ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Agility

Operations and Maintenance Guide

Product Version: 2102, Internal: V3.5.0

Document Version: 20210719

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

1.1. Overview

This topic describes the management system of the Apsara Uni-manager Operations service.

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.

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 service. The following figure shows the management framework of the Apsara Uni-manager Operations system.

Apsara Uni-manager Operations System

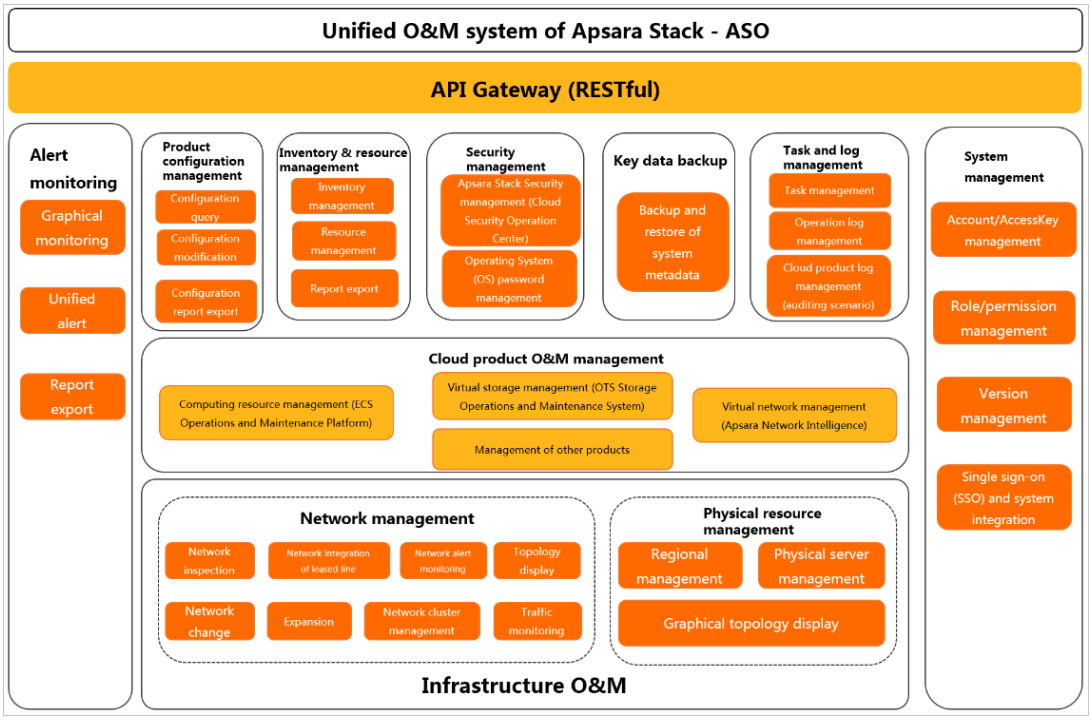

The Apsara Uni-manager Operations Console is a unified and intelligent O&M platform. In the Apsara Uni-manager Operations system, cloud operations is classified into the following layers: infrastructure, cloud service, and business operations. The management framework of the Apsara Uni-manager Operations service 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 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 operations processes to build an improved operations service management platform.

Based on ITIL and ISO 20000, the management framework of the Apsara Uni-manager Operations system uses management supporting 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.

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).

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.

The Apsara Uni-manager Operations system 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 service management system and can provide different levels of operations services based on the service agreements.

1.2. Get started

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.

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

1. Log on to the Apsara Uni-manager Operations Console as an administrator.
2. Create a role. For more information, see Manage roles.
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 **Appendix > OAM**, grant permissions to the role, and then grant the role to the corresponding operations account as an administrator.

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 URL of the Apsara Uni-manager Operations Console is in the following format: *region-id*.*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 (*region-id*.*ops*.asconsole.*intranet-domain-id*.com). Then, press the Enter key.

   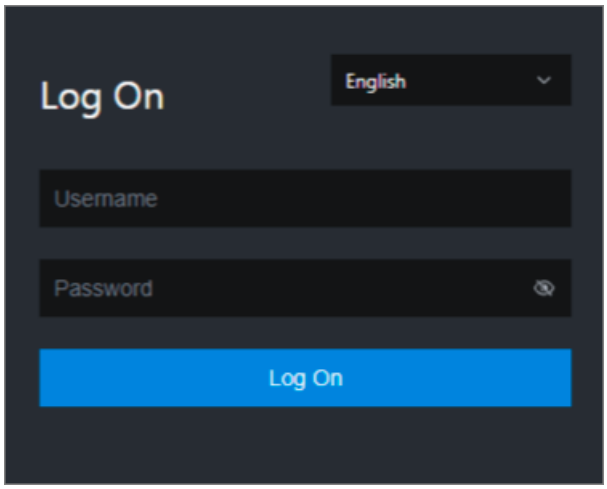

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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
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.

NO.	Section	Description
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.

1.2.4. 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.

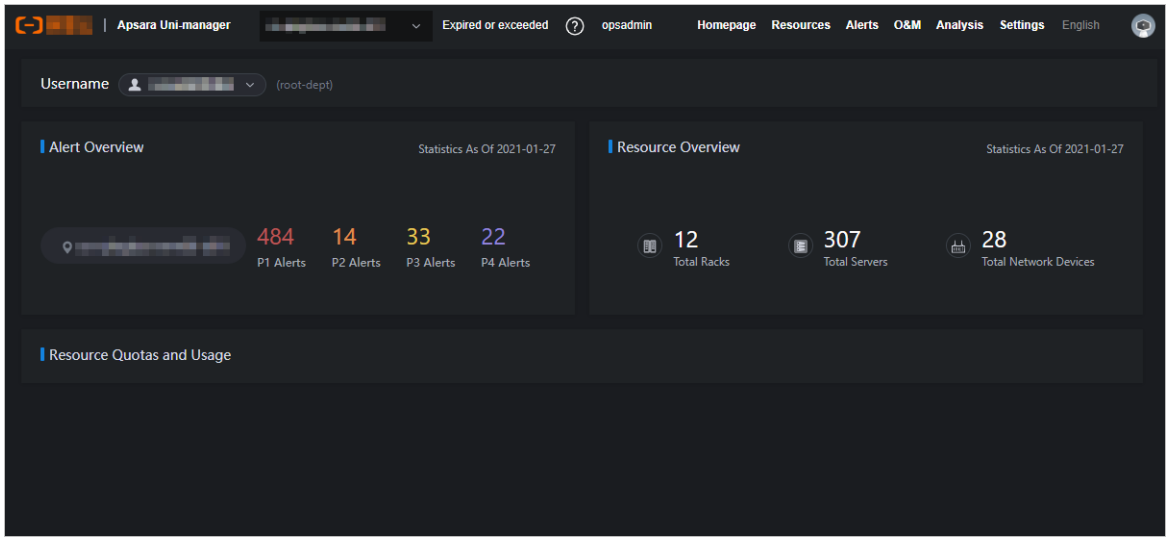

- In the **opsadmin** section, select a user. The department to which the user belongs is displayed on the right.
- In the **Alert Overview** section, view the total number of alerts at the P1, P2, P3, and P4 levels.
- In the **Resource Overview** section, view the total number of racks, servers, and 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 quantity, used and available resources, and usage.

Alibaba Cloud service	Statistical metric
RDS	CPU (Core)
	Disk (GB)
	Memory (GB)
OSS	Storage Capacity (GB)

1.3. Settings

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.
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.

1.3.2. Security policies

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 > 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 prohibited logon IP addresses. Click **OK**. If you select blacklist for logon settings, the specified IP addresses are prohibited from logging on to the Apsara Uni-manager Operations Console. If you select whitelist for logon settings, the specified IP addresses are allowed to log on to the Apsara Uni-manager Operations Console. For more information about operations on logon settings, 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.

1.3.2.2. Physical server passwords

The Server Password module allows you to configure and manage passwords, as well as query the historical passwords of all physical servers in the Apsara Stack environment.

Context

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

- The system automatically collects information of all the servers in the Apsara Stack environment.
- The server password is 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 appears. 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 filter 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 Passwords** tab, select multiple servers.
     b. Click **Batch Update** in the lower part of the tab.
     c. In the dialog box that appears, specify **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. Take note of the following items:

- **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**.

1.3.3. Log clearance

The Log Clearance 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.

1.3.3.1. Import container or server log cleanup rules

If you have configured log cleanup rules on your computer, you can batch import multiple cleanup rules for 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.
- If a container belongs to a service whose service roles are deployed on the PaaS platform, logs of the container cannot be cleaned 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 **Log Clearance > Rules**.
4. Click the **Containers** or **Servers** tab.
5. Click **Import**.
6. Select the XLS or XLSX files that you want to import and click **Open**. You can import multiple log clearance rules.

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

1.3.3.2. Export container or physical server log cleanup rules

You can batch export multiple container or physical server log cleanup rules.

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 Clearance > 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 service 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 service role in the list and select the corresponding search result. In subsequent queries, the system shows all available options in the drop-down list.

1.3.3.3. 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 Clearance > Rules**.
4. Click the **Containers** or **Servers** tab.
5. (Optional)In the upper part of the tab, select the product, service, and service role, and then click **Search** to query the cleanup rules that meet the filter conditions.

   > **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 service 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 utilization and specify whether to automatically clean up logs that match the cleanup rule.

   > **Note** The maximum disk utilization 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**.

1.3.3.4. 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 Clearance > Rules**.
4. Click the **Containers** or **Servers** tab.
5. (Optional)In the upper part of the tab, select the product, service, and service role, and click **Search** to query cleanup rules that meet the filter conditions.

   > **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 service 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 delete and click **Delete** in the **Actions** column.

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

> **Note** When you delete a log cleanup rule, execution plans corresponding to the rule are not deleted right away. Existing execution plans are cleaned up at 02:00 every day and new execution plans are generated based on the current cleanup rules.

1.3.3.5. Obtain the usage information of containers or physical servers

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

Solution 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 Clearance > 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 service role, and click **Search**. In the search results, find the container or physical server for which you want to query disk usage information. Click **Obtain Watermark** in the **Actions** column to obtain the disk usage information of the container or physical server.

     > **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 service 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 **Obtain Watermarks** 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 Clearance > Rules**.
4. Click the **Containers** or **Servers** tab.
5. (Optional)In the upper part of the page, select a product, service, and service role, and click **Search**.

> **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 service 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 Plan** in the **Actions** column. The **Execution Plan** page appears.

1.3.3.6. 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.

Solution 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 Clearance > Plans**.
4. Click the **Containers** or **Servers** 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 service role, and click **Search**. In the search results, find the container or physical server for which you want to clean up logs and click **Clear** in the **Actions** column.

     > **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 service 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 **Clear Logs** 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.

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 Clearance > Rules**.
4. Click the **Containers** or **Servers** tab.
5. (Optional)In the upper part of the tab, select a product, service, and service role, and click **Search**.

> **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 service 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 result, find the cleanup rule of the container or physical servers for which you want to clean up logs and click **Execution plan** in the **Actions** column. The **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 **Clear** in the **Actions** column to clean up the logs of a single container or physical server.
   - Select multiple containers or physical servers and click **Clear Logs** 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.

1.3.3.7. 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

Existing execution plans are cleaned up at 02:00 every day and new execution plans are generated 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 Clearance > Rules**.
4. Click the **Containers** or **Servers** 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 service role, and click **Search**. In the search results, find the cleanup rule for which you want to set automatic cleanups and turn on **Automatically delete**. 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 service 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 **Automatically delete**.

- Select multiple cleanup rules and click **Turn on automatic cleanup**. The system cleans up the container or physical server logs that meet the selected cleanup rules.

  To disable the automatic cleanup feature, click **Turn off automatic cleanup**.

1.3.3.8. 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.

Solution 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 Clearance > Records**.
4. (Optional)In the upper part of the page, select a product, service, and service role, and click **Search**.

   > **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 service 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 **Cleanup details** column to view the detailed cleanup information.

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 Clearance > Plans**.
4. Click the **Containers** or **Servers** tab.
5. (Optional)In the upper part of the tab, select a product, service, and service role, and click **Search**.

> **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 service 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 result, find the execution plan for which you want to view the cleanup records and click **Clear 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 **Cleanup details** column to view the detailed cleanup information.

1.3.4. System settings

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.

1.3.4.1. User management

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

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 Manage roles.

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 delete and click **Delete** in the **Actions** column. In the message that appears, click **OK**.

  > **Note** Deleted users are displayed on the **Locked Instances** tab. To restore a deleted user, click the **Locked Instances** 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, select the user to which you want to attach a logon policy and click **Bind Logon Policy** in the lower 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**.

  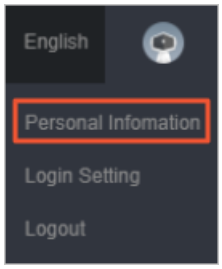

  On the **User Profile** page, you can also change the password that the current user uses to log on to the Apsara Uni-manager Operations Console. For more information about how to change the logon password, see Change the logon password.

- Configure logon settings

Move the pointer over the profile picture in the upper-right corner of the page and click **Logon Setting**. On the **Login 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.

1.3.4.2. 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**.

1.3.4.3. 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**.

1.3.4.4. Menu management

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

1.3.4.4.1. 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.

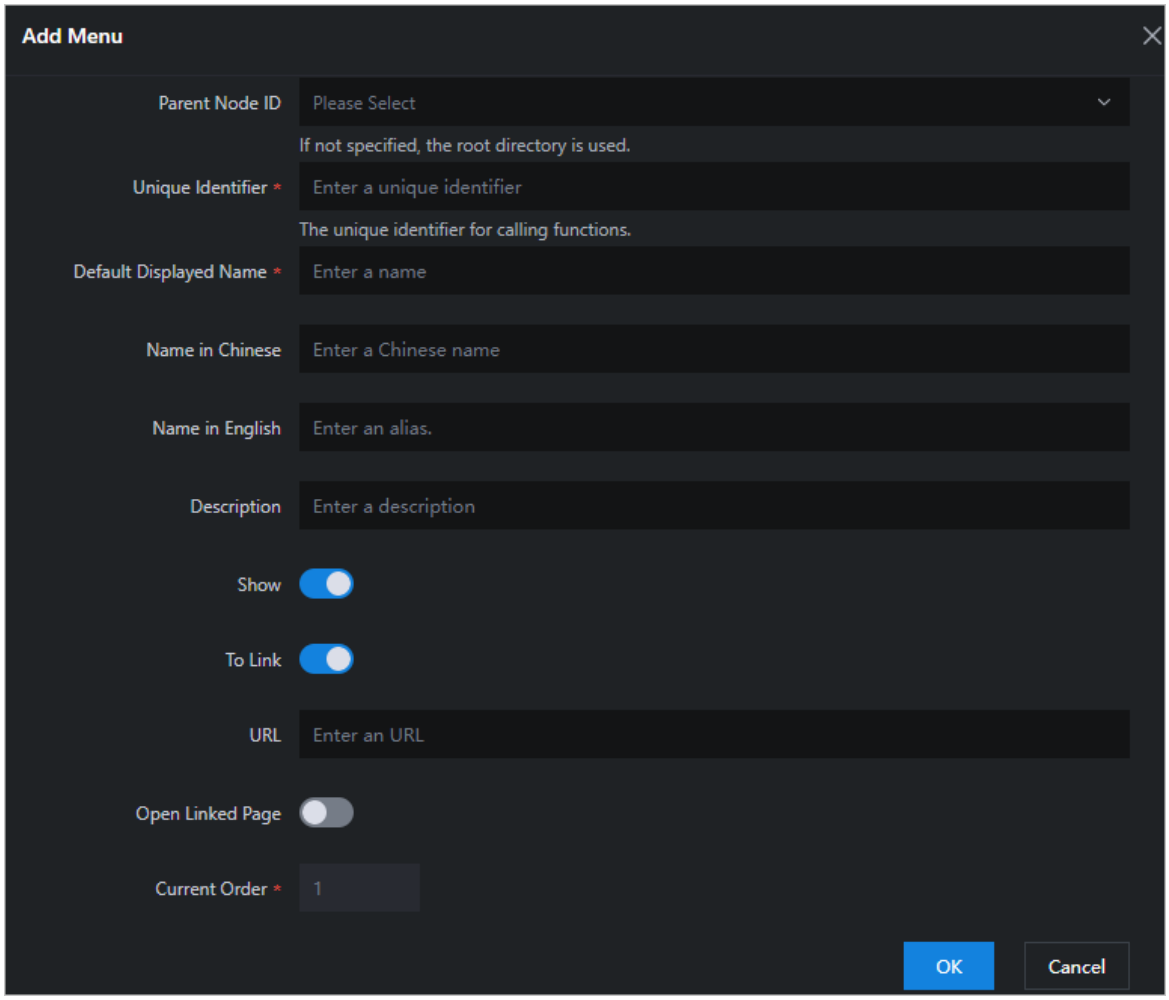

The following table describes the parameters.

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 of only letters and can be 5 to 20 characters in length.
Default Displayed Name	The default display name of the menu.
Name in Chinese	The menu name in Chinese. In the Chinese language environment, if the Chinese name of the menu is specified, the default display name of the menu is the specified Chinese name.
Name in English	The menu name in English. In the English language environment, if the English name of the menu is specified, the default display name of the menu is the specified English name.
Description	The description of the menu.
Show	Specifies whether to show the menu after it is added. You can turn on or off **Show**. By default, Show is turned on.

Parameter	Description
To Link	Specifies whether to go to another page when you click the menu. 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 menu. ◦ 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 menu. You can turn on or off **Open Linked Page**. By default, Open Linked Page is turned 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.

1.3.4.4.2. 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.

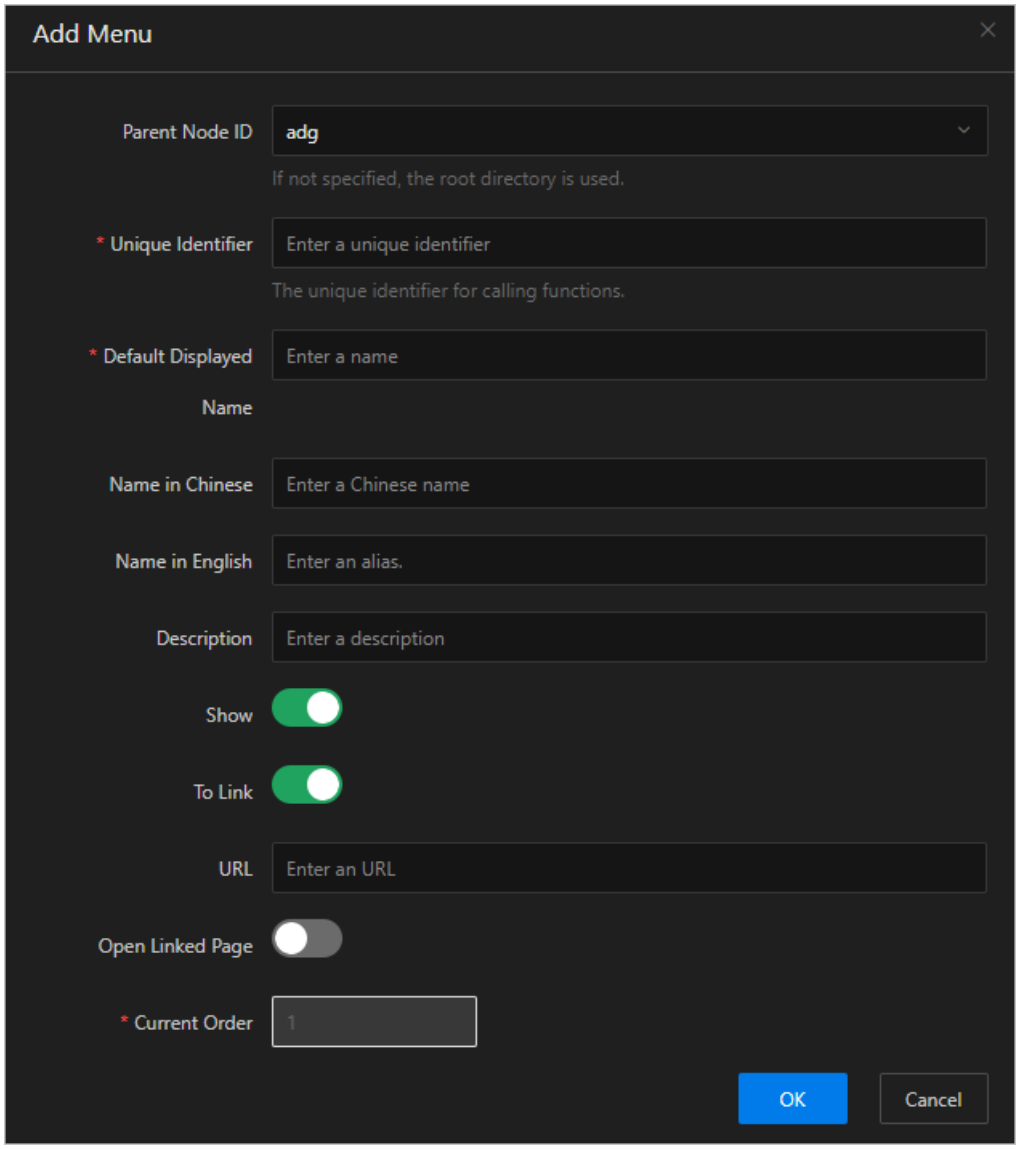

The following table describes the parameters.

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 Displayed 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.

Parameter	Description
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.
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, Open Linked Page is turned 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.

1.3.4.4.3. 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**.

1.3.4.4.4. 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.

1.3.4.4.5. 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 delete and click **Delete** in the

Actions column.

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

1.3.4.5. 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 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.

1.3.4.6. 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**, **Department Leader**, and **Department Role**, and click **OK**. Then, you can view the created department in the left catalog tree.

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

   - Modify a department

     In the left catalog tree, select the department that you want to modify and click **Modify Department**. In the **Modify Department** dialog box, set **Department Name**, **Department Leader**, and **Department 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**. In the message that appears, click **OK**.

1.3.4.7. 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.

1.3.4.8. 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.

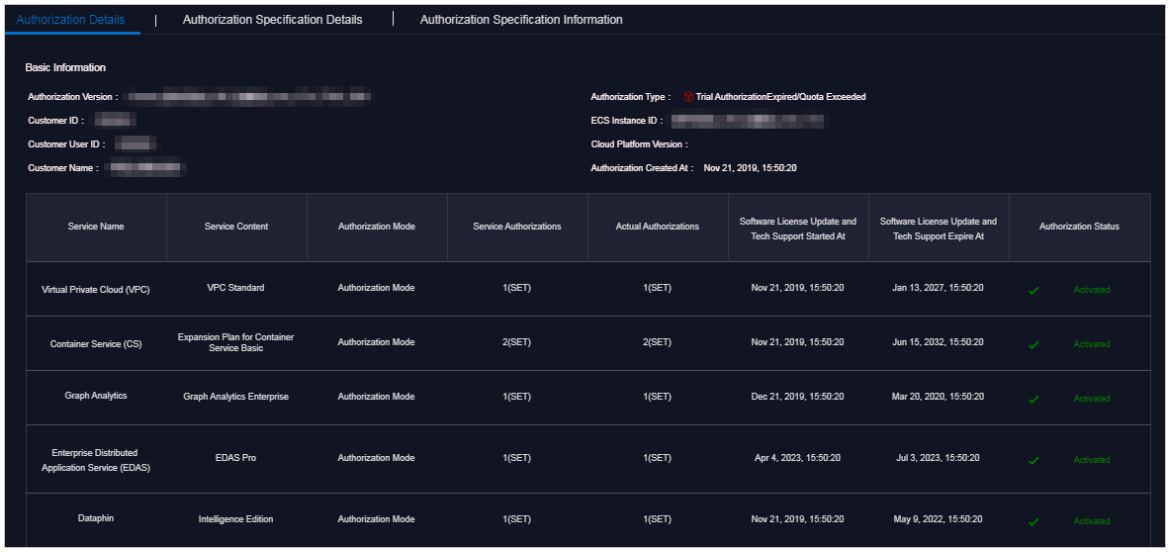

4. Perform the following operations to view the authorization information.

   > **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
Authorization Version	You can use the BP number in the version to associate a project or contract. Take note of the following items: ■ **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 statuses are available: ■ **Not Activated** ■ **Expire Soon** ■ **Activated** ■ **Expired** ■ **Expired/Quota Exceeded**
Customer information	The information about the customer, including the customer name, customer ID, and customer user ID.
ECS Instance ID	The ECS instance ID in the deployment planner of the field environment.
Cloud Platform Version	The Apsara Stack version of the current cloud platform.
Authorization Created At	The start time of the authorization.
Licensing Details of Apsara Stack Products (IDC Level)	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 Specification 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.
Specification Name	The specification name of an authorized service.
Specifications	The total number of current authorizations of a specification for a service.
Specification Quota	The authorization quota of a specification for a service.
Specification Status	The current authorization status of a specification for a service.

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

  Set **Licensing Specification Level**, **Region ID** or **IDC 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 points as well as the average number of specifications within the specified time range.

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

1.3.4.9. Multi-cloud management

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

1.3.4.9.1. 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 another data center and click **OK**.

Parameter	Description
Name	The name of another data center.
Console link	The console link of another data center. Make sure that the console link is valid. Otherwise, an error message is returned.

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.

1.3.4.9.2. Modify the name of a data center

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

Procedure

1. 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. (Optional)In the **Name** search box, enter the data center name that you want to modify and click **Search**.
5. Find the data center that you want to modify and click **Modify** in the **Actions** column.
6. In the dialog box that appears, modify the name of the data center and click **OK**.

1.3.5. Personal Settings

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

1.3.5.1. 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**.

1.3.5.2. 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 may be 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 create 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**.

1.4. Resources

1.4.1. Products

1.4.1.1. Product overview

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

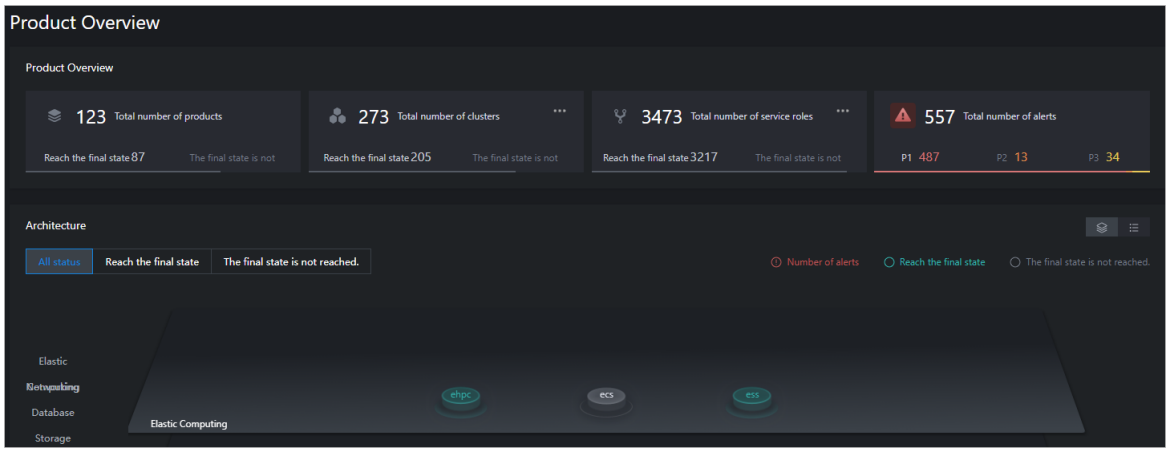

- In the **Product Overview** section, view the total numbers of products, clusters, service roles, and alerts as well as the final-state status of these products, clusters, and service roles.
  - Click the icon next to **Total number of clusters** and view the information about the clusters on the **Cluster list** page.
  - Click the icon next to **Total number of service roles** and view the information about the service roles on the **Service role** page.

- In the **Architecture** section, view the product information.
  - Click **All status**, **Reach the final state**, or **The final state is not reached** 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, service roles, machines, and alerts in a list form.

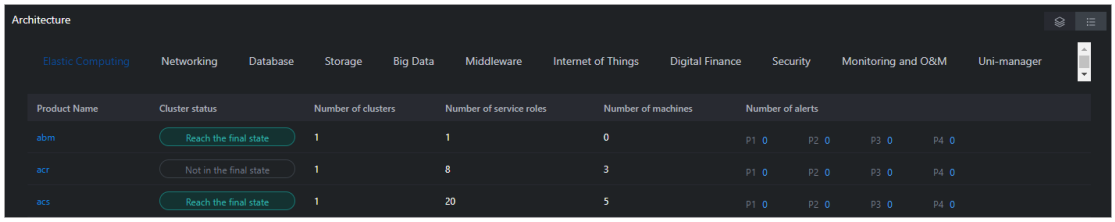

  - Click a product type in the upper product type list to view the information about the products of that type.
  - Click a product name in the Product Name column to view the details of the product on the **Product Details** page.

