructure Management Framework console and view its details.

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.

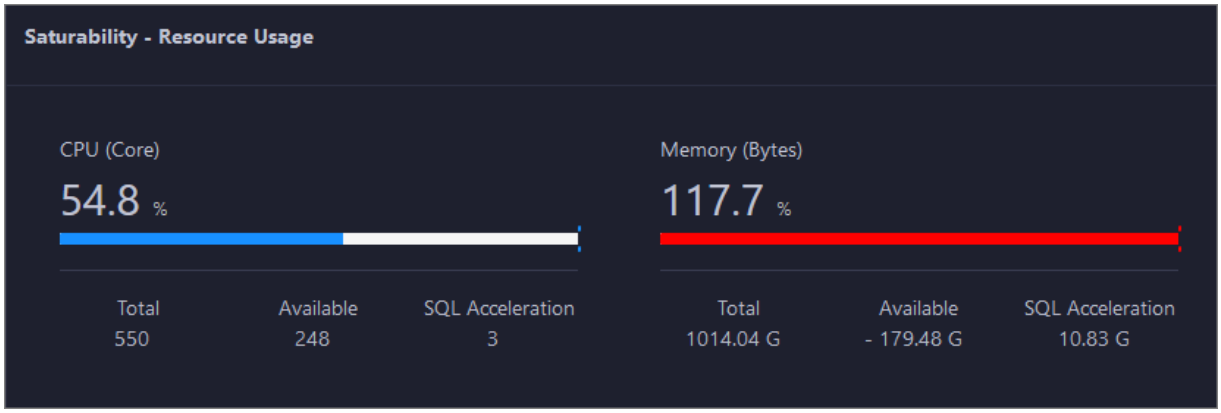

View the trend charts of CPU and memory usage

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.

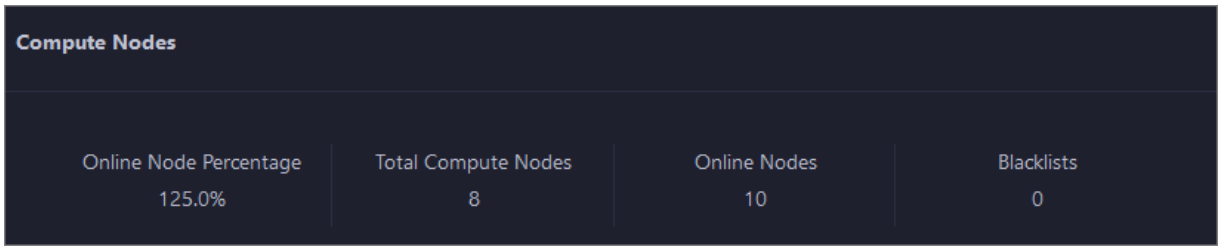

Service instances

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.

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.

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.

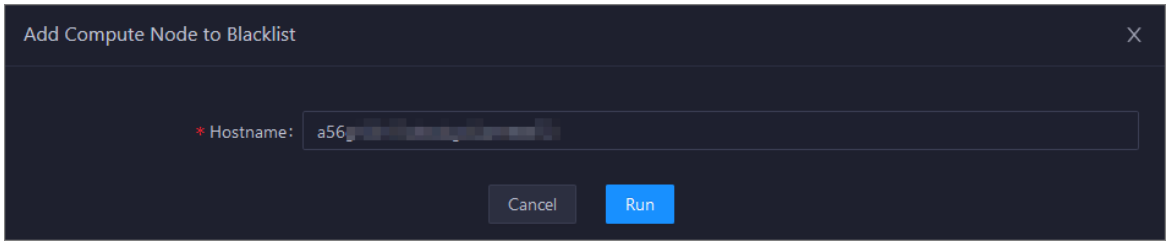

The value of the **Hostname** parameter is automatically filled. You do not need to specify a value for this parameter.

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

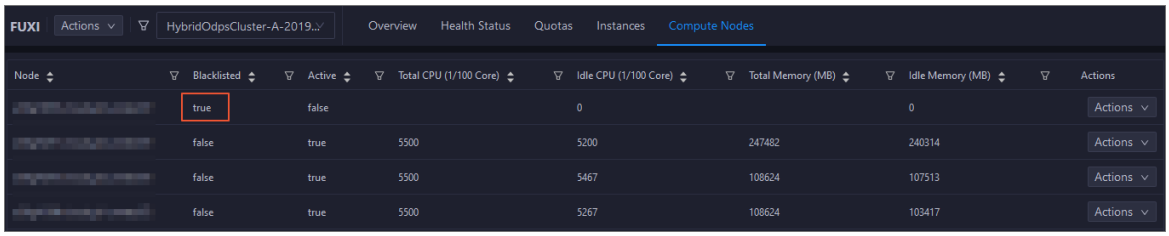

6.7.1.4.5.3. Service O&M for Apsara Distributed File System

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.

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.

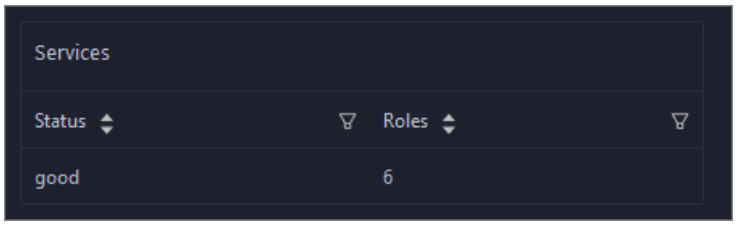

Roles

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.

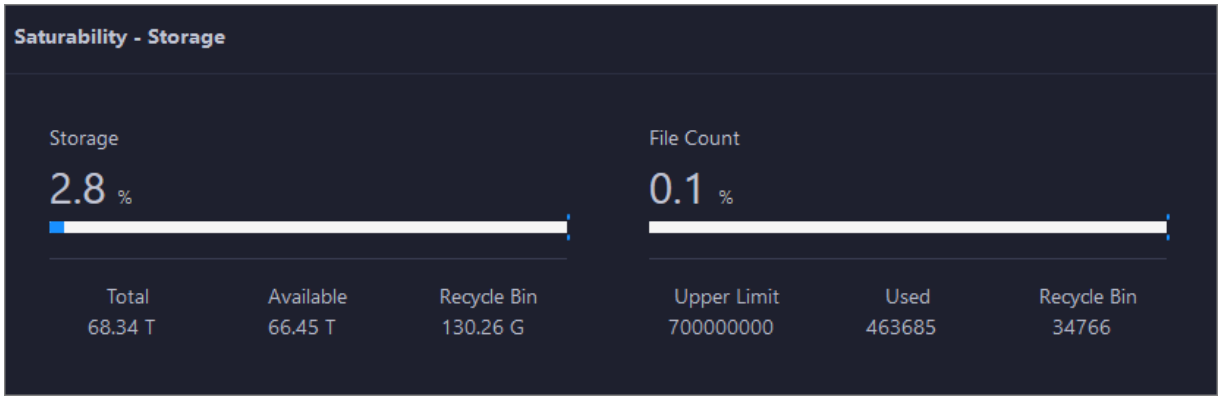

Storage Trend and File Count Trend

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.

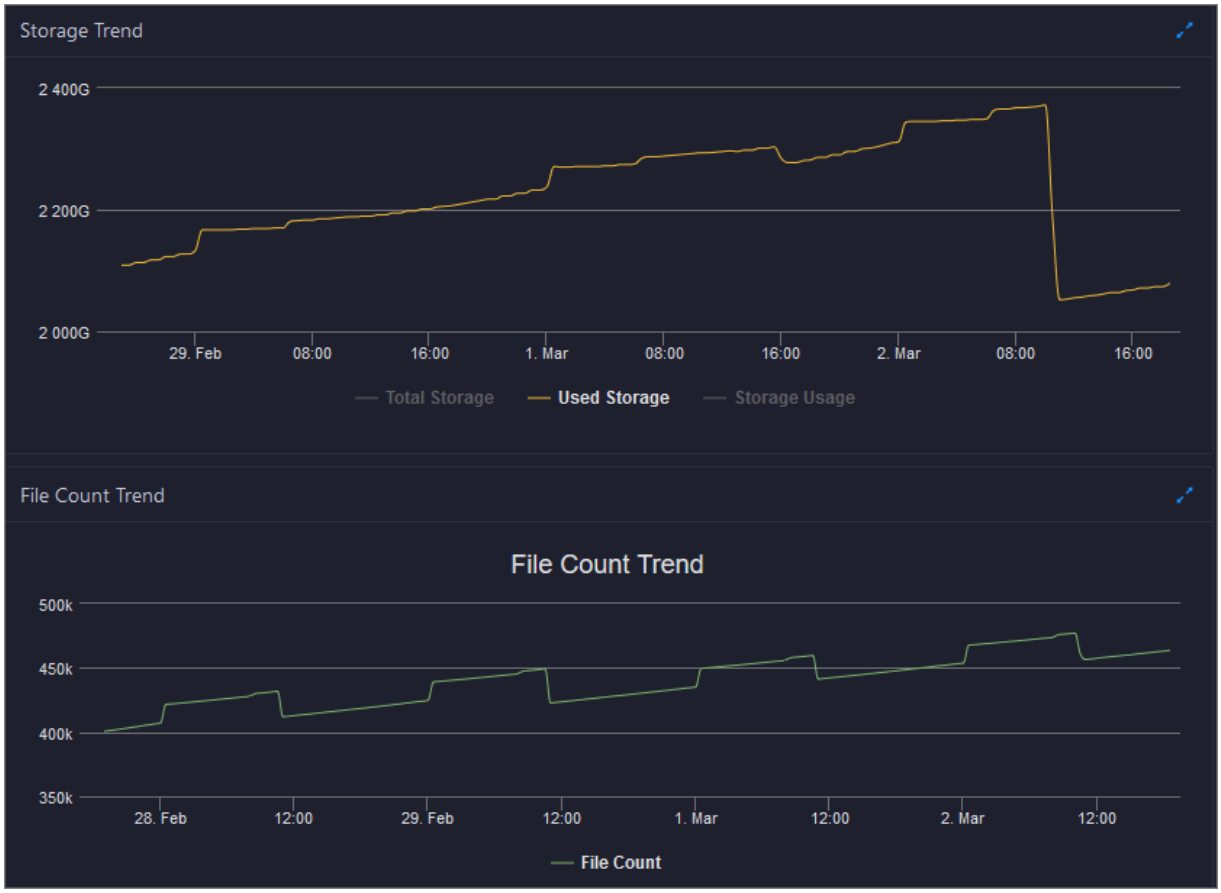

In the upper-right corner of the chart, click the icon to zoom in the chart. The following figure shows an enlarged chart of storage usage.

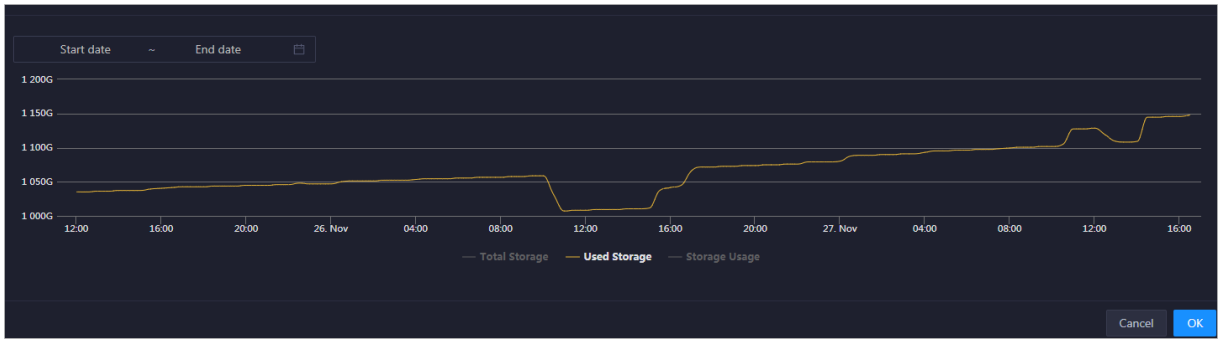

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

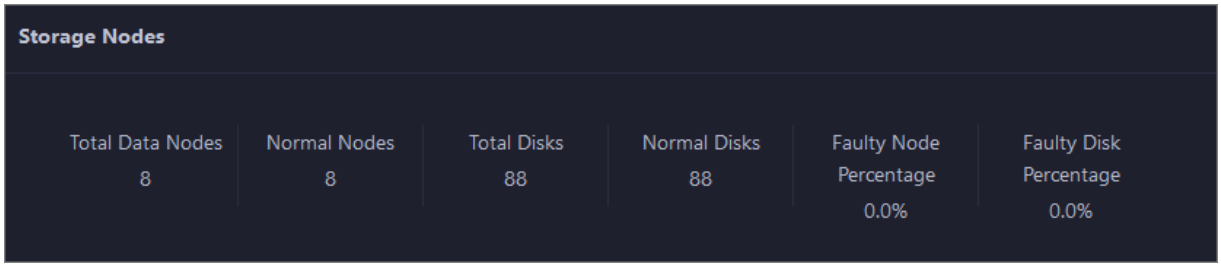

Service roles

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.

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.

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.

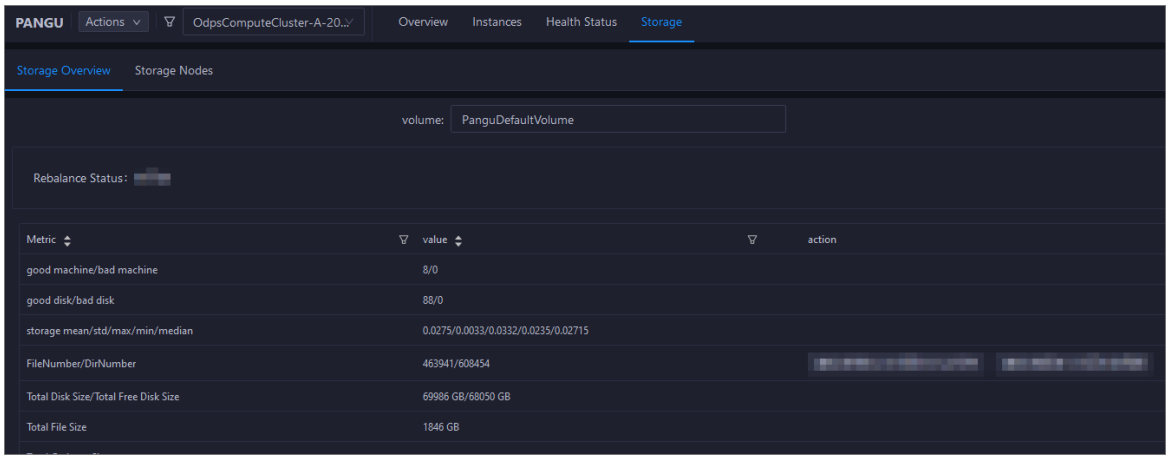

   The **Storage Overview** page displays whether data rebalancing is enabled, key metrics and their values, suggestions to handle exceptions, and rack specifications of Apsara Distributed File System.

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 **Set Node Status to Shutdown** panel, click **Run**. A message appears, indicating that the action has been submitted.

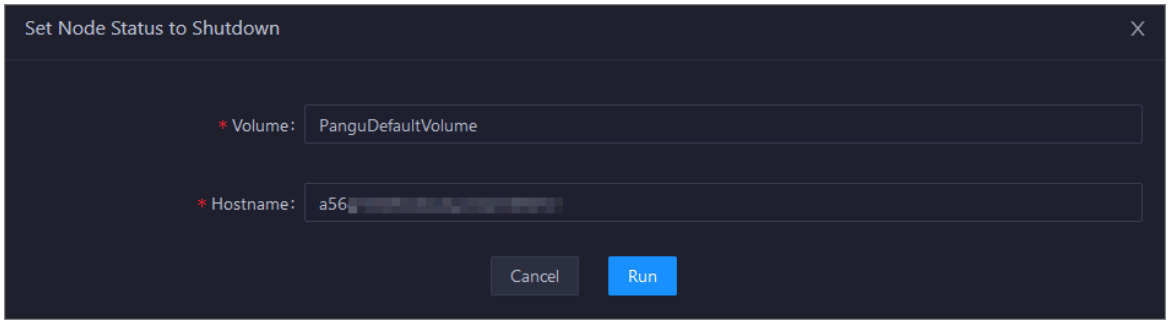

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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 **Set Disk Status to Error** panel, set the **Diskid** parameter.

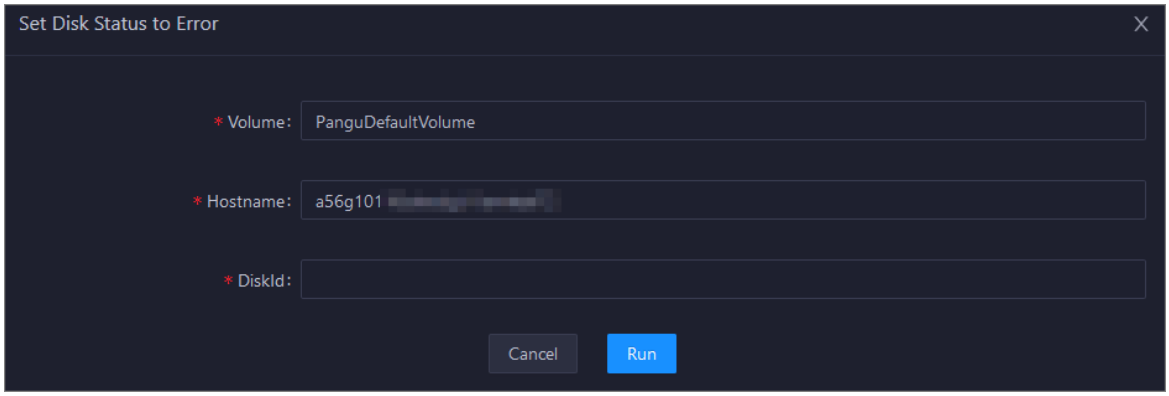

The values of the **Volume** and **Hostname** parameters are automatically filled based on the selected storage node. You do not need to specify values for the parameters.

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

Clear the recycle bin of Apsara Distributed File System

Apsara Big Data Manager (ABM) allows you to clear the recycle bin of Apsara Distributed File System to release storage space.

Prerequisites

Your Apsara Big Data Manager 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.

   Click **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.

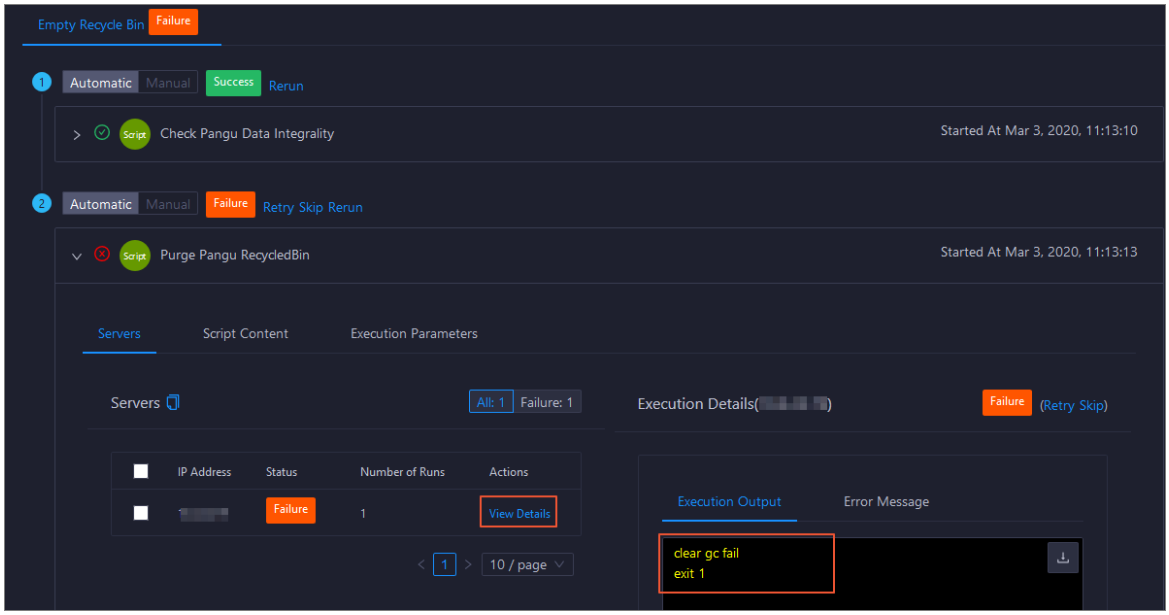

You can view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

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**.