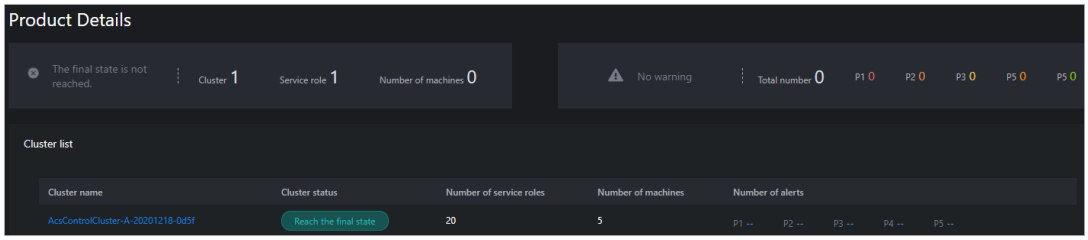

    a. Click a cluster name in the Cluster Name column to view the details of the cluster on the **Cluster Details** page.

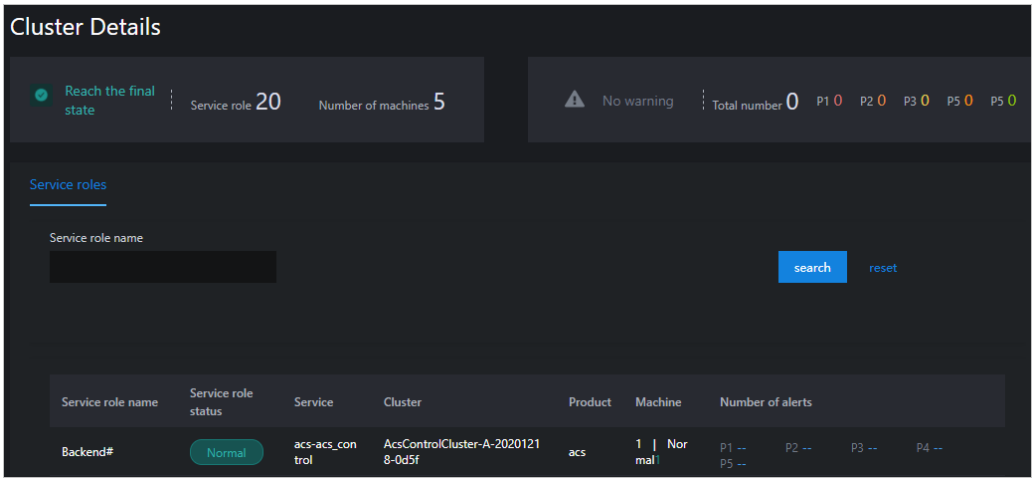

    b. Enter a service role name in the **Service Role Name** search box and click **Search** to view the details about the service role.

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

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

1.4.1.2. Clusters

You can view the status and alerts of all the deployed clusters and their service 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.

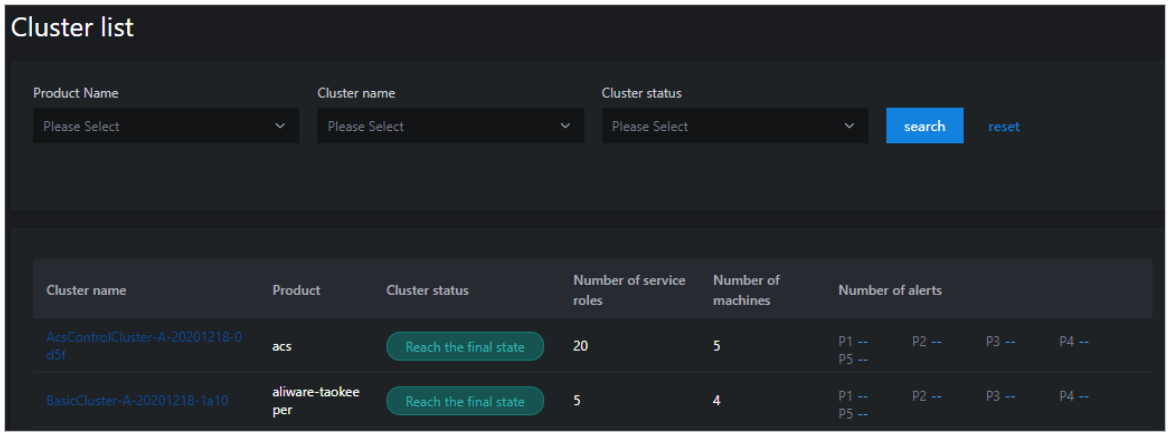

5. (Optional) Click **Reset** to clear the filter conditions.
6. Click a cluster name in the Cluster Name column to view the details of the cluster on the **Cluster Details** page.

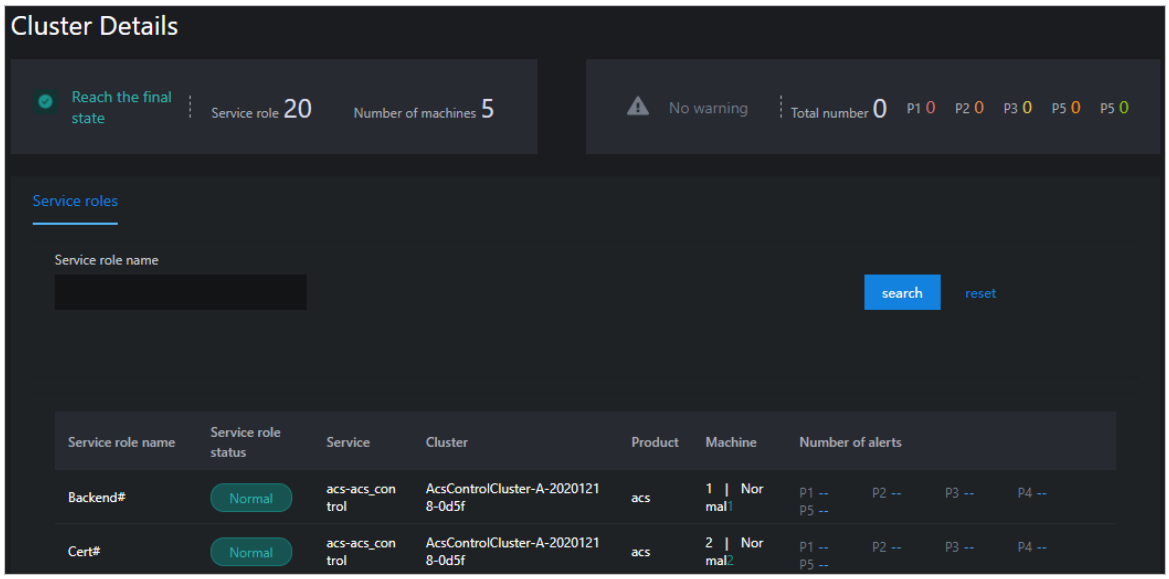

7. Enter a service role name in the **Service Role Name** search box and click **Search** to view the details about the service role.
8. (Optional) Click **Reset** to clear the filter conditions.

1.4.1.3. Service roles

On the Service role page, you can view the service 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 **Service Roles**.
4. On the **Service role** 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 **Service role status** drop-down list. Then, click **Search** and view the search results displayed below.

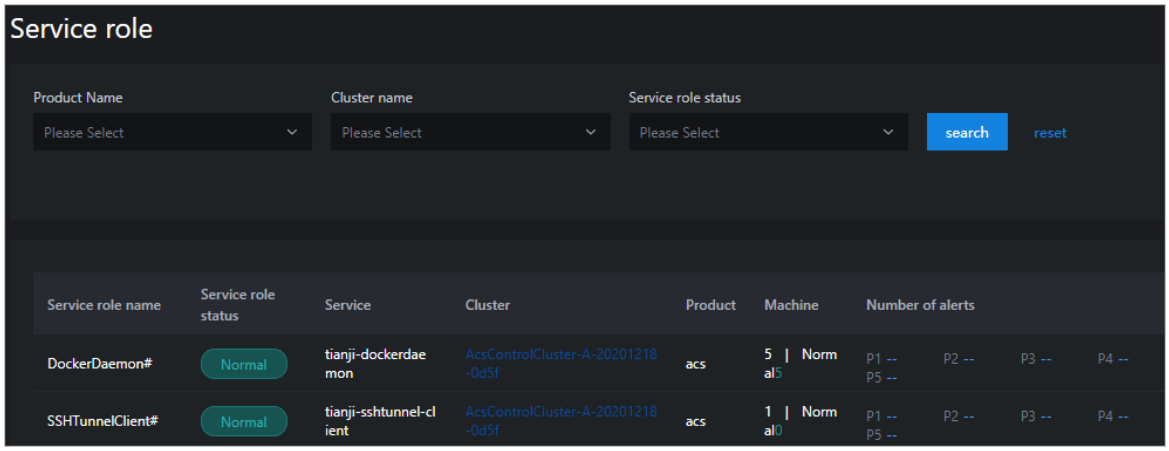

5. (Optional) Click **Reset** to clear the filter conditions.
6. Click a cluster name in the Cluster column to view the details of the cluster on the **Cluster Details** page.

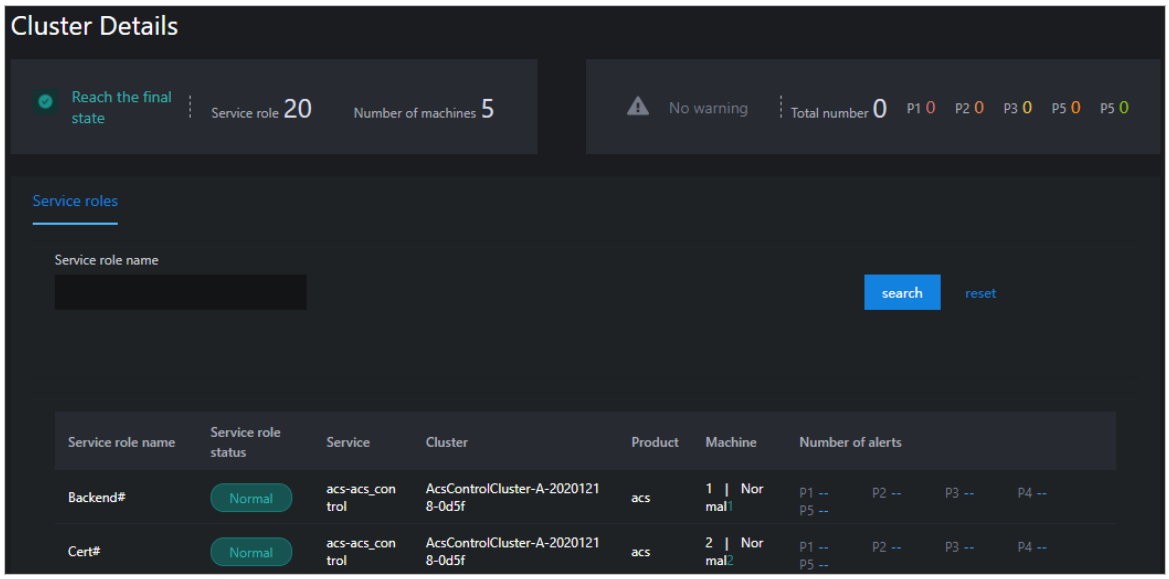

7. Enter a service role name in the **Service Role Name** search box and click **Search** to view the details about the service role.
8. (Optional) Click **Reset** to clear the filter conditions.

1.4.2. Data centers

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

1.4.2.1. View the 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.

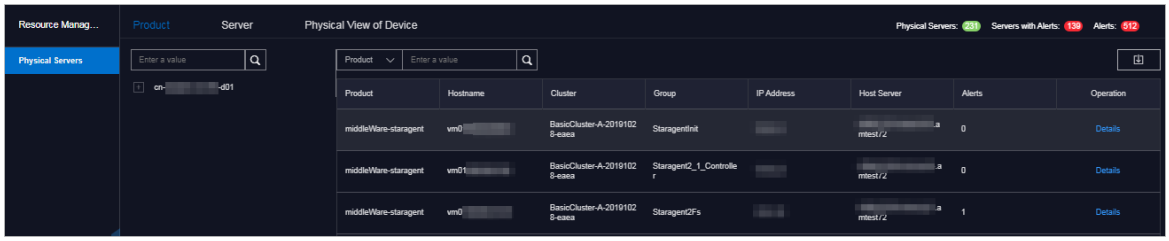

Product tab

1. On the Product tab, perform the following operations to view the 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.

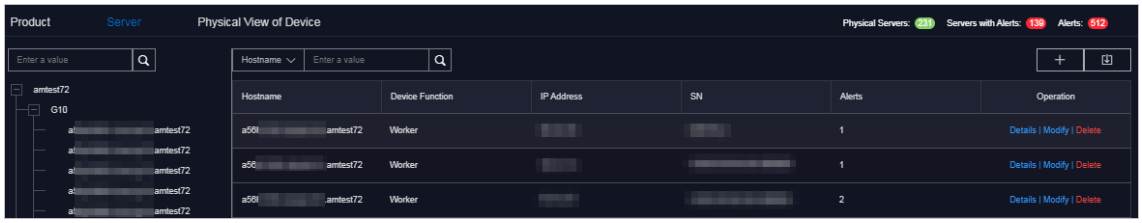

3. To view the details of a physical server, enter the hostname, IP address, device role, or serial number (SN) in the right-side search box and press the Enter key.
4. Find the physical server that you want to view and click **Details** in the **Operation** column. On the **Physical Server Details** page, view the basic information, monitoring details, and alert information of the physical server.

   You can switch 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.

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.

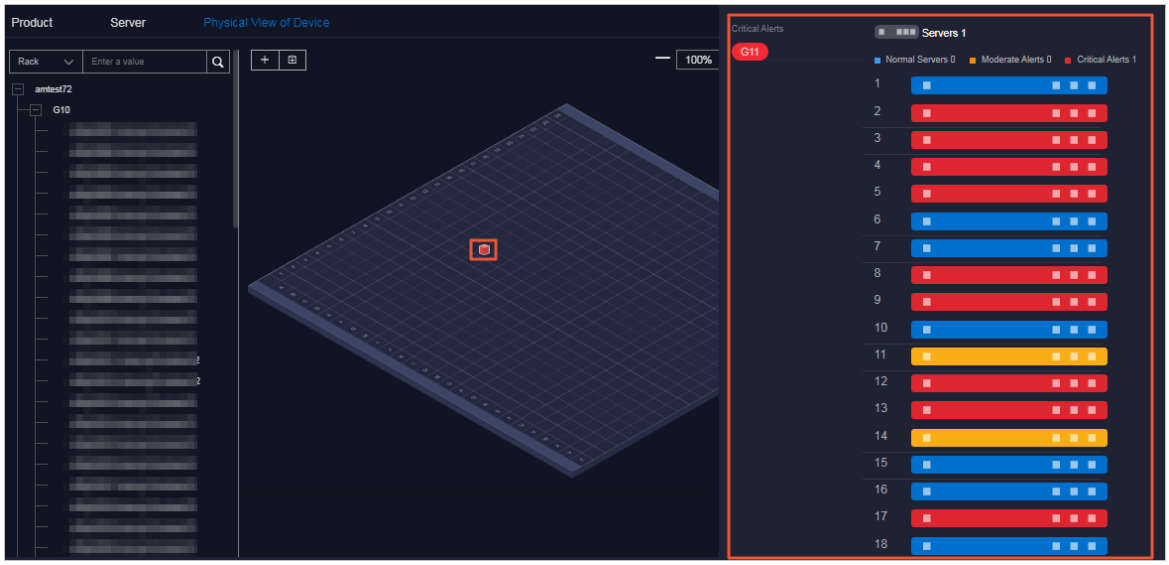

3. To view the details of a physical server, perform the following operations:
   i. Find the physical server that you want to view in the left-side hierarchy tree or right-side rack graph of the tab.
   ii. In the rack details panel that appears, click the color block corresponding to a server to view the basic information of the server.
   iii. Click **Details** in the **Operation** row of the basic information.

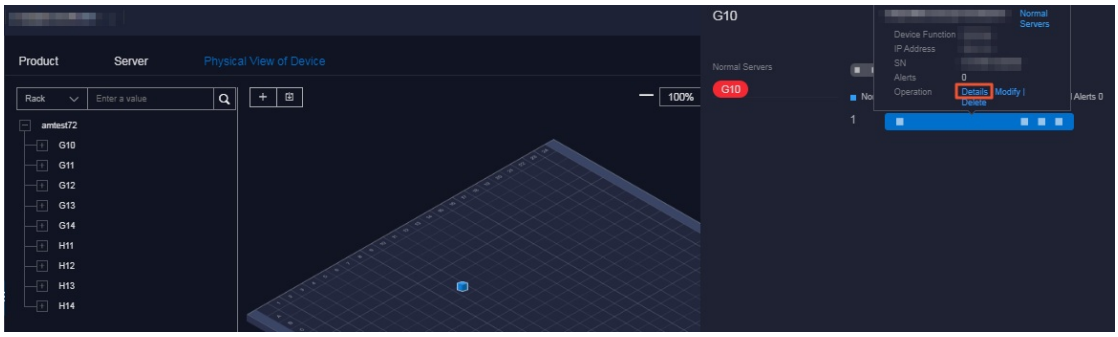

iv. On the **Physical Server Details** page, view the basic information, monitoring details, and alert information of the physical server.

You can switch 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.

1.4.2.2. Add physical servers

Operations personnel can add the existing physical servers in an environment to the Apsara Uni-manager Operations Console.

Procedure

1. Log on to the Apsara Uni-manager Operations Console.
2. In the top navigation bar, choose **Resources > Data Centers**.
3. Click the **Server** or **Physical View of Device** tab.
4. In the upper-right corner of the **Server** tab or the upper-left corner of the **Physical View of Device** tab, click the icon.
5. In the **Add Physical Server** pane, configure the parameters.

   The following table describes the parameters.

Parameter	Description
Zone	The zone where the physical server that you want to add is located.
Data Center	The data center where the physical server is located.
Rack	The rack where the physical server is located.
Room	The room where the physical server is located.
Physical Server Name	The name of the physical server.
Memory	The memory size of the target physical server.

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

6. Click **OK**.

1.4.2.3. Modify a physical server

You can modify the physical server information in the Apsara Uni-manager Operations Console when the information is changed in the Apsara Stack environment.

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**.

Server tab

1. Click the **Server** tab.
2. (Optional)In the right-side search box, search for the physical server that you want to modify by hostname, IP address, device role, or serial number (SN).
3. Find the physical server that you want to modify and click **Modify** in the **Operation** column.

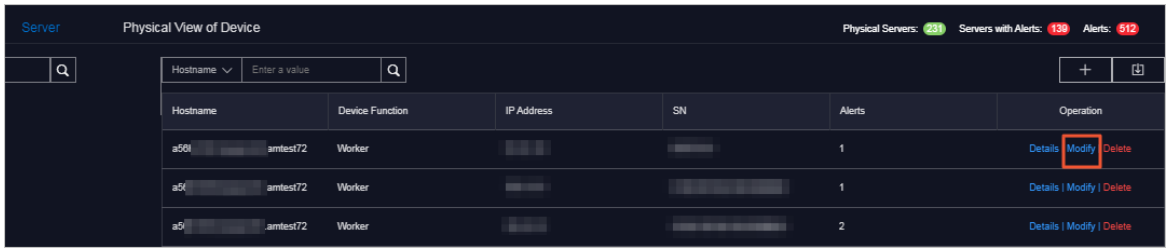

4. In the **Modify Physical Server** panel, modify the physical server information.

   You can modify the following physical server information: zone, data center, rack, room, physical server name, memory size, disk size, number of CPU cores, rack group, server type, server role, serial number, operating system template, and IP address.

5. Click **OK**.

Physical View of Device tab

1. Click the **Physical View of Device** tab.
2. Expand the left-side hierarchy tree by selecting an IDC and a rack in sequence to find the physical server that you want to modify.

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

3. In the rack details panel, click the color block corresponding to a server to view the basic information of the server.
4. Click **Modify** in the **Operation** row of the basic information.

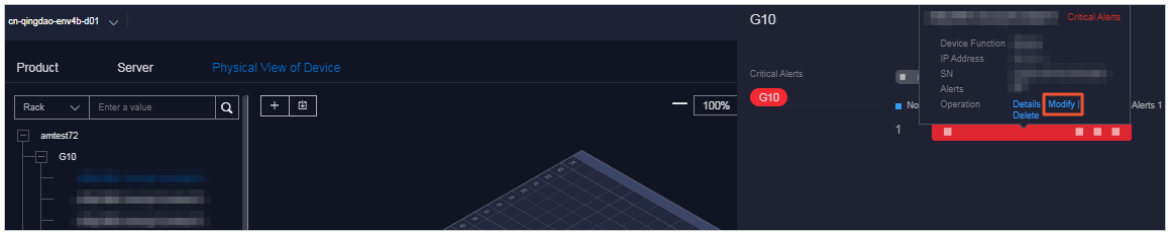

5. In the **Modify Physical Server** panel, modify the physical server information.

   You can modify the following physical server information: zone, data center, rack, room, physical server name, memory size, disk size, number of CPU cores, rack group, server type, server role, serial number, operating system template, and IP address.

6. Click **OK**.

1.4.2.4. Export the 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.

1.4.2.5. Delete a physical server

This topic describes how to delete physical servers that no longer need to be monitored.

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**.

Server tab

1. Click the **Server** tab.
2. (Optional)In the right-side search box, search for the physical server that you want to delete by hostname, IP address, device role, or serial number (SN).
3. Find the physical server and click **Delete** in the **Operation** column.
4. In the message that appears, click **OK**.

Physical View of Device tab

1. Click the **Physical View of Device** tab.
2. Expand the left-side hierarchy tree by selecting an IDC and a rack in sequence to find the physical server that you want to delete.

   > **Note** You can also expand the left-side hierarchy tree by selecting an IDC and a rack in sequence to find the physical server that you want to delete.

3. In the rack details panel that appears, click the color block corresponding to a server to view the basic information of the server.
4. Click **Delete** in the **Operation** row of the basic information.

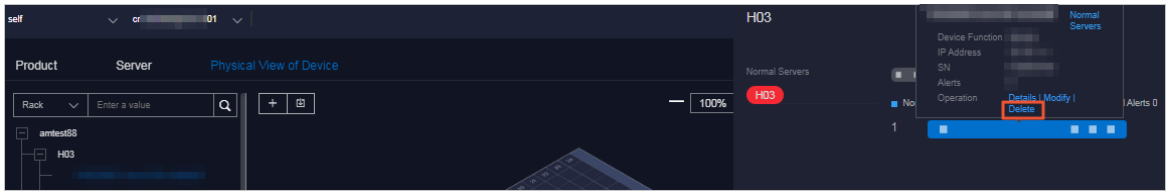

5. In the message that appears, click **OK**.

1.5. Alerts

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.

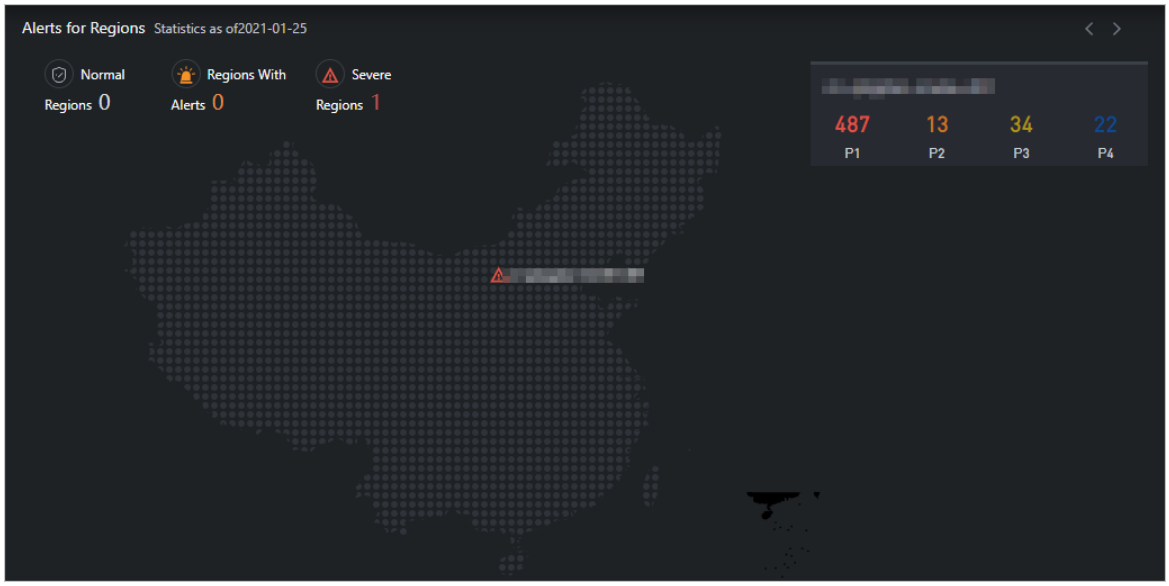

> **Note** P1, P2, P3, and P4 have the following meanings:
>
> - P1: urgent alerts
> - P2: major alerts
> - P3: minor alerts
> - P4: reminder alerts

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

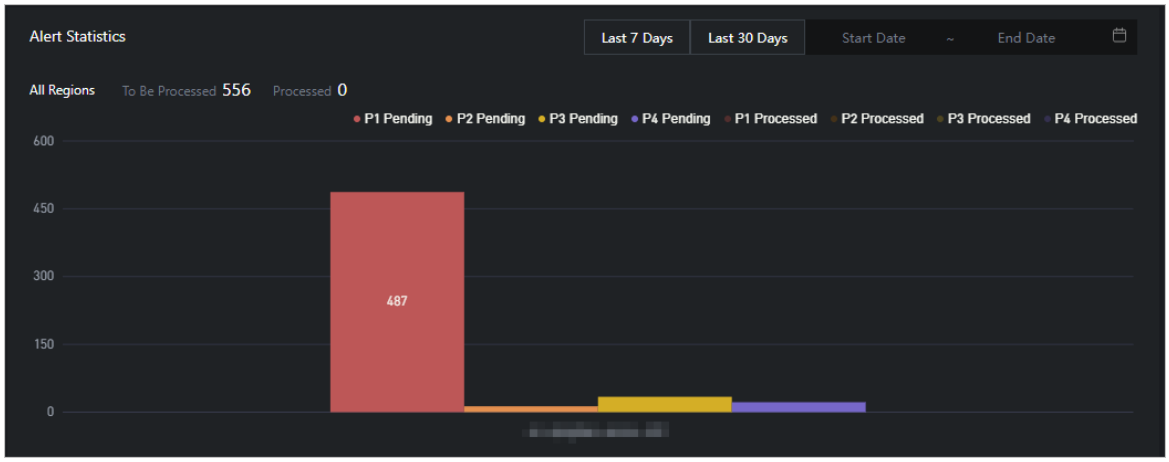

- Click **Last 7 Days**, **Last 30 Days**, or specify **Start Date** and **End Date** to view the statistical data of alert handling within the period.
- Move the pointer over a column chart, and the corresponding statistical data is displayed.
- Select the required options from the **Region**, **Priority Level**, and **Status** drop-down lists, and click **search** to view the alert details.

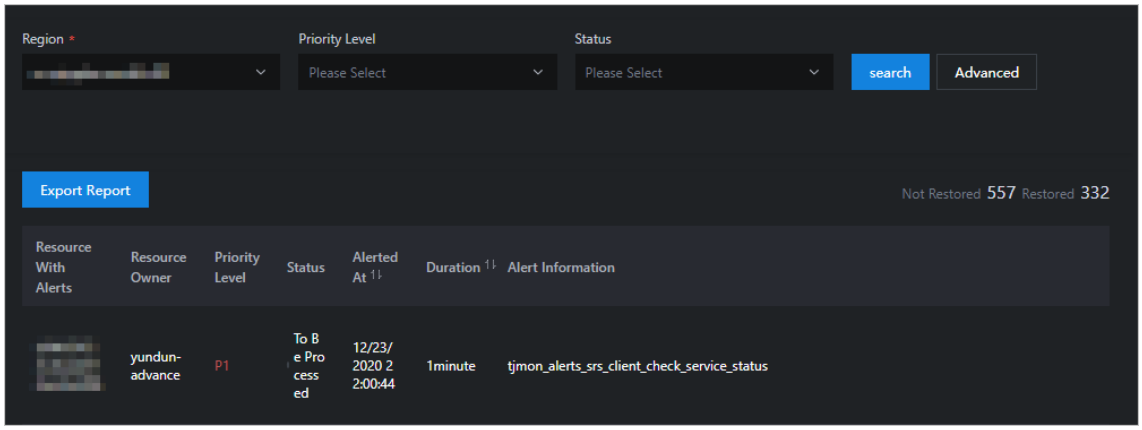

- Specify more filter conditions to query alerts.
  - Click **Advanced**. Select the required options from the **Region**, **Priority Level**, **Status**, and **Resource** drop-down lists, and click **Search** to view the alert details.

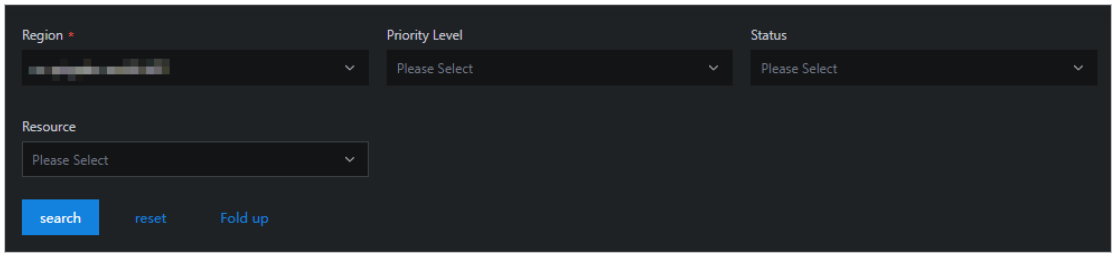

  - (Optional) Click **Reset** to clear the filter conditions.
  - (Optional) Click **Fold up** to hide the **Resource** option.

5. Click **Export Report** to download the alert details to your computer.

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.

     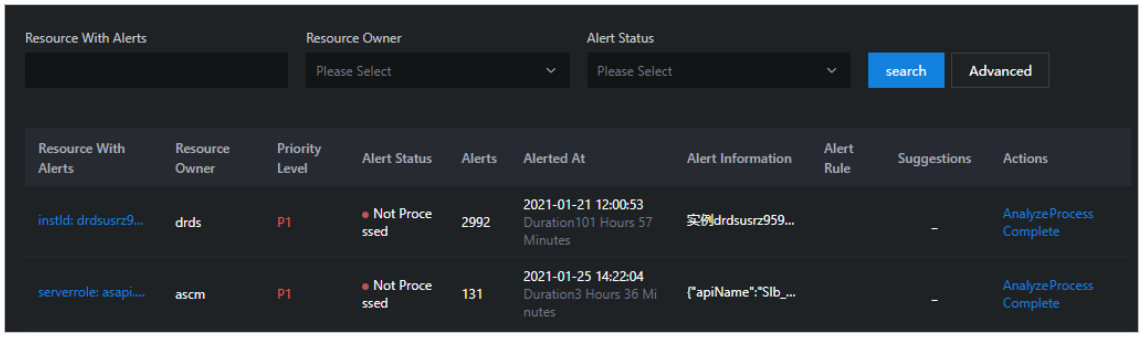

   - Specify more filter conditions to view the alerts.
     a. Click **Advanced**. Enter an alert resource ID in the **Resource With Alerts** search box, select the required options from the **Resource Owner**, **Alert Status**, and **Alert Level** drop-down lists, and then select a start date and an end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

     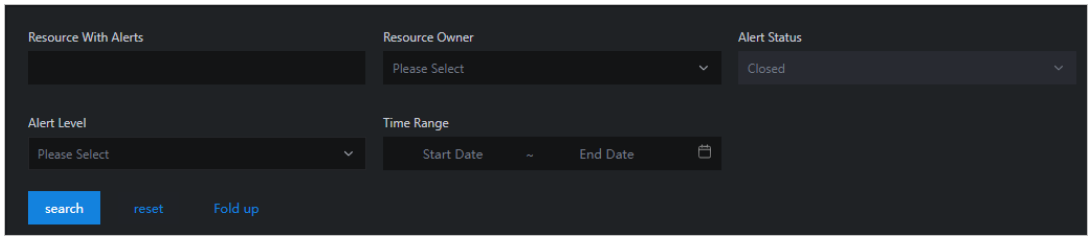

     b. (Optional) Click **reset** to clear the filter conditions.
     c. (Optional) Click **Fold up** to hide the **Alert Level** and **Time Range** options.
5. Analyze the alert information.

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.

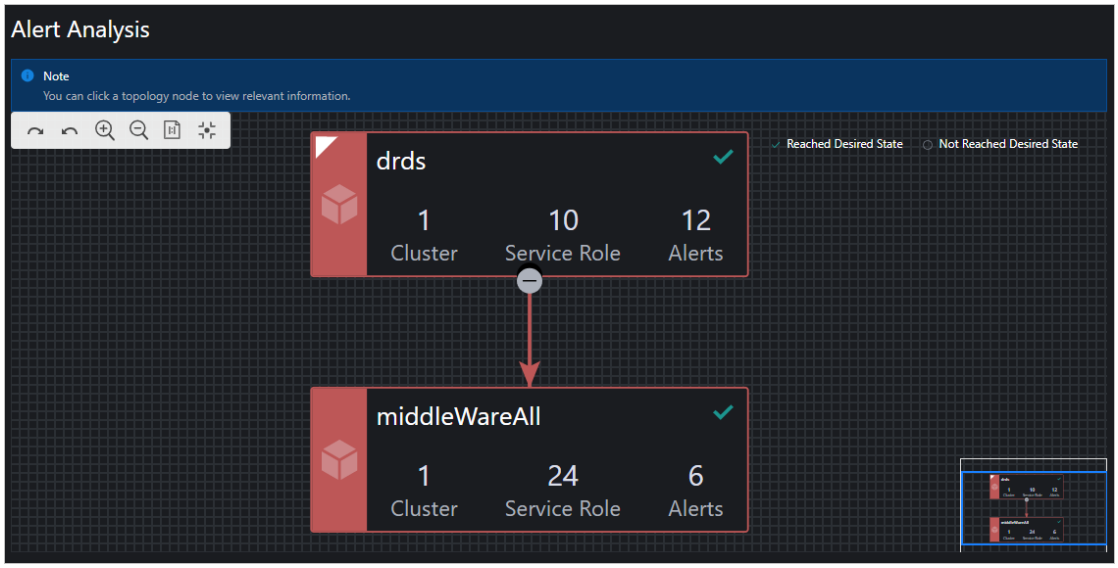

> **Note** Icons:
>
> - : indicates the server.
> - : indicates the cluster
> - : Indicates the service.
> - : indicates the service role.

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.

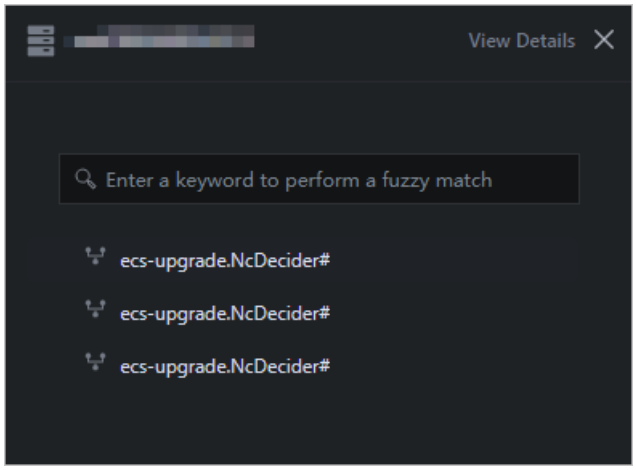

iii. Click **View Details** in the upper-right corner to view the details of the node.

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

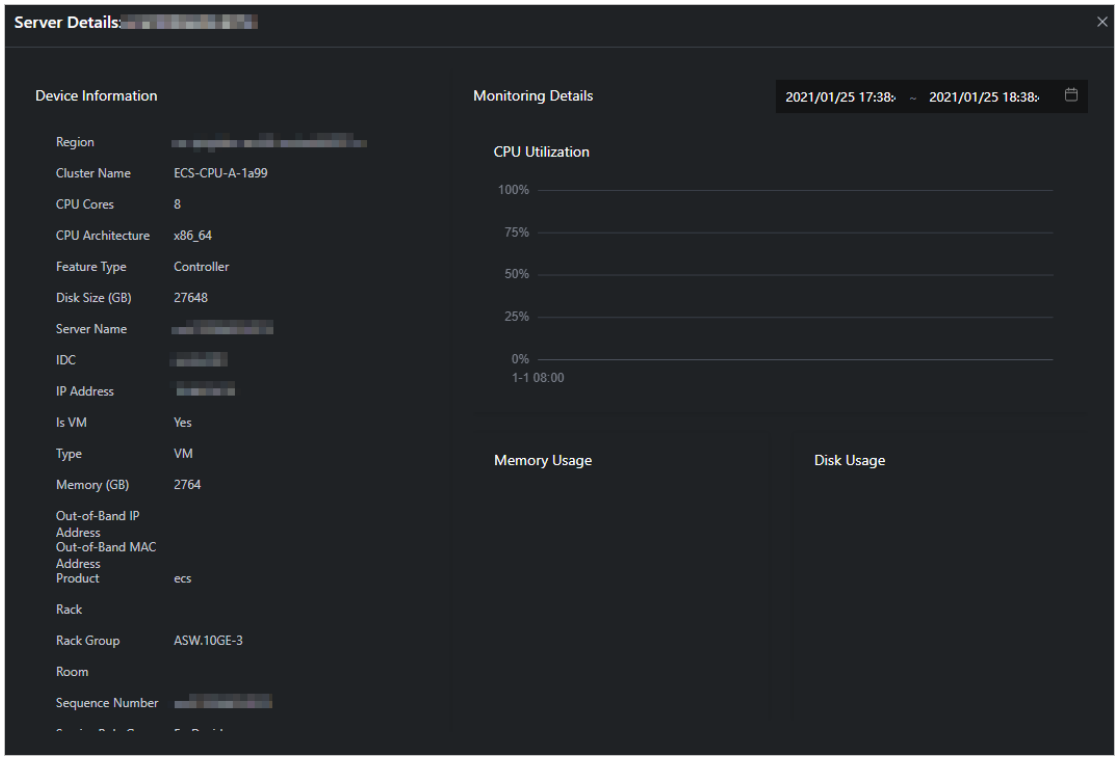

iv. In the node details panel, click an element related to this node to view the details of this element.

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

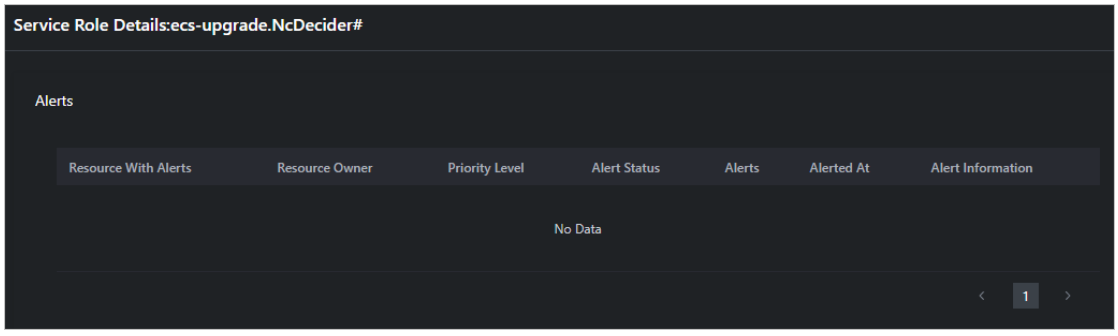

6. Click **Treatment** in the Actions column corresponding to the alert information. In the message that appears, click **OK** to mark the alert as being processed.
7. Click **Complete** in the Actions column corresponding to the alert information. In the message that appears, click **OK** to mark the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

1.5.3. Alert settings

1.5.3.1. Policy management

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

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**.

1.5.3.1.2. 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**.

1.5.3.1.3. 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.

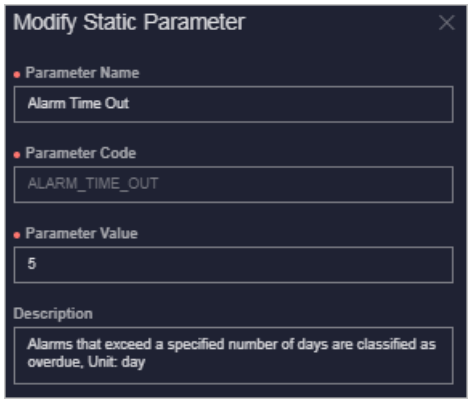

Parameter	Description
Parameter Name	Enter a parameter name related to the configuration.

Parameter	Description
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**.

1.5.3.2. 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.

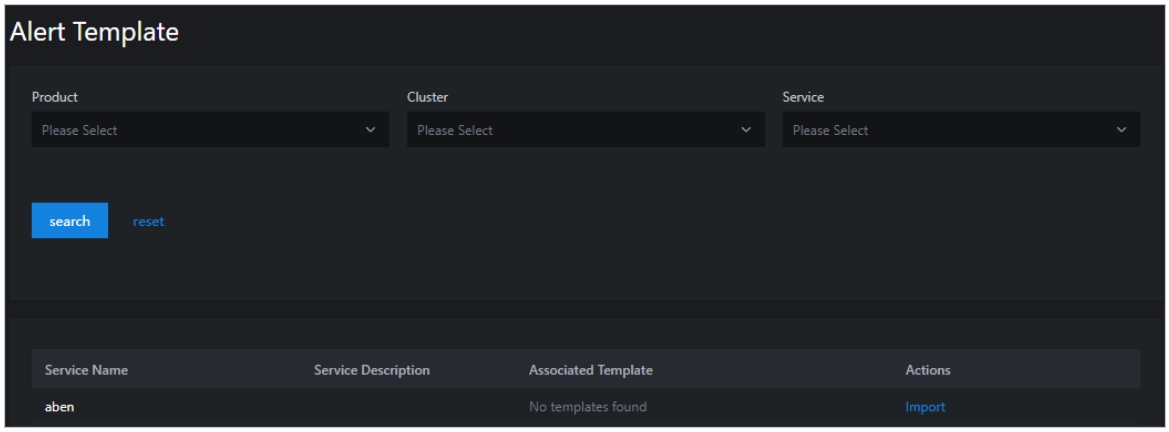

5. (Optional) Click **Reset** to clear the filter conditions.
6. Download Alert Templates.

> **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**.

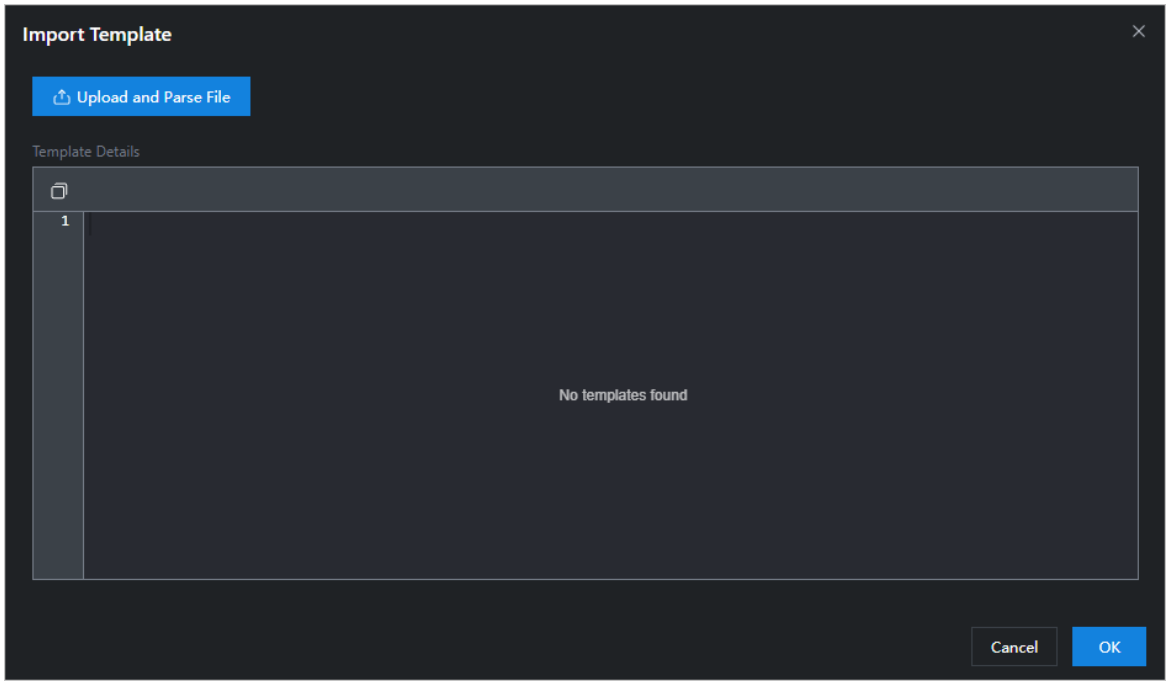

1.5.3.3. Notification management

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**.

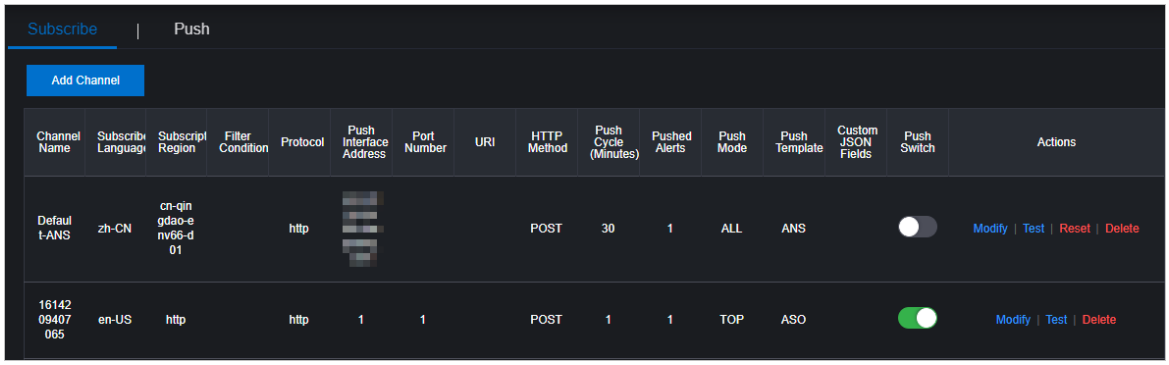

4. On the **Subscribe** tab, click **Add Channel**.
5. In the **Add Subscription** panel, configure the parameters described in the following table.

Parameter	Description
Channel Name	The name of the subscription channel.

Parameter	Description
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.
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.

Parameter	Description
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.

1.5.3.4. Alert masking

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

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.

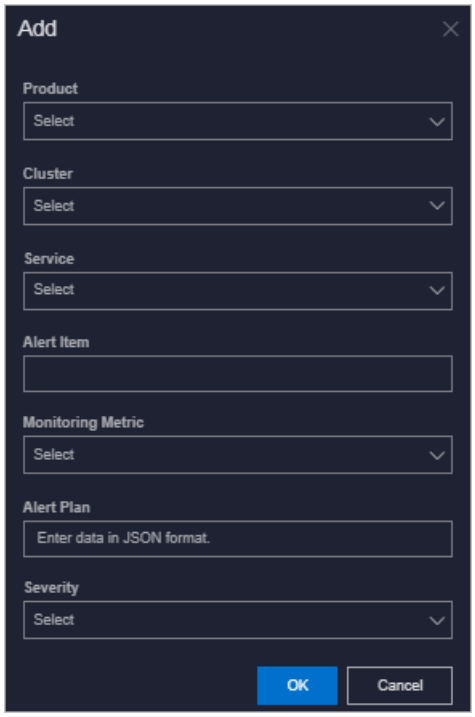

Parameter	Description
Product	Optional. The product to which the alerts to be masked belong.
Cluster	Optional. The cluster to which the alerts to be masked belong.
Service	Optional. The service to which the alerts to be masked belong.
Alert Item	Optional. The name of the alerts to be masked. **Note** It may take an extended period of time if the number of alerts is large when you configure **Alert Item**.
Monitoring Metric	Optional. The monitoring metric to which the alerts to be masked belong.
Alert Plan	Optional. Details about the alerts to be masked. Example: {"serverrole":"ecs-yaochi.ServiceTest#","machine":"vm01001******","level":"error"}

Parameter	Description
Severity	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**.

1.5.3.4.2. 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.

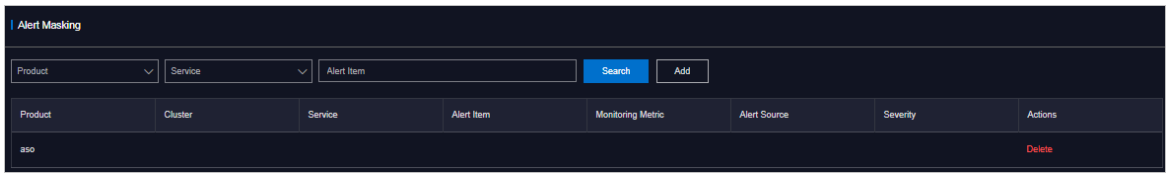

6. In the message that appears, click **OK**.

Result

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

1.6. O&M

1.6.1. 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.

1.6.1.1. 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**.

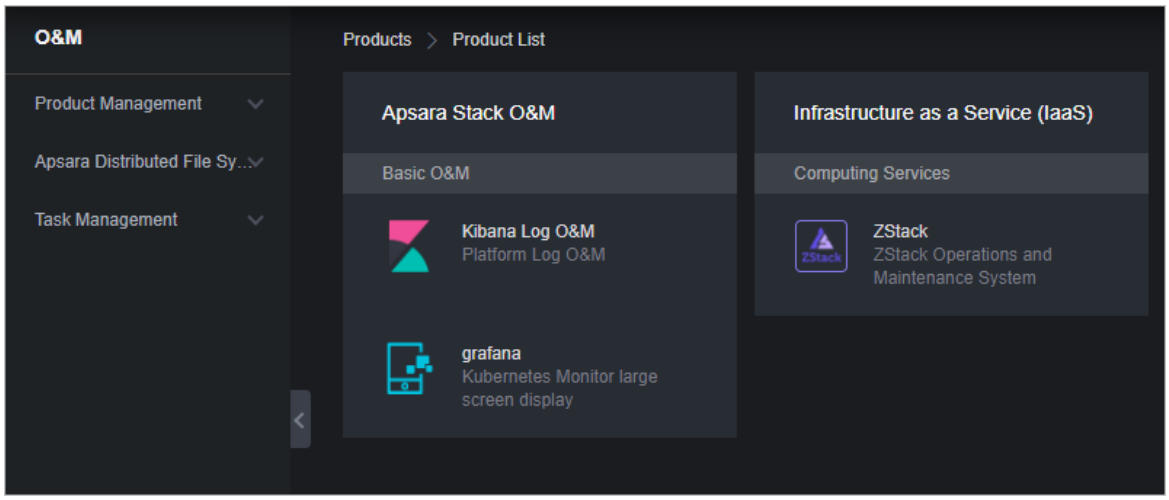

4. On the **Product List** page, you can view the O&M icons of different products and ISV icons based on your permissions.

1.6.1.2. ISV access settings

The ISV Access Settings module allows you to configure, modify, and delete the ISV access information.

1.6.1.2.1. 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.

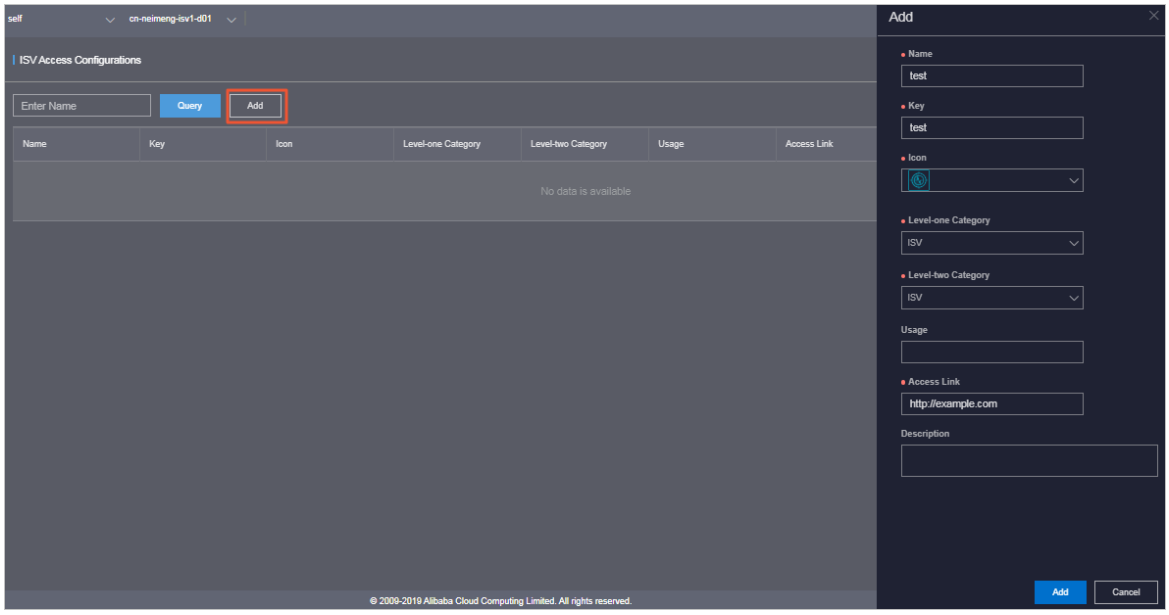

The following table describes the parameters.

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.

Parameter	Description
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.

1.6.1.2.2. 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**.

1.6.1.2.3. 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.

1.6.2. Apsara Distributed File System Management

1.6.2.1. Apsara Distributed File System

1.6.2.1.1. 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.

     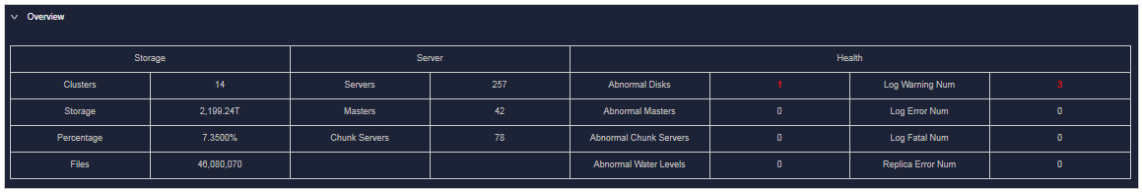

   - **Heatmap of Health**

     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.

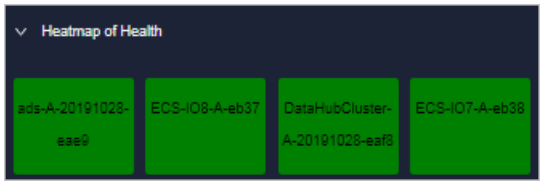

- **Data of Top 5 Services**

  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.

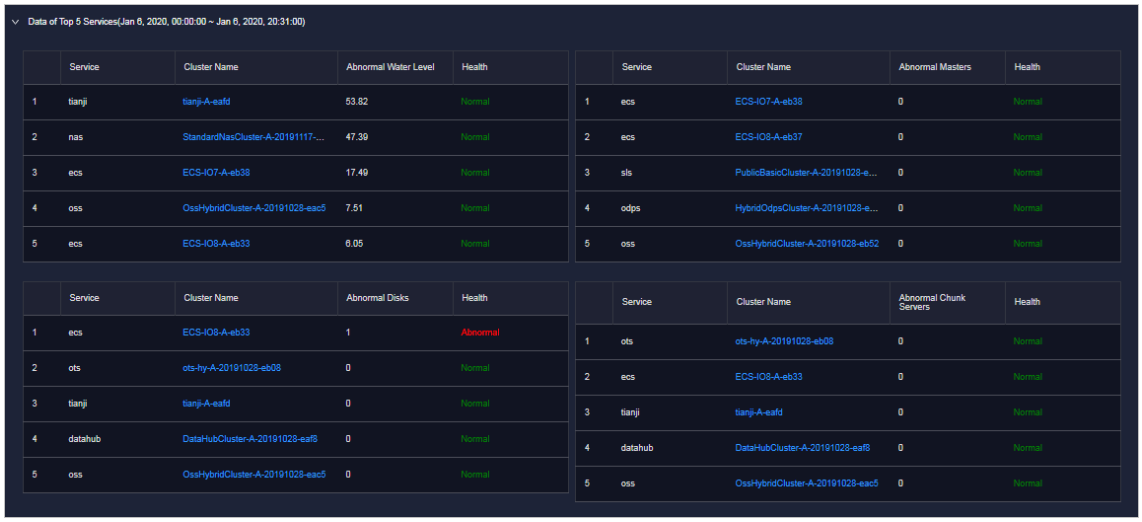

1.6.2.1.2. Clusters

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.

- **Alert Monitor**

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

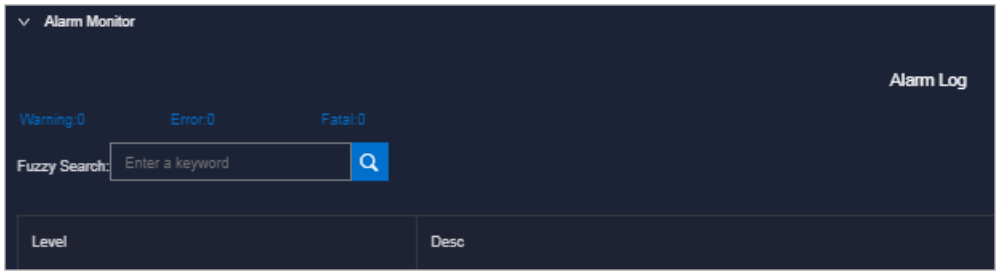

- **Replica**

  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.