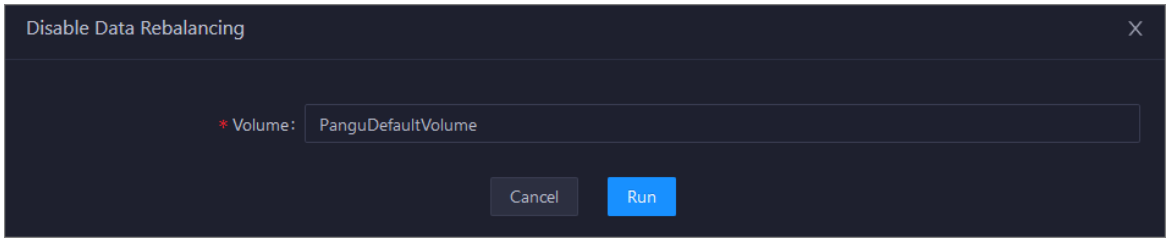

4. Click **Run**. A message appears, indicating that the request has been submitted.
5. View the execution status.

   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**.

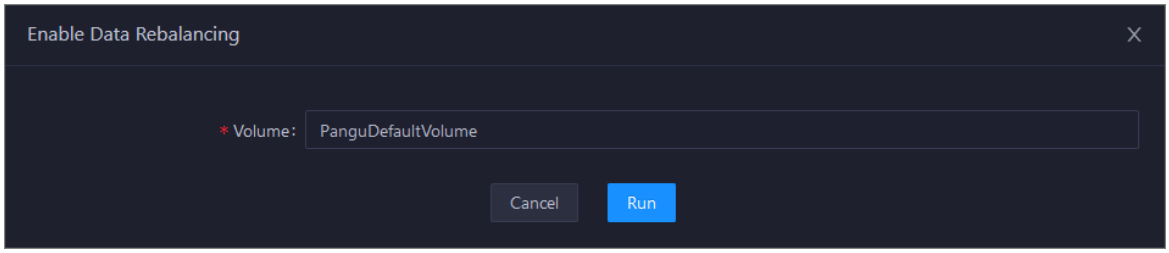

4. Click **Run**. A message appears, indicating that the request has been submitted.
5. View the execution status.

   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.

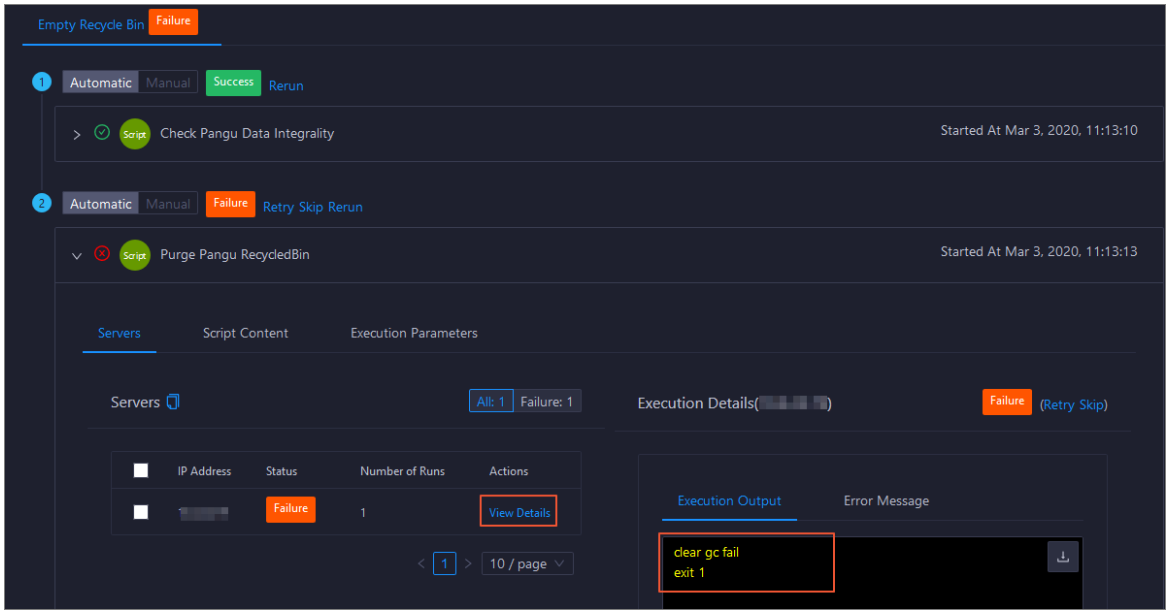

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

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.

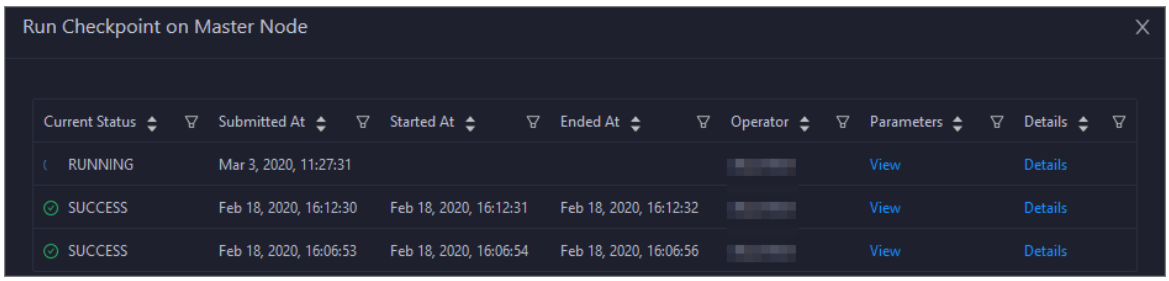

In the Current Status column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. If the status is FAILED, click **Details** in the Details column to locate the failure cause.

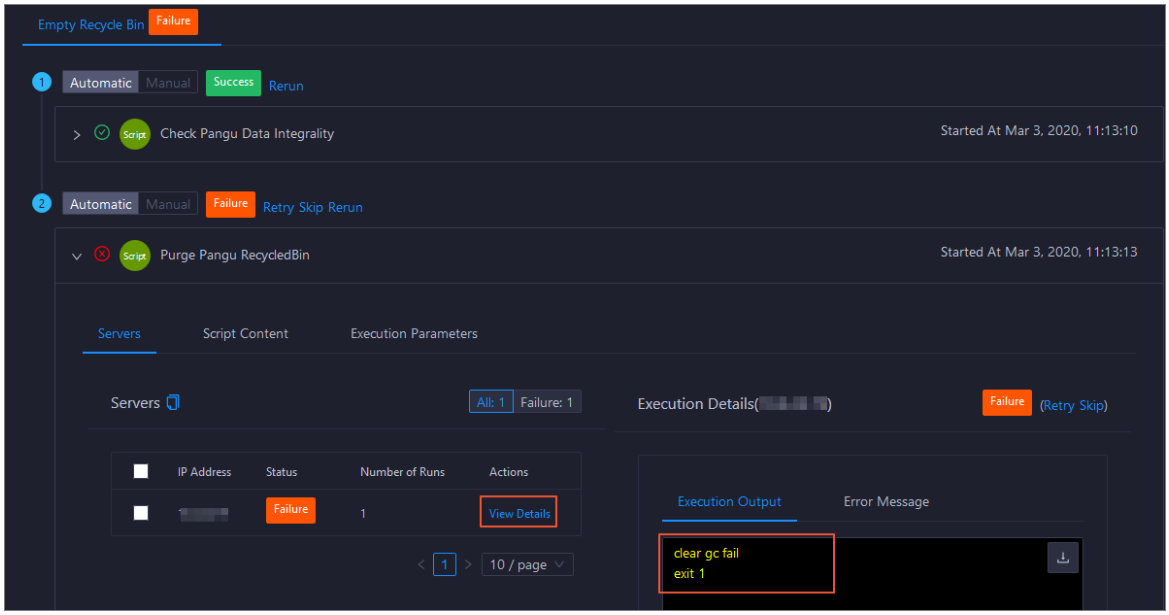

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

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.

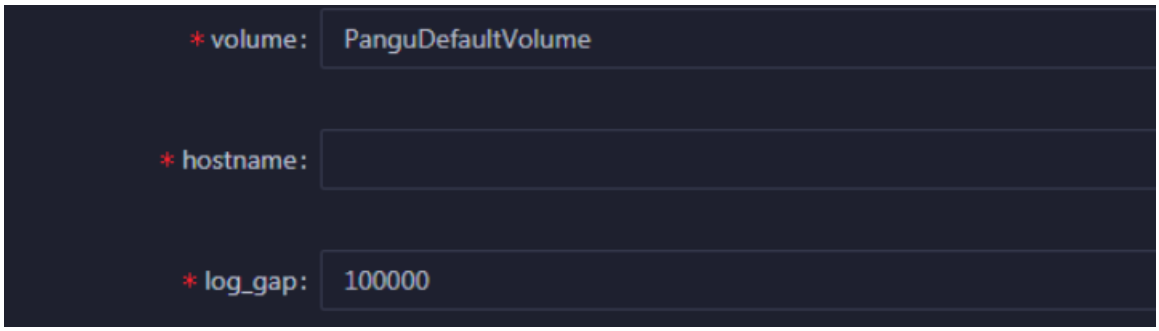

You must set the following parameters in this step:

- **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.

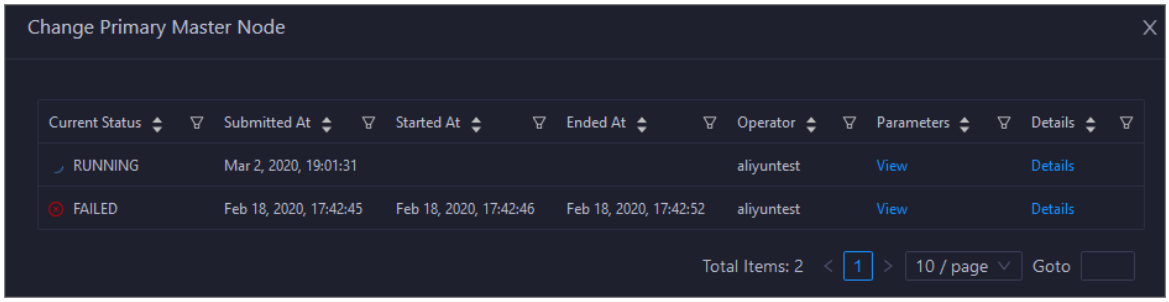

   In the Current Status column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

5. If the status is FAILED, click **Details** in the Details column to locate the failure cause.

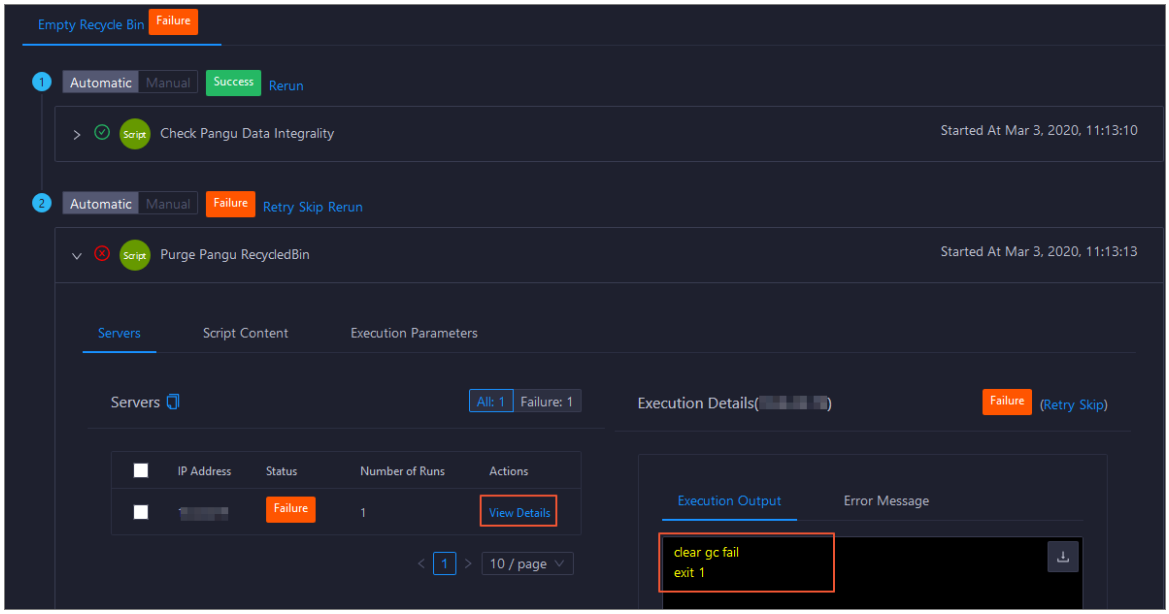

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

6.7.1.4.6. Cluster O&M

6.7.1.4.6.1. 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.

6.7.1.4.6.2. 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.

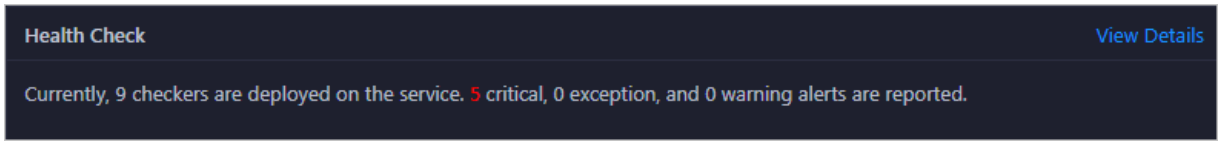

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Cluster health.

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.

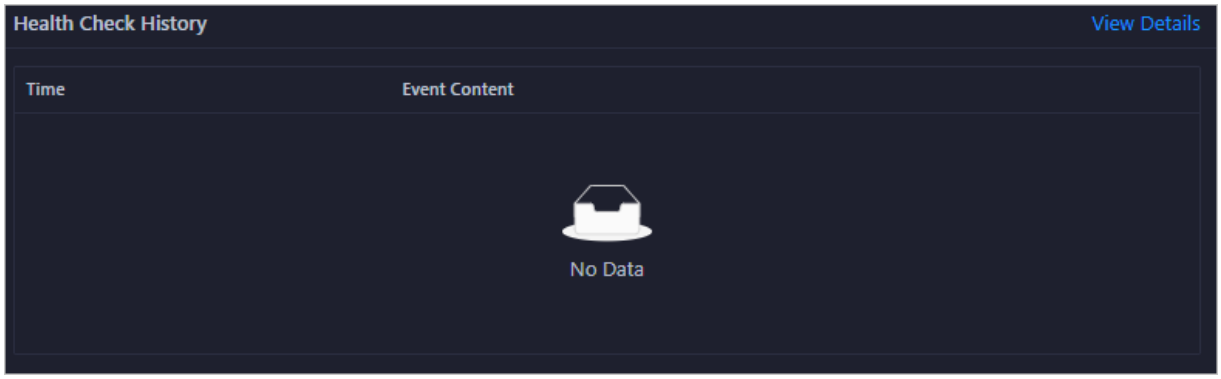

You can click the event content of a check to view the anomalous items.

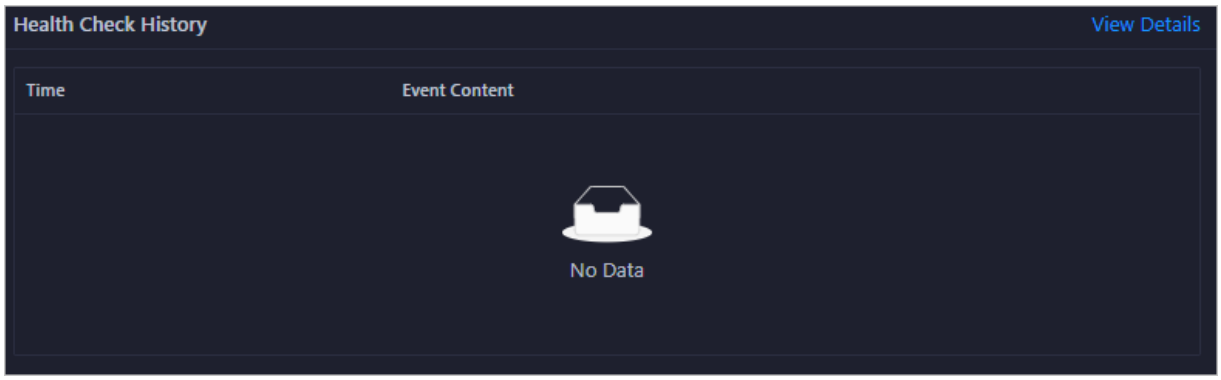

CPU

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.

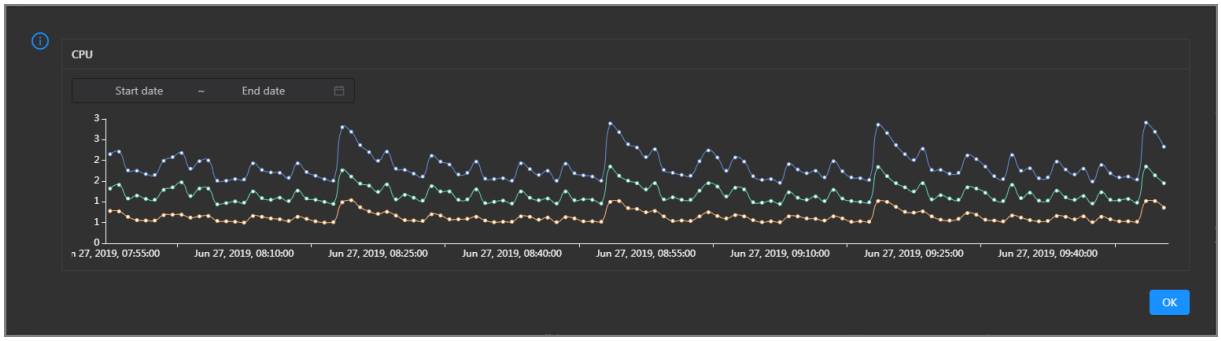

DISK

This chart shows the trend lines of the storage usage on the */*, */boot*, */home/admin*, and */home* directories for the cluster over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

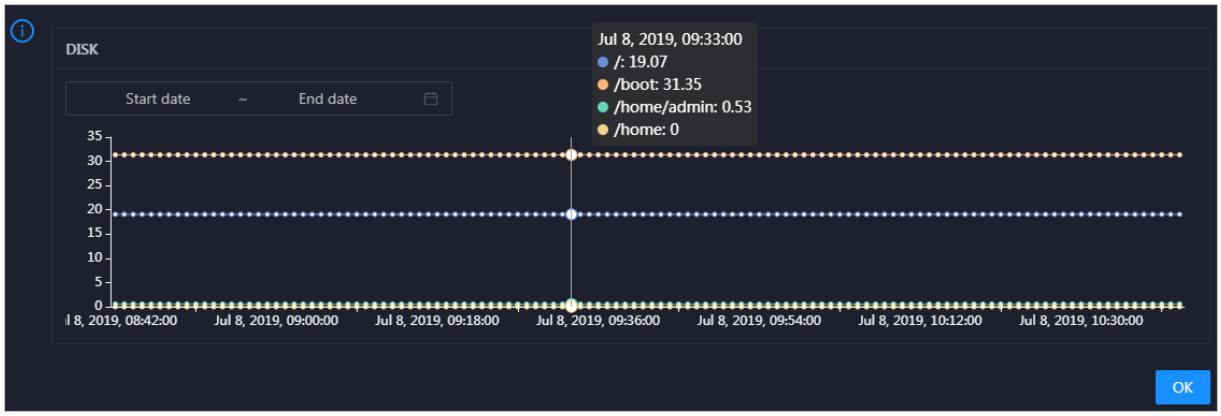

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

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.

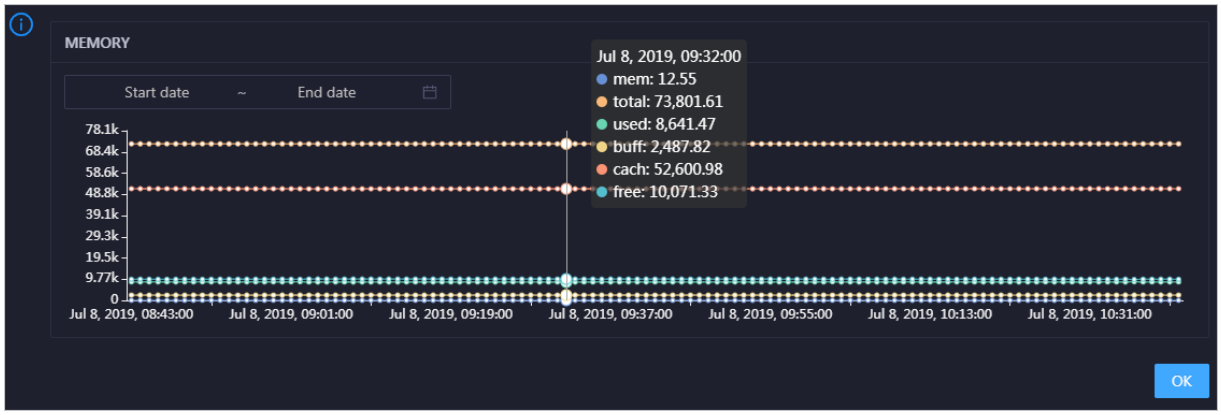

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

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.