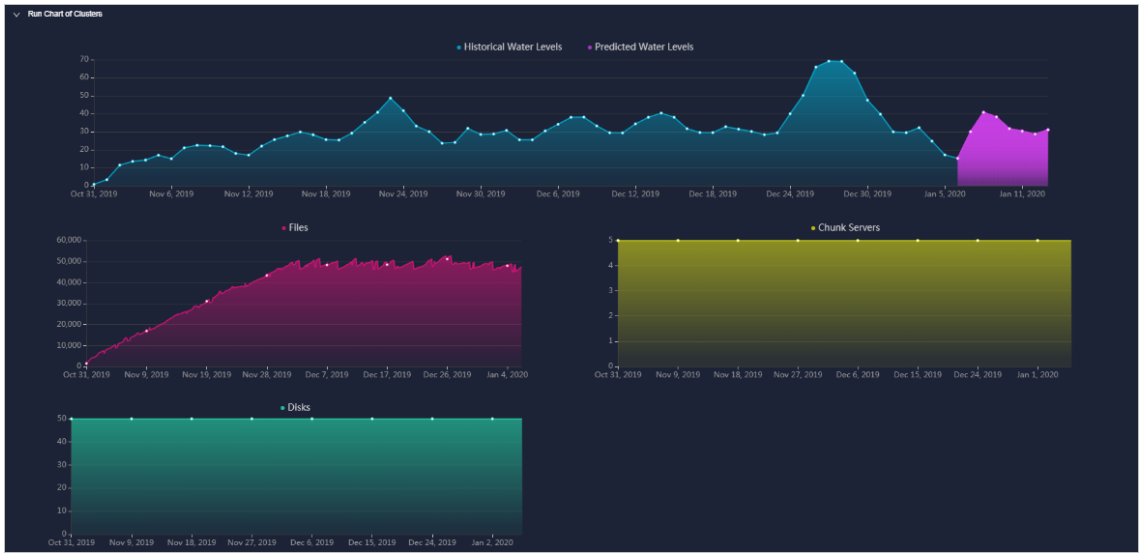

- **Rack Information**

  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.

1.6.2.1.3. 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.

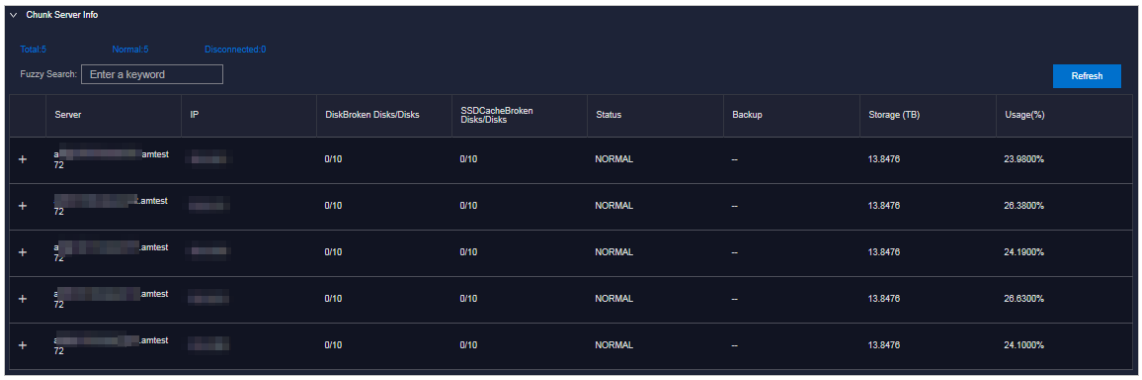

1.6.2.1.4. Operations and maintenance

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.

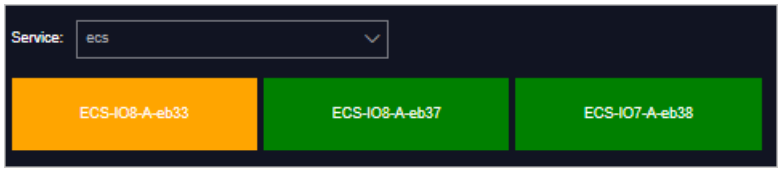

5. Move the pointer over a cluster name to view the service name, server name, and IP address of the cluster.

1.6.2.1.5. 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.

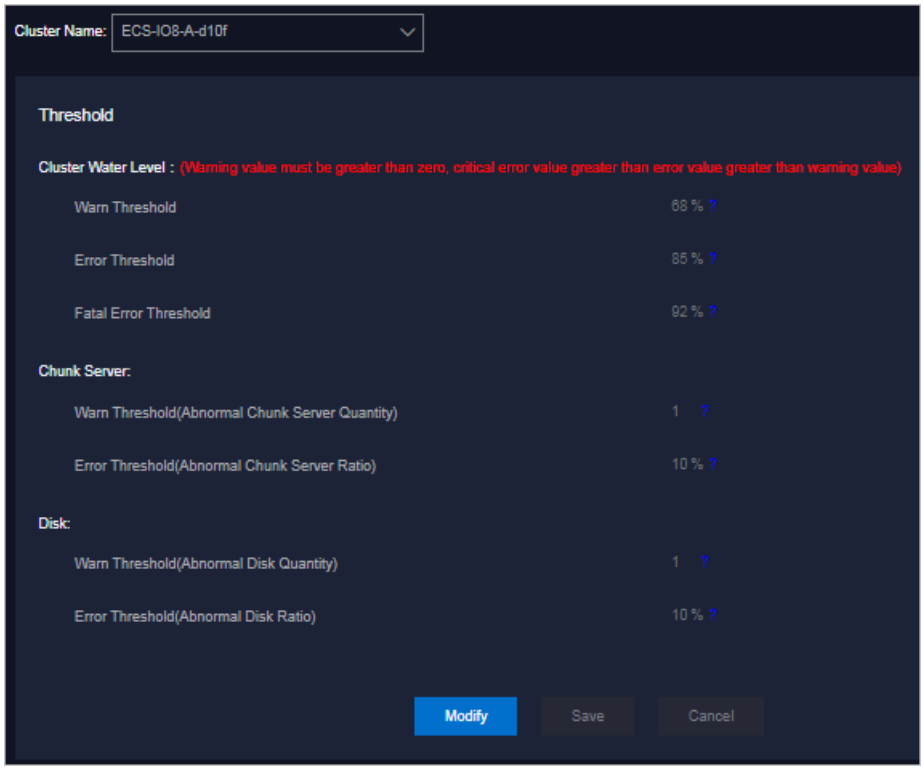

The following table describes the parameters.

Parameter	Description

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.
	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%.

Parameter		Description
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**.

1.6.2.2. EBS

1.6.2.2.1. 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**, **Abnormal Block GcWorker**, 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.

1.6.2.2.2. Block master nodes 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.

     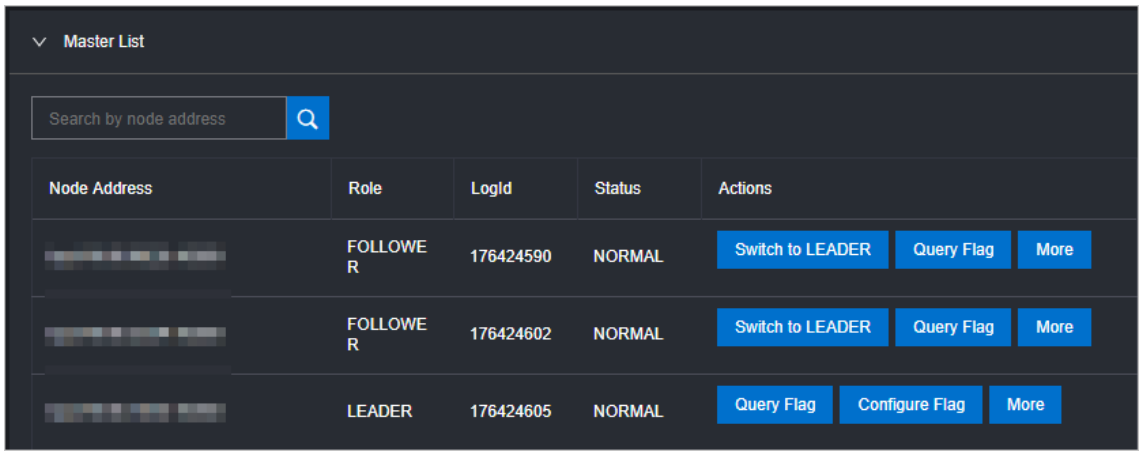

   - Switch to LEADER

     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 customized 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.

1.6.2.2.3. 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 customized flag value.

Parameter	Description
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 has been disabled.
UPGRADE	The node has been upgraded.
RECOVERY	The node has been 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.

1.6.2.2.4. SnapShotServer

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

Procedure

1. 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.

     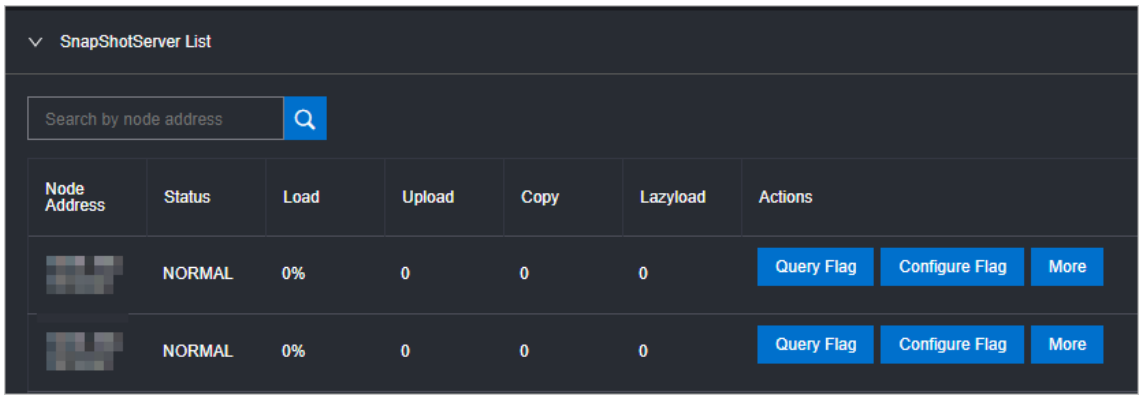

   - Query a flag

     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 customized falg 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	Indicates that the node is running normally.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	Indicates that the node is disabled.

- Query the version information

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

1.6.2.2.5. 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.

     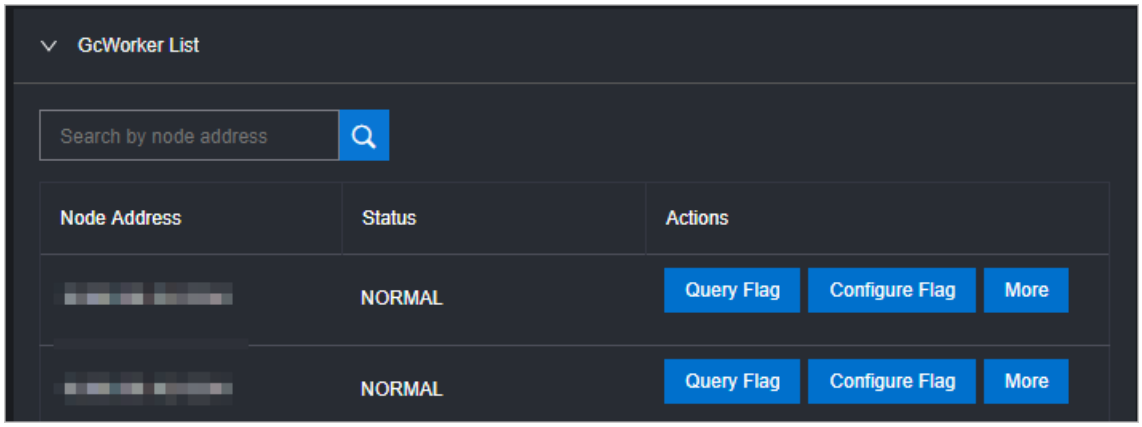

   - Query a flag

     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_blockgcworke r_flag.json* file.
flag_value	The customized 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 > Configure 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	Indicates that the node is running normally.
DISCONNECTED	Indicates that the node is disconnected.
OFFLOADING	Indicates that the node is being disabled.
OFFLOADED	Indicates that the node 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.

1.6.2.2.6. Device operations

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

Procedure

1. 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 > 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. Perform the following operations:
   - View the device list

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

     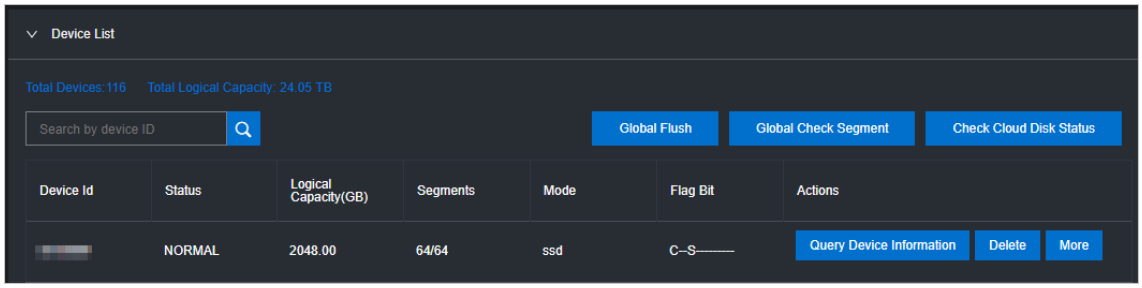

   - Check global segments

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

   - Check disk status

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

   - Query device information

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

   - Delete a device

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

     After you delete the device, its status becomes **DELETING**, and the device is unavailable. You are not allowed to perform operations such as enabling the device or modifying the configurations.

   - 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 its normal state.

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

   - Enable a device

     In the device list, find a device and choose **More > Enable** in the **Actions** column. In the dialog box that appears, configure the required parameters and click **Submit**.

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

     The following table describes the parameters for enabling a device.

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

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

The following table describes the parameters for disabling a device.

Parameter	Description
client_ip	Specifies the client IP address of the disk to be disabled. If the client IP address is not specified, the IP address of the local server is used.
token	Specifies the token used to disable the device. The token is configured when the device is enabled. You can query the token by running the **dev -query** command on a server in the EBS cluster.
open_ver	Specifies the current openversion of the device if the client IP address is not specified. If a client IP address is specified, you do not need to specify the openversion. You can query the openversion by running the **dev -query** command on a server in the EBS cluster.

- Flush

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

The following table describes the parameters.

Parameter	Description
segment	Select the segment to be flushed. If you do not specify this parameter, all segments are flushed.
ifnsw	Specifies whether to flush the index file. Valid values: ■ **0**: specifies to flush the index file. ■ **1**: specifies not to flush the index file.
dfnsw	Specifies whether to flush the data files. Valid values: ■ **0**: specifies to flush the data files. ■ **1**: specifies not to flush the data files.

- Global flush

  You can perform the 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**. Then, the transaction logs of all the disks or segments in the current cluster are flushed.

- Query the configuration status

  In the device list, find a device and choose **More > Query Configuration Status** in the **Actions** column corresponding to the device. 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 of the queried device information.

- Modify device configurations

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

  In the device list, find a device and choose **More > Modify Device Configurations** in the **Actions** column corresponding to the device. In the dialog box that appears, modify the 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	Specifies a data compression algorithm. Valid values: ■ 0: specifies not to use data compression algorithms. ■ 1: specifies to use the snappy data compression algorithm. ■ 2: specifies to use the lz4 data compression algorithm.
ec	Specifies whether to enable the EC storage mode. Default value: disable. Valid values: ■ enable ■ disable
data_chunks	Specifies the number of data chunks. Default value: 8.
parity_chunks	Specifies the number of parity chunks. Default value: 3.
packet_bits	Specifies the size of a single data block in ec mode. Default value: 15.
copy	Specifies the number of data replicas. Default value: 3.
storage_mode	Specifies the storage mode of the disk.
cache	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	Specifies the data storage name.
simsuppress	Specifies whether to enable the delay simulation feature. Default value: disable. Valid values: ■ enable ■ disable
baselatency	Specifies the basic latency. Default value: 300.
consumespeed	Specifies the processing speed. Default value: 256 bit/μs.
lat80th	Specifies the quantile jitter control of the latency as 80%.
lat90th	Specifies the quantile jitter control of the latency as 90%.
lat99th	Specifies the quantile jitter control of the latency as 99%.

- Query segment information

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

- Check a segment

  In the device list, find a device and choose **More > Check Segment** in the **Actions** column corresponding to the device. In the dialog box that appears, select the segment to be checked and click **Submit**.

1.6.2.2.7. 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**.

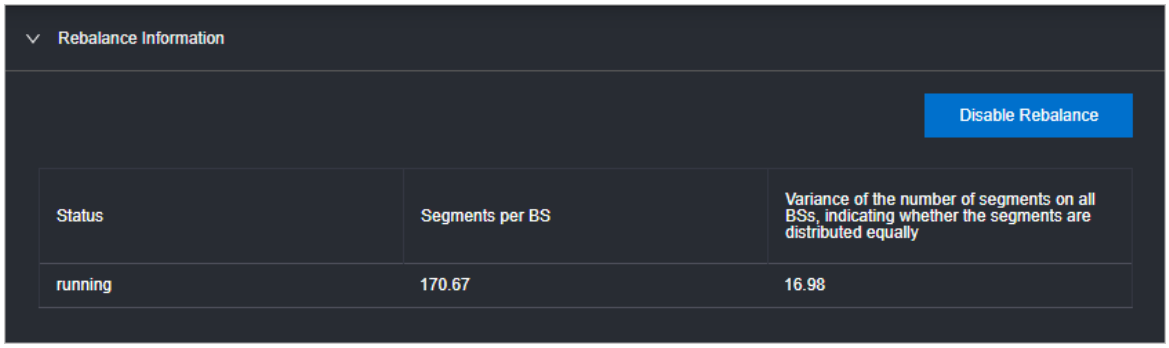

1.6.2.3. miniOSS

The miniOSS module provides features such as monitoring dashboard, user management, permission and quota management, array monitoring, and system management.

1.6.2.3.1. Monitoring dashboard

This topic describes how to view the overview, bucket usage heatmap, user quota usage heatmap, usage trends, and network traffic trends of miniOSS in the system, and download logs 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 **Apsara Distributed File System Management > miniOSS > Dashboard**.
4. On the page that appears, view the following information:
   - **Overview**

     This section shows the bucket information, user information, and health information of miniOSS.

     In the **Health** section, if the value of **Abnormal Bucket Watermark** or **Abnormal User Quota Watermarks** is greater than 0, the corresponding value is displayed in red.

     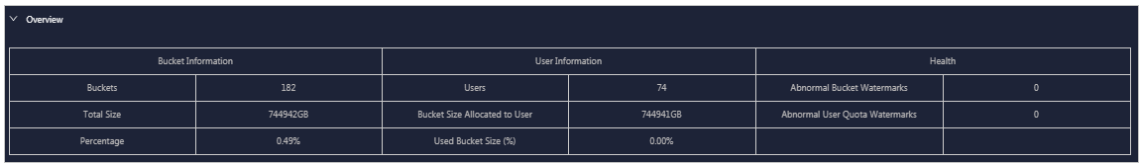

   - **Bucket Watermark Heatmap**

     This section shows the bucket capacity usage.

     The number of sections in **Bucket Watermark Heatmap** is the same as the value of Buckets in **Overview**. Buckets are displayed in different colors based on their usage status.

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

     Move the pointer over a bucket section to view the usage of the bucket.

     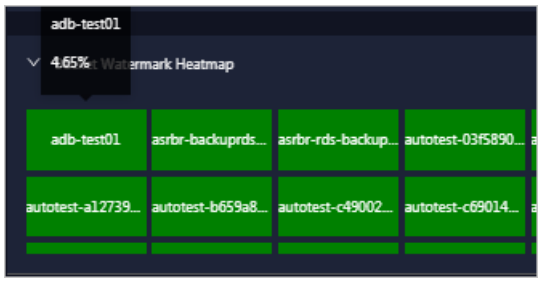

   - **User Quota Watermark Heatmap**

     This section shows the user quota usage information.

     User quota usage = Used capacity of all buckets of the user/Total capacity of all buckets of the user. Buckets are displayed in different colors based on their usage status.

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

     Move the pointer over a section to view the percentage of capacity used by all buckets of a user.

- **Watermark Trend**

  This section shows the historical and predicted data usage of a user or bucket. Usage represents the disk usage, and usage of a user indicates the disk usage of buckets.

  Usage trend data is based on scheduled tasks in the system. The system stores or updates data every 30 minutes.

  Select a bucket or user from the drop-down list to view the corresponding usage trends.

  > **Note** You can enter a keyword of a **Bucket Name** or **Username** to perform a fuzzy search.

  The top 10 data in terms of the bucket usage is displayed on the right. Click a bucket name in the top 10 data to view the usage trends of the bucket on the left.

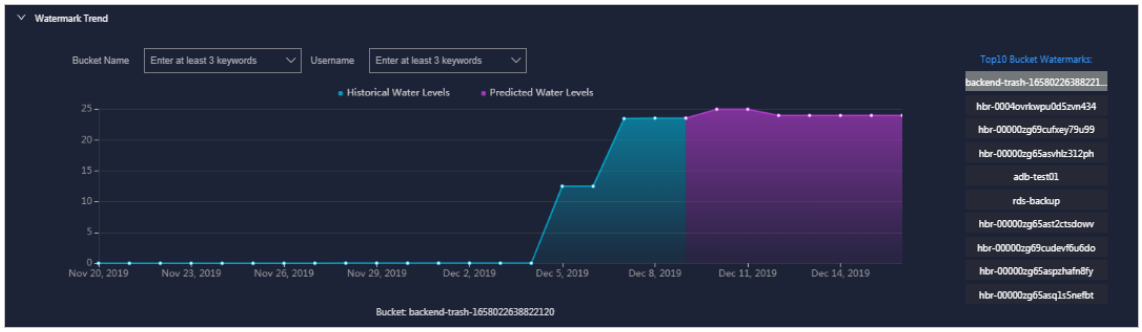

- **Network Traffic Trend**

  This section shows the daily network traffic data of miniOSS in the last month, including the normal network traffic, abnormal network traffic, average weekly network traffic, and average monthly network traffic.

  Network traffic trends are displayed in different colors based on status.

  - Green indicates that the network traffic is normal.
  - Yellow indicates that the network traffic is abnormal.
  - Orange indicates the average weekly network traffic.
  - Blue indicates the average monthly network traffic.

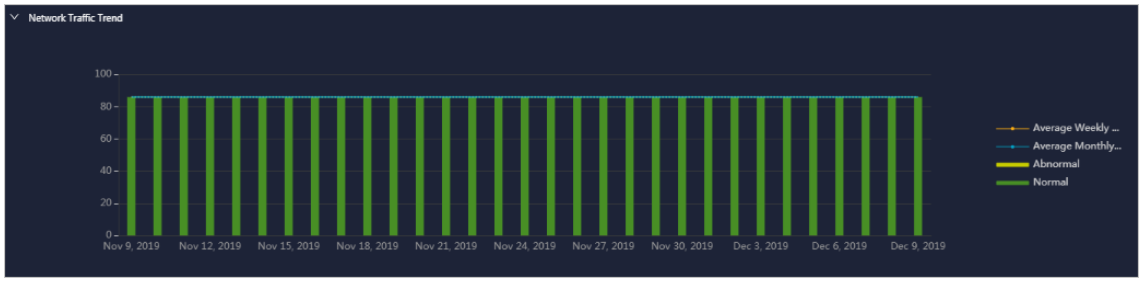

5. (Optional)In the **Download Log** section, click **Download Log Package** and use the download URL to download logs to your computer for subsequent review and analysis.

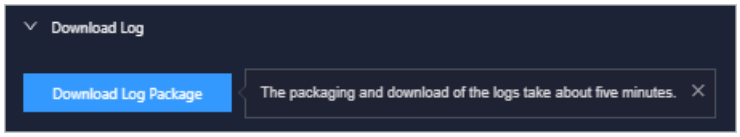

1.6.2.3.2. User management

The User Management module has the User List, Bucket List of User, and Network Traffic Control sections. You can use this module to view user information, network traffic bandwidth, and the list of user 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 **Apsara Distributed File System Management > miniOSS > User Management**.
4. On the page that appears, perform the following operations:
   - View user information

     By default, all users including the administrator and common users are displayed in the list.

     Click **View All** in the **Actions** column corresponding to a user. In the dialog box that appears, view the AccessKey secret of the user.

     In the **User List** section, enter three characters of the username, such as def, and press the Enter key or click the Search icon. This section shows information such as the username, user role, Accesskey, and network traffic bandwidth of the users that meet the search condition.

     > **Note** After the search is complete, click **Refresh** to view all the users in the list.

     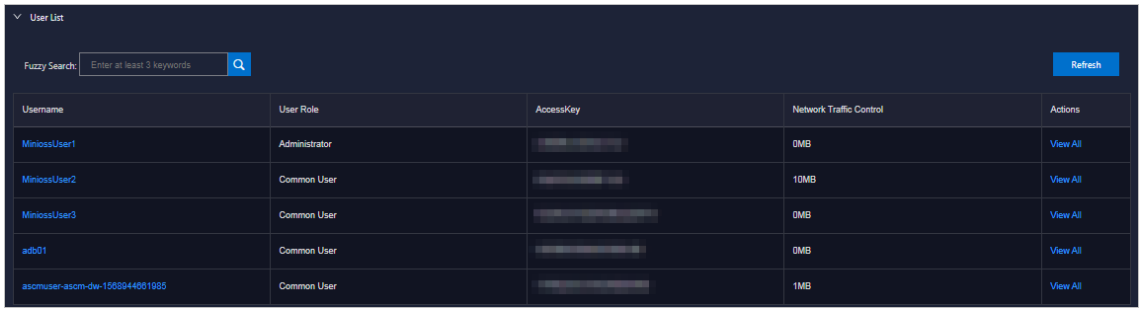

   - View the bucket information of a user

     In the **User List** section, click a username. In the **Bucket List of User** section, view the corresponding bucket information, including the bucket name, bucket ACL, user ACL, quota, and bucket creation time of the user.

     In the **Bucket List of User** section, enter five characters of the bucket name, such as atest, and press the Enter key or click the Search icon. The information of the buckets that meet the search condition is displayed.

     > **Note** After the search is complete, click **Refresh** to view all the information of all buckets.

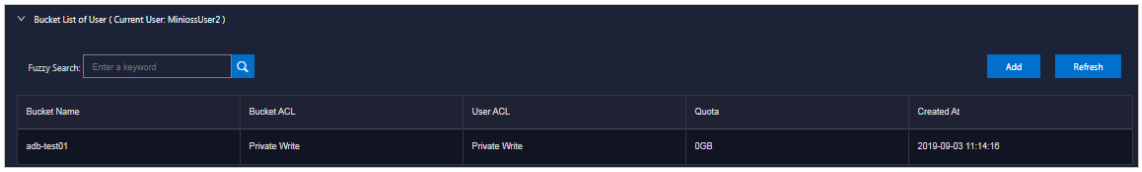

- Add a bucket for a common user

  > **Notice** You can add a bucket only for a common user, instead of for an administrator.

  In the **User List** section, find the common user for whom you want to add a bucket and click the username. In the **Bucket List of User** section, click **Add**. In the dialog box that appears, enter the quota and click **OK**.

  > **Note** Enter an integer from 0 to 4094 as the quota. Unit: GB.

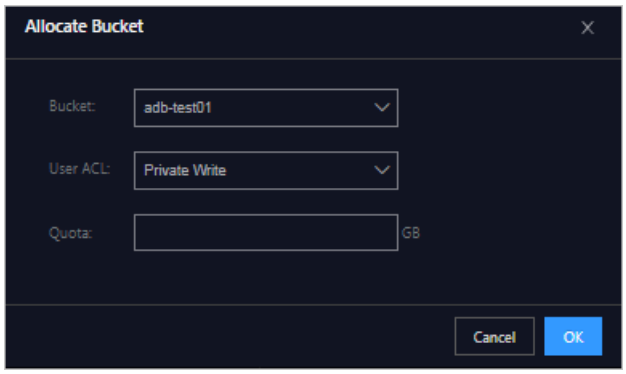

- View the network traffic bandwidth of a user

  In the **User List** section, find the user whose network traffic bandwidth you want to view and click the username. In the **Network Traffic Control** section, view the network traffic bandwidth of the user.

- Modify the network traffic bandwidth of a user

  In the **User List** section, find the user whose network traffic bandwidth you want to modify and click the username. In the **Network Traffic Control** section, click **Modify** to modify the network traffic bandwidth of the user. Then, click **Save**. The traffic bandwidth value must be 0 or a positive integer.

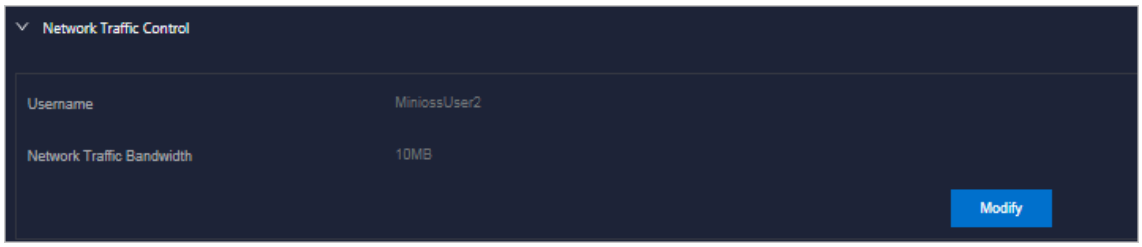

1.6.2.3.3. Permission and quota management

The Permission/Quota Management module allows you to view the bucket list and user lists of buckets, as well as add, modify, and delete 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 **Apsara Distributed File System Management > miniOSS > Permission/Quota Management**.
4. Perform the following operations:
   - View bucket information

     By default, all buckets are displayed in the bucket list. In the **Bucket List** section, you can view basic information such as the names, ACLs, quotas, and network traffic bandwidth for all buckets.

     In the upper-left corner of the section, enter a keyword of the bucket name in the search box and press the Enter key or click the Search icon. The information of the buckets that meet the search conditions is displayed.

     > **Note** After the search is complete, click **Refresh** to view all the buckets in the list.

   - Add a bucket

     In the **Bucket List** section, click **Add**. In the dialog box that appears, enter the bucket name and click **OK**. The bucket name must be 3 to 63 characters in length. It can contain only lowercase letters, digits, hyphens (-), and cannot start or end with a hyphen (-).

   - Modify a bucket

     In the **Bucket List** section, click **Modify** in the **Actions** column corresponding to a bucket. In the dialog box that appears, modify the bucket ACL, quota, and network traffic bandwidth, and click **OK**.

   - Delete a bucket

     In the **Bucket List** section, click **Delete** in the **Actions** column corresponding to a bucket. In the message that appears, click **OK**.

   - View the user information of the user to which a bucket belongs

     In the **Bucket List** section, click a bucket name to view the user information related to the bucket in the **User List of Bucket** section.

     In the upper-left corner of the section, enter a keyword of the username in the search box and press the Enter key or click the Search icon. The information of the users that meet the search conditions is displayed.

     > **Note** After the search is complete, click **Refresh** to view the information of all the users.

1.6.2.3.4. Array monitoring

The Array Monitoring module allows you to view the running status 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 **Apsara Distributed File System Management > miniOSS > Array Monitoring**.
4. View the running status of each device.

   By default, you can view the status information of all devices.

   Device types include rack, controller, hard drive, rack battery, power supply unit (PSU), fan, FC port, iSCSI port, SAS port, USB port, cluster node, cluster system, block storage, volume information, storage pool information, host, NFS service, CIFS service, FTP service, and file system.

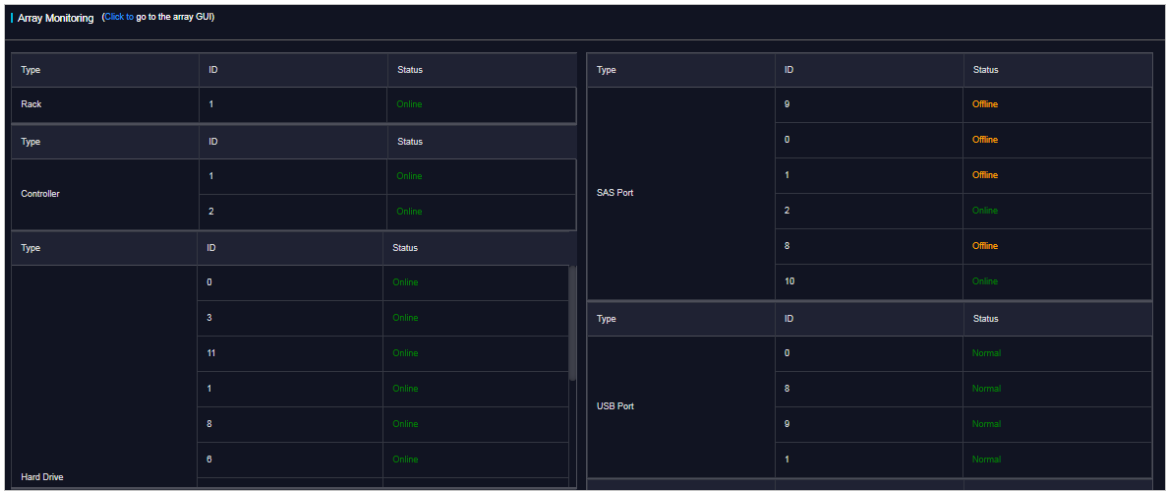

   Values in different colors indicate different states:

   - Green indicates that the device is online or is running normally.
   - Yellow indicates that the device is offline.
   - Red indicates that the device has an exception.

   In the upper part of the page, click **Click** to go to the array GUI.

1.6.2.3.5. System management

The System Management module allows you to modify the bucket watermark threshold and user watermark threshold.

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 > miniOSS > System Management**.
4. Perform the following operations:
   - Modify the bucket usage threshold
     a. In the **Bucket Watermark Threshold** section, click **Modify**.
     b. Enter values greater than 0 and less than or equal to 100 in the Warning Value, Error Value, and Fatal Error Value fields. Make sure that the warning value is less than the error value and the error value is less than the fatal error value.
     c. Click **Save**.

- Modify the user usage threshold
  a. In the **User Watermark Threshold** section, click **Modify**.
  b. Enter values greater than 0 and less than or equal to 100 in the Warning Value, Error Value, and Fatal Error Value fields. Make sure that the warning value is less than the error value and the error value is less than the fatal error value.
  c. Click **Save**.

1.6.3. Task Management

The system allows you to run operations scripts on the cloud platform, which reduces your actions by using command lines, lowers misoperations, and improves the security and stability of the cloud platform.

1.6.3.1. 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.

1.6.3.2. 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.

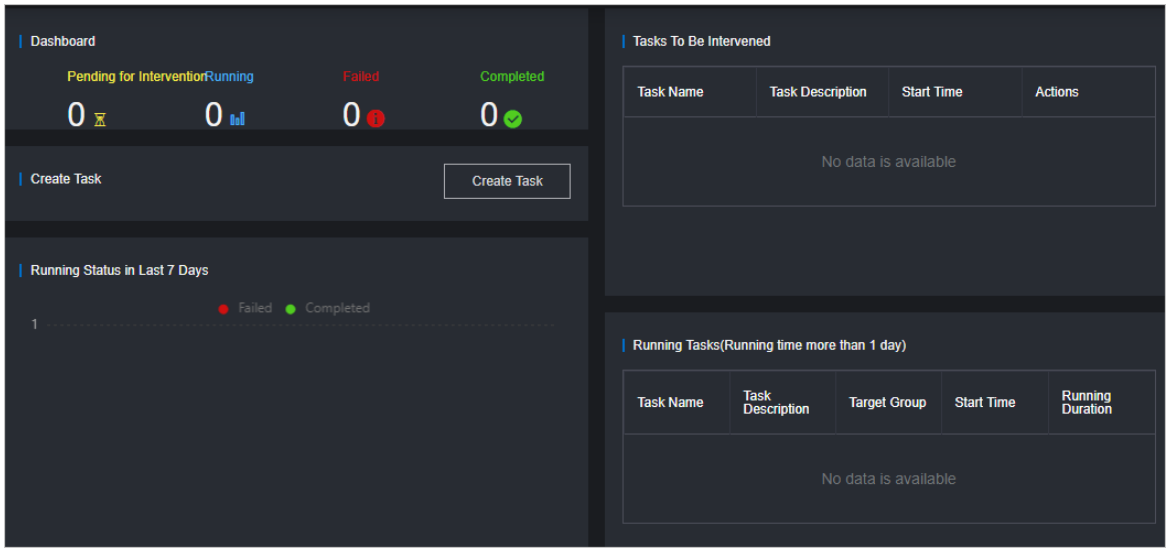

4. Perform the following operations:
   - In the **Dashboard** section, view the number of tasks that are in the **Pending for Intervention**, **Running**, **Failed**, or **Completed** state.

     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.

1.6.3.3. Create a task

You can make regular modifications as tasks to run 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.

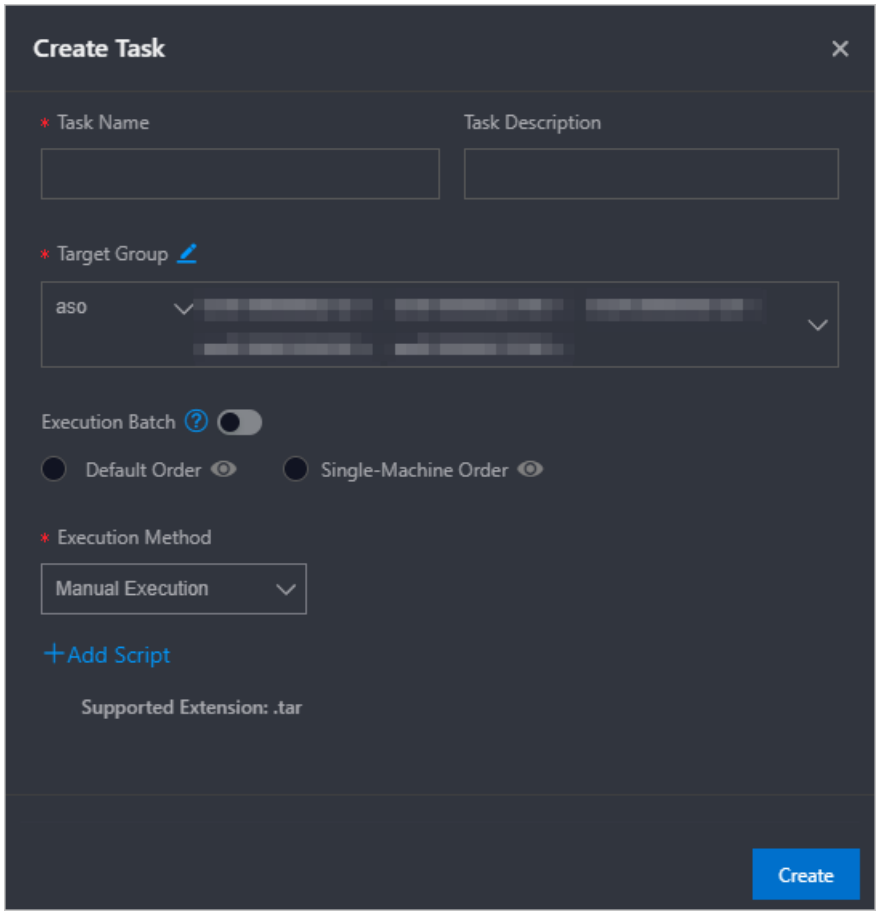

Parameter	Description
Task Name	The name of the O&M task.
Task Description	The description of the O&M task.
Target Group	The task target. You can use one of the following methods to configure the target group: ○ Select a product. Enter the VM or physical machine in the field and press the Enter key. 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**.

Parameter	Description
Execution Batch	Optional. This option appears after you specify the target group. If **Execution Batch** is not specified, **Target Group** is displayed in the **Target Group** column, which can be viewed by choosing **Task Management > 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 options: ○ **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 rest machines in batch 5. You can change the number of machines in each batch. By default, if the number of machines is greater than 10, the machines are allocated to different batches, with one machine in batch 1, three machines in batch 2, five machines in batch 3, N/3-1 (an integer) machines in batch 4, N/3-1 (an integer) machines in batch 5. You can change the number of machines in each batch. ○ **Single-Machine Order**: By default, each batch has one machine. You can change the number of machines in each batch.
Execution Method	If **Execution Batch** is specified, **Execution Method** can be set only to **Manual Execution**. If **Execution Batch** is not enabled, you can select one of the following execution methods: ○ **Manual Execution**: You must manually start the task. 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.
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.

1.6.3.4. 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)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 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.

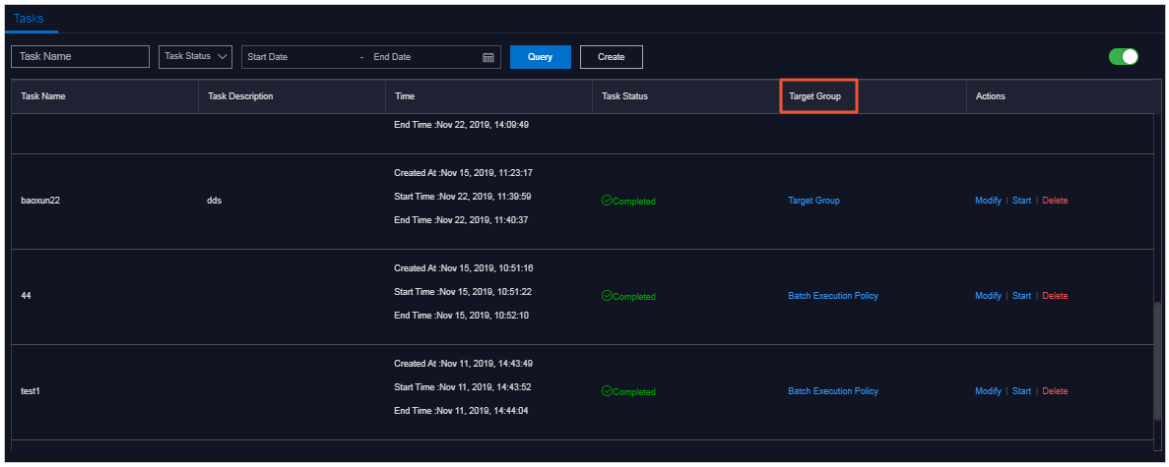

6. In the dialog box that appears, view the task execution status based on the machine color. Click a machine name to view the task execution results on it.

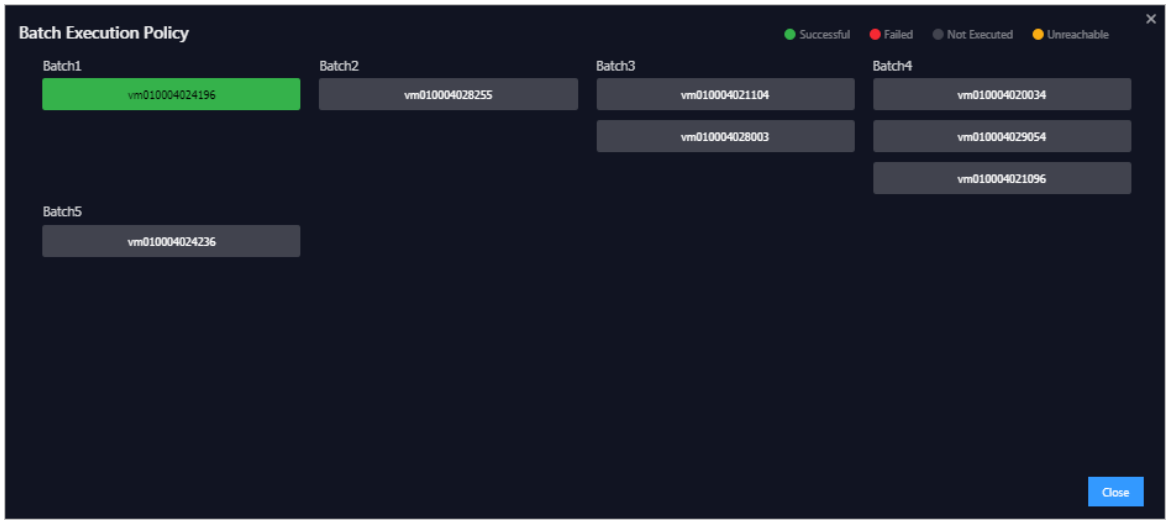

1.6.3.5. Start a task

If you select **Manual Execution** when you create a task, you must manually start the task after it is created.

Procedure

1. 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.

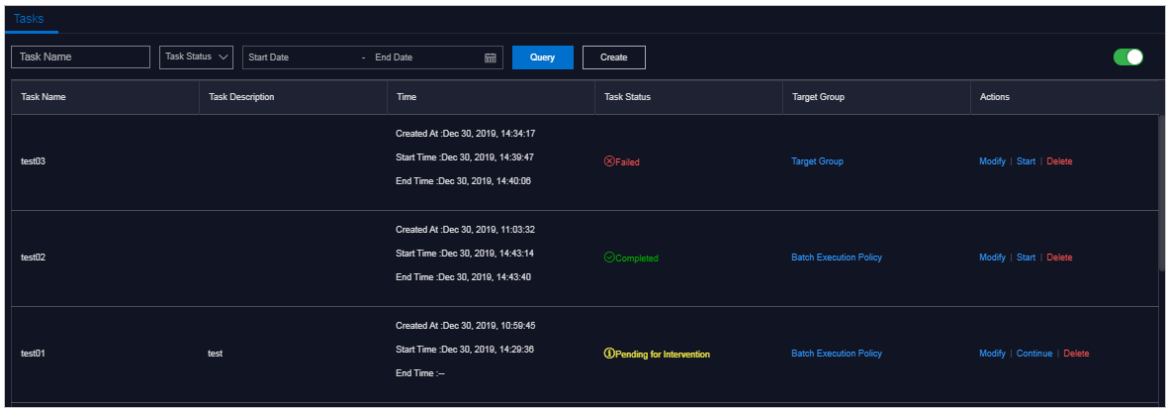

1.6.3.6. Delete a task

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**.

1.6.3.7. 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.

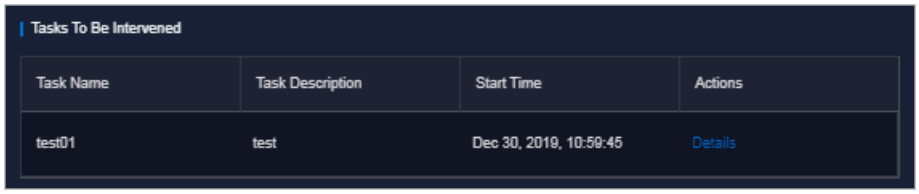

5. On the **Task Details** tab, check the information and click **Continue** for the task to continue.

1.6.3.8. Configure an XDB backup task

The XDB Backup module allows you to configure XDB data backups without using command lines. You can configure and modify backup tasks on the XDB Backup page to regularly back up platform data and back up data in real time.

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 > XDB Backup**.
4. On the **XDB Backup** page, configure the parameters for the XDB backup task.

Parameter	Description
Task Name	The name of the XDB backup task. By default, the name is **xdbBackup** and cannot be modified.
Task Description	The description of the XDB backup task.

Parameter	Description
Target Group	Required. The target of the XDB backup task. You can use one of the following methods to configure the target group: ○ Select the **product**, **cluster**, **service**, **server role**, and **virtual machine (VM) or physical machine** in sequence. ○ Select a product. Enter the VM or physical machine in the field and 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. You can set **Execution Batch** to one of the following options: ○ **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 rest machines in batch 5. You can change the number of machines in each batch. By default, if the number of machines is greater than 10, the machines are allocated to different batches, with one machine in batch 1, three machines in batch 2, five machines in batch 3, N/3-1 (an integer) machines in batch 4, and N/3-1 (an integer) machines in batch 5,until all of the machines are allocated. N is the total number of servers in the cluster. You can change the number of machines in each batch. ○ **Single-Machine Order**: By default, each batch has one machine. You can change the number of machines in each batch. If **Execution Batch** is not specified, **Execution Batch** is disabled by default, and **Target Group** is displayed in the **Target Group** column in the task list of the **Task Management > Task Management** page. If **Execution Batch** is specified and saved, **Execution Batch** is automatically enabled, and **Batch Execution Policy** is displayed in the **Target Group** column.

Parameter	Description
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**: You must manually start the task. When **Manual Execution** is selected, you must click **Start** in the **Actions** column to run the task after the task is created. ○ **Scheduled Execution**: You must select the execution time. The task automatically runs when the execution time is reached. ○ **Regular Execution**: You must select the time interval and times to run the task. If the execution condition is met, the task is run again. ○ **Advanced**: You must enter the crontab expression to configure the command to run the task periodically. For example, `0 20 20 * * ?` indicates that the task runs at 20:20 every day.
Execution Scripts	By default, the system automatically loads the XDB backup script.

5. Click **Create**.

   You can view the created XDB task in the task list of the **Task Management > Task** page. The system automatically runs the XDB backup task when the task execution condition is met. If **Execution Method** of the XDB backup task is set to **Manual Execution**, you must start the backup task based on the procedures described in **O&M tools > Task management > Start a task**.

   > **Note** After the XDB backup task is created, you can click **Modify** in the lower part of the **XDB Backup** page to modify the information of the backup task.

   After the XDB backup task is completed, O&M engineers can view the backup file of each instance in the */alidata/xdb-backup/instance name* directory on the backup server. The backup file name is in the following format: instance name-timestamp (specific to day).tar. The temporary backup information in the */alidata/xdb-backup-tmp* directory of the temporary backup folder is automatically deleted.

1.6.4. Apsara Infrastructure Management Framework O&M

1.6.4.1. Old console

1.6.4.1.1. What is Apsara Infrastructure Management Framework?

This topic describes what Apsara Infrastructure Management Framework is from the aspects of core functions and basic concepts.

1.6.4.1.1.1. Overview

Apsara Infrastructure Management Framework is a distributed data center management system, which manages applications on clusters containing multiple machines and provides the basic functions such as deployment, upgrade, expansion, contraction, and configuration changes.

Apsara Infrastructure Management Framework also supports monitoring data and analyzing reports, which facilitates users to perform one-stop Operation & Maintenance (O&M) control. Based on the Apsara Infrastructure Management Framework services, automated O&M is implemented in the large-scale distributed environment, which greatly improves the operations efficiency and enhances the system availability.

Apsara Infrastructure Management Framework is mainly composed of TianjiMaster and TianjiClient. Apsara Infrastructure Management Framework installs TianjiClient as the agent on machines it manages. Then, TianjiMaster accepts and issues the upper-layer instructions to TianjiClient for execution. In the upper layer, Apsara Infrastructure Management Framework is divided into different components based on different functions, and then provides API server and portal for external use.

Core functions

- Network initialization in data centers
- Server installation and maintenance process management
- Deployment, expansion, and upgrade of cloud products
- Configuration management of cloud products
- Automatic application for cloud product resources
- Automatic repair of software and hardware faults
- Basic monitoring and business monitoring of software and hardware

1.6.4.1.1.2. Basic concepts

Before using Apsara Infrastructure Management Framework, you must know the following basic concepts for a better understanding.

project

A collection of clusters, which provides service capabilities for external entities.

cluster

A collection of physical machines, which logically provides services and is used to deploy project software.

- A cluster can only belong to one project.
- Multiple services can be deployed on a cluster.

service

A set of software, which provides relatively independent functions. A service is composed of one or more server roles and can be deployed on multiple clusters to form multiple sets of services and provide the corresponding service capabilities. For example, Apsara Distributed File System, Job Scheduler, and Apsara Name Service and Distributed Lock Synchronization System are all services.

service instance

A service that is deployed on a cluster.

server role

One or more indivisible deployment units into which a service can be divided based on functions. A server role is composed of one or more specific applications. If a service is deployed on a cluster, all the server roles of the service must be deployed to machines of this cluster. Multiple server roles, such as PanguMaster and TianjiClient, can be deployed on the same server.

server role instance

A server role that is deployed on a machine. A server role can be deployed on multiple machines.

application

Applications correspond to each process-level service component in a server role and each application works independently. The application is the minimum unit for deployment and upgrade in Apsara Infrastructure Management Framework, and can be deployed to each machine. Generally, an application is an executable software or Docker container.

If a server role is deployed on a machine, all applications in the server role must be deployed to this machine.

rolling

Each time when a user updates configurations, Apsara Infrastructure Management Framework upgrades services and modifies the cluster configurations based on the updated configurations. This process is called rolling.

service configuration template

Some configurations are the same when services are deployed on clusters. A service configuration template can be created to quickly write the same configurations to different clusters. The service configuration template is basically used for large-scale deployment and upgrade.

associated service template

A *template.conf* file that exists in the configurations. This file declares the service configuration template and its version, of which the configuration is used by the service instance.

final status

If a cluster is in this status, all hardware and software on each of its machines are normal and all software are in the target version.

dependency

The dependency between server roles in a service defines that server roles with dependencies run tasks or are upgraded based on the dependency order. For example, if A depends on B, B is upgraded first. A starts to be downloaded after B is downloaded successfully, and upgraded after B is successfully upgraded. (By default, the dependency does not take effect for configuration upgrade.)

upgrade

A way of aligning the current status with the final status of a service. After a user submits the version change, Apsara Infrastructure Management Framework upgrades the service version to the target version. With the server role as the processing unit, upgrade aims to update the versions of all machines to the target version.

At the beginning, the final status and current status of the cluster are the same. When a user submits the change, the final status is changed, whereas the current status is not. A rolling task is generated and has the final status as the target version. During the upgrade, the current status is continuously approximating to the final status. Finally, the final status and the current status are the same when the upgrade is finished.

1.6.4.1.2. 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 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 URL of the Apsara Uni-manager Operations Console is in the following format: *region-id*.*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 (*region-id*.*ops*.asconsole.*intranet-domain-id*.com). Then, press the Enter key.

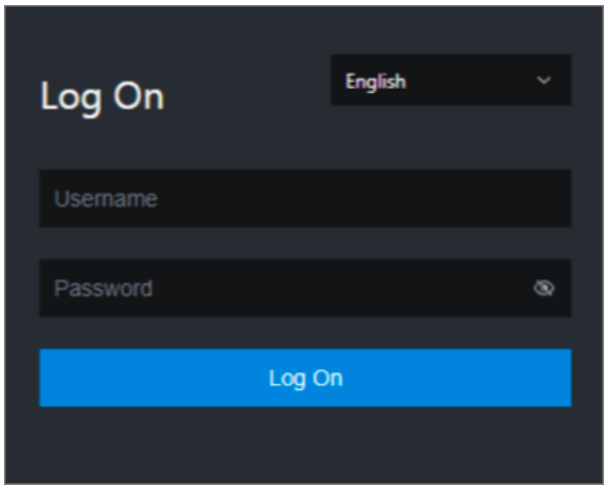

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

1.6.4.1.3. Web page introduction

Before performing Operation & Maintenance (O&M) operations on Apsara Infrastructure Management Framework, you must have a general understanding of the Apsara Infrastructure Management Framework page.

1.6.4.1.3.1. Instructions for the homepage

After you log on to the Apsara Infrastructure Management Framework console, the homepage appears. This topic describes the basic operations and functions on the homepage.

Log on to Apsara Infrastructure Management Framework. The homepage appears, as shown in Homepage of the Apsara Infrastructure Management Framework console.

Homepage of the Apsara Infrastructure Management Framework console

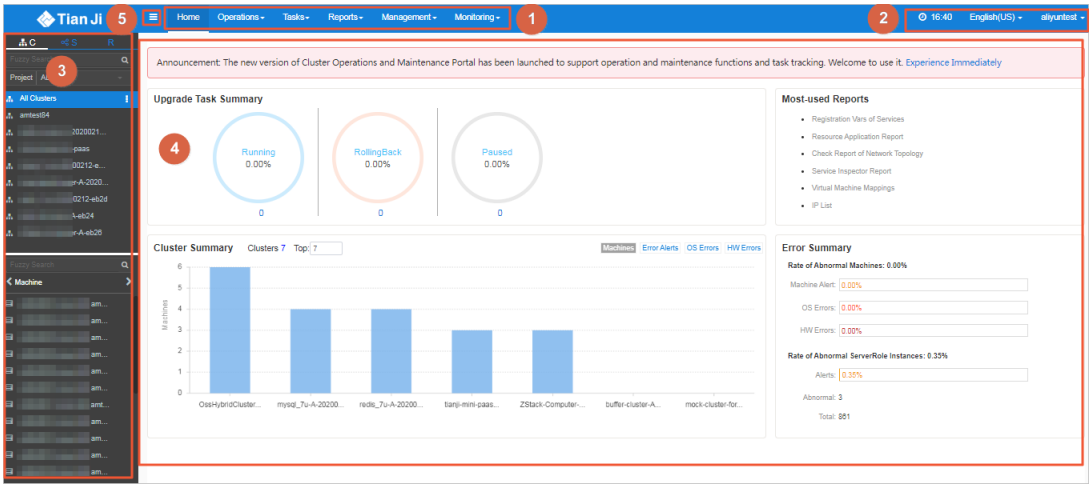

Description of functional sections describes the functional sections on the homepage.

Description of functional sections

Section		Description

①	Top navigation bar	• **Operations**: the quick entrance to operations and maintenance (O&M) operations and their objects. This menu consists of the following submenus:
		  ◦ **Cluster Operations**: allows you to use the project permissions to perform O&M and management operations on clusters. For example, you can view the cluster status.
		  ◦ **Service Operations**: allows you to use the service permissions to manage services. For example, you can view the service list.
		  ◦ **Machine Operations**: allows you to perform O&M and management operations on machines. For example, you can view the machine status.
		• **Tasks**: Rolling tasks are generated when you modify the configurations in the system. This menu allows you to view the running tasks, task history, and deployment of clusters, services, and server roles in all projects.
		• **Reports**: allows you to view monitoring data in tables and find specific reports by using fuzzy search.
		• **Monitoring**: monitors metrics during system operations and sends alert notifications for abnormal conditions. This menu allows you to view the alert status, modify alert rules, and search alert history.