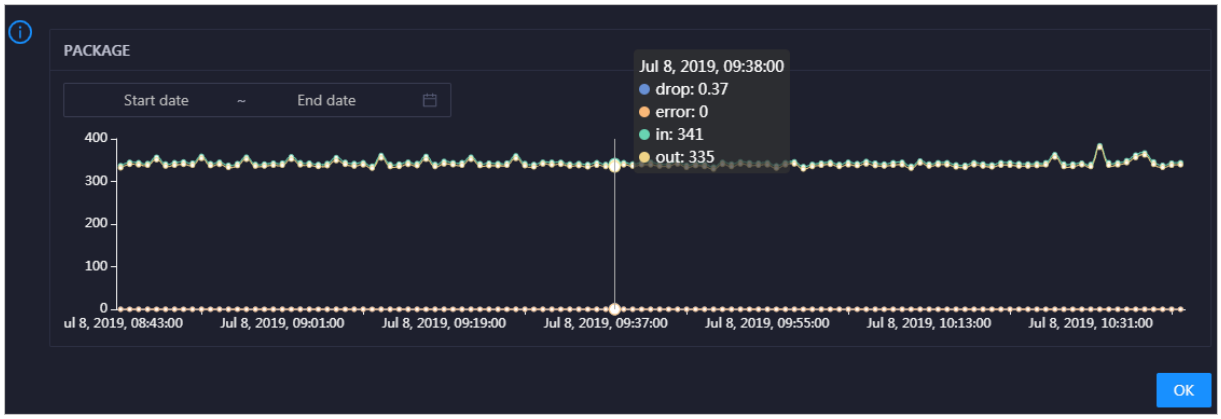

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

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.

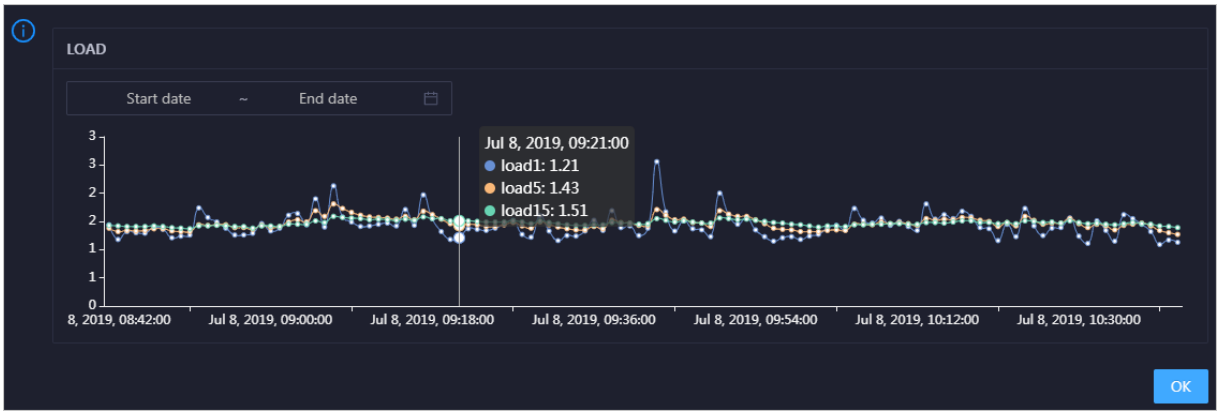

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

6.7.1.4.6.3. 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 represented in different colors. You must handle the alerts in a timely manner, especially the **CRITICAL** and **WARNING** alerts.

View checker details

1. On the Health Status tab, click **Details** in the Actions column of a checker. On the Details page, view checker details.

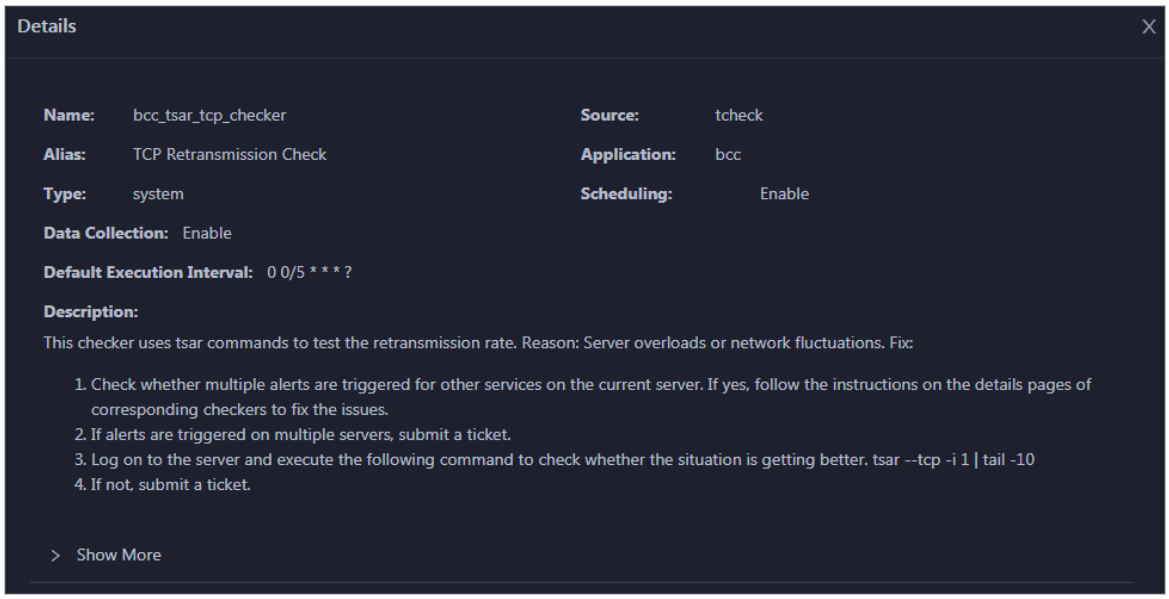

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

2. Click **Show More** to view more information about the checker.

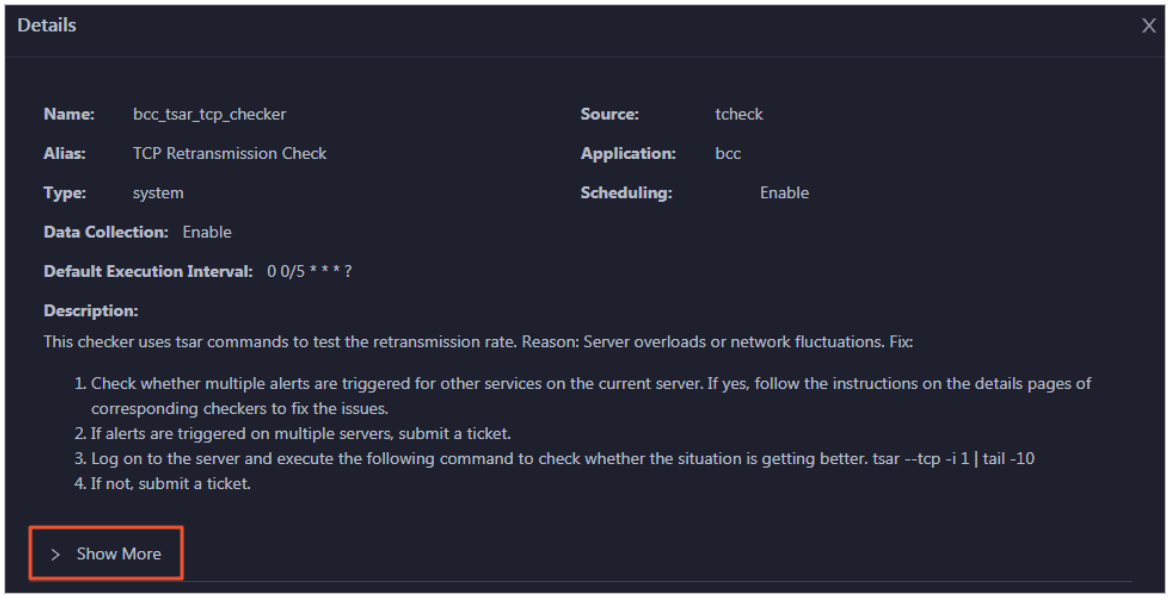

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

View the hosts for which alerts are reported and causes for the alerts

You can view the check history and check results of a checker on a host.

1. On the Health Status tab, click **+** to expand a checker for which alerts are reported. You can view all hosts where the checker is run.

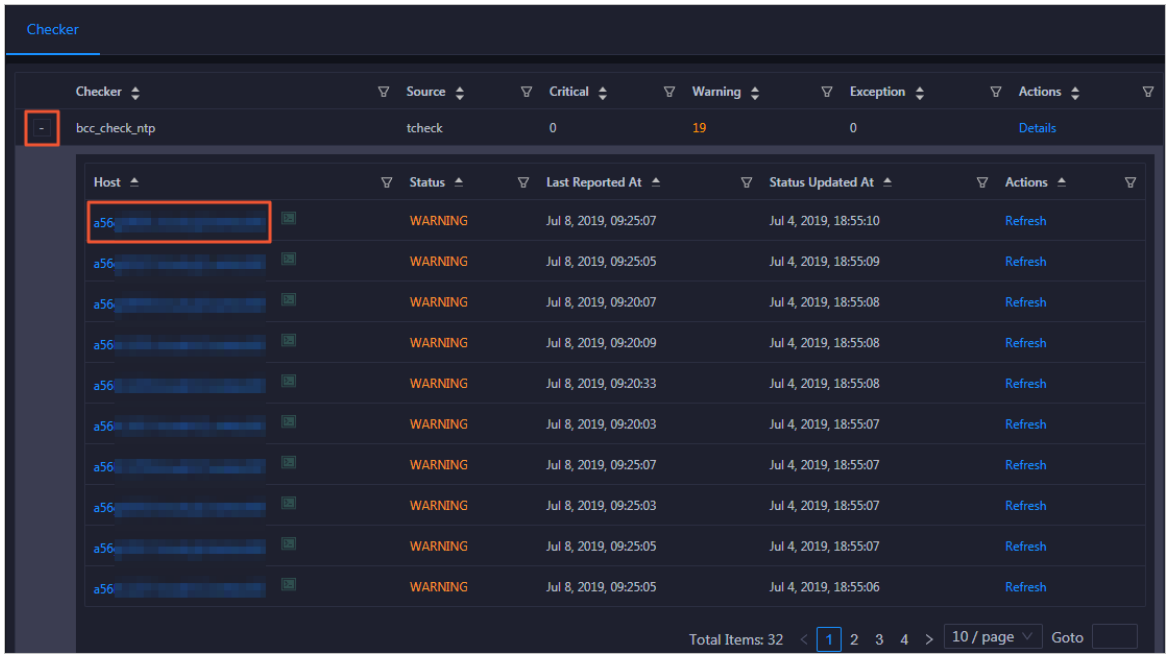

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

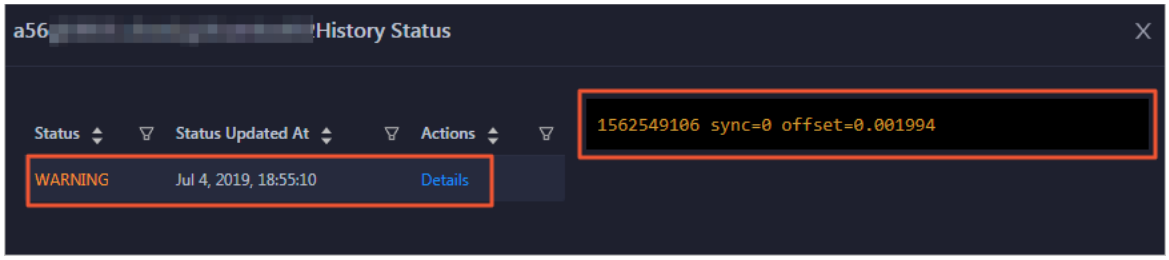

Clear alerts

On the Health Status tab, click **Details** in the Actions column of a checker for which alerts are reported. On the Details page, view the schemes to clear alerts.

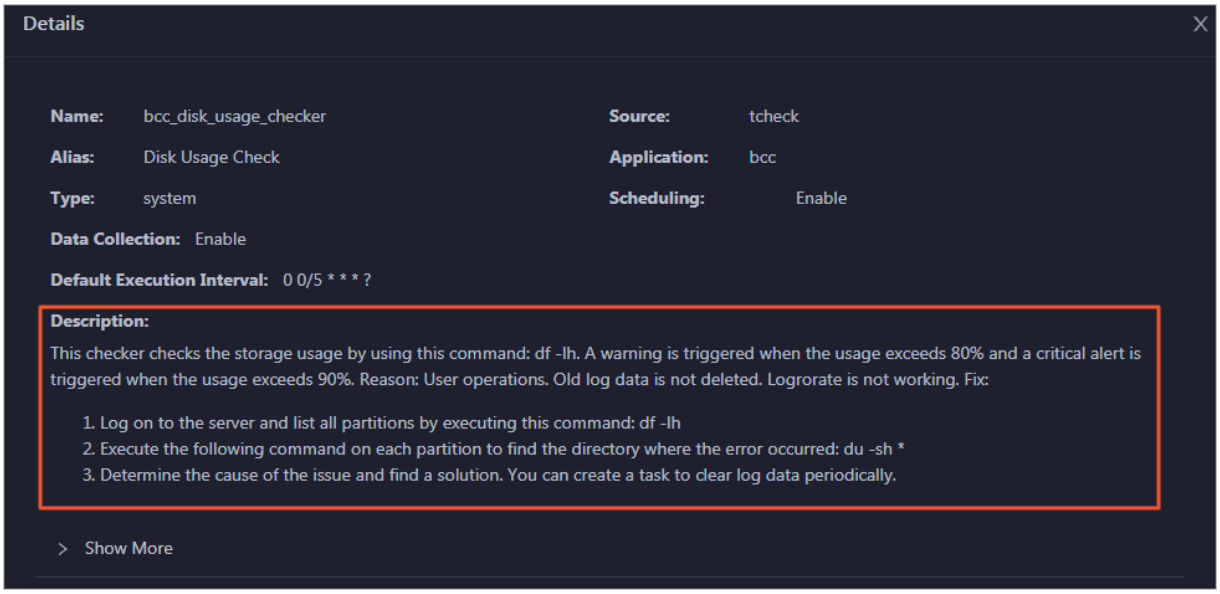

Log on to a host

You may need to log on to a host to handle alerts or other issues that occurred on the host.

1. On the Health Status tab, click + to expand a checker for which alerts are reported.

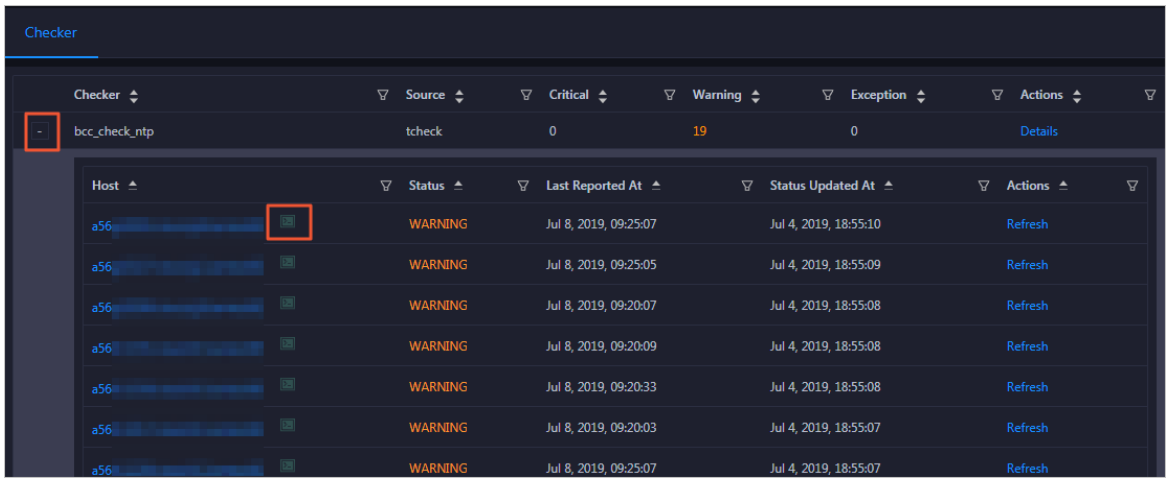

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

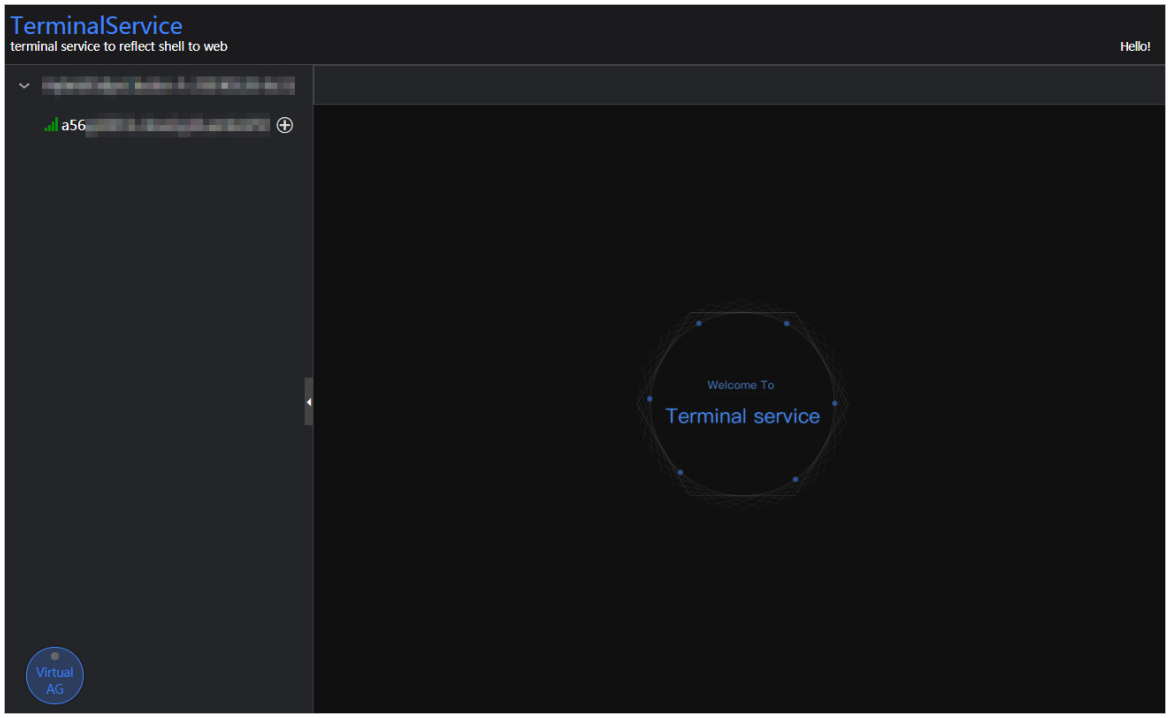

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

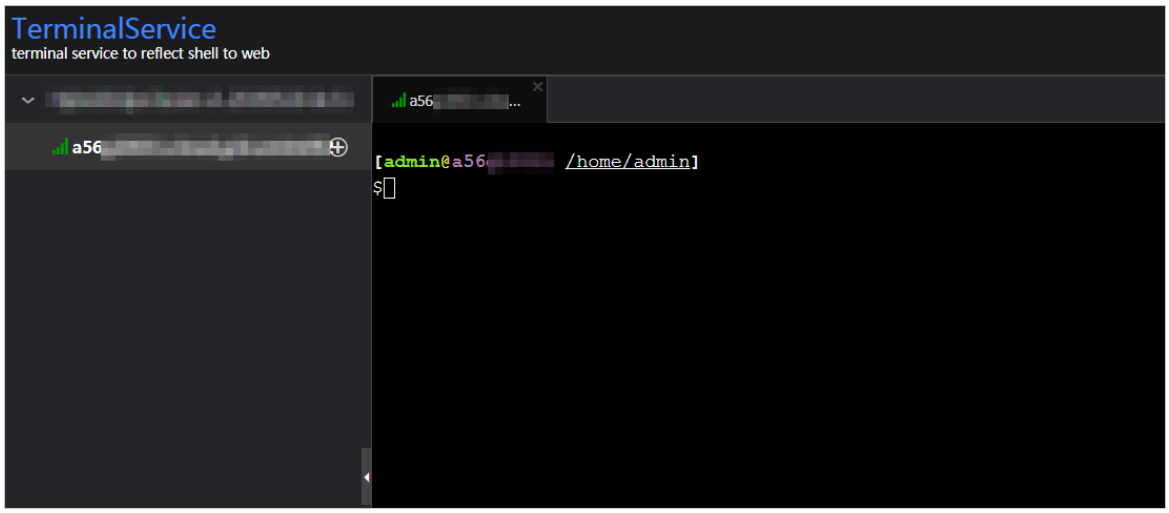

Run a checker again

After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. This way, you can check whether the alert is cleared.