Section		Description

②	Upper-right buttons	•  :
		  ○ **TJDB Synchronization Time**: the time when the data on the current page is generated.
		  ○ **Final Status Computing Time**: the time when the desired-state data on the current page is calculated.
		  The system processes data as fast as it can after the data is generated. Latency exists because Apsara Infrastructure Management Framework is an asynchronous system. Time information helps explain why data on the current page is generated and determine whether the system experiences an error.
		• 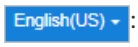 : the current display language of the console. You can select another language from the drop-down list.
		•  : your logon account. You can select **Logout** from the drop-down list to log out of your account.

③	Left-side navigation pane	In the left-side navigation pane, you can view the logical architecture of Apsara Infrastructure Management Framework. The tabs allow you to view details and perform operations. For more information, see Introduction on the left-side navigation pane.
④	Workspace	The workspace shows a summary of tasks and other information. • **Upgrade Task Summary**: shows the numbers and proportions of running, rolling back, and suspended upgrade tasks. • **Cluster Summary**: shows the numbers of machines, error alerts, operating system errors, and hardware errors in each cluster. • **Error Summary**: shows metric values about the rate of abnormal machines and the rate of abnormal server role instances. • **Most-used Reports**: shows links of common statistical reports.
⑤	Show/hide button	If you do not need to use the left-side navigation pane, click this button to hide the pane and enlarge the workspace.

1.6.4.1.3.2. Instructions for the left-side navigation pane

The left-side navigation pane contains three tabs: **C** (cluster), **S** (service), and **R** (report). This topic describes how to use the tabs to view information.

Cluster

You can search for clusters in a project and their information such as the cluster status, cluster operations and maintenance (O&M), service desired state, and logs by fuzzy match.

On the **C** tab of the left-side navigation pane, you can perform the following operations:

- Enter a cluster name or a part of a cluster name in the search box to filter clusters.
- Select a project from the **Project** drop-down list to view all clusters in the project.
- Move the pointer over the icon next to a cluster and select menu items to perform corresponding operations on the cluster.

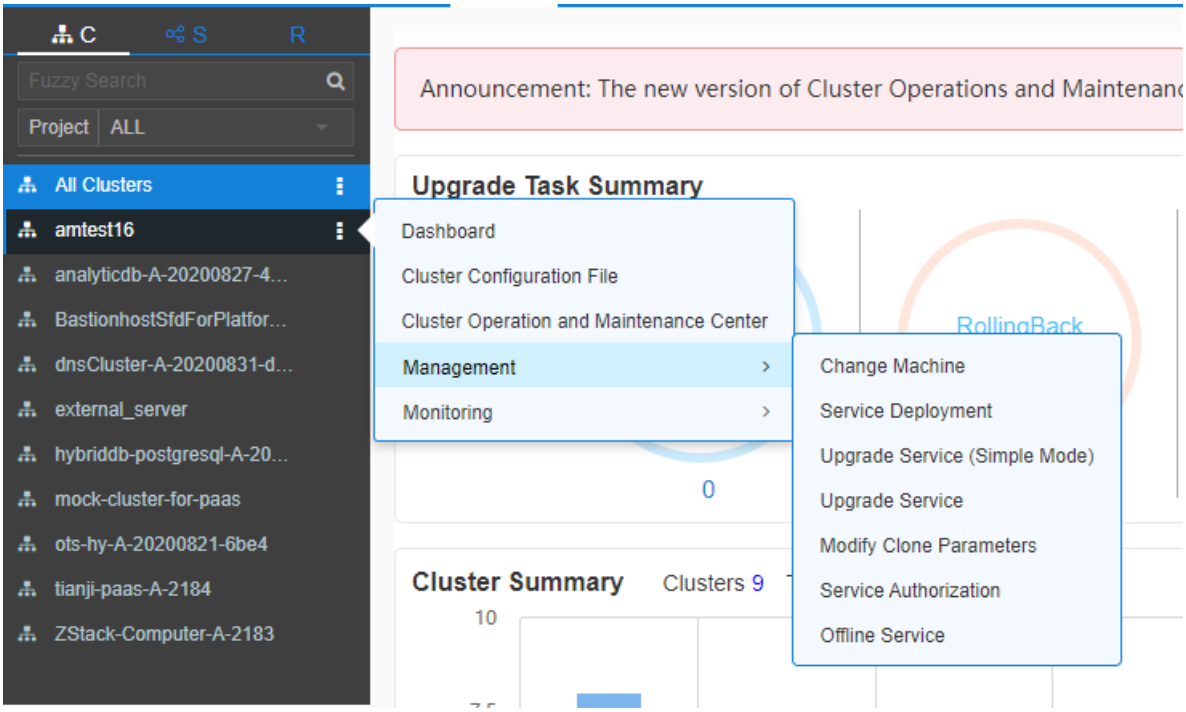

- Click a cluster. All machines and services within the cluster are displayed in the lower part of the left-side navigation pane. Move the pointer over the icon next to a machine or service on the **Machine** or **Service** tab and select menu items to perform corresponding operations on the machine or service.

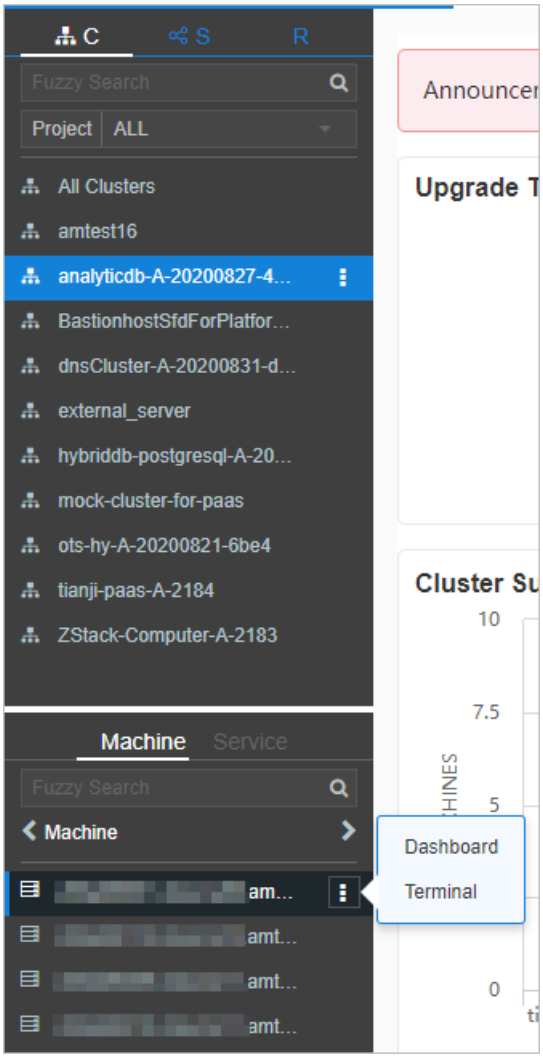

- Click the **Machine** tab. Double-click a machine to view information about all server roles on the machine. Double-click a server role to view applications, and then double-click an application to view log files.
- Click the **Service** tab. Double-click a machine to view information about all server roles on the machine. Double-click a server role to view machines, double-click a machine to view applications, and then double-click an application to view log files.
- Double-click a log file. Move the pointer over the log file, click the icon next to the log file, and then click **Download** to download the log file.

  Alternatively, move the pointer over a log file and click **View** next to the log file. The time-ordered log details are displayed on the **Log Viewer** page. You can search for log details by keyword.

Service

You can search for services and view information about services and service instances by fuzzy match.

On the **S** tab of the left-side navigation pane, you can perform the following operations:

- Enter a service name or a part of a service name in the search box to filter services.
- Move the pointer over the icon next to a service and select menu items to perform corresponding operations on the service.

- Click a service. All service instances within the service are displayed in the lower part of the left-side navigation pane. Move the pointer over the icon next to a service instance and select menu items to perform corresponding operations on the service instance.

Report

You can search for reports by fuzzy match and view report details.

On the **R** tab of the left-side navigation pane, you can perform the following operations:

- Enter a report name or a part of a report name in the search box to filter reports.
- Click **All Reports** or **Favorites**. Corresponding groups are displayed in the lower part of the left-side navigation pane. Double-click a group to view all reports in the group. Double-click a report to view details of the report.

1.6.4.1.4. Cluster operations

This topic describes the actions about cluster operations.

1.6.4.1.4.1. View configuration information of a cluster

This topic describes how to view the basic information, deployment plan, and configuration information of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Cluster Operations**.

   The **Cluster Operations** page contains the following information:

   - Cluster

     The name of a cluster. Click a cluster name to go to the Cluster Dashboard page. For more information, see View dashboard information of a cluster.

   - Scale-Out/Scale-In

     The numbers of machines and server roles that are scaled in and out. Click a number to go to the Cluster Operation and Maintenance Center page. For more information, see View information of the cluster O&M center.

   - Abnormal Machine Count

     The number of machines that are not in the Good state within a cluster. Click the number to go to the Cluster Operation and Maintenance Center page. For more information, see View information of the cluster O&M center.

   - Final Status of Normal Machines

     Specifies whether a cluster has reached the desired state. Select **Clusters not Final** above the cluster list to view all clusters that have not reached the desired state. Click a link in the column to view desired state information. For more information, see View the desired state of a service.

   - Rolling

Specifies whether rolling tasks are running within a cluster. Select **Rolling Tasks** above the cluster list to view all clusters that have rolling tasks. Click rolling in the column to view rolling tasks. For more information see View rolling tasks.

3. (Optional)Select a project from the drop-down list or enter a cluster name to search for the cluster.
4. Click the cluster name or click **Cluster Configuration** in the **Actions** column to go to the **Cluster Configuration** page.

   Cluster configuration description describes the parameters on the **Cluster Configuration** page.

   Cluster configuration description

Section	Parameter	Description
Basic Information	**Cluster**	The name of the cluster.
	Project	The project to which the cluster belongs.
	Clone Switch	○ **Pseudo-clone**: The system is not cloned when a machine is added to the cluster. ○ **Real Clone**: The system is cloned when a machine is added to the cluster.
	Machines	The number of machines included in the cluster. Click View Clustering Machines to view the list of machines.
	Security Verification	The access control among processes. By default, security verification is disabled in non-production environments. You can enable or disable security verification based on your business requirements.
	Cluster Type	○ RDS ○ NETFRAME ○ T4: a type of cluster that renders special configurations for the mixed deployment of e-commerce ○ Default
Deployment Plan	**Service**	The service that is deployed within the cluster.
	Dependency Service	The service on which the current service depends.
	Service Information	The service that you want to view. Select a service from the drop-down list to view its configuration information.
	Service Template	The template that is used by the service.

Section	Parameter	Description
Service Information	Monitoring Template	The monitoring template that is used by the service.
	Machine Mappings	The machines where server roles of the service are deployed.
	Software Version	The version of the software that is included in server roles of the service.
	Availability Configuration	The percentage of availability configuration for server roles of the service.
	Deployment Plan	The deployment plan of server roles of the service.
	Configuration Information	The configuration file that is used for the service.
	Role Attribute	The server roles and their parameter information.

5. Click **Operation Logs** in the upper-right corner to view version differences. For more information about operation logs, see View operation logs.

1.6.4.1.4.2. View dashboard information of a cluster

This topic describes how to view the basic information and related statistics of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the **Cluster Dashboard** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and select **Dashboard**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click the name of the target cluster.
3. View all information about the cluster on the **Cluster Dashboard** page. The following table describes the information that you can view, such as basic information, desired state information, rolling tasks, dependencies, resources, virtual machine (VM) mappings, and monitoring status.

Parameter	Description

Parameter	Description
Basic Cluster Information	The basic information about the cluster. ○ **Project Name**: the name of the project. ○ **Cluster Name**: the name of the cluster. ○ **IDC**: the data center to which the cluster belongs. ○ **Final Status Version**: the latest version of the cluster. ○ **Cluster in Final Status**: specifies whether the cluster has reached the desired state. ○ **Machines Not In Final Status**: the number of machines that have not reached the desired state. ○ **Real/Pseudo Clone**: specifies whether the system is cloned when a machine is added to the cluster. ○ **Expected Machines**: the number of machines that are expected within the cluster. ○ **Actual Machines**: the number of machines that are deployed in the current environment. ○ **Machines Not Good**: the number of machines that are not in the Good state within the cluster. ○ **Actual Services**: the number of services that are deployed within the cluster. ○ **Actual Server Roles**: the number of server roles that are deployed within the cluster. ○ **Cluster Status**: specifies whether the cluster is starting or shutting down machines.
Machine Status Overview	The status of machines within the cluster.
Machines In Final State	The distribution of machines where services are deployed, based on whether the machines have reached the desired state.
Load-System	The statistics chart of the cluster system load.
CPU-System	The statistics chart of the CPU load.
Mem-System	The statistics chart of the memory load.
Disk_Usage-System	The statistics chart of the disk usage.
Traffic-System	The statistics chart of the system traffic.
TCP State-System	The statistics chart of the CPU request status.
TCP Retrans-System	The statistics chart of the CPU retransmission traffic.
Disk_IO-System	The statistics chart of the disk I/O information.

Parameter	Description
Service Instances	The service instances that are deployed within the cluster and their desired state information. ○ **Service Instance**: the service instance that is deployed within the cluster. ○ **Final Status**: specifies whether the service instance has reached the desired state. ○ **Expected Server Roles**: the number of server roles that are expected to deploy in the service instance. ○ **Server Roles in Final Status**: the number of server roles that have reached the desired state in the service instance. ○ **Server Roles Going Offline**: the number of server roles that are being unpublished from the service instance. ○ Actions: Click **Details** to go to the **Service Instance Information Dashboard** page. For more information about the service instance dashboard, see View the service instance dashboard.
Upgrade Tasks	The upgrade tasks within the cluster. ○ **Cluster Name**: the name of the cluster. ○ **Type**: the type of the upgrade task. Valid values: app and config. app indicates version upgrade, and config indicates configuration change. ○ **Git Version**: the change version of the upgrade task. ○ **Description**: the description of the change. ○ **Rolling Result**: the result of the upgrade task. ○ **Submitted By**: the user who submits the change. ○ **Submitted At**: the time when the change is submitted. ○ **Start Time**: the time when rolling starts. ○ **End Time**: the time when the upgrade task ends. ○ **Time Used**: the time consumed for the upgrade. ○ Actions: Click **Details** to go to the **Rolling Task** page. For more information about rolling tasks, see View rolling tasks.
Cluster Resource Request Status	○ **Version**: the version of the resource request. ○ **Msg**: the error message. ○ **Begintime**: the time when the resource request analysis starts. ○ **Endtime**: the time when the resource request analysis ends. ○ **Build Status**: the build status of resources. ○ **Resource Process Status**: the resource request status of the version.

Parameter	Description
Cluster Resource	○ **Service**: the name of the service. ○ **Service Role**: the name of the server role. ○ **App**: the name of the application of the server role. ○ **Name**: the name of the resource. ○ **Type**: the type of the resource. ○ **Status**: the status of the resource request. ○ **Error_Msg**: the error message. ○ **Parameters**: the parameters of the resource. ○ **Result**: the result of the resource request. ○ **Res**: the ID of the resource. ○ **Reprocess Status**: the request status of AnyTunnel VIP addresses. ○ **Reprocess Msg**: the error message reported when AnyTunnel VIP addresses are requested. ○ **Reprocess Result**: the request result of AnyTunnel VIP addresses. ○ **Refer Version List**: the version that uses the resource.
VM Mappings	The VMs within the cluster. VM information is displayed only when VMs are deployed within the cluster. ○ **VM**: the hostname of the VM. ○ **Currently Deployed On**: the hostname of the physical machine where the VM is deployed. ○ **Target Deployed On**: the hostname of the physical machine where you expect to deploy the VM.
Service Dependencies	The dependency configuration of service instances and server roles within the cluster, and the desired state information of dependency services or server roles. ○ **Service**: the name of the service. ○ **Server Role**: the name of the server role. ○ **Dependent Service**: the service on which the server role depends. ○ **Dependent Server Role**: the server role on which the server role depends. ○ **Dependent Cluster**: the cluster where the dependency server role is deployed. ○ **Dependency in Final Status**: specifies whether the dependency server role has reached the desired state.

1.6.4.1.4.3. View information of the cluster O&M center

This topic describes how to view the status and statistics of services and machines within a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the **Cluster Operation and Maintenance Center** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and select **Cluster Operation and Maintenance Center**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and choose **Monitoring > Cluster Operation and Maintenance Center** in the Actions column.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click the name of the target cluster. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**.
3. View information on the **Cluster Operation and Maintenance Center** page.

Parameter	Description
SR not in Final Status	All server roles that have not reached the desired state within the cluster. Click the number to view the list of server roles. Click a server role to view information of machines where the server role is deployed.
Running Tasks	Specifies whether rolling tasks are running within the cluster. Click **Rolling** to go to the **Rolling Task** page. For more information about rolling tasks, see View rolling tasks.
Head Version Submitted At	The time when the HEAD version is submitted. Click the time to view details.
Head Version Analysis	The status of desired state analysis. During desired state analysis, Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to specific change contents. Desired state analysis can be in one of the following states: ○ **Preparing**: No new version is detected. ○ **Waiting**: The latest version has been detected, but the analysis module has not started. ○ **Doing**: The application to be changed is being analyzed. ○ **done**: The desired state analysis succeeds. ○ **Failed**: The desired state analysis fails to parse change contents. Apsara Infrastructure Management Framework can obtain change contents of server roles in the latest version only when the desired state analysis is in the **done** state. Click a state to view related information.
Service	The service deployed within the cluster. Select a service from the drop-down list.

Parameter	Description
Server Role	The server role of a service within the cluster. Select a server role from the drop-down list. **Note** After you select a service and a server role, machines that are related to the service or the server role are displayed.
Total Machines	The total number of machines within the cluster or machines where the selected server roles are deployed.
Scale-in Scale-out	The numbers of machines and server roles that are scaled in and out.
Abnormal Machines	The numbers of machines in an abnormal state for the following reasons: ○ **Ping Failed**: the number of machines that experience ping_monitor errors because TianjiMaster cannot ping the machines. ○ **No Heartbeat**: the number of machines that experience TianjiClient or network errors because TianjiClient does not report data on a regular basis. ○ **Status Error**: the number of machines that experience critical or fatal errors. Resolve problems based on alert information.
Abnormal Services	The number of machines that have abnormal services. The following rules are used to check whether a service has reached the desired state: ○ Each server role on the machine is in the GOOD state. ○ The actual version of each application of each server role on the machine is consistent with the HEAD version. ○ Before the Image Builder builds an application of the HEAD version, Apsara Infrastructure Management Framework cannot obtain the value of the HEAD version, and the desired state of the service is unknown. This process is called change preparation. The desired state of the service cannot be obtained when the preparation process is in progress or if the preparation fails.

Parameter	Description
Machines	All machines within the cluster or machines where the selected server roles are deployed. ○ Click the Machine Search search box. In the dialog box that appears, enter one or more machines. Fuzzy match and batch search are supported. ○ Click the name of a machine to view its physical information in the Machine Information dialog box. Click **DashBoard** to go to the **Machine Details** page. For more information about machine details, see View the machine dashboard. ○ Move the pointer over the **Final Status** or **Final SR Status** column and click **Details** to view the machine status and system service information, as well as status information and error messages of server roles on the machine. ○ Before you filter machines by service and service role, move the pointer over the **Running Status** column and click **Details** to view status information and error messages of the machine. After you filter machines by service and service role, move the pointer over the **SR Running Status** column and click **Details** to view status information and error messages of server roles on the machine. ○ Click **Error**, **Warning**, or **Good** in the **Monitoring Statistics** column to view machine and server role metrics. ○ Click **Terminal** in the **Actions** column to log on to the machine and perform operations. ○ Click **Machine Operation** in the **Actions** column to perform reboot, out-of-band reboot, or reclone operations on the machine.

1.6.4.1.4.4. View the desired state of a service

This topic describes how to check whether a service within a cluster has reached the desired state and how to view desired state details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the **Service Final Status Query** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and choose **Monitoring > Service Final Status Query**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and choose **Monitoring > Service Final Status Query** in the Actions column.
3. View information on the **Service Final Status Query** page.

Parameter	Description
Project Name	The project to which the cluster belongs.

Parameter	Description
Cluster Name	The name of the cluster.
Head Version Submitted At	The time when the HEAD version is submitted.
Head Version Analysis	The status of desired state analysis. During desired state analysis, Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to specific change contents. Desired state analysis can be in one of the following states: ○ **Preparing**: No new version is detected. ○ **Waiting**: The latest version has been detected, but the analysis module has not started. ○ **Doing**: The application to be changed is being analyzed. ○ **done**: The desired state analysis succeeds. ○ **Failed**: The desired state analysis fails to parse change contents. Apsara Infrastructure Management Framework can obtain change contents of server roles in the latest version only when the desired state analysis is in the **done** state.
Cluster Rolling Status	Specifies whether the cluster has reached the desired state. If a rolling task is running, its task information is displayed.
Cluster Machine Final Status Statistics	The status of all machines within the cluster. Click **View Details** to go to the **Cluster Operation and Maintenance Center** page and view machine details. For more information about the operations and maintenance (O&M) center, see View the cluster operation and maintenance center.
Final Status of Cluster SR Version	The desired state of services within the cluster. **Note** This section includes only the services that have not reached the desired state due to version inconsistency or status exceptions. For other services that fail to reach the desired state due to machine errors, see desired state information of machines within the cluster.
Final Status of SR Version	The number of machines that have not reached the desired state. The number is displayed if server roles have rolling tasks.

1.6.4.1.4.5. View operations logs

This topic describes how to view differences between Git versions from operation logs.

Procedure

1. Log on to Apsara Infrastructure Management Framework.

2. Use one of the following methods to go to the **Cluster Operation Logs** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over the icon next to the target cluster and choose **Monitoring > Operation Logs**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and choose **Monitoring > Operation Logs** in the Actions column.
3. On the **Cluster Operation Logs** page, click **Refresh** in the upper-right corner to view the Git version, description, submission information, and task status.
4. (Optional)On the **Cluster Operation Logs** page, view differences between versions.
   i. Find the target operation log and click **View Release Changes** in the **Actions** column.
   ii. On the **Version Difference** page, configure the following parameters:
      - **Select Base Version**: Select a basic version.
      - **Configuration Type**: Select **Extended Configuration** or **Cluster Configuration**. **Extended Configuration** allows you to view differences between the merging results of cluster and template configurations. **Cluster Configuration** allows you to view differences between cluster configurations.
   iii. Click **Obtain Difference**.

   Difference files are displayed.
   iv. Click each difference file to view its difference details.

1.6.4.1.5. Service operations

This topic describes the actions about service operations.

1.6.4.1.5.1. View the service list

The service list allows you to view the list of all services and the related information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Service Operations**.
3. View the information on the **Service Operations** page.

Item	Description
Service	The service name.
Service Instances	The number of service instances in the service.
Service Configuration Templates	The number of service configuration templates.

Item	Description
Monitoring Templates	The number of monitoring templates.
Service Schemas	The number of service configuration validation templates.
Actions	Click **Management** to view the service instances, service templates, monitoring templates, monitoring instances, service schemas, and detection scripts.

1.6.4.1.5.2. View dashboard information of a service instance

This topic describes how to view the basic information and related statistics of a service instance.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, click the **S** tab.
3. (Optional)Enter a service name in the search box to search for the service.
4. Click the service name to view service instances of the service.
5. Move the pointer over the icon next to the target service instance and select **Dashboard**.
6. View information on the **Service Instance Information Dashboard** page.

Parameter	Description

Parameter	Description
Service Instance Summary	The basic information about the service instance. ○ **Cluster Name**: the name of the cluster where the service instance is deployed. ○ **Service Name**: the name of the service to which the service instance belongs. ○ **Actual Machines**: the number of machines that are deployed in the current environment. ○ **Expected Machines**: the number of machines that are expected for the service instance. ○ **Target Total Server Roles**: the number of server roles that are expected for the service instance. ○ **Actual Server Roles**: the number of server roles that are deployed in the current environment. ○ **Template Name**: the name of the service template that is used by the service instance. ○ **Template Version**: the version of the service template that is used by the service instance. ○ **Schema**: the name of the service schema that is used by the service instance. ○ **Monitoring System Template**: the name of the Monitoring System template that is used by the service instance.
Server Role Statuses	The status of server roles in the service instance.
Machine Statuses for Server Roles	The status of machines where server roles are deployed.
Service Monitoring Information	○ **Monitored Item**: the name of the metric. ○ **Level**: the level of the metric. ○ **Description**: the description of the metric. ○ **Updated At**: the time when the data is updated.
Service Alert Status	○ **Alert Name** ○ **Instance Information** ○ **Alert Start** ○ **Alert End** ○ **Alert Duration** ○ **Severity Level** ○ **Occurrences**: the number of occurrences of the alert.

Parameter	Description
Server Role List	○ **Server Role** ○ **Current Status** ○ **Expected Machines** ○ **Machines In Final Status** ○ **Machines Going Offline** ○ **Rolling Task Status** ○ **Time Used**: the time that is used for the execution of rolling tasks. ○ **Actions**: Click **Details** to go to the View the server role dashboard page.
Service Alert History	○ **Alert Name** ○ **Alert Time** ○ **Instance Information** ○ **Severity Level** ○ **Contact Group**
Service Dependencies	The dependency configuration of service instances and server roles, and the desired state information of dependency services or server roles. ○ **Server Role**: the name of the server role. ○ **Dependent Service**: the service on which the server role depends. ○ **Dependent Server Role**: the server role on which the server role depends. ○ **Dependent Cluster**: the cluster where the dependency server role is deployed. ○ **Dependency in Final Status**: specifies whether the dependency server role has reached the desired state.

1.6.4.1.5.3. View the server role dashboard

The server role dashboard allows you to view the statistics of a server role.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, click the S tab.
3. (Optional)Enter the service name in the search box. Services that meet the search condition are displayed.
4. Click a service name and then service instances in the service are displayed in the lower-left corner.
5. Move the pointer over at the right of a service instance and then select **Dashboard**.
6. In the **Server Role List** section of the **Service Instance Information Dashboard** page, click **Details** in the **Actions** column.

7. View the information on the **Server Role Dashboard** page.

Item	Description
Server Role Summary	Displays the basic information of the server role as follows: ○ **Project Name**: the name of the project to which the server role belongs. ○ **Cluster Name**: the name of the cluster to which the server role belongs. ○ **Service Instance**: the name of the service instance to which the server role belongs. ○ **Server Role**: the server role name. ○ **In Final Status**: whether the server role reaches the final status. ○ **Expected Machines**: the number of expected machines. ○ **Actual Machines**: the number of actual machines. ○ **Machines Not Good**: the number of machines whose status is not Good. ○ **Machines with Role Status Not Good**: the number of server roles whose status is not Good. ○ **Machines Going Offline**: the number of machines that are going offline. ○ **Rolling**: whether a running rolling task exists. ○ **Rolling Task Status**: the current status of the rolling task. ○ **Time Used**: the time used for running the rolling task.
Machine Final Status Overview	The statistical chart of the current status of the server role.
Server Role Monitoring Information	○ **Updated At**: the time when the data is updated. ○ **Monitored Item**: the name of the monitored item. ○ **Level**: the level of the monitored item. ○ **Description**: the description of the monitored item.

Item	Description
Machine Information	○ **Machine Name**: the hostname of the machine. ○ **IP**: the IP address of the machine. ○ **Machine Status**: the machine status. ○ **Machine Action**: the action that the machine is performing. ○ **Server Role Status**: the status of the server role. ○ **Server Role Action**: the action that the server role is performing. ○ **Current Version**: the current version of the server role on the machine. ○ **Target Version**: the expected version of the server role on the machine. ○ **Error Message**: the exception message. ○ **Actions**: ■ Click **Terminal** to log on to the machine and perform operations. ■ Click **Restart** to restart the server roles on the machine. ■ Click **Details** to go to the **Machine Details** page. For more information about the machine details, see View the machine dashboard. ■ Click **Machine System View** to go to the **Machine Info Report** page. For more information about the machine info report, see Machine info report. ■ Click **Machine Operation** to restart, out of band restart, or clone the machine again.
Server Role Monitoring Information of Machines	○ **Updated At**: the time when the data is updated. ○ **Machine Name**: the machine name. ○ **Monitored Item**: the name of the monitored item. ○ **Level**: the level of the monitored item. ○ **Description**: the description of the monitored item.
VM Mappings	The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster. ○ **VM**: the hostname of the virtual machine. ○ **Currently Deployed On**: the hostname of the physical machine where the virtual machine is currently deployed. ○ **Target Deployed On**: the hostname of the physical machine where the virtual machine is expected to be deployed.

Item	Description
Service Dependencies	The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role. ○ **Dependent Service**: the service on which the server role depends. ○ **Dependent Server Role**: the server role on which the server role depends. ○ **Dependent Cluster**: the cluster to which the dependent server role belongs. ○ **Dependency in Final Status**: whether the dependent server role reaches the final status.

1.6.4.1.6. Machine operations

This topic describes the actions about machine operations.

1.6.4.1.6.1. View the machine dashboard

The machine dashboard allows you to view the statistics of a machine.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, click the **C** tab.
3. (Optional)On the **Machine** tab in the lower-left corner, enter the machine name in the search box. Machines that meet the search condition are displayed.
4. Move the pointer over at the right of a machine and then select **Dashboard**.
5. On the **Machine Details** page, view all the information of this machine. For more information, see the following table.