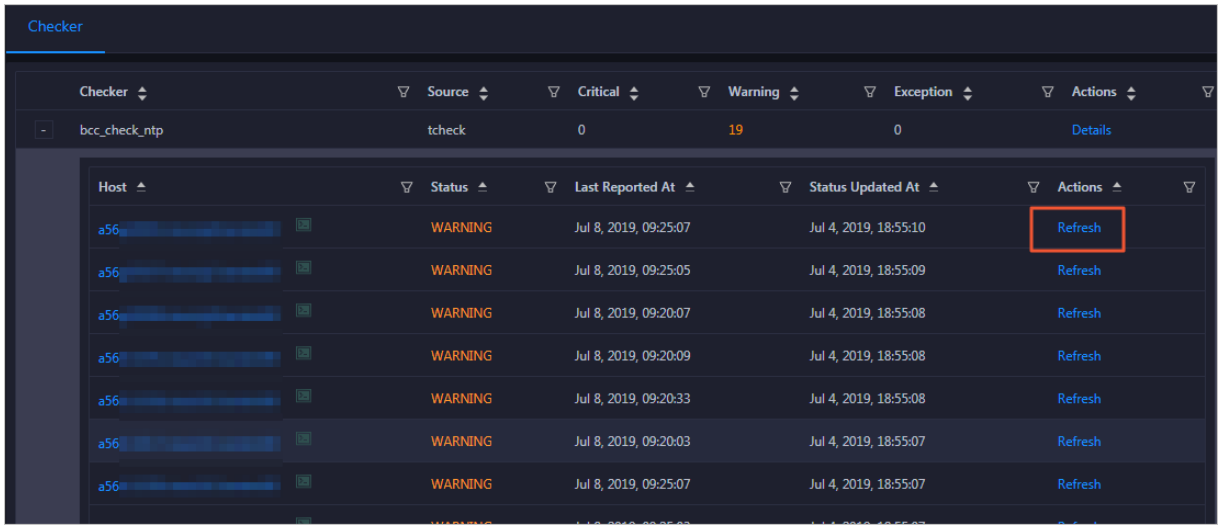

6.7.1.4.6.4. Cluster hosts

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.

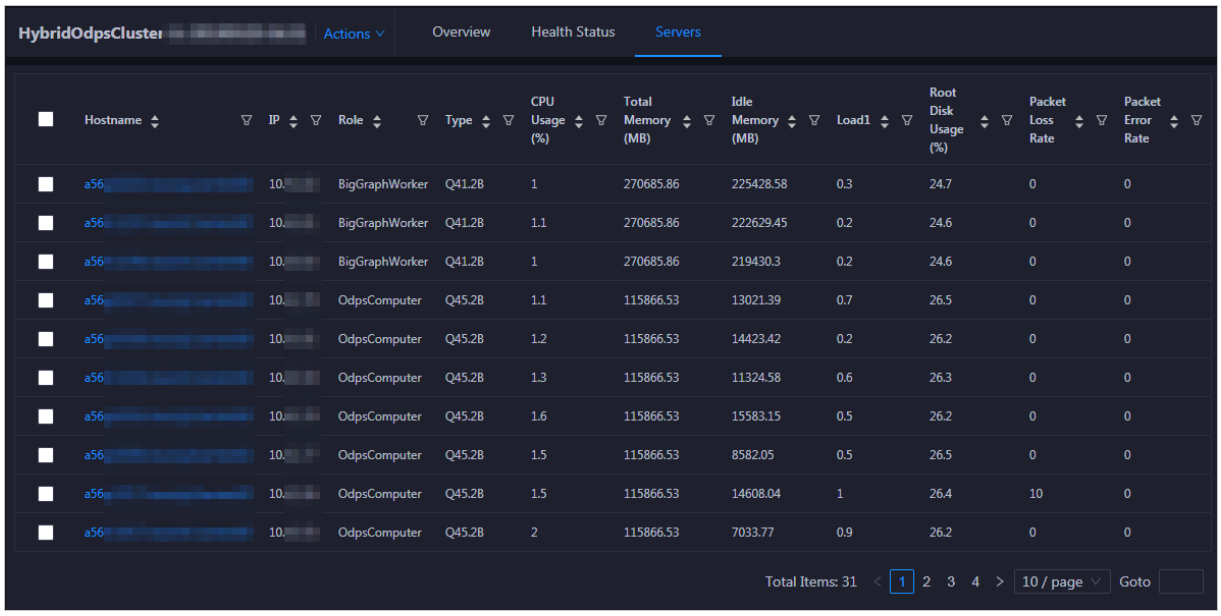

To view more information about a host, click the name of the host. The Overview tab of the Hosts page appears. For more information, see Host overview.

6.7.1.4.6.5. 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.

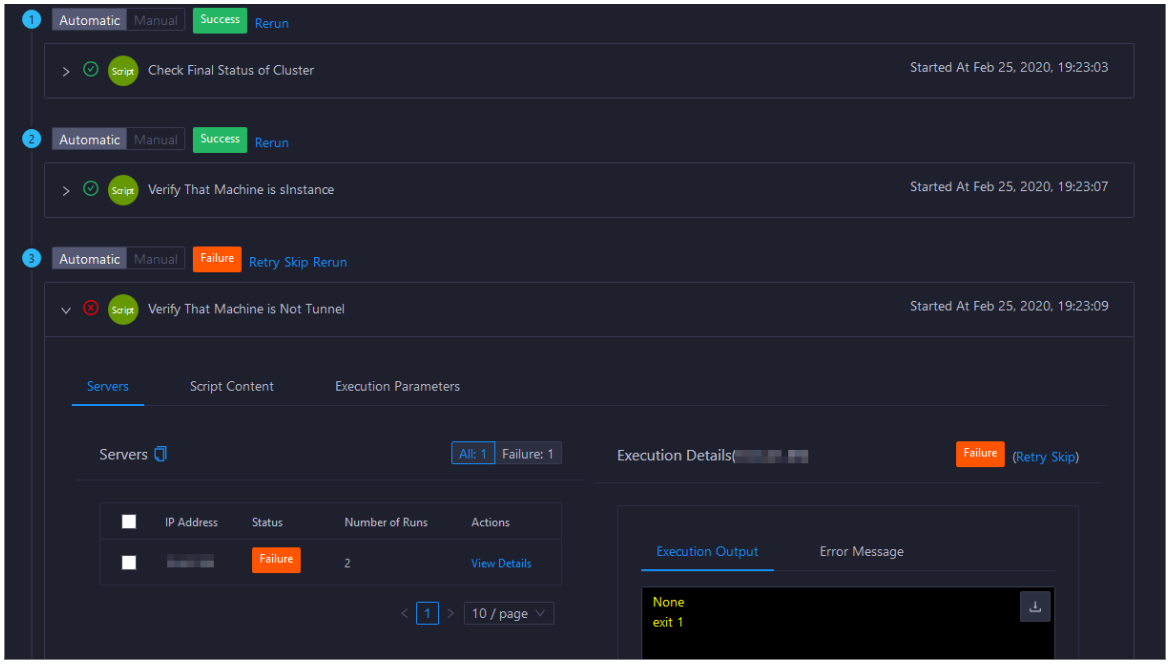

You can also view the parameter settings, host details, script, and execution parameters to locate the failure cause.

6.7.1.4.6.6. 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.

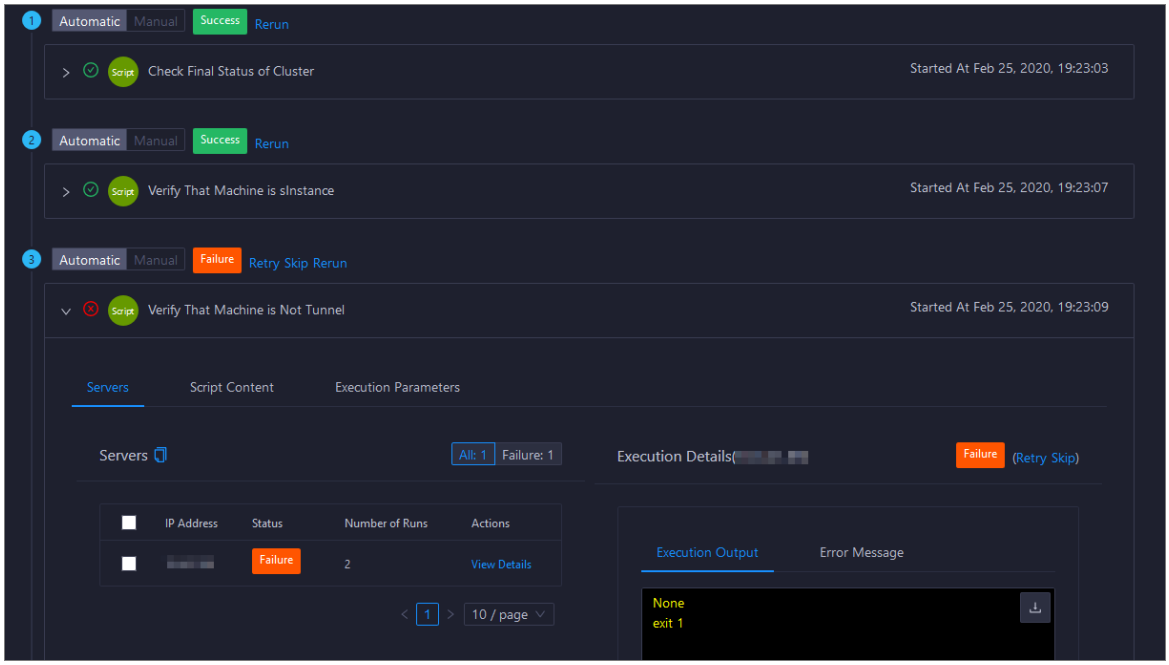

You can also view the parameter settings, host details, script, and execution parameters to locate the failure cause.

6.7.1.4.6.7. Delete topics from a smoke testing project

Apsara Big Data Manager 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** 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.
   - 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 appears in the Execution Details section in the lower-right corner. The execution result includes the time when the job was run and the IP address of the host.

6.7.1.4.6.8. Reverse parse request ID

Apsara Big Data Manager allows you to reverse parse request ID 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** 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.
   - 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 appears in the Execution Details section in the lower-right corner. The execution result includes the time when the job was run and the IP address of the host.

6.7.1.4.7. Host O&M

6.7.1.4.7.1. 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.

6.7.1.4.7.2. 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.

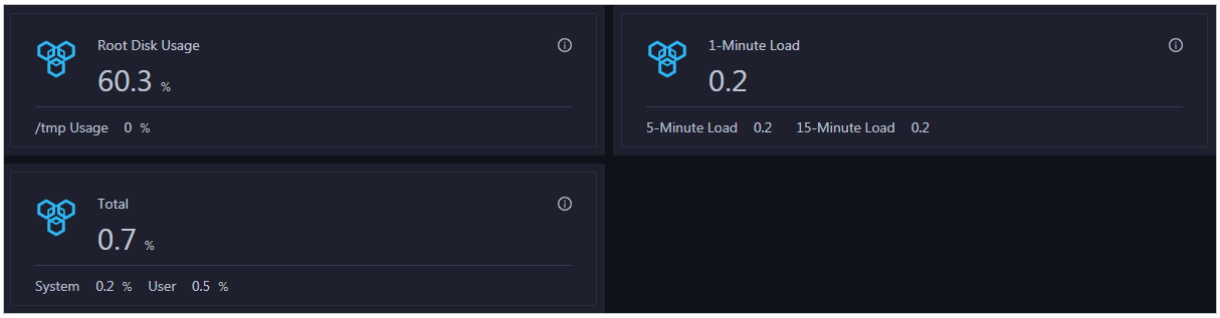

CPU

The CPU chart shows the trend lines of the total CPU utilization (cpu), CPU utilization for executing code in kernel space (sys), and CPU utilization for executing code in user space (user) of the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

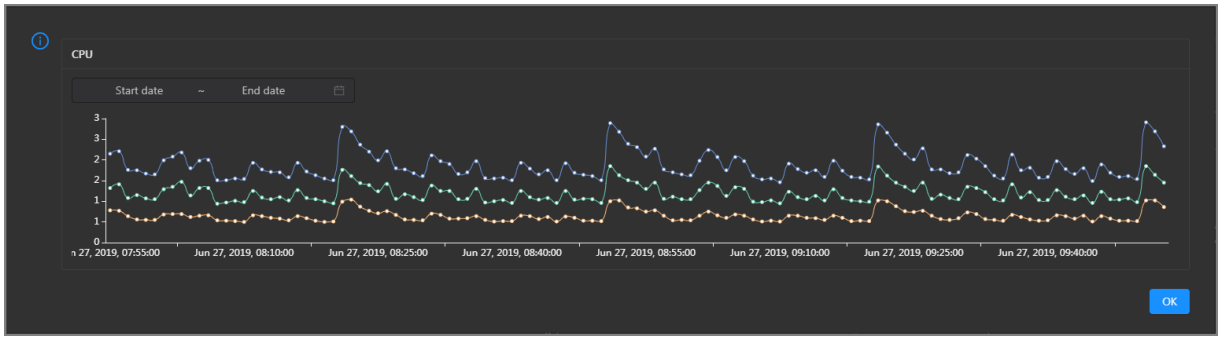

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

The DISK chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

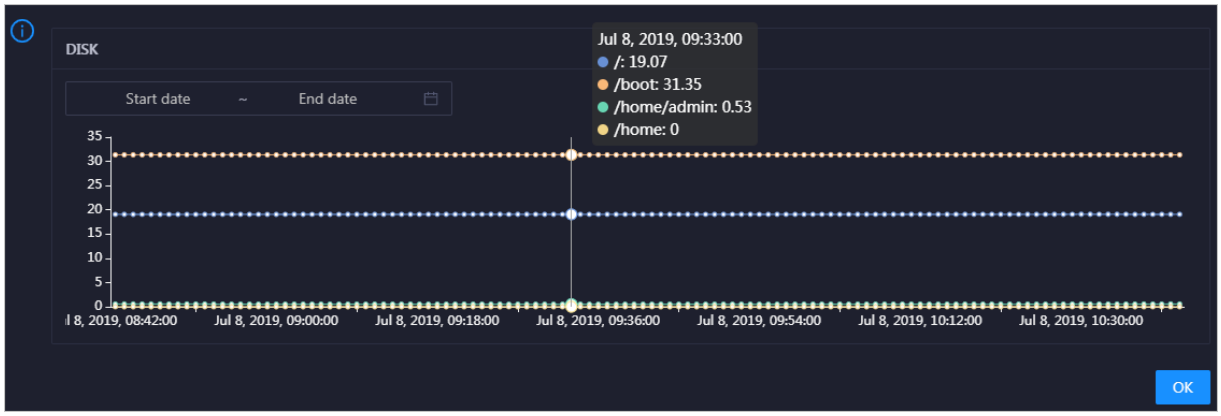

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

The MEMORY chart shows the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

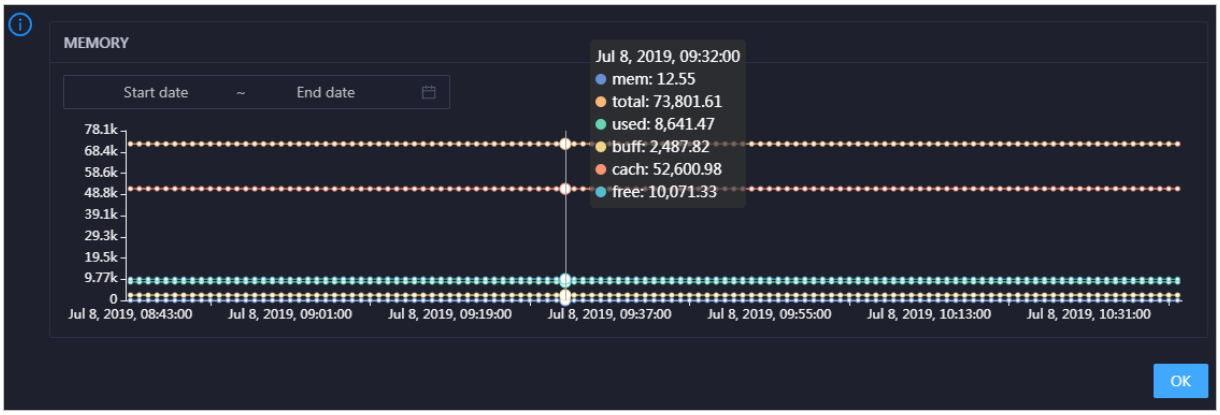

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

The LOAD chart shows the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

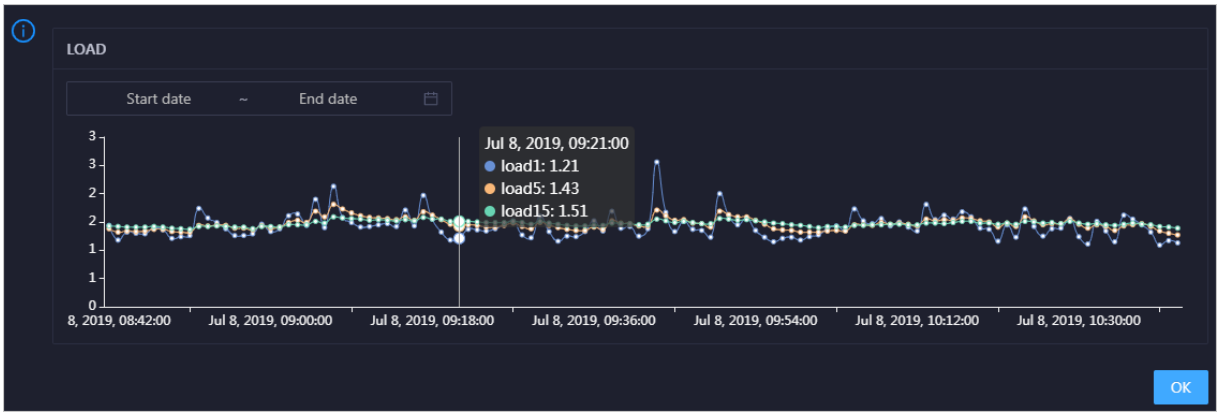

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

The PACKAGE chart shows the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

In the upper-right corner of the chart, click the icon to zoom in the chart.

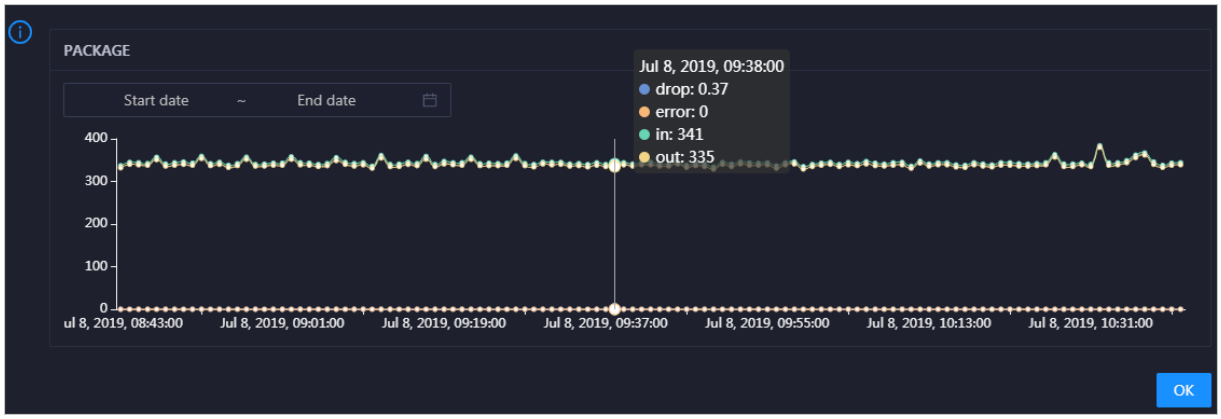

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click in the upper-right corner of the chart to zoom in the chart.