Item	Description
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
DISK Usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-System	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
DISK IO-System	The statistical table of the disk input and output.

Item	Description
Machine Summary	○ **Project Name**: the name of the project to which the machine belongs. ○ **Cluster Name**: the name of the cluster to which the machine belongs. ○ **Machine Name**: the machine name. ○ **SN**: the serial number of the machine. ○ **IP**: the IP address of the machine. ○ **IDC**: the data center of the machine. ○ **Room**: the room in the data center where the machine is located. ○ **Rack**: the rack where the machine is located. ○ **Unit in Rack**: the location of the rack. ○ **Warranty**: the warranty of the machine. ○ **Purchase Date**: the date when the machine is purchased. ○ **Machine Status**: the running status of the machine. ○ **Status**: the hardware status of the machine. ○ **CPUs**: the number of CPUs for the machine. ○ **Disks**: the disk size. ○ **Memory**: the memory size. ○ **Manufacturer**: the machine manufacturer. ○ **Model**: the machine model. ○ **os**: the operating system of the machine. ○ **part**: the disk partition.
Server Role Status of Machine	The distribution of the current status of all server roles on the machine.
Machine Monitoring Information	○ **Monitored Item**: the name of the monitored item. ○ **Level**: the level of the monitored item. ○ **Description**: the description of the monitored contents. ○ **Updated At**: the time when the monitoring information is updated.

Item	Description
Machine Server Role Status	○ Service Instance ○ Server Role ○ Server Role Status ○ Server Role Action ○ Error Message ○ Target Version ○ Current Version ○ Actual Version Update Time ○ Actions: ■ Click Details to go to the Server Role Dashboard page. For more information about the server role dashboard, see View the server role dashboard. ■ Click Restart to restart the server roles on the machine.
Application Status in Server Roles	○ Application Name: the application name. ○ Process Number ○ Status: the application status. ○ Current Build ID: the ID of the current package version. ○ Target Build ID: the ID of the expected package version. ○ Git Version ○ Start Time ○ End Time ○ Interval: the interval between the time when Apsara Infrastructure Management Framework detects that the process exits and the time when Apsara Infrastructure Management Framework repairs the process. ○ Information Message: the normal output logs. ○ Error Message: the abnormal logs.

1.6.4.1.7. Monitoring center

You can view the alert status, alert rules, and alert history in the monitoring center.

1.6.4.1.7.1. Modify an alert rule

You can modify an alert rule based on the actual business requirements.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Service Operations**.

3. (Optional)Enter the service name in the search box.
4. Find the service and then click **Management** in the **Actions** column.
5. Click the **Monitoring Template** tab.
6. Find the monitoring template that you are about to edit and then click **Edit** in the **Actions** column.
7. Configure the monitoring parameters based on actual conditions.
8. Click **Save Change**.

   Wait about 10 minutes. The monitoring instance is automatically deployed. If the status becomes Successful and the deployment time is later than the modified time of the template, the changes are successfully deployed.

1.6.4.1.7.2. View the status of a monitoring instance

After a monitoring instance is deployed, you can view the status of the monitoring instance.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Operations > Service Operations**.
3. (Optional)Enter the service name in the search box.
4. Find the service and then click **Management** in the **Actions** column.
5. Click the **Monitoring Instance** tab.

   In the **Status** column, view the current status of the monitoring instance.

1.6.4.1.7.3. View the alert status

This topic describes how to view the alerts related to different services and the alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert Status**.
3. (Optional)Search for an alert by service name, cluster name, alert name, or alert time range.
4. View alert details on the **Alert Status** page. The following table describes the related parameters.

Parameter	Description
Service	The name of the service.
Cluster	The name of the cluster where the service is deployed.
Instance	The name of the monitored instance. Click the name of an instance to view the alert history of the instance.
Alert Status	Two alert states are available, which are Normal and Alerting.

Parameter	Description
Alert Level	Alerts are divided into five levels in descending order of severity: ○ P0: an alert that has been cleared ○ P1: an urgent alert ○ P2: a major alert ○ P3: a minor alert ○ P4: a reminder alert
Alert Name	The name of the alert. Click the name of an alert to view alert rule details.
Alert Time	The time when the alert is triggered and how long the alert lasts.
Actions	Click **Show** to view the data before and after the alert time.

1.6.4.1.7.4. View alert rules

This topic describes how to view alert rules.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert Rules**.
3. (Optional)Search for alert rules by service name, cluster name, or alert name.
4. View alert rules on the **Alert Rules** page. The following table describes the related parameters.

Parameter	Description
Service	The name of the service.
Cluster	The name of the cluster where the service is deployed.
Alert Name	The name of the alert.
Alert Conditions	The conditions that trigger the alert.
Periods	The frequency at which the alert rule is executed.
Alert Contact	The groups and members to notify when the alert is triggered.
Status	The status of the alert rule. ○ **Running**: Click it to stop the alert rule. ○ **Stopped**: Click it to execute the alert rule.

1.6.4.1.7.5. View the alert history

This topic describes how to view the historical alerts related to different services and the alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert History**.
3. (Optional)Search for an alert by service name, cluster name, or alert time range.
4. View the alert history on the **Alert History** page. The following table describes the related parameters.

Parameter	Description
Service	The name of the service to which the alert belongs.
Cluster	The name of the cluster where the service is deployed.
Alert Instance	The name of the instance where the alert is triggered.
Status	Two alert states are available, which are Normal and Alerting.
Alert Level	Alerts are divided into five levels in descending order of severity: ○ P0: an alert that has been cleared ○ P1: an urgent alert ○ P2: a major alert ○ P3: a minor alert ○ P4: a reminder alert
Alert Name	The name of the alert. Click the name of an alert to view alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members to notify when the alert is triggered.
Actions	Click **Show** to view the data before and after the alert time.

1.6.4.1.8. Tasks and deployment summary

This topic describes how to view rolling tasks, running tasks, history tasks, and deployment summary on Apsara Infrastructure Management Framework.

1.6.4.1.8.1. View rolling tasks

This topic describes how to view rolling tasks and their status.

Procedure

1. Log on to Apsara Infrastructure Management Framework.

2. In the top navigation bar, choose **Operations > Cluster Operations**.
3. Select **Rolling Tasks** to view all clusters that have rolling tasks.
4. Click **rolling** in the **Rolling** column.
5. On the **Rolling Task** page. view the change task information and change details.

   Change task parameters

Parameter	Description
Change Version	The source version of the rolling task.
Description	The description of the change.
Head Version Analysis	The status of desired state analysis. During desired state analysis, Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to specific change contents. Desired state analysis can be in one of the following states: ○ **Preparing**: No new version is detected. ○ **Waiting**: The latest version has been detected, but the analysis module has not started. ○ **Doing**: The application to be changed is being analyzed. ○ **done**: The desired state analysis succeeds. ○ **Failed**: The desired state analysis fails to parse change contents. Apsara Infrastructure Management Framework can obtain change contents of server roles in the latest version only when the desired state analysis is in the **done** state.
Blocked Server Role	The server role that is blocked by dependencies in the rolling task.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.
Actions	Click **View Difference** to go to the **Version Difference** page. For more information, see View operation logs. Click **Stop** to terminate the rolling task. Click **Pause** to suspend the rolling task.

   Change details parameters

Parameter	Description
Service Name	The name of the service that has changes.

Parameter	Description
Status	The current status of the service. The rolling status of a service is an aggregation result of rolling statuses of multiple server roles. Services can be in one of the following states: ○ **succeeded**: A task succeeds. ○ **blocked**: A task is blocked. ○ **failed**: A task fails.
Server Role Status	The status of the server role. Click **>** to the left of a service name to view the rolling task status of each server role in the service. Server roles can be in one of the following states: ○ **Downloading**: A task is being downloaded. ○ **Rolling**: A rolling task is in progress. ○ **RollingBack**: A rolling task fails and is performing rollback.
Depend On	The services on which the service depends, or the server roles on which the server role depends.
Actions	Click **Stop** to terminate the change of the server role. Click **Pause** to suspend the change of the server role.

1.6.4.1.8.2. View running tasks

This topic describes how to view running tasks.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > Running Tasks**.
3. (Optional)Search for running tasks by cluster name, server role name, task status, task submitter, Git version, or time range.
4. Find the target task, move the pointer over the **Rolling Task Status** column, and then click **View Tasks** to go to the **Rolling Task** page. For more information about rolling task details, see View rolling tasks.

1.6.4.1.8.3. View historical tasks

This topic describes how to view historical tasks.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > History Tasks**.
3. (Optional)Search for historical tasks by cluster name, Git version, submitter, or time range.
4. Find the target task and click **Details** in the **Actions** column to go to the **Rolling Task** page. For more information about rolling task details, see View rolling tasks.

1.6.4.1.8.4. View the deployment summary

On the **Deployment Summary** page, you can view the deployment conditions of clusters, services, and server roles in all projects on Apsara Infrastructure Management Framework.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > Deployment Summary**.
   - View the deployment status and the duration of a certain status for each project.
     - Gray: wait to be deployed. It indicates that some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed.
     - Blue: being deployed. It indicates that the project has not reached the final status for one time yet.
     - Green: has reached the final status. It indicates that all clusters in the project have reached the final status.
     - Orange: not reaches the final status. It indicates that a server role does not reach the final status for some reason after the project reaches the final status for the first time.
   - Configure the global clone switch.
     - **normal**: Clone is allowed.
     - **block**: Clone is forbidden.
   - Configure the global dependency switch.
     - **normal**: All configured dependencies are checked.
     - **ignore**: The dependency is not checked.
     - **ignore_service**: None of the service-level dependencies, including the server role dependencies across services, are checked, and only the server role-level dependencies are checked.
3. Click the **Deployment Details** tab to view the deployment details.

   For more information, see the following table.

Item	Description

Item	Description
Status Statistics	The general statistics of deployment conditions, including the total number of projects that are currently available. Click each status to display the projects in the corresponding status in the list. The projects have five deployment statuses: ○ **Final**: All the clusters in the project have reached the final status. ○ **Deploying**: The project has not reached the final status for one time yet. ○ **Waiting**: Some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed. ○ **Non-final**: A server role does not reach the final status for some reason after the project reaches the final status for the first time. ○ **Inspector Warning**: An error is detected on service instances in the project during the inspection.
Start Time	The time when Apsara Infrastructure Management Framework starts the deployment.
Progress	The proportion of server roles that reach the final status to all the server roles in the current environment.
Deployment Status	The time indicates the deployment duration for the following statuses: **Final**, **Deploying**, **Waiting**, and **Inspector Warning**. The time indicates the duration before the final status is reached for the **Non-final** status. Click the time to view the details.
Deployment Progress	The proportion of clusters, services, and server roles that reach the final status to the total clusters, services, and server roles in the project. Move the pointer over the blank area at the right of the data of roles and then click **Details** to view the deployment statuses of clusters, services, and server roles. The deployment statuses are indicated by icons, which are the same as those used for status statistics.
Resource Application Progress	**Total** indicates the total number of resources related to the project. ○ **Done**: the number of resources that have been successfully applied for. ○ **Doing**: the number of resources that are being applied for and retried. The number of retries (if any) is displayed next to the number of resources. ○ **Block**: the number of resources whose applications are blocked by other resources. ○ **Failed**: the number of resources whose applications failed.
Inspector Error	The number of inspection alerts for the current project.

Item	Description
Monitoring Information	The number of alerts generated for the machine monitor and the machine server role monitor in the current project.
Dependency	Click the icon to view the project services that depend on other services, and the current deployment status of the services that are depended on.

1.6.4.1.9. Reports

The system allows you to search for and view reports based on your business needs, and add commonly used reports to your favorites.

1.6.4.1.9.1. View reports

The **Reports** menu allows you to view the statistical data.

Context

You can view the following reports on Apsara Infrastructure Management Framework.

- System reports: default and common reports in the system.
- All reports: includes the system reports and custom reports.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You can go to the report list in the following three ways:
   - In the top navigation bar, choose **Reports > System Reports**.
   - In the top navigation bar, choose **Reports > All Reports**.
   - In the left-side navigation pane, click the **R** tab. Move the pointer over at the right of **All Reports** and then select **View**.

   See the following table for the report descriptions.

Item	Description
Report	The report name. Move the pointer over next to **Report** to search for reports by report name.
Group	The group to which the report belongs. Move the pointer over next to **Group** to filter reports by group name.
Status	Indicates whether the report is published.
Public	Indicates whether the report is public.
Created By	The person who creates the report.

Item	Description
Published At	The published time and created time of the report.
Actions	Click **Add to Favorites** to add this report to your favorites. Then, you can view the report by choosing **Reports > Favorites** in the top navigation bar or moving the pointer over at the right of **Favorites** on the **R** tab in the left-side navigation pane and then selecting **View**.

3. (Optional)Enter the name of the report that you are about to view in the search box.
4. Click the report name to go to the corresponding report details page.

   For more information about the reports, see Appendix.

1.6.4.1.9.2. Add a report to favorites

You can add common reports to favorites. Then, find them quickly on the **Favorites** page.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You can go to the report list in the following three ways:
   - In the top navigation bar, choose **Reports > System Reports**.
   - In the top navigation bar, choose **Reports > All Reports**.
   - In the left-side navigation pane, click the **R** tab. Move the pointer over at the right of **All Reports** and then select **View**.
3. (Optional)Enter the name of the report that you are about to add to favorites in the search box.
4. At the right of the report, click **Add to Favorites** in the **Actions** column.
5. In the displayed **Add to Favorites** dialog box, enter tags for the report.
6. Click **Add to Favorites**.

1.6.4.1.10. Metadata operations

In this version, you can use only command lines to perform metadata operations.

1.6.4.1.10.1. Common parameters

Common parameters consist of the common request parameters and the common response parameters.

Common request parameters

Common request parameters are request parameters that you must use when you call each API.

Parameter descriptions

Name	Type	Required	Description
Action	String	Yes	The API name. For more information about the valid values, see APIs on the control side and APIs on the deployment side.

Common response parameters

Each time you send a request to call an API, the system returns a unique identifier, regardless of whether the call is successful.

Parameter descriptions

Name	Type	Required	Description
RequestID	String	Yes	The request ID. The request ID is returned, regardless of whether the API call is successful.
Code	String	No	The error code.
Message	String	No	The reason of failure, which appears when the API call fails.
Result	The type varies with the request, which is subject to the returned result of the specific API.	No	The request result, which appears when the API call is successful.

> **Note**
>
> - If the API call is successful, RequestID is returned and the HTTP return code is 200.
> - If the API call fails, RequestID, Code, and Message are returned and the HTTP return code is 4xx or 5xx.

Instance types

```
{
  "rds.mys2.small":{
    "cpu":2,
    "memory":4096,
    "disk":51200,
    "max_connections":60
  },
  "rds.mys2.mid":{
    "cpu":4,
    "memory":4096,
    "disk":51200,
    "max_connections":150
  },
  "rds.mys2.standard":{
    "cpu":6,
    "memory":4096,
    "disk":51200,
    "max_connections":300
  },
  "rds.mys2.large":{
    "cpu":8,
    "memory":7200,
    "disk":102400,
    "max_connections":600
  },
  "rds.mys2.xlarge":{
    "cpu":9,
    "memory":12000,
    "disk":204800,
    "max_connections":1500
  },
  "rds.mys2.2xlarge":{
    "cpu":10,
    "memory":20000,
    "disk":512000,
    "max_connections":2000
  }
}
```

1.6.4.1.10.2. Make API requests

This topic describes how to make API requests from the control and deployment sides.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Machine Operations**.
3. Select a project from the drop-down list or enter a cluster or machine name to search for the machine.
4. Make API requests.

- Make API requests from the control side
  a. Find the machine and click **Terminal** in the **Actions** column to log on to the machine.
  b. On the command line, enter the following command and press the Enter key to obtain the value of intranet-domain:

```
grep 'intranet-domain' /cloud/app/tianji/TianjiClient#/service_manager/current/conf.global/kv.json
```

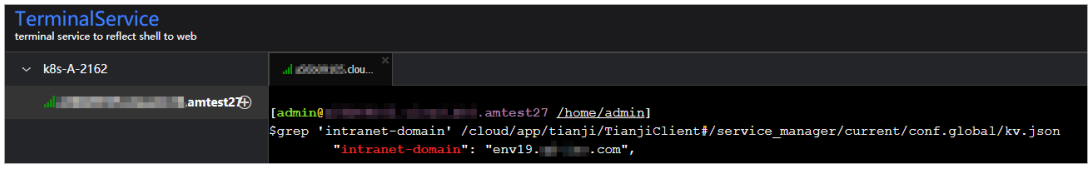

  c. Use one of the following methods to make API requests from the control side. ListInstance is used in the example.
     - GET request

```
curl 'xdb-master.xdb.{intranet-domain}:15678?Action=ListInstance'
```

     - POST request

```
curl 'xdb-master.xdb.{intranet-domain}:15678' -X POST -d '{"Action":"ListInstance"}'
```

- Make API request from the deployment side
  a. Find the machine and record the IP address in the Hostname column.
  b. Use one of the following methods to make API requests from the deployment side. CheckState is used in the example.

     Assume that the IP address of the machine is 127.0.XX.XX.

     - GET request

```
curl '127.0.XX.XX:18765?Action=CheckState&Port=3606'
```

     - POST request

```
curl '127.0.XX.XX:18765' -X POST -d '{"Action":"CheckState","Port":3606}'
```

1.6.4.1.10.3. APIs on the control side

DescribeInstance

You can call this operation to query an instance.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeInstance.

Parameter	Type	Required	Description
InstanceName	String	Yes	The name of the instance.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
InstanceID	Integer	Yes	The ID of the instance.
InstanceName	String	Yes	The name of the instance.
Domain	String	Yes	The domain name.
Port	Integer	Yes	The port of the instance.
PaxosPort	Integer	Yes	The communication port between instance nodes.
InstanceDir	String	Yes	The directory of the instance.
Level	String	Yes	The instance type.
User	String	Yes	The username.
Password	String	Yes	The password.
Config	String	No	The custom my.cnf configuration of the instance in the JSON format.
LeaderIP	String	No	The IP address of the leader node.
ActionName	String	Yes	The name of the operation.
ActionStatus	String	Yes	The status of the operation.
Description	String	Yes	The description of the instance.

Parameter	Type	Required	Description
IsDeleted	Integer	No	Indicates whether the instance was deleted. Valid values: 0 and 1. 0 indicates no, and 1 indicates yes.
NodeList	[]NodeInfo	Yes	The information of the instance node.

The following table describes the parameters of NodeInfo.

Parameter	Type	Required	Description
InstanceID	Integer	Yes	The ID of the instance.
InstanceName	String	Yes	The name of the instance.
IP	String	Yes	The IP address of the instance node.
NodeID	Integer	Yes	The ID of the instance node.
ActionName	String	Yes	The name of the operation.
ActionStatus	String	Yes	The status of the operation.
Description	String	Yes	The description of the instance.
IsDeleted	Integer	No	Indicates whether the instance was deleted. Valid values: 0 and 1. 0 indicates no, and 1 indicates yes.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DescribeInstance&InstanceName=xdb-meta'
```

Sample success responses

```
{
  "Result": {
    "ActionName": "",
    "Level": "rds.mys2.standard",
    "InstanceID": 1,
    "LeaderIP": "10.39.XX.XX",
    "Config": "{}",
    "Description": "",
    "ActionStatus": "",
    "Domain": "xdb-meta.xdb.env8c-inc.com",
    "PaxosPort": 11606,
    "InstanceName": "xdb-meta",
    "User": "xdb",
    "Password": "xdb",
    "Port": 3606,
    "IsDeleted": 0,
    "InstanceDir": "/apsarapangu/disk1/xdb/xdb_instance_3606",
    "NodeList": [
      {
        "ActionStatus": "",
        "ActionName": "",
        "Description": "",
        "InstanceID": 1,
        "IP": "10.38.XX.XX",
        "InstanceName": "xdb-meta",
        "NodeID": 1,
        "IsDeleted": 0
      },
      {
        "ActionStatus": "",
        "ActionName": "",
        "Description": "",
        "InstanceID": 1,
        "IP": "10.39.XX.XX",
        "InstanceName": "xdb-meta",
        "NodeID": 2,
        "IsDeleted": 0
      },
      {
        "ActionStatus": "",
        "ActionName": "",
        "Description": "",
        "InstanceID": 1,
        "IP": "10.39.XX.XX",
        "InstanceName": "xdb-meta",
        "NodeID": 3,
        "IsDeleted": 0
      }
    ]
  },
  "RequestID": "3CFCBA07-3D87-4A99-B8C1-E861A7D1A573"
}
```

ListInstance

You can call this operation to list the basic information of instances.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ListInstance.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
InstanceNames	String	Yes	The list of one or more instance names.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=ListInstance'
```

Sample success responses

```
{
  "Result": {
    "InstanceNames": [
      "xdb-meta",
      "xdb-instance-1",
      "xdb-instance-2",
      "xdb-instance-3"
    ]
  },
  "RequestID": "A921B8C7-C833-417C-B46A-E0CE129EBE48"
}
```

CreateInstance

You can call this operation to create an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CreateInstance.
InstanceName	String	Yes	The name of the instance.
User	String	Yes	The username.
Password	String	Yes	The password.
Level	String	Yes	Instance types
Config	String	No	The custom my.cnf configuration of the instance in the JSON format. The key must be the same as the value of the field in my.cnf. The value must be of the String type.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=CreateInstance&InstanceName=xdb-instance-1&User=admin&password=xdb&Level=rds.mys2.small'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "8BCB3B39-6140-459F-B283-F83C03ADC3CA"
}
```

DeleteInstance

You can call this operation to delete an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DeleteInstance.
InstanceName	String	Yes	The name of the instance.

Response parameters

Common response parameters

Examples

Sample requests

```
curl '127.0.XX.XX:15678? Action=DeleteInstance&InstanceName=xdb-instance-1'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "9C40CCB3-4FAB-4242-9B87-792E8154E5CD"
}
```

RestartInstance

You can call this operation to restart an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to RestartInstance.
InstanceName	String	Yes	The name of the instance.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=RestartInstance&InstanceName=xdb-instance-2'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "47277A23-5FFE-4A46-B65F-E6F2569F44E5"
}
```

UpgradeInstance

You can call this operation to upgrade the minor version of an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to UpgradeInstance.
InstanceName	String	Yes	The name of the instance.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=UpgradeInstance&InstanceName=xdb-instance-2'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "95E8B098-B04A-4BCA-BEBE-DA1D11BBAD4A"
}
```

DescribeTaskProgress

You can call this operation to query the task progress.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeTaskProgress.
RequestID	String	Yes	The ID of the request.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
Progress	String	Yes	The task progress of the instance. Valid values: pending, doing, done, and failed.
Description	String	Yes	The description of the task progress.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DescribeTaskProgress&RequestID=47277A23-5FFE-4A46-B65F-E6F2569F44E5'
```

Sample success responses

```
{
  "Result": {
    "Progress": "done",
    "Description": "Success"
  },
  "RequestID": "AC535130-F40E-4D45-BC05-0F45C8473346"
}
```

ChangeLeaderTo

You can call this operation to change the leader role of an instance to another node.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ChangeLeaderTo.
InstanceName	String	Yes	The name of the instance.
IP	String	Yes	The IP address of the machine where the new leader node is deployed.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=ChangeLeaderTo&InstanceName=xdb-instance-1&IP=10.39.XX.XX'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "37638DE5-14C1-4D2E-984F-FEA1F29C9F84"
}
```

ModifyInstanceLevel

You can call this operation to modify the instance type. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ModifyInstanceLevel.
InstanceName	String	Yes	The name of the instance.
Level	String	Yes	The new instance type.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=ModifyInstanceLevel&InstanceName=xdb-instance-1&Level=rds.mys2.mid'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "21B91211-BB09-4665-835D-9471A6F07F24"
}
```

DescribeLeader

You can call this operation to query the leader node information of an instance.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeLeader.
InstanceName	String	Yes	The name of the instance.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
LeaderIP	String	Yes	The IP address of the leader node.
Port	Integer	Yes	The port of the instance.
User	String	Yes	The username.
Password	String	Yes	The password.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DescribeLeader&InstanceName=xdb-meta'
```

Sample success responses

```
{
  "Result": {
    "LeaderIP": "10.27.XX.XX",
    "Password": "xdb",
    "Port": 3606,
    "User": "xdb"
  },
  "RequestID": "2F05EE81-DC47-478E-9CA9-9AE8CA809151"
}
```

RecreateNode

Recreates an instance node.

Description

Uses other available nodes to recreate an instance node by backup and recovery. This is an asynchronous task. You can view the task result by calling the DescribeTaskProgress API based on the RequestID in the responses.

Request parameters

For more information about the common request parameters, see Common request parameters.

Name	Type	Required	Description
Action	String	Yes	The parameter specified by the system. Value: RecreateNode
InstanceName	String	Yes	The instance name.
IP	String	Yes	The IP address of the instance node to be recreated.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=RecreateNode&InstanceName=xdb-instance-1&IP=10.39.XX.XX'
```

Sample responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "7F079E11-1DE9-4148-A9FA-683E4C58F9C2"
}
```

CreateDatabase

You can call this operation to create a database and a user.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CreateDatabase.
InstanceName	String	Yes	The name of the instance.
DBName	String	Yes	The name of the database.
User	String	Yes	The username.
Password	String	Yes	The password.

Response parameters

Common response parameters

Examples

Sample requests

```
curl '*db-master.*db.env8c-inc.com:15678? Action=CreateDatabase&InstanceName=*db-instance-1&DBNam
e=***&User=***&Password=***_password'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "A2BEF74F-5C3A-4CEF-A2B8-C14C71E36569"
}
```

DeleteDatabase

You can call this operation to delete a database.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DeleteDatabase.
InstanceName	String	Yes	The name of the instance.
DBName	String	Yes	The name of the database to be deleted.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DeleteDatabase&InstanceName=xdb-instance-1&DBName=xdb'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "23F75A0A-B1D6-4341-BD5B-1A5F3FD45848"
}
```

DeleteUser

You can call this operation to delete a user.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DeleteUser.
InstanceName	String	Yes	The name of the instance.
User	String	Yes	The username.

Parameter	Type	Required	Description
Host	String	No	The range of IP addresses of hosts to which the user logs on. If you do not specify this parameter, the user is deleted from all the IP addresses.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DeleteUser&InstanceName=xdb-instance-1&User=admin&Host=10.39.XX.XX'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "6A82AFF6-2B4D-48EF-868D-BBA54667D846"
}
```

1.6.4.1.10.4. APIs on the deployment side

CheckHealth

You can call this operation to check whether an instance node assumes a leader role and is able to read and write data.

Description

You can call this operation to check whether an instance node assumes a leader role. An instance node is considered to be healthy only when it assumes a leader role and is able to read and write data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CheckHealth.
Port	Integer	Yes	The port of the instance node.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
health	Boolean	Yes	The health status.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=CheckHealth&Port=3606'
```

Sample success responses

```
{
  "Result": {
    "health": true
  },
  "RequestID": "304B69CE-1566-4E87-B618-233F40238FFF"
}
```

```
{
  "Message":"{\"health\": false}",
  "Code":"NodeNotHealth",
  "RequestID":"E939DB9B-4337-4B1C-8680-F62BEDD645DC"
}
```

CheckState

You can call this operation to check whether the status of an instance node is normal.