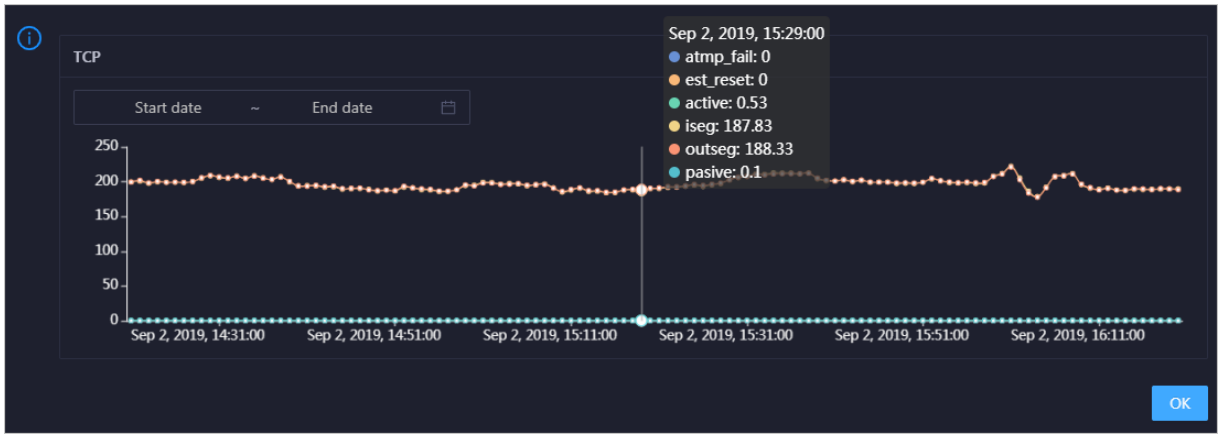

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This chart displays the trend line of the average usage of the root disk (/) for the host over time.

Click in the upper-right corner of the chart to zoom in the chart.

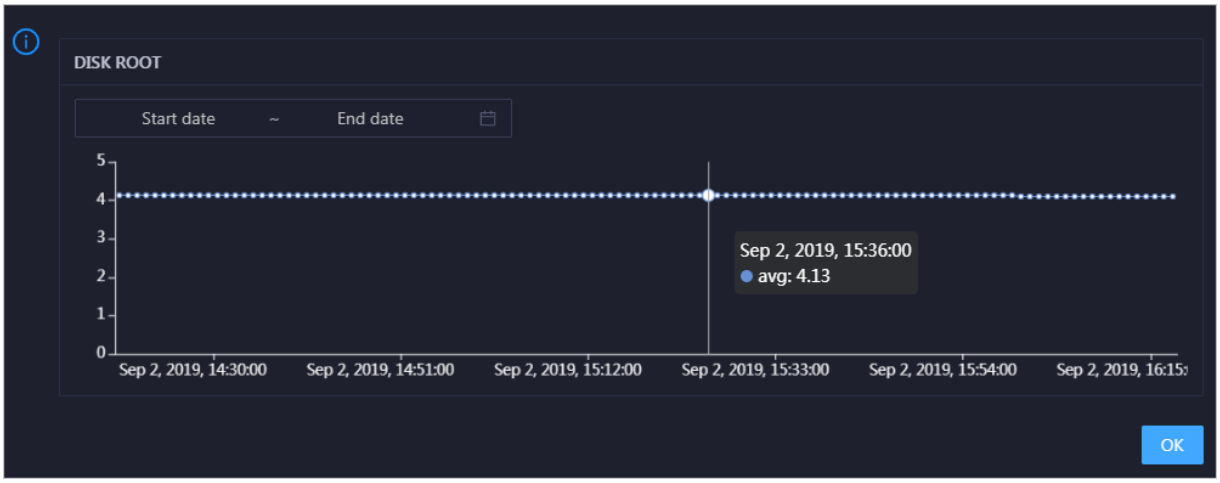

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Click **View Details** to go to the 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.

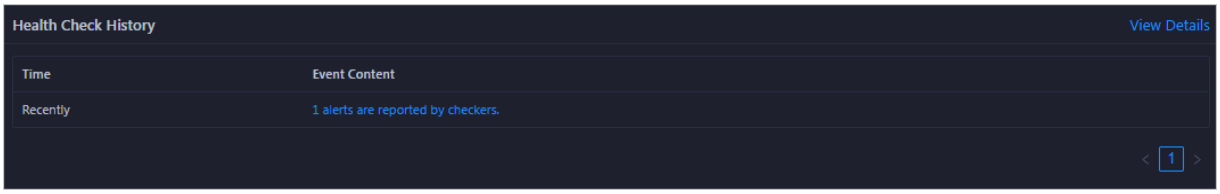

Click **View Details** to go to the Host health page. On this page, you can view the health check details.

You can click the event content of a check to view the exception items.

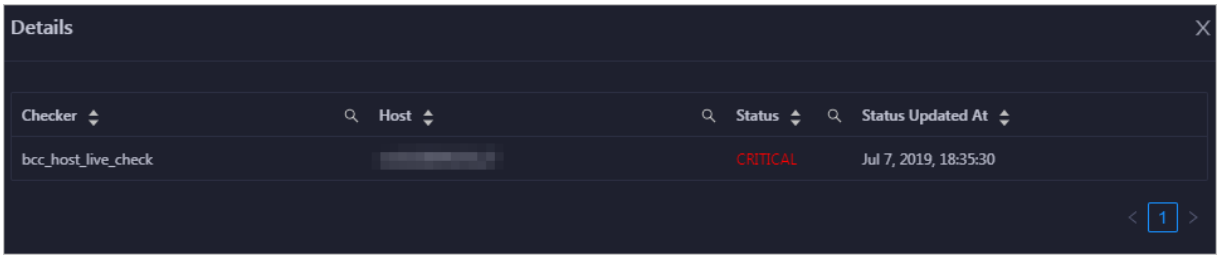

6.7.1.4.7.3. Host charts

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.

6.7.1.4.7.4. 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.

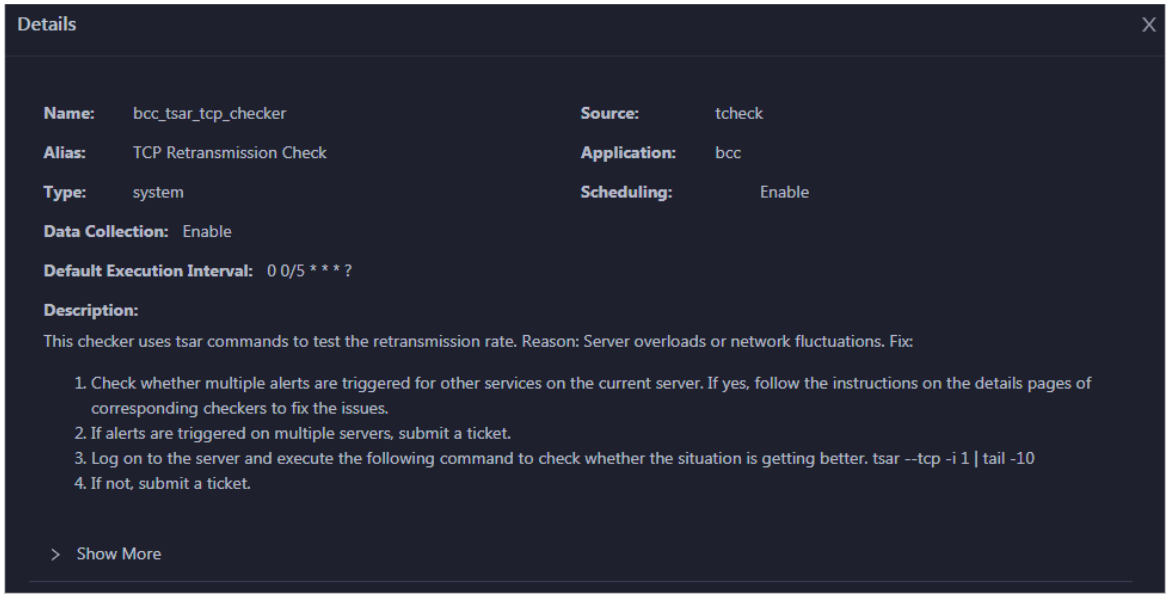

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click **Show More** at the bottom to view more information about the checker.

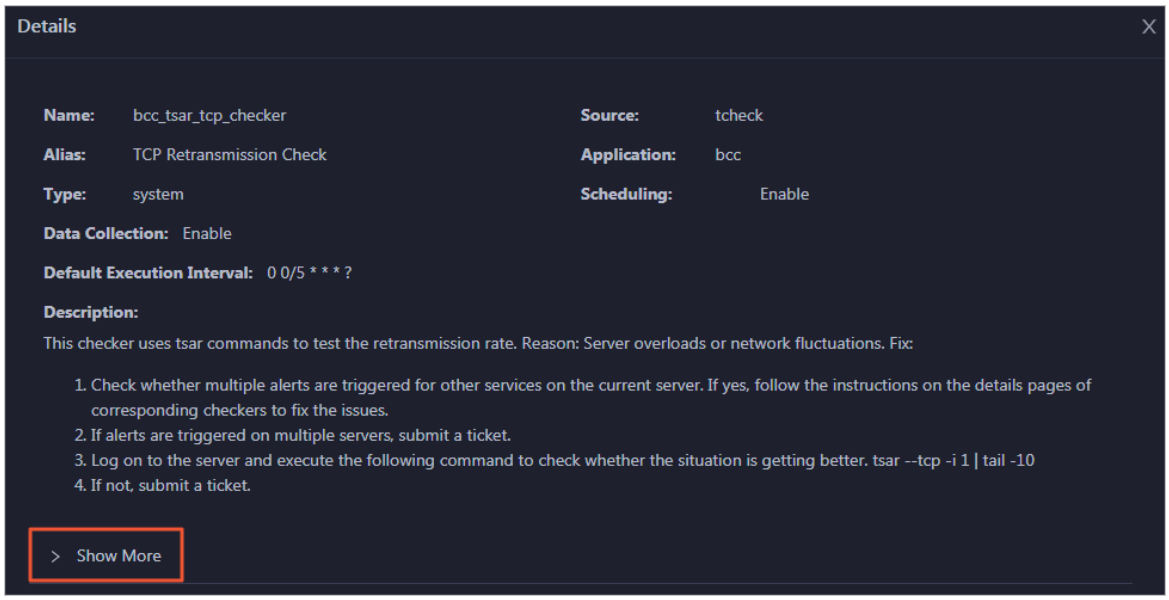

You can view information about the execution script, execution target, default threshold, and mount point for data collection.

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.

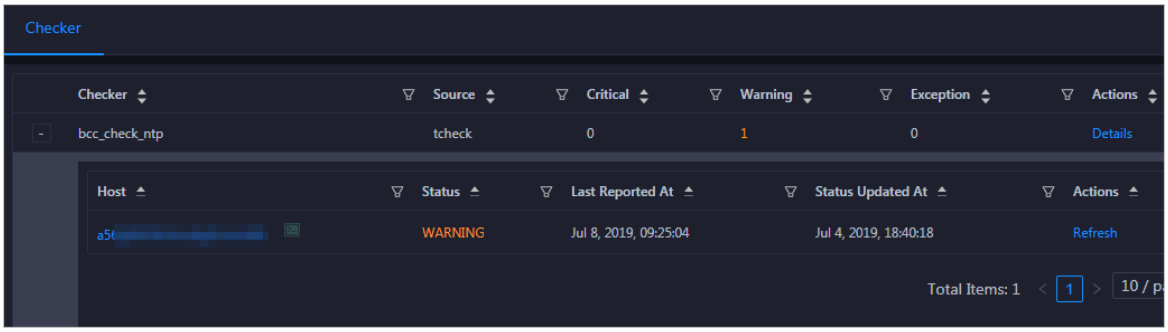

2. Click the hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

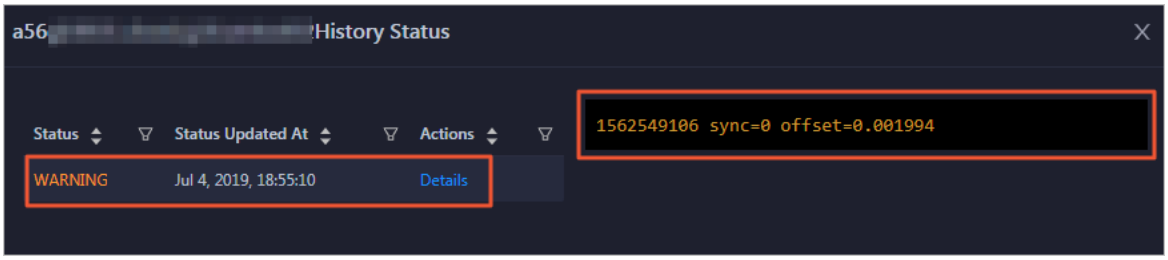

Clear alerts

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.

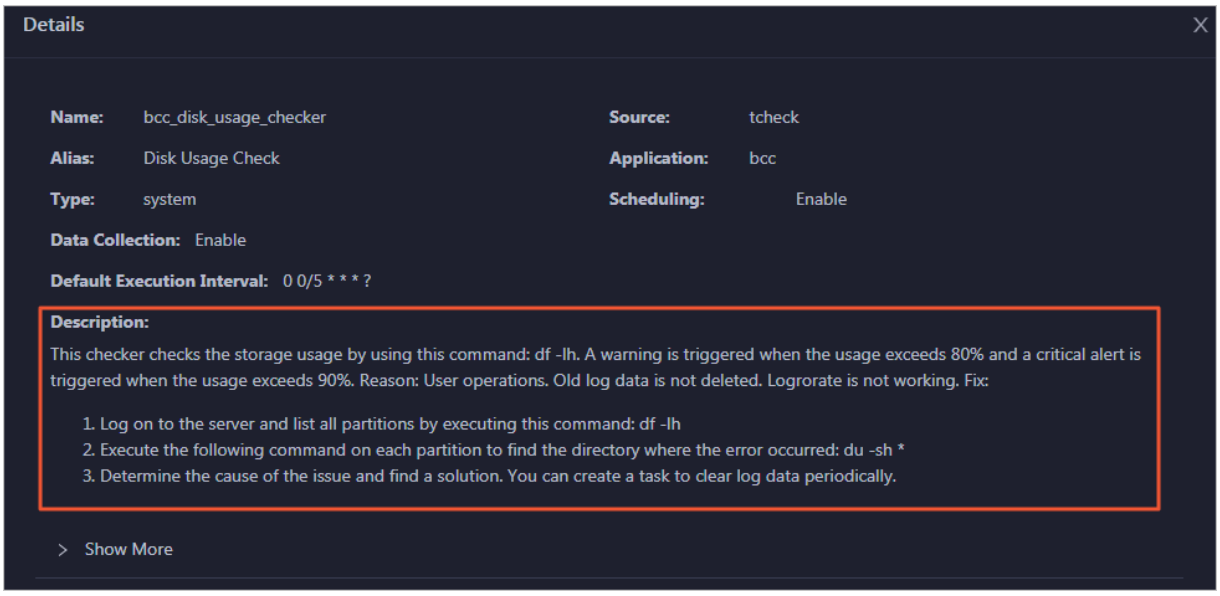

Log on to a host

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.

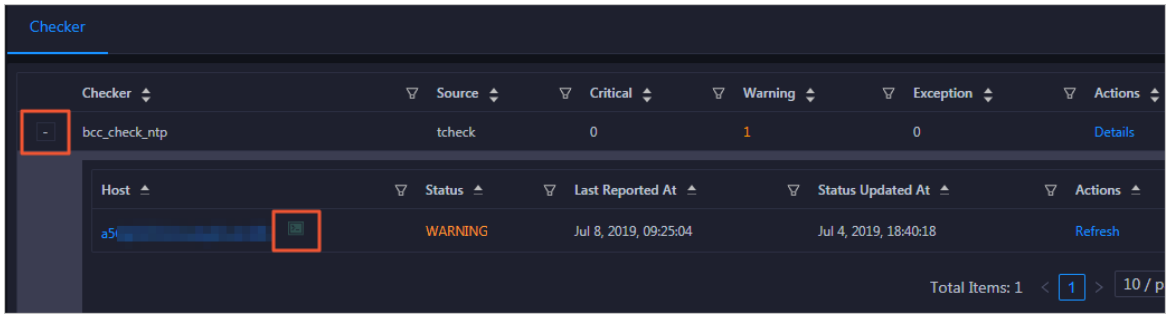

2. Click the **Log On** icon of a host. The **TerminalService** page appears.

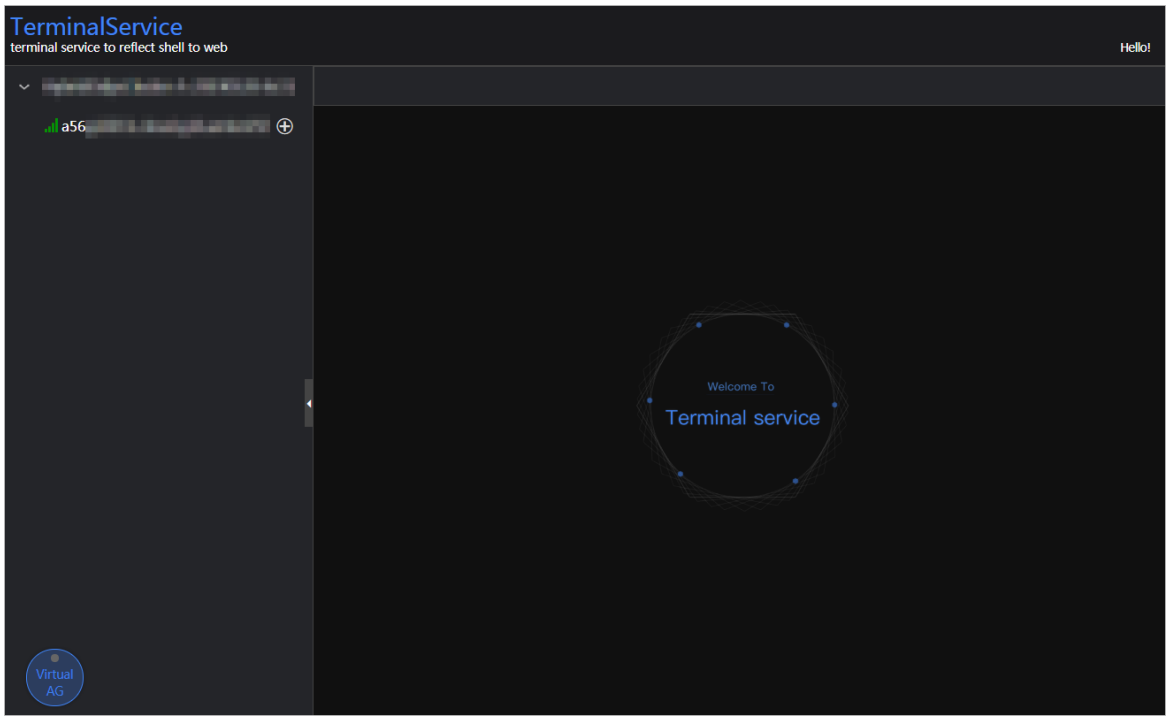

3. On the **TerminalService** page, click the hostname on the left to log on to the host.

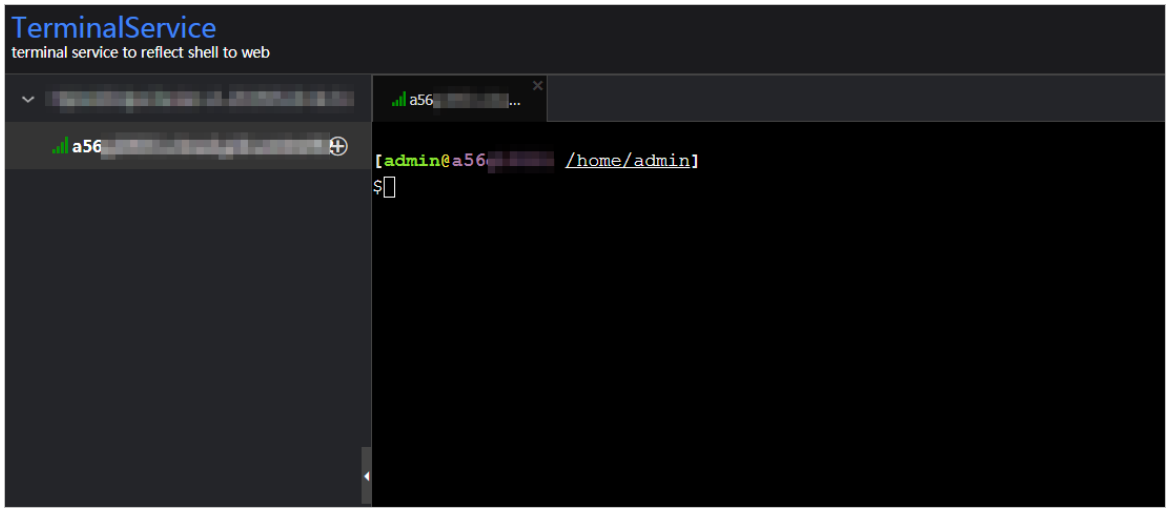

Run a checker again

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.

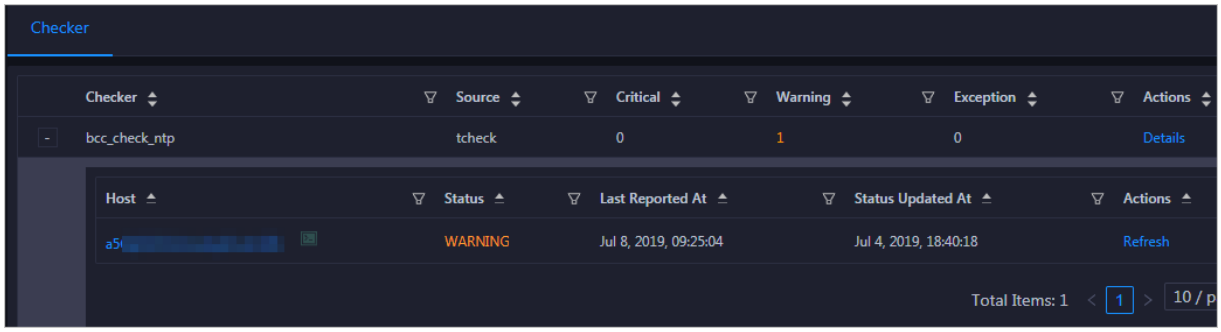

6.7.1.4.7.5. Host services

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.

6.7.1.5. 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.

6.7.1.6. Appendix

6.7.1.6.1. Installation environment

Operation system: AliOS5U7-x86-64

Template: Bigdata

6.7.1.6.2. 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

6.7.1.6.3. 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**.

7.Appendix

7.1. Operation Access Manager (OAM)

7.1.1. OAM

7.1.1.1. Introduction to OAM

This topic describes the features and permission model of Operation Administrator Manager (OAM).

Overview

OAM is a centralized permission management platform in the Apsara Uni-manager Operations Console. OAM uses a simplified role-based access control (RBAC) model. Administrators can use OAM to assign roles to O&M personnel who are then granted the corresponding operation permissions on O&M systems.

OAM permission model

In RBAC, administrators do not directly grant system operation permissions to users. Instead, they create a role set that can be associated with sets of users and permissions. Each role has a set of permissions. When a role is assigned to a user, the user is granted all the operation permissions of that role. This way, administrators only need to assign a role to a user when they create the user, eliminating the need to grant permissions to the user. In addition, changes in role permissions occur less often than changes in user permissions, which leads to simplified user permission management and reduced system overheads.

The following figure shows the OAM permission model.

OAM permission model

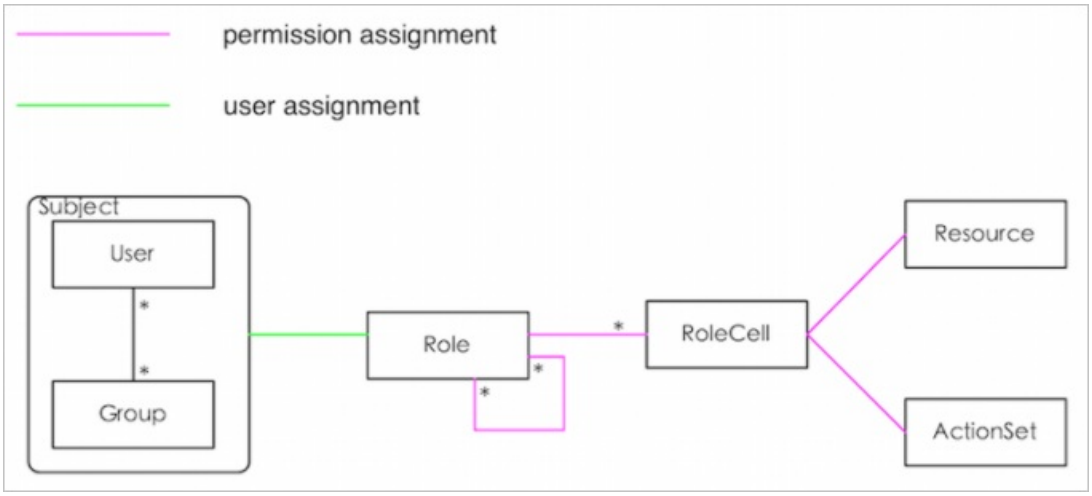

7.1.1.2. Usage instructions

Before you use OAM, you must understand the following basic terms about permission management.

subject

The operators of the access control system. OAM has two types of subjects: user and group.

user

The administrators and operators of O&M systems.

group

A collection of users.

role

The core of the role-based access control (RBAC) system.

Typically, a role can be regarded as a collection of permissions. One role can include multiple role cells or roles.

role hierarchy

In OAM, one role can include other roles to form role hierarchy.

role cell

The specific description of a permission. A role cell consists of resources, action sets, and grant options.

resource

The description of an authorized object. For more information about the resources of O&M platforms, see **Operation permissions on O&M platforms**.

action set

The description of authorized actions. An action set can include multiple actions. For more information about the actions of O&M platforms, see **Operation permissions on O&M platforms**.

grant option

The maximum number of grants in the cascaded grant, 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 Grant Option to 5 when Administrator A grants a permission to Administrator B, the permission can be granted another five times at most. When Administrator B grants the permission to Administrator C, the value of Grant Option cannot be greater than 4. If Grant Option 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** OAM does not support the cascaded revocation for cascaded grant. Therefore, Administrator C and Operator D still have the permission even if the permission is revoked for Administrator B.

7.1.1.3. Quick Start

By completing the steps in this guide, you will learn how to create and assign roles for O&M.

7.1.1.3.1. Log on to OAM

This topic describes how to log on to OAM.

Prerequisites

- The endpoint of the Apsara Uni-manager Operations Console and the username and password used to log on to the console are obtained from the deployment personnel or an administrator.

  The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id*.ops.console.*intranet-domain-id*.

- A browser is available. We recommend that you use Google Chrome.

Procedure

1. Open your Chrome browser.
2. In the address bar, enter the endpoint of the Apsara Uni-manager Operations Console. Press the Enter key.

   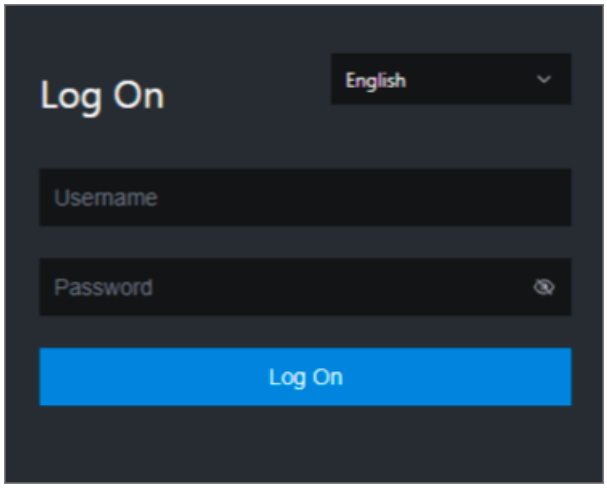

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password used to log on to the Apsara Uni-manager Operations Console from the deployment personnel or an administrator.

   When you log on to the Apsara Uni-manager Operations Console for the first time, you must change the password of your username.

   For security reasons, your password must meet the following requirements:

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %
   - The password must be 10 to 20 characters in length.

4. Click **Log On**.
5. In the top navigation bar, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. In the **Apsara Stack O&M** section, click **OAM**.

7.1.1.3.2. Create a group

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, set **Group Name** and **Description**.
4. Click **Confirm**.

   After the group is created, it is displayed on the **Owned Groups** page.

7.1.1.3.3. Add a group member

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 **Group Member** section, click **Add Member** in the upper-right corner.

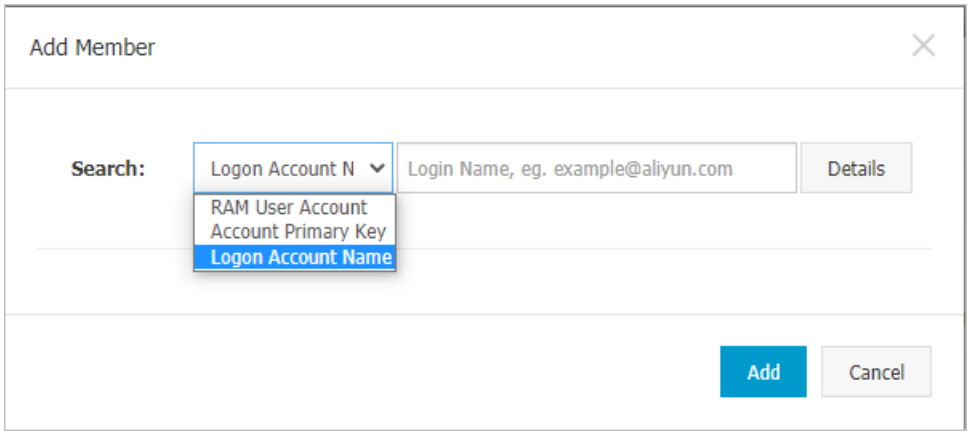

5. Select a search mode, enter the corresponding information, and then click **Details**. Details of the specified user are displayed.

   Three search modes are available:

   - **RAM User Account**: Enter a RAM user in the format of RAM user@Apsara Stack tenant account ID to search for the RAM user.
   - **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.

7.1.1.3.4. Add a group role

You can add roles to an existing group.

Prerequisites

- The role to be added is created. For more information, see Create a role.
- 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**.

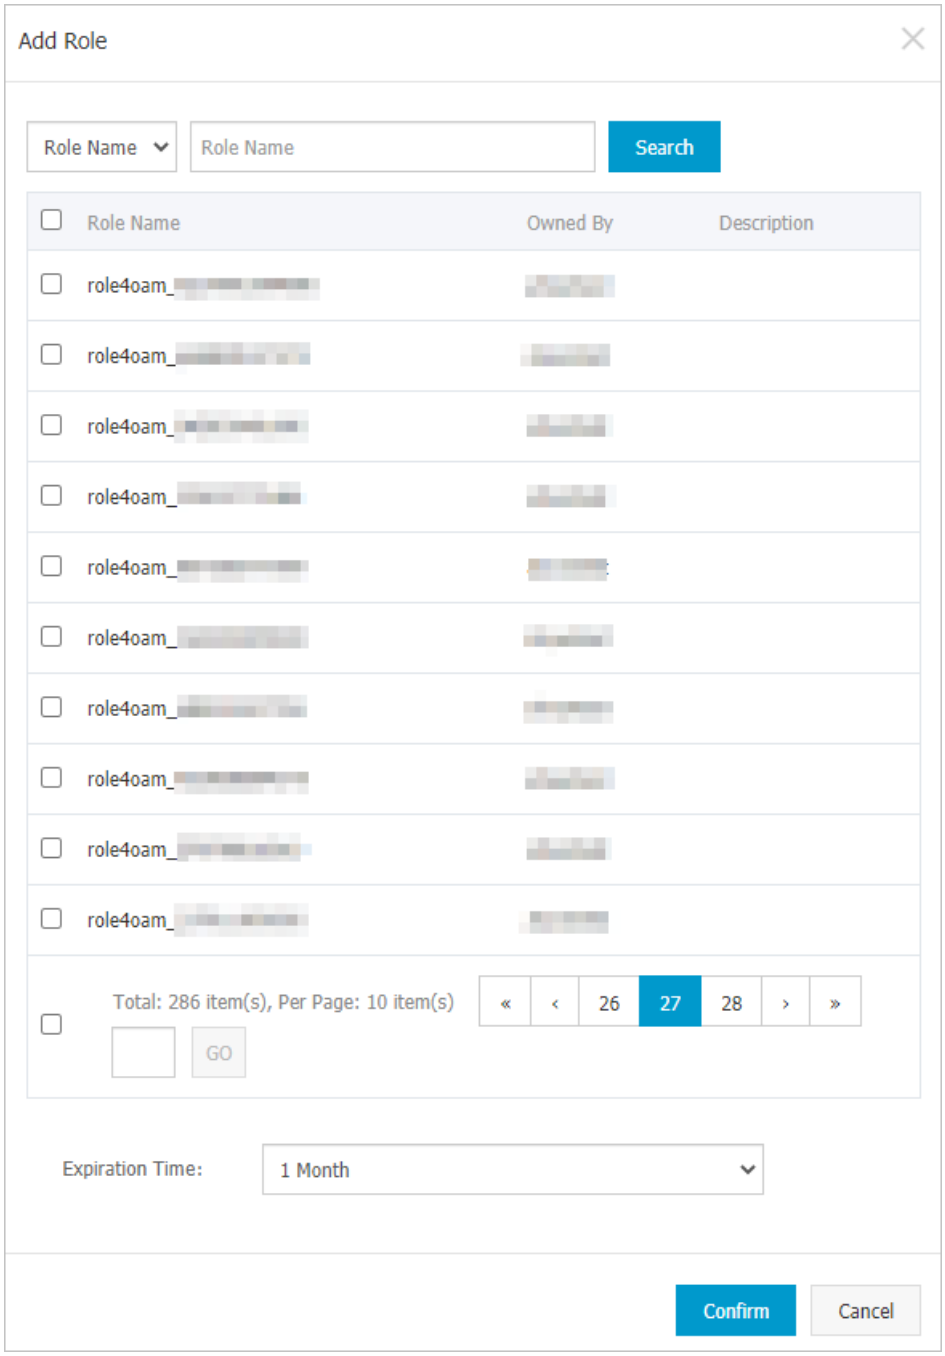

5. Search for roles by **Role Name**. Select one or more roles and set Expiration Time.
6. Click **Confirm**.

   To remove a role from a group, find the role in **Role List** and click **Remove** in the **Actions** column.

7.1.1.3.5. Create a role

This topic describes how to create a role.

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.

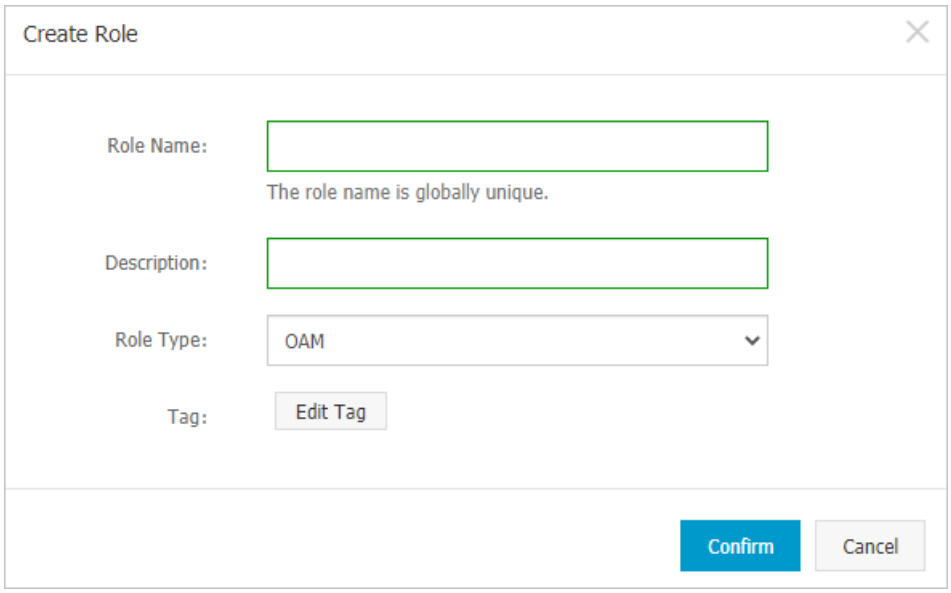

4. In the Create Role dialog box, set **Role Name**, **Description**, and **Role Type**.
5. (Optional)Add tags for the role. Tags can be used to search for roles.
   i. Click **Edit Tag**.

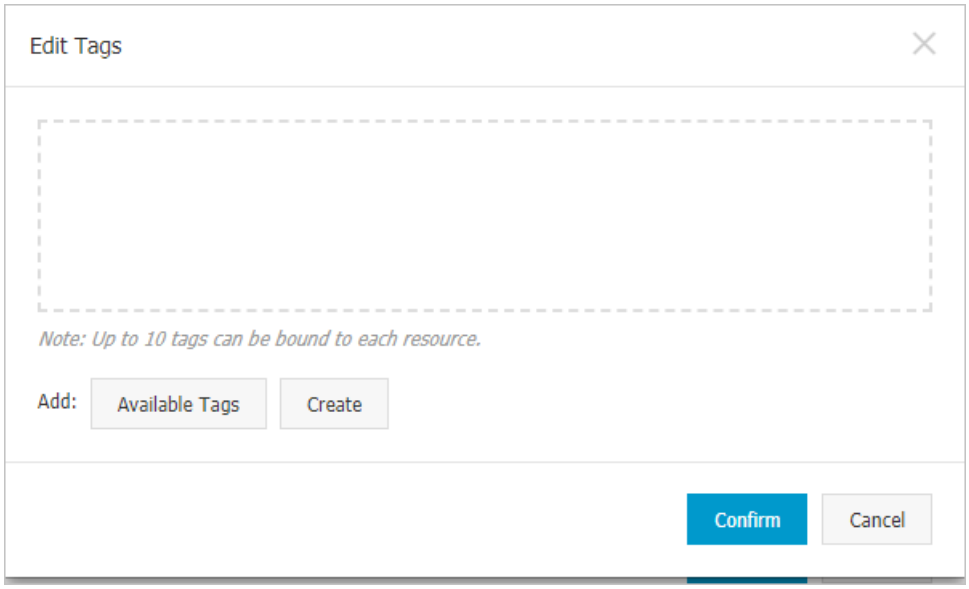

   ii. In the **Edit Tags** dialog box, click **Create**.

iii. Set **Key** and **Value** for the tag and click **Confirm**.

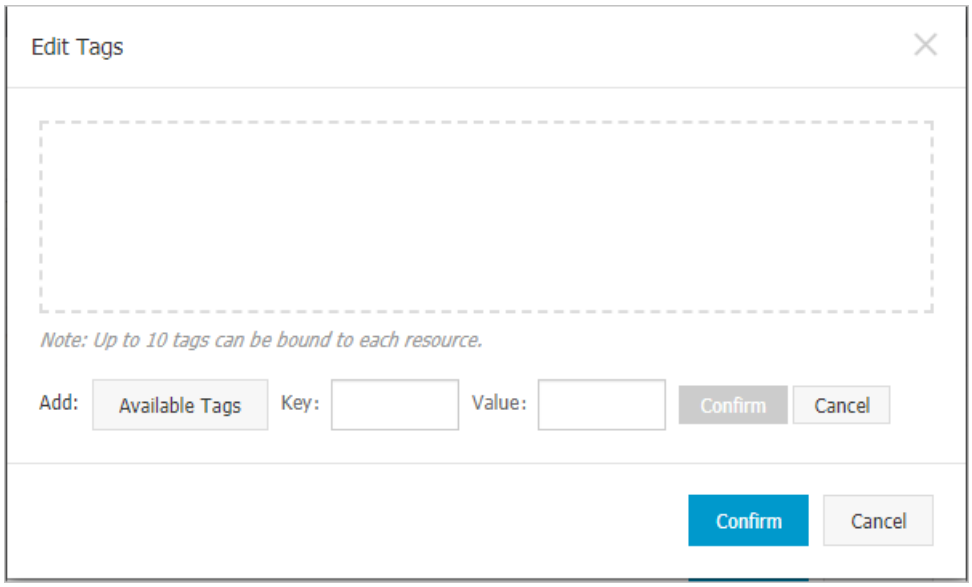

iv. Repeat the preceding step to create more tags.

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.

7.1.1.3.6. Add an inherited role 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, search for roles by **Role Name**. Select one or more roles.

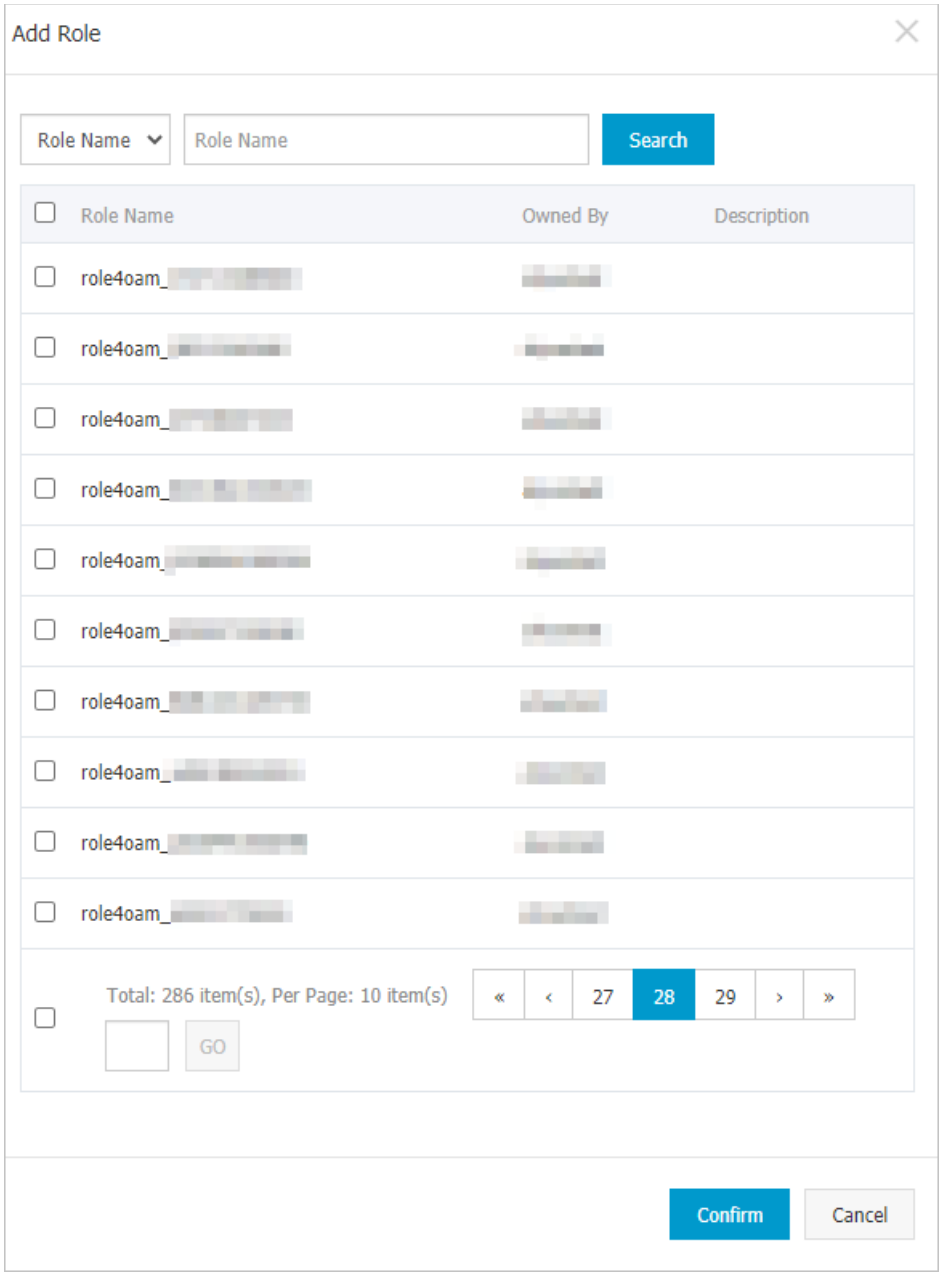

6. Click **Confirm**.

7.1.1.3.7. Add a resource to a role

After you create a role, you must add the permissions on resources to the 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**.

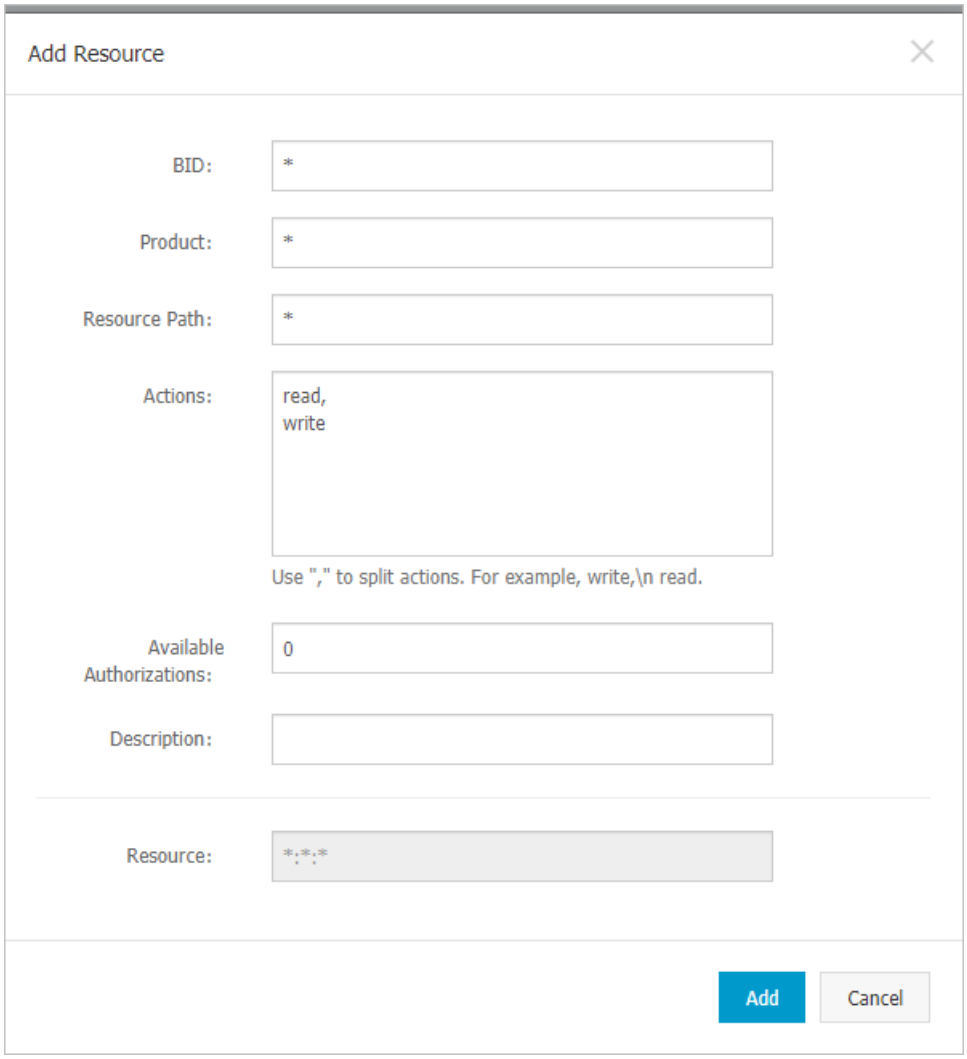

6. In the Add Resource dialog box, configure the parameters. For more information, see Parameters.

Parameters

Parameter	Description
BID	The deployment region ID.
Product	The cloud service to be added. Example: rds. **Note** The cloud service name must be lowercase. For example, enter **rds** instead of **RDS**.
Resource Path	The resources of the cloud service. For more information about resources of the O&M platforms, see **Operation permissions on O&M platforms**.
Actions	An action set, which can contain multiple actions. For more information about actions on the O&M platforms, see **Operation permissions on O&M platforms**.

Parameter	Description
Available Authorizations	The maximum number of grants in cascaded grant, 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**.

7.1.1.3.8. Assign a role to authorized users

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 Uni-manager Operations Console. For more information about how to create a user group, see Create a group.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role that you want to assign to a 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.

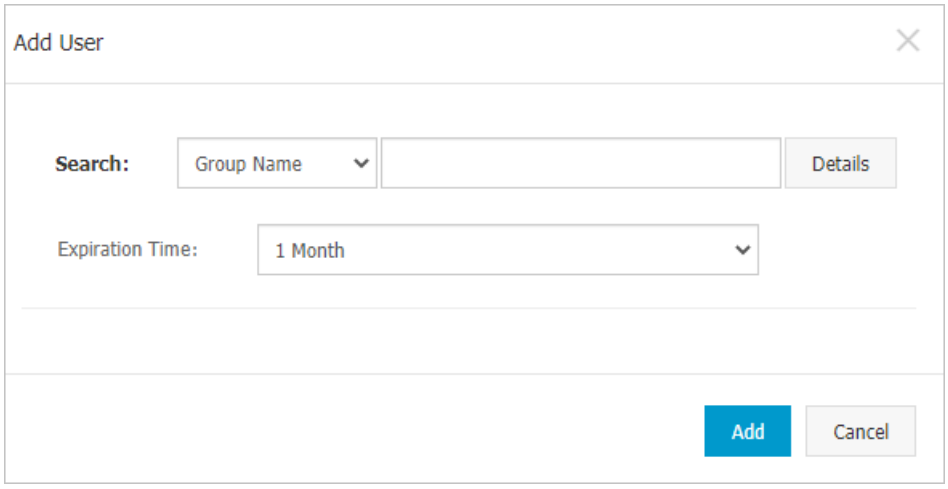

6. Select a search mode and enter corresponding information to search for the user to which you want to assign the role.

   Four search modes are available:

   - **RAM User Account**: Enter a RAM user in the format of *RAM user@Apsara Stack tenant account ID* to search for the RAM user.
   - **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 a group.

7. Set Expiration Time.

   When the specified expiration time is due, the user no longer has the permissions of the role. To grant permissions to the user again, 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.

7.1.1.4. Manage groups

Group management allows you to view, modify, and delete groups.

7.1.1.4.1. 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, modify the group name and description.
6. Click **Confirm**.

7.1.1.4.2. View group role details

You can view the information about 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 and click **Manage** in the **Actions** column.
4. In **Role List** section, click **Details** in the Actions column corresponding to the 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 resources to this role, see Add a resource 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.

7.1.1.4.3. Delete a group

You can delete a group that is 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 that you want to delete and click **Remove** in the **Actions** column. In the message that appears, click **OK**.

7.1.1.4.4. 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.

7.1.1.5. Manage roles

Role management allows you to view, modify, transfer, and delete roles.

7.1.1.5.1. 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** search box and click **Search** to search for roles that meet the search condition.

   > **Note** If the role that you want to search for has a tag, you can click **Tag** and select a tag key to search for the role based on the tag.

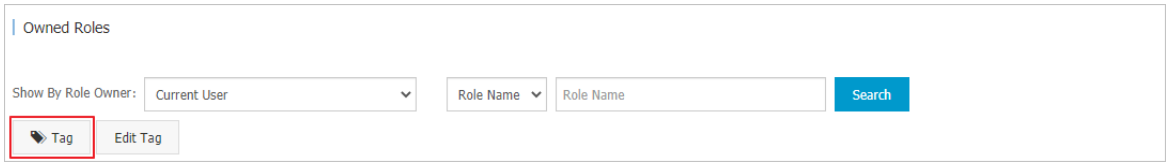

7.1.1.5.2. Modify role information

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, set **Role Name**, **Description**, **Role Type**, and **Tag**.
6. Click **Confirm**.

7.1.1.5.3. 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.

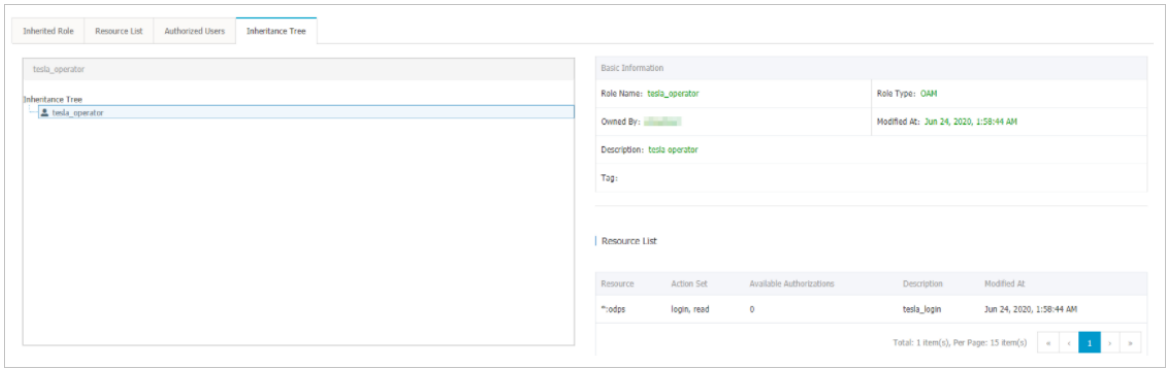

7.1.1.5.4. Transfer a role

You can transfer a role to other users or groups based on your 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, set the search conditions and search for the roles that you want to transfer.
4. Select one or more roles in the search results and click **Transfer** in the lower-left corner.
5. In the **Transfer** dialog box, 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**: Enter a RAM user in the format of RAM user@Apsara Stack tenant account ID to search for the RAM user.
   - **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.

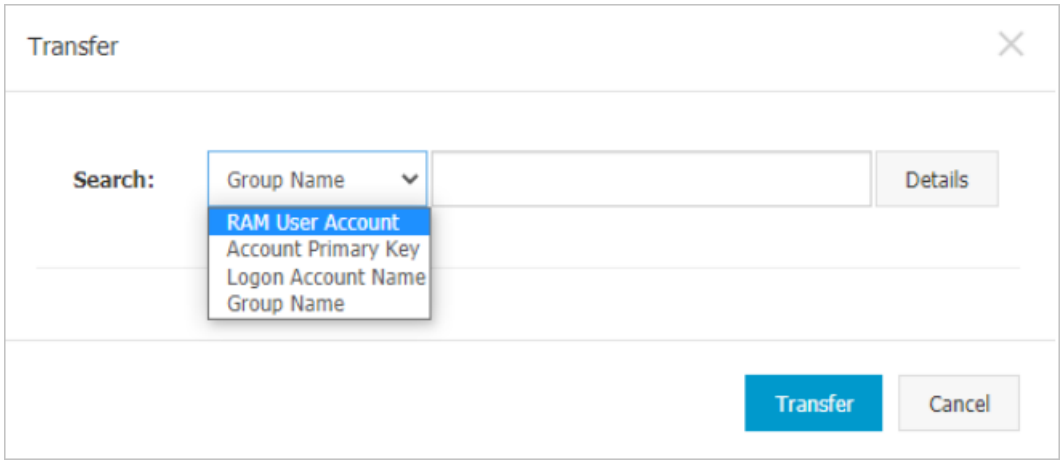

6. Click **Transfer**.

7.1.1.5.5. Delete a role

You can delete roles that are no longer needed.

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. Find the role that you want to delete and click **Delete** in the Actions column. In the message that

appears, click **OK**.

7.1.1.5.6. View assigned roles

You can view the roles assigned to you and the 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, 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.

7.1.1.5.7. View all roles

You can view all the roles in 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.

7.1.1.6. 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.

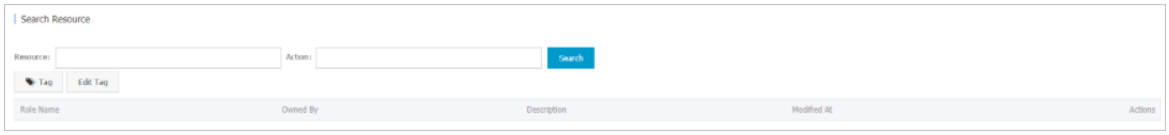

4. Click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

7.1.1.7. 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, 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, 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 (,).

7.1.1.8. Default roles and permissions

7.1.1.8.1. Default roles and their functions

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

7.1.1.8.1.1. Default roles of OAM

This topic describes the default roles of Operation Administrator Manager (OAM) and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
Super administrator	Has all permissions of the root user.	*:*	*	10

7.1.1.8.1.2. Default roles of Tablestore Operations and Maintenance System

This topic describes the default roles of Tablestore Operations and Maintenance System and the corresponding grant options.

Tablestore Operations and Maintenance System is an O&M platform for Tablestore.

The following table describes the default roles of Tablestore Operations and Maintenance System and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
Public permissions on Tablestore Operations and Maintenance System	Has basic permissions on Tablestore Operations and Maintenance System. This role is required for granting.	*:ots:*	["*"]	0

7.1.1.8.1.3. Default roles of Apsara Infrastructure Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding grant options.

Apsara Infrastructure Management Framework is a distributed data center management system. It can manage applications within clusters that include multiple machines and provide basic features such as deployment, upgrade, scale-in, scale-out, and configuration change.

The following table describes the default roles of Apsara Infrastructure Management Framework and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
Tianji_Project read-only	Has read-only permissions on Apsara Infrastructure Management Framework projects, which allows you to view the configurations and status of all projects and clusters.	*:tianji:projects	["read"]	0

Role	Description	Resource	Actions	Grant Option
Tianji_Project administrator	Has all permissions on Apsara Infrastructure Management Framework projects, which allows you to view and modify the configurations and status of all projects and clusters.	*:tianji:projects	["*"]	0
Tianji_Service read-only	Has read-only permissions on 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 permissions on 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 permissions on Apsara Infrastructure Management Framework data centers, which allows you to view and modify the information of data centers.	*:tianji:idcs	["*"]	0

Role	Description	Resource	Actions	Grant Option
Tianji administrator	Has all permissions on Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastructure Management Framework configurations.	*:tianji	["*"]	0
Tianji_Report_group_RO	Has read-only permissions on Apsara Infrastructure Management Framework reports.	*:tianjireport:groups:*	["read"]	0
Tianji_Report_group_RW	Has read and write permissions on Apsara Infrastructure Management Framework reports.	*:tianjireport:groups:*	["*"]	0

7.1.1.8.1.4. Default roles of Webapp-rule

This topic describes the default roles of Webapp-rule and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
Webapp-rule O&M administrator	Has all permissions on Webapp-rule projects, which allows you to view, modify, add, and delete all the configurations and status information.	26842:webapp-rule:*	["read", "write"]	0

Role	Description	Resource	Actions	Grant Option
Webapp-rule read-only	Has read-only permissions on Webapp-rule projects, which allows you to view all the configurations and status information.	26842:webapp-rule:*	["read"]	0

7.1.1.8.1.5. Default roles of Grandcanal

This topic describes the default roles of Grandcanal and the corresponding grant options.

Grandcanal is an internally distributed process development framework that allows developers to assemble, retry, roll back, and manually intervene with processes. O&M engineers also can use Grandcanal to manually intervene with the corresponding processes.

The following table describes the default roles of Grandcanal and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
grandcanal.ADMIN	Has permissions of the Grandcanal administrator, which allows you to query workflows and activities, and retry, roll back, terminate, and restart workflows.	26842:grandcanal	["write","read"]	0
grandcanal.Reader	Has read-only permissions on Grandcanal, which allows you to perform only read operations.	26842:grandcanal	["read"]	0

7.1.1.8.1.6. Default roles of Tianjimon

This topic describes the default roles of Tianjimon and the corresponding grant options.

Tianjimon is the monitoring module of Apsara Infrastructure Management Framework and monitors the physical machines and services deployed based on Apsara Infrastructure Management Framework.

The following table describes the default roles of Tianjimon and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option

Role	Description	Resource	Actions	Grant Option
Tianjimon O&M	Has all the permissions on Tianjimon, which allows you to perform basic monitoring and O&M operations.	26842:tianjimon:*	["*"]	0

7.1.1.8.1.7. 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	Description	Resource	Actions	Grant Option
Rtools administrator	Has all permissions in the Rtools console.	26842:drds:rtools:*	*	0

7.1.1.8.1.8. Default roles of the Apsara Uni-manager Operations Console

This topic describes the default roles of the Apsara Uni-manager Operations Console and the corresponding grant options.

The Apsara Uni-manager Operations Console 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 the Apsara Uni-manager Operations Console and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
		: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
	Has permissions			

Role	Description	Resource	Actions	Grant Option
System administrator	to manage platform nodes, physical devices, and virtual resources, back up, restore, and migrate product data, and query and back up system logs.	*: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
		: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
Security officer	Has permissions to manage permissions, security polices, and network security, and review and analyze security logs and activities of security auditors.	*:aso:auth:*	["read","write"]	0
		:aso:plat-access:	["read","write"]	0
		:aso:twoFactorAuth:	["read","write"]	0
Security auditor	Has 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
	Has permissions	*:aso:api-adapter:*	["read"]	0
		:aso:backup:	["read"]	0
		:aso:cmdb:	["read"]	0
		:aso:doc:	["read"]	0
		:aso:fullview:	["read","write"]	0

Role	Description	Resource	Actions	Grant Option
Product O&M officer	to perform O&M operations such as data import and export, modification, configuration, upgrade, and troubleshooting coordination.	*: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
Common O&M officer	Has permissions to perform daily health checks and query service status, inventory information, and 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

Role	Description	Resource	Actions	Grant Option
Duty observer	Has permissions to view and monitor the dashboard and monitor system alerts.	*:aso:doc:*	["read"]	0
		:aso:fullview:	["read"]	0

7.1.1.8.1.9. 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	Description	Resource	Actions	Grant Option
PaaS_Operation_Manager	Has all the permissions on the PaaS console.	*:paas-ops:*	["*"]	0

7.1.1.8.1.10. 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	Description	Resource	Actions	Grant Option
ocp_readonly	Has 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

Role	Description	Resource	Actions	Grant Option
ocp_dev	Has all permissions on OCP, but does not have the grant permission.	*:oceanbase:role:ocp_dev	["access"]	0

7.1.1.8.1.11. 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	Description	Resource	Actions	Grant Option
Apsara Stack Security administrator	Has all permissions on Apsara Stack Security, which allows you to manage data in all Apsara Stack Security modules.	*:yundun-luban:*	["*"]	0
Apsara Stack Security viewer	Has read permissions on Apsara Stack Security, which allows you to read data in all Apsara Stack Security modules.	*:yundun-luban:*	["read"]	0

7.1.1.8.1.12. 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	Description	Resource	Actions	Grant Option
Instance querier	Has permissions to query various instance resources.	*:qitian:instance:*	["read","create","delete","update"]	0
		:qitian:user:	["read","create","delete","update"]	0
Product O&M personnel	Has permissions to use the features under the "Products" menu of Apsara Network Intelligence.	*:qitian:product:*	["read","create","delete","update"]	0
		:qitian::*	["read","create","delete","update"]	0
R&D and O&M personnel	Has permissions to use the features under the "System" menu of Apsara Network Intelligence.	*:qitian:system:*	["read","create","delete","update"]	0
		:qitian::*	["read","create","delete","update"]	0

7.1.1.8.1.13. Default roles of CDS

This topic describes the default roles of Content Delivery Service (CDS) and the corresponding grant options.

CDS is a content distribution service that allows you to access content cached in a cache server that is closest to you when you visit a website. This reduces the response time of your request and network latency, and reduces loads on the website server.

The following table describes the default roles of CDS and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
Cds_admin	Has permissions to perform universal O&M and storage operations.	26842:cds:management:*	["read", "write"]	0
Cds_super_admin	Has permissions of the super user.	26842:cds:management:*	["read", "write", "super_read", "super_write"]	0

7.1.1.8.1.14. Default roles of Config Logic Quarantine Fram

This topic describes the default roles of Config Logic Quarantine Fram and the corresponding grant options.

Config Logic Quarantine Fram is the O&M platform of Apsara Distributed File System and provides features such as cluster list, cluster overview, monitoring data statistics, and screen monitoring of Distributed System Performance Monitor.

The following table describes the default roles of Config Logic Quarantine Fram and the corresponding grant options.

Role	Description	Resource	Actions	Grant Option
panshi_admin	Has permissions of the administrator.	*:panshi	*	10

7.1.1.8.1.15. Default roles of ECS

This topic describes the default roles of Elastic Compute Service (ECS) and the corresponding grant options.

ECS is a simple and efficient cloud computing service that provides elastic processing capabilities. ECS enables you to build stable and secure applications and perform O&M in a more efficient manner and helps lower IT costs.

Role	Description	Resource	Actions	Grant Option
ECS_administrator	Has all permissions of the ECS administrator.	26842:ecs	["*"]	0

Role	Description	Resource	Actions	Grant Option
ECS_read-only	Has all permissions to perform read operations.	26842:ecs	["inner_getAllUrls","inner_getCurrentUser","inner_getAccountByIdkp","inner_getIdkpByAccount","inner_allErrorCode","inner_getOptions","vm_describe","vm_export","vm_describeMountedSnapshots","region_describeRegions","group_queryVms","group_queryAcls","group_describe","disk_describe","monitor_deviceIOStat","monitor_vmMonitor","monitor_deviceIOBlock","monitor_deviceLatency","nc_queryAvailableNcs","snapshot_describe","vnc_generateUrl","iso_queryAvailableIsos","iso_queryMountedIso","image_describe"]	0

7.1.1.8.1.16. Default roles of SLB

This topic describes the default roles of Server Load Balancer (SLB) and the corresponding grant options.

SLB is a load balancing service that distributes network traffic across multiple Elastic Compute Service (ECS) instances. SLB distributes network traffic to backend servers to improve the responsiveness and availability of your applications. You can use SLB to prevent service interruption caused by single points of failure (SPOFs) and improve the availability of applications. SLB can be used with ECS instances and Enterprise Distributed Application Service (EDAS) applications.

Role	Description	Resource	Actions	Grant Option
SLB_administrator	Has permissions to manage all data in SLB.	26842:slb	["read","create","delete","update"]	0
SLB_read-only	Has permissions to read all data in SLB.	26842:slb	["read"]	0

7.1.1.8.1.17. Default roles of VPC

This topic describes the default roles of Virtual Private Cloud (VPC) and the corresponding grant options.

VPC allows you to build a custom, private network on Apsara Stack. VPCs are logically isolated from each other. You can create and manage instances, such as Elastic Compute Service (ECS) instances, Server Load Balancer (SLB) instances, and ApsaraDB RDS instances, in your VPCs.

Role	Description	Resource	Actions	Grant Option
VPC_administrator	Has permissions to manage all data in VPC.	26842:vpc	["read","create","delete","update"]	0
VPC_read-only	Has permissions to read all data in VPC.	26842:vpc	["read"]	0

7.1.1.8.1.18. Default roles of MaxCompute

This topic describes the default roles of MaxCompute and the corresponding grant options.

MaxCompute is an industry-leading distributed big data processing platform developed by Alibaba Cloud.

Role	Description	Resource	Actions	Grant Option
mcadmin	Has all permissions of the MaxCompute substation administrator.	26842:bcc:/api/ias/	["*"]	0
		26842:bcc:/api/bccapi/odps/	["*"]	0
		26842:bcc:/api/tflow/	["*"]	0
		26842:bcc:/api/bccapi/sysadmin/	["*"]	0
mc administrator	Has all permissions to perform operations in the MetaCenter console.	26842:drds:mc:*	["*"]	0

7.1.1.8.1.19. Default roles of AnalyticDB

This topic describes the default roles of AnalyticDB and the corresponding grant options.

AnalyticDB is a real-time data warehousing service that can process petabytes of data with high concurrency and low latency. It is fully compatible with the MySQL protocol and SQL:2003 syntax and can perform instant multi-dimensional analysis and business exploration for terabytes of data within milliseconds. AnalyticDB offers flexible computing for massive data volumes, rapid response, and low costs to drive the big data business revolution.

Role	Description	Resource	Actions	Grant Option
adadmin	Has all permissions of the AnalyticDB administrator.	26842:bcc:/api/ias/	["*"]	0
		26842:bcc:/api/bccapi/ads/	["*"]	0
		26842:bcc:/api/tflow/	["*"]	0
		26842:bcc:/api/bccapi/sysadmin/	["*"]	0

7.1.1.8.1.20. Default roles of StreamCompute

This topic describes the default roles of StreamCompute and the corresponding grant options.

StreamCompute is a streaming big data analysis platform that runs in Apsara Stack. The platform provides tools to analyze streaming data in real time on the cloud.

Role	Description	Resource	Actions	Grant Option
scadmin	Has all permissions of the StreamCompute administrator.	26842:bcc:/api/ias/	["*"]	0
		26842:bcc:/api/bccapi/galaxy/	["*"]	0
		26842:bcc:/api/tflow/	["*"]	0
		26842:bcc:/api/bccapi/sysadmin/	["*"]	0

7.1.1.8.1.21. Default roles of DataWorks

This topic describes the default roles of DataWorks and the corresponding grant options.

DataWorks is a platform that provides offline data processing, analysis, and mining capabilities. DataWorks supports fully hosted scheduling of tasks by time and dependency. It allows tens of millions of tasks to be performed on time with maximum accuracy based on DAG relationships every day. DataWorks offers visual task monitoring and management tools and displays global conditions in the DAG format when tasks are running.

Role	Description	Resource	Actions	Grant Option
dwadmin	Has all permissions of the DataWorks administrator.	26842:bcc:/api/ias/	["*"]	0
		26842:bcc:/api/bccapi/base/	["*"]	0
		26842:bcc:/api/tflow/	["*"]	0
		26842:bcc:/api/bccapi/sysadmin/	["*"]	0

7.1.1.8.1.22. Default roles of Big Data Manager

This topic describes the default roles of Big Data Manager and the corresponding grant options.

Big Data Manager is an O&M platform tailored for big data services. Big Data Manager supports O&M for big data services from various aspects, such as business, service, cluster, and host.

Role	Description	Resource	Actions	Grant Option
bcc_admin	Has all permissions of the super administrator of the Big Data Manager backend.	26842:bcc:/api/product/odps/	["*"]	0
		26842:bcc:/api/product/ads/	["*"]	0
		26842:bcc:/api/product/dataworks/	["*"]	0
		26842:bcc:/api/product/apsara/	["*"]	0
		26842:bcc:/api/product/minirds/	["*"]	0
		*:bcc	["*"]	0
		26842:bcc:/api/product/streamcompute/	["*"]	0
		26842:bcc:/api/product/minilvs/	["*"]	0
bcc_admin_ads	Has all permissions of the AnalyticDB administrator of the Big Data Manager backend.	26842:bcc:/api/product/ads/	["*"]	0

Role	Description	Resource	Actions	Grant Option
bcc_admin_odps	Has all permissions of the MaxCompute administrator of the Big Data Manager backend.	26842:bcc:/api/product/odps/	["*"]	0
bcc_admin_dataworks	Has all permissions of the DataWorks administrator of the Big Data Manager backend.	26842:bcc:/api/product/dataworks/	["*"]	0
bcc_admin_streamcompute	Has all permissions of the StreamCompute administrator of the Big Data Manager backend.	26842:bcc:/api/product/streamcompute/	["*"]	0
bcc_admin_dataapp	Has all permissions of the Data Application administrator of the Big Data Manager backend.	26842:bcc:/api/product/elasticsearch/	["*"]	0
		26842:bcc:/api/product/asap/	["*"]	0
		26842:bcc:/api/product/datahub/	["*"]	0
		26842:bcc:/api/product/dataphin/	["*"]	0
		26842:bcc:/api/product/datav/	["*"]	0
		26842:bcc:/api/product/dtboost/	["*"]	0
		26842:bcc:/api/product/es/	["*"]	0
		26842:bcc:/api/product/iplus/	["*"]	0
		26842:bcc:/api/product/pai/	["*"]	0
		26842:bcc:/api/product/quickbi/	["*"]	0

Role	Description	Resource	Actions	Grant Option
bcc_admin_biggraph	Has all permissions of the BigGraph administrator of the Big Data Manager backend.	26842:bcc:/api/product/biggraph/	["*"]	0
bcc_account_admin	Has all permissions of the account administrator of the Big Data Manager backend.	26842:bcc:/api/account/	["*"]	0

7.1.1.8.1.23. Default roles of ApsaraDB RDS

This topic describes the default roles of ApsaraDB RDS and the corresponding grant options.

ApsaraDB RDS is a stable and reliable online database service that supports elastic scaling. ApsaraDB RDS is built on top of the Apsara distributed operating system and high-performance storage technologies and supports a wide range of engines, such as MySQL, SQL Server, PostgreSQL, and Postgres Plus Advanced Server (PPAS, highly compatible with Oracle). ApsaraDB RDS provides a portfolio of solutions to disaster recovery, data backup, fault recovery, business monitoring, and data migration to facilitate database operations and maintenance.

Role	Description	Resource	Actions	Grant Option
ROLE_CONTROLLER	Has all permissions of the ApsaraDB RDS controller.	26842:rds	["TASK_START_STEP","TASK_CLOSE"]	0

7.1.1.8.2. Operation permissions on O&M platforms

This topic describes the operation permissions on O&M platforms.

7.1.1.8.2.1. Permissions on Apsara Infrastructure Management Framework

This topic describes the operation permissions on Apsara Infrastructure Management Framework.

Resource	Operation	Description
*:tianji:services:[sname]:tjmontemplates:[tmplname]	delete	Deletes a monitoring template.

Resource	Operation	Description
*:tianji:services:[sname]:tjmontemplates:[tmplname]	write	Creates a monitoring template.
*:tianji:services:[sname]:templates:[tmplname]	write	Creates a service template.
*:tianji:services:[sname]:templates:[tmplname]	delete	Deletes a service template.
*:tianji:services:[sname]:serviceinstances:[siname]:tjmontemplate	read	Obtains a monitoring template.
*:tianji:services:[sname]:serviceinstances:[siname]:tssessions	terminal	Creates a remote service.
*:tianji:services:[sname]:serviceinstances:[siname]:template	write	Updates a service template reference.
*:tianji:services:[sname]:serviceinstances:[siname]:template	delete	Deletes a service template.
*:tianji:services:[sname]:serviceinstances:[siname]:template	read	Obtains a service template.
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	delete	Deletes a service template tag.
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	write	Adds a service template tag.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	Obtains a service resource.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	write	Modifies a machine.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	read	Obtains a machine.

Resource	Operation	Description
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	delete	Deletes a machine.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	read	Obtains a machine role.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	delete	Batch deletes machines.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	write	Modifies a machine role.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:resources	read	Obtains a service resource.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	Obtains Apsara Infrastructure Management Framework logs.
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles	read	Obtains a service role.
*:tianji:services:[sname]:serviceinstances:[siname]:schema	write	Sets a service specification.
*:tianji:services:[sname]:serviceinstances:[siname]:schema	delete	Deletes a service specification.
*:tianji:services:[sname]:serviceinstances:[siname]:rollings:[version]	write	Modifies an upgrade task.
*:tianji:services:[sname]:serviceinstances:[siname]:rollings	read	Lists upgrade tasks.

Resource	Operation	Description
*:tianji:services:[sname]:serviceinstances:[siname]:resources	read	Obtains an instance resource.
*:tianji:services:[sname]:serviceinstances:[siname]:machines:[machine]	read	Obtains all the machine roles.
*:tianji:services:[sname]:serviceinstances:[siname]	write	Deploys a service instance.
*:tianji:services:[sname]:serviceinstances:[siname]	read	Obtains service configurations.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:name	read	Obtains a list of machine service files.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:download	read	Obtains the information about downloading a machine service file.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:content	read	Obtains the content of a machine service file.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:filelist	read	Obtains a list of machine files.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:dockerlogs	read	Obtains container logs.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:debuglog	read	Obtains machine debugging information.
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps	read	Obtains a list of machine services.

Resource	Operation	Description
*:tianji:services:[sname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	Obtains the information about a container.
*:tianji:services:[sname]:schemas:[schemaname]	write	Modifies a service specification.
*:tianji:services:[sname]:schemas:[schemaname]	delete	Deletes a service specification.
*:tianji:services:[sname]:resources	read	Obtains a service resource.
*:tianji:services:[sname]	delete	Deletes a service.
*:tianji:services:[sname]	write	Creates a service.
*:tianji:projects:[pname]:machinebuckets:[bname]:machines:[machine]	read	Obtains machine information.
*:tianji:projects:[pname]:machinebuckets:[bname]:machines	read	Obtains a list of machines.
*:tianji:projects:[pname]:machinebuckets:[bname]	write	Creates a machine pool.
*:tianji:projects:[pname]:machinebuckets:[bname]	write	Modifies a machine pool.
*:tianji:projects:[pname]:machinebuckets:[bname]	delete	Deletes a machine pool.
*:tianji:projects:[pname]:machinebuckets:[bname]	read	Obtains a list of machines.
*:tianji:projects:[pname]:machinebuckets	read	Obtains a list of machine pools.
*:tianji:projects:[pname]:projects:[pname]:clusters:[cname]:tssessions:[tssessionname]:tsses	terminal	Updates a remote connection.

Resource	Operation	Description
*:tianji:projects:[pname]:projects:[pname]:clusters:[cname]:tssessions	terminal	Creates a remote connection.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tjmontemplate	read	Obtains a service monitoring instance.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	delete	Deletes a service monitoring instance.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	write	Sets a service template.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	read	Obtains a service template.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	write	Adds a service product tag.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	delete	Deletes a service product tag.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	Obtains a role resource.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:name	read	Obtains a list of machine service files.

Resource	Operation	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:download	read	Obtains the information about downloading a machine service file.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:content	read	Obtains the content of a machine service file.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:filelist	read	Obtains a list of machine files.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:dockerlogs	read	Obtains container logs.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:debuglog	read	Obtains machine debugging information.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps	read	Obtains a list of machine files.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	read	Obtains role information.

Resource	Operation	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	write	Modifies machine role information.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	delete	Deletes a machine role.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	write	Modifies machine role information.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	delete	Batch deletes machine roles.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	read	Obtains the information about all machine services.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:resources	read	Obtains a service resource.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	Obtains Apsara Infrastructure Management Framework logs.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	Obtains information about the container group.

Resource	Operation	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles	read	Obtains a service instance role.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	delete	Deletes a service specification.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	write	Sets a service specification.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:resources	read	Obtains an instance resource.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	delete	Deletes a service instance.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	write	Creates a service instance.
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	read	Obtains service instance configurations.
*:tianji:projects:[pname]:clusters:[cname]:rollings:[version]	write	Modifies an upgrade task.
*:tianji:projects:[pname]:clusters:[cname]:rollings	read	Obtains a list of upgrade tasks.
*:tianji:projects:[pname]:clusters:[cname]:resources	read	Obtains a cluster resource.
*:tianji:projects:[pname]:clusters:[cname]:quota	write	Sets a cluster quota.
*:tianji:projects:[pname]:clusters:[cname]:machinesinfo	read	Obtains machine information.

Resource	Operation	Description
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	read	Obtains all the machine roles.
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	write	Configures a machine operation.
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	delete	Deletes a machine operation.
*:tianji:projects:[pname]:clusters:[cname]:machines	write	Modifies a machine cluster.
*:tianji:projects:[pname]:clusters:[cname]:difflist	read	Obtains a list of edition differences.
*:tianji:projects:[pname]:clusters:[cname]:diff	read	Obtains the content of an edition difference.
*:tianji:projects:[pname]:clusters:[cname]:deploylogs:[version]	read	Obtains the content of a cluster deployment log.
*:tianji:projects:[pname]:clusters:[cname]:deploylogs	read	Obtains a list of cluster deployment logs.
*:tianji:projects:[pname]:clusters:[cname]:builds:[version]	read	Obtains the information about a build task.
*:tianji:projects:[pname]:clusters:[cname]:builds	read	Obtains a list of build tasks.
*:tianji:projects:[pname]:clusters:[cname]	write	Modifies a cluster.
*:tianji:projects:[pname]:clusters:[cname]	delete	Deletes a cluster.
*:tianji:projects:[pname]:clusters:[cname]	read	Obtains cluster configurations.
*:tianji:projects:[pname]:clusters:[cname]	write	Deploys a cluster.
*:tianji:projects:[pname]	write	Creates a project.
*:tianji:projects:[pname]	delete	Deletes a project.

Resource	Operation	Description
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	Creates a slot.
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	Sets slot properties.
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	delete	Deletes a slot.
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	Sets rack properties.
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	Creates a rack.
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	delete	Deletes a rack.
*:tianji:idcs:[idc]:rooms:[room]	write	Creates a room.
*:tianji:idcs:[idc]:rooms:[room]	delete	Deletes a room.
*:tianji:idcs:[idc]:rooms:[room]	write	Sets room properties.
*:tianji:idcs:[idc]	delete	Deletes a data center.
*:tianji:idcs:[idc]	write	Sets data center properties.
*:tianji:idcs:[idc]	write	Creates a data center.

7.1.1.8.2.2. 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

7.1.1.8.2.3. Permissions on Grandcanal

This topic describes the operation permissions on Grandcanal.

Resource	Action	Description
26842:grandcanal	read	Queries the activity details and summary of a workflow.
26842:grandcanal	write	Restarts, retries, rolls back, and terminates a workflow.

7.1.1.8.2.4. 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

7.1.1.8.2.5. 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.