Description

You can call this operation to check whether the status of an instance node is normal. Two scenarios are available:

- The node assumes a leader role and is able to read and write data.
- The node assumes a follower role and is able to read data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CheckState.
Port	Integer	Yes	The port of the instance node.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
IP	String	Yes	The IP address of the instance node.
Port	Integer	Yes	The port of the instance.
Role	String	Yes	The role of the instance node.
Writeable	String	Yes	Indicates whether the instance node can write data.
Readable	String	Yes	Indicates whether the instance node can read data.
State	String	Yes	The status of the instance node. If the instance node is normal, *GOOD* is returned. Otherwise, *ERROR* is returned.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=CheckState&Port=3606'
```

Sample success responses

```
{
  "Result": {
    "Readable": true,
    "State": "GOOD",
    "Role": "Follower",
    "Port": 3606,
    "IP": "10.39.XX.XX"
  },
  "RequestID": "45A59426-46D3-4709-8DD6-CD9F243336E0"
}
```

DescribeNodeStatus

You can call this operation to query the status of an instance node.

Description

You can call this operation to query the status of an instance node. A leader node must be able to read and write data. A follower node must be able to read data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeNodeStatus.
Port	Integer	Yes	The port of the instance node. This parameter must be specified if the instance is in single_machine mode.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
IP	String	Yes	The IP address of the instance node.
Port	Integer	Yes	The port of the instance node.
Role	String	Yes	The role of the instance node.
Writeable	String	Yes	Indicates whether the node can write data.
Readable	String	Yes	Indicates whether the node can read data.
ConnectionCount	Integer/String	Yes	The number of connections. If this parameter value cannot be obtained, unknown is returned.
MaxConnectionCount	Integer/String	Yes	The maximum number of connections. If this parameter value cannot be obtained, unknown is returned.

Parameter	Type	Required	Description
ConnectionPercent	Integer/String	Yes	The usage of connections. If this parameter value cannot be obtained, unknown is returned.
QPS	Integer/String	Yes	The queries per second (QPS). If this parameter value cannot be obtained, unknown is returned.
CpuPercent	Integer/String	Yes	The CPU utilization. If this parameter value cannot be obtained, unknown is returned.
MemoryPercent	Integer/String	Yes	The memory usage. If this parameter value cannot be obtained, unknown is returned.
DiskPercent	Integer/String	Yes	The disk usage. If this parameter value cannot be obtained, unknown is returned.
State	String	Yes	The status of the instance node. If the instance node is normal, GOOD is returned. Otherwise, ERROR is returned.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=DescribeNodeStatus&Port=3606'
```

Sample success responses

```
{
  "Result": {
    "CpuPercent": 2.74,
    "IP": "10.39.XX.XX",
    "Readable": true,
    "MemoryPercent": 56.13,
    "State": "GOOD",
    "Role": "Follower",
    "MaxConnectionCount": 500,
    "ActiveThreadCount": 34,
    "Writeable": false,
    "ConnectionCount": 37,
    "DiskPercent": 3.0,
    "ConnectionPercent": 7.4,
    "QPS": 15.95,
    "Port": 3606
  },
  "RequestID": "D18328B1-78A9-4F3E-BB2E-B27AB7683C19"
}
```


ListNode

You can call this operation to query the basic information of instance nodes.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ListNode.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
Nodes	String	Yes	The list of one or more instance node names.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=ListNode'
```

Sample success responses

```
{
  "Result": {
    "Nodes": [
      "xdb-instance-1",
      "xdb-instance-2",
      "xdb-instance-3",
      "xdb-meta"
    ]
  },
  "RequestID": "3F7BB536-FA3F-4597-A3DF-E5830F5A3A21"
}
```


BackupNode

You can call this operation to back up the data of an instance node and transfer the data to a specified location. An nc command is used to specify a port for receiving the data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to BackupNode.
Port	Integer	Yes	The port of the instance.
TargetIP	String	Yes	The IP address of the location to which you want to transfer data.
TargetPort	Integer	Yes	The port of the location to which you want to transfer data.

Response parameters

Common response parameters

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=BackupNode&Port=3606&TargetIP=10.39.XX.XX&TargetPort='
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "6A82AFF6-2B4D-48EF-868D-BBA54667D846"
}
```

1.6.4.1.11. Appendix

1.6.4.1.11.1. 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.

1.6.4.1.11.2. 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.

Item	Description
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.

1.6.4.1.11.3. Machine info report

This report displays the statuses of machines and server roles on the machines.

Machine Status

Displays all the machines currently managed by Apsara Infrastructure Management Framework and their corresponding statuses. In the **Global Filter** section at the top of the page, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists, and then click **Filter** on the right to filter the data.

Item	Description
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The machine status.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status.
Status Description	The description about the machine status.

Expected Server Role List

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The name of the expected server role on the machine.

Abnormal Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

Server Role Version and Status on Machine

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.
Server Role Status	The status of the server role.
Target Version	The expected version of the server role on the machine.
Current Version	The current version of the server role on the machine.
Status Description	The description about the status.
Error Message	The exception message of the server role.

Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.
Monitored Item	The name of the monitored item.

Item	Description
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

1.6.4.1.11.4. 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.

Item	Description
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.
Action Name	The ongoing action of the server role.
Action Status	The status of the action.

1.6.4.1.11.5. Machine RMA approval pending list

Some Apsara Infrastructure Management Framework actions (such as restart) on machines and server roles can be triggered by users, but this type of actions must be reviewed and approved. This report is used to process the actions that must be reviewed and approved.

Machine

Displays the basic information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
State	The running status of the machine.
Action Name	The action on the machine.
Action Status	The status of the action on the machine.
Actions	The approval button.

Machine Serverrole

Displays the information of server roles on the pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
Serverrole	The server role name.
State	The running status of the server role.
Action Name	The action on the server role.
Action Status	The status of the action on the server role.
Actions	The approval button.

Machine Component

Displays the hard disk information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
Component	The hard disk on the machine.
State	The running status of the hard disk.
Action Name	The action on the hard disk.
Action Status	The status of the action on the hard disk.
Actions	The approval button.

1.6.4.1.11.6. Registration vars of services

This report displays values of all service registration variables.

Item	Description
Service	The service name.
Service Registration	The service registration variable.
Cluster	The cluster name.
Update Time	The updated time.

1.6.4.1.11.7. Virtual machine mappings

Use the global filter to display the virtual machines of a specific cluster.

Displays the information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.

Item	Description
Project	The project name.
Cluster	The cluster name.
VM	The hostname of the virtual machine.
Currently Deployed On	The hostname of the physical machine on which the virtual machine is currently deployed.

Item	Description
Target Deployed On	The hostname of the physical machine on which the virtual machine is expected to be deployed.

1.6.4.1.11.8. Service inspector report

Use the global filter to display the service inspection reports of a specific cluster.

Service Inspector: Data is available only for services with inspection configured.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Description	The contents of the inspection report.
Level	The level of the inspection report.

1.6.4.1.11.9. Resource application report

In the **Global Filter** section, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists and then click **Filter** on the right to display the corresponding resource application data.

Change Mappings

Item	Description
Project	The project name.
Cluster	The cluster name.
Version	The version where the change occurs.
Resource Process Status	The resource application status in the version.
Msg	The exception message.
Begintime	The start time of the change analysis.
Endtime	The end time of the change analysis.

Changed Resource List

Item	Description
Res	The resource ID.
Type	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.
Parameters	The resource parameters.
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
APP	The application of the server role.
Name	The resource name.
Type	The resource type.
Status	The resource application status.
Parameters	The resource parameters.
Result	The resource application result.
Res	The resource ID.
Reprocess Status	The status of the interaction with Business Foundation System during the VIP resource application.
Reprocess Msg	The error message of the interaction with Business Foundation System during the VIP resource application.
Reprocess Result	The result of the interaction with Business Foundation System during the VIP resource application.
Refer Version List	The version that uses the resource.

Item	Description
Error Msg	The exception message.

1.6.4.1.11.10. 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.
Description	The description of the alert.
Updated At	The update time of the alert.

1.6.4.1.11.11. Relationship of service dependency

This report displays the dependencies among server roles. Use the global filter to display the data of a specific cluster in the list.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.

Item	Description
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.

1.6.4.1.11.12. Check report of network topology

This report checks if network devices and machines have wirecheck alerts.

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Network Instance	The name of the network device.
Level	The alert level.
Description	The description about the alert information.

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Machine Name	The server (machine) name.
Level	The alert level.
Description	The description about the alert information.

1.6.4.1.11.13. Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Status	The running status of the machine.
Clone Progress	The progress of the current clone process.

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
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.

1.6.4.1.11.14. Auto healing/install approval pending report

The list structure is the same as the machine RMA approval pending list, whereas this view is used for the approval during the installation. For more information, see Machine RMA approval pending list.

1.6.4.1.11.15. Machine power on or off statuses of clusters

After a cluster starts or shuts down machines, you can view the related information in this report.

Cluster Running Statuses

If a cluster is starting or shutting down machines, the corresponding data is available in this list. No data indicates that no cluster has machines shut down.

Item	Description
Project	The project name.
Cluster	The cluster name.
Action Name	The startup or shutdown action that is being performed by the cluster.
Action Status	The status of the action.

Server Role Power On or Off Statuses

Displays the power on or off statuses of server roles in the cluster selected in the **Cluster Running Statuses** section.

Select a row in the **Cluster Running Statuses** section to display the information of the corresponding cluster in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
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.

1.6.4.2. New console

1.6.4.2.1. Introduction to Apsara Infrastructure Management Framework

This topic describes the features and terms of Apsara Infrastructure Management Framework.

1.6.4.2.1.1. 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.

1.6.4.2.1.2. 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.

1.6.4.2.1.3. 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.

1.6.4.2.2. 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 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.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 endpoint *region-id.ops.asconsole.intranet-domain-id*.com. Press the Enter key.

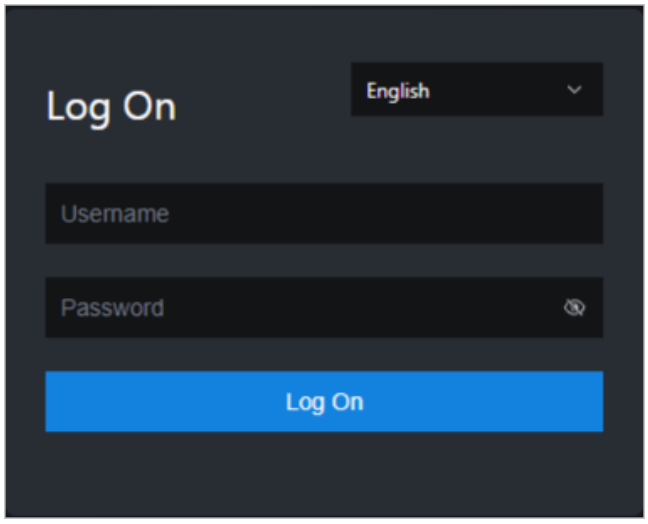

> **Note**
>
> You can select a language from the drop-down list in the upper-right corner of the page.

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 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**.

1.6.4.2.3. 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

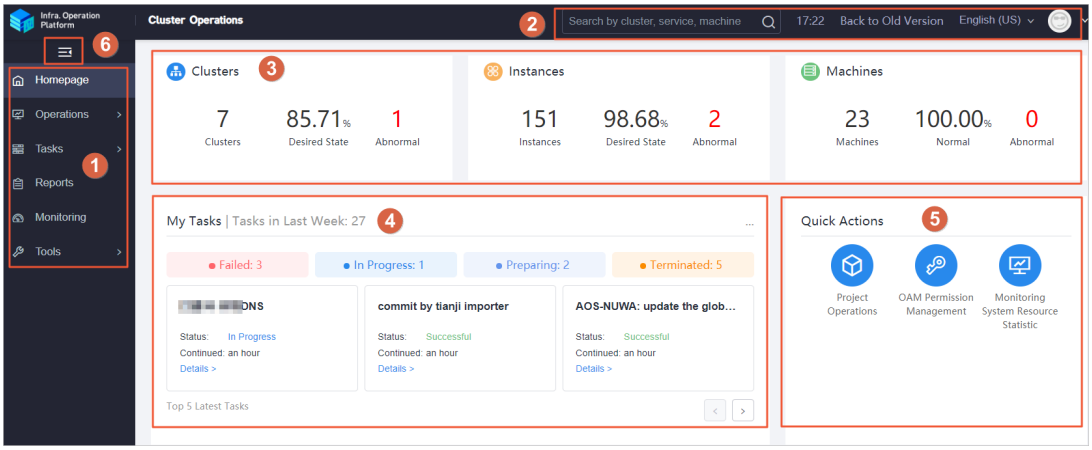

The following table describes the functional sections on the homepage.

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. • **Monitoring System Resource Statistics**: allows you to go to the Grafana console of Monitoring System. The Grafana console displays the running data of Monitoring System and facilitates your O&M operations. **Note** **Monitoring System Resource Statistics** is displayed only when Monitoring System is deployed in the environment.
⑥	Show/hide button	Allows you to expand or collapse the left-side navigation pane to narrow or enlarge the workspace.

1.6.4.2.4. Operations

1.6.4.2.4.1. 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**.

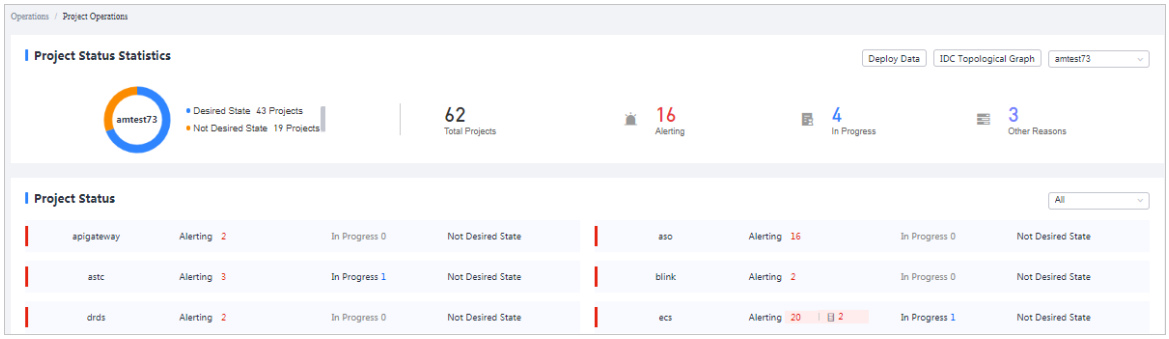

3. On the Project Operations page, perform the following operations:
   - Query a project

     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.

1.6.4.2.4.2. 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**.

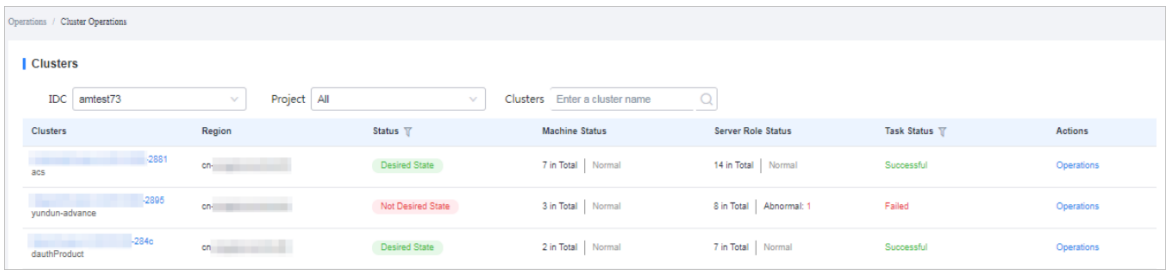

The following table describes the information displayed in the cluster list.

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.
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.

Parameter	Description
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. 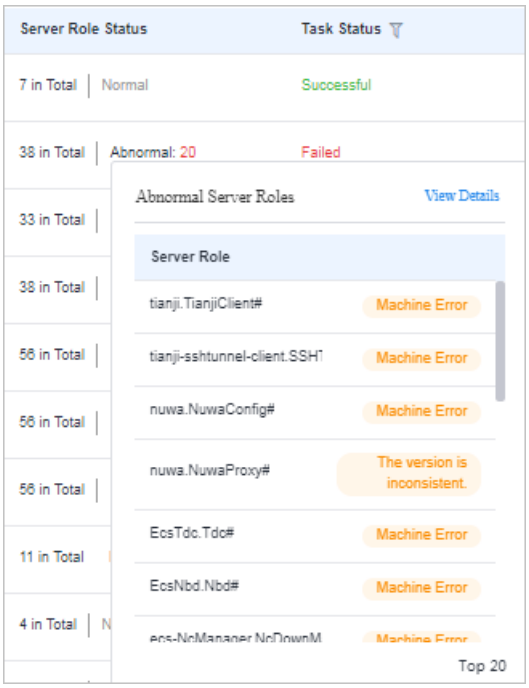
Task Status	The status of the task related to the cluster. Click the icon to filter clusters. Click the task status to view the task details.
Actions	The available operations. Click **Operations** to go to the **Cluster Details** page.

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.
	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.

Section	Parameter	Description
	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.
	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.

Section	Parameter	Description
	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.

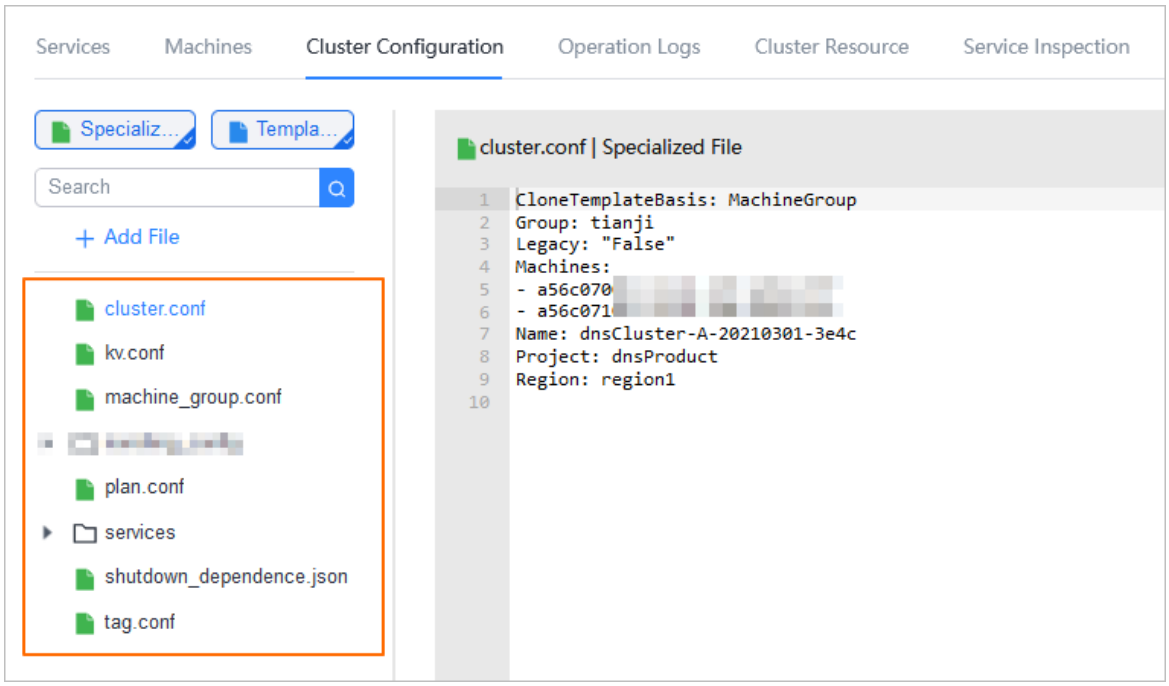

The following table describes configuration files and folders of a cluster.

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.

Parameter	Description
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**.

   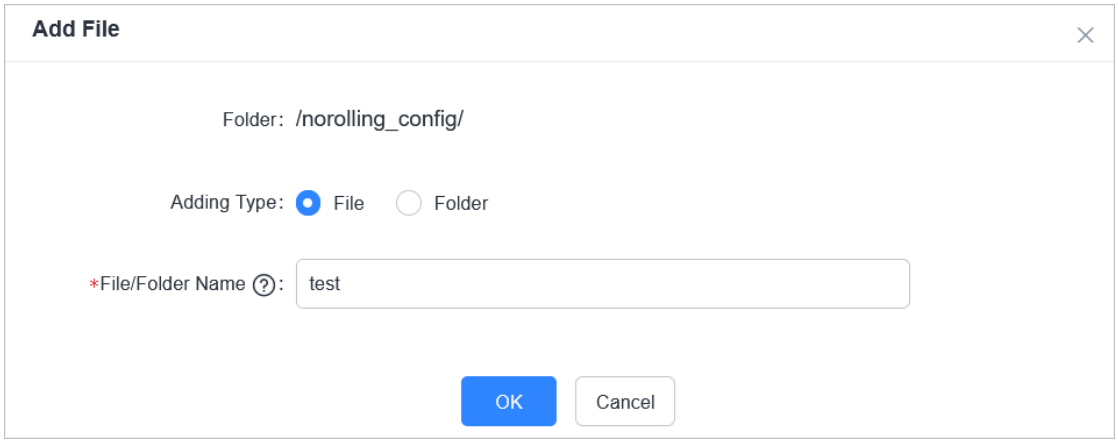

   > **Note** After you enter a folder name and click **OK**, the folder is added.

ii. Enter configuration file information into the **Cluster File** text editor. Click **Preview and Submit**.

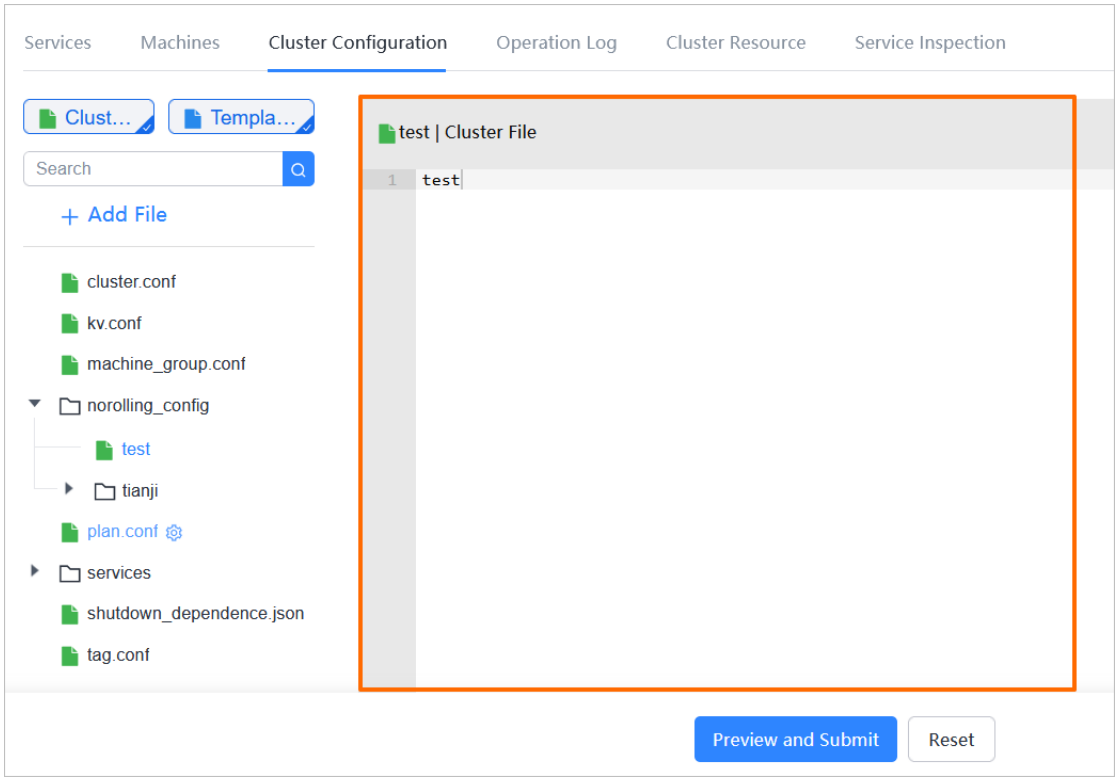

iii. In the **Confirm and Submit** dialog box, enter **Description** and click **Submit**.

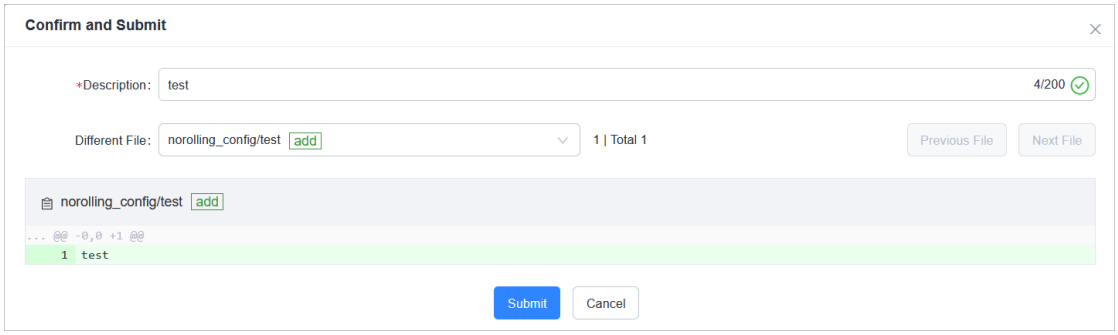

The configuration file is added. You can click the **Operation Logs** tab to view related records.

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.

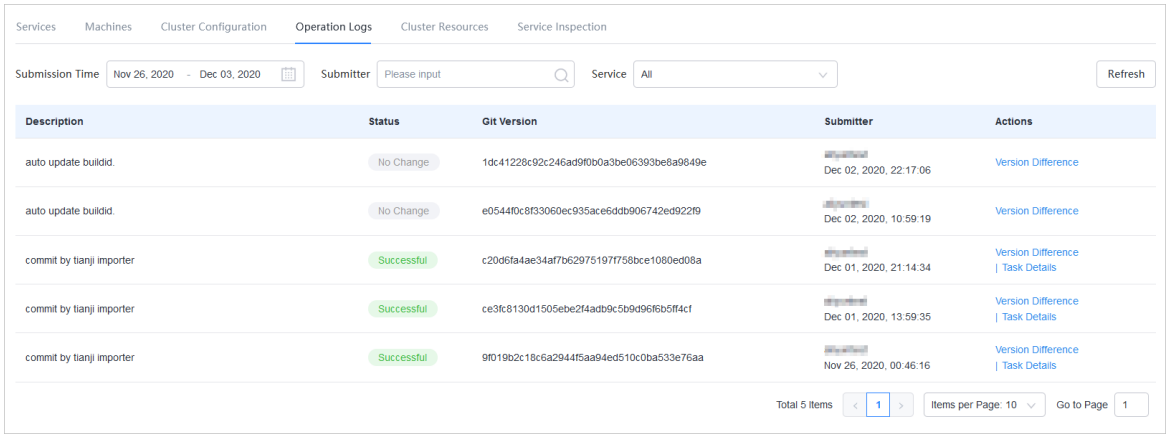

3. View the version differences on the **Operation Logs** tab.
   i. Find the operation log that you want to view and click **Version Difference** in the **Actions** column.
   ii. Set **Configuration Type** to **Show Configuration** or **Cluster Configuration**.
      - **Show Configuration**: displays 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.

1.6.4.2.4.3. 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**.

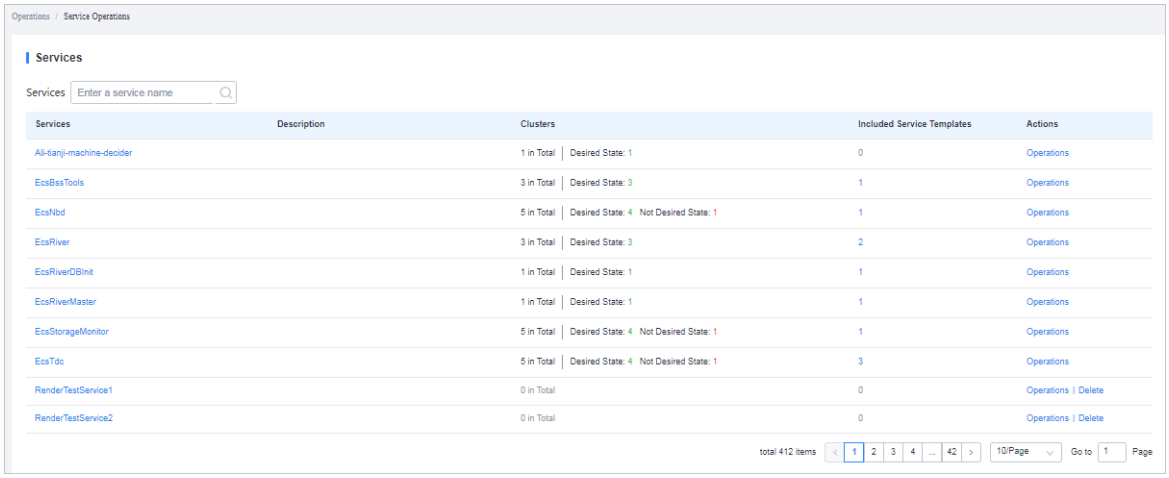

The following table describes the information displayed in the service list.

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.

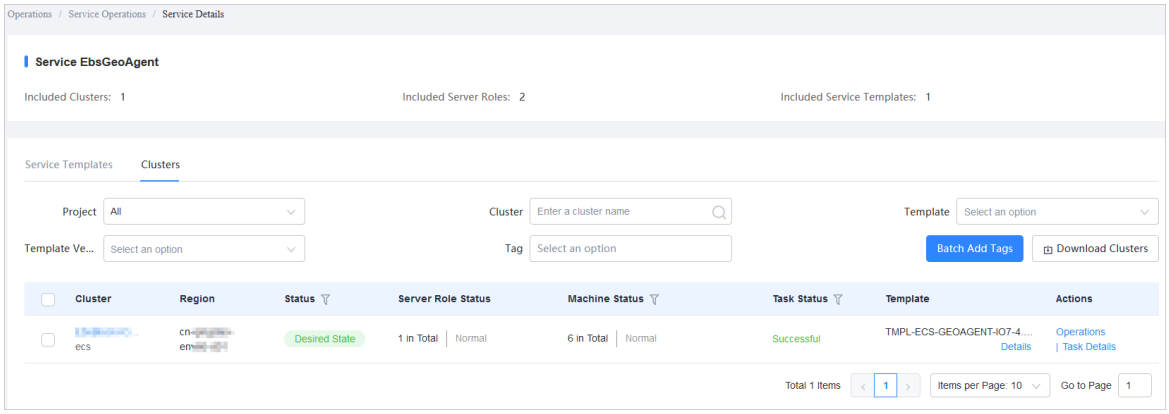

6. Select the server role that you want to view.
   - Click the **Machines** tab to view details of the server role.

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 block hardware alerts.

Background information

You must block hardware alerts in the following scenarios:

- Alerts are triggered improperly 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.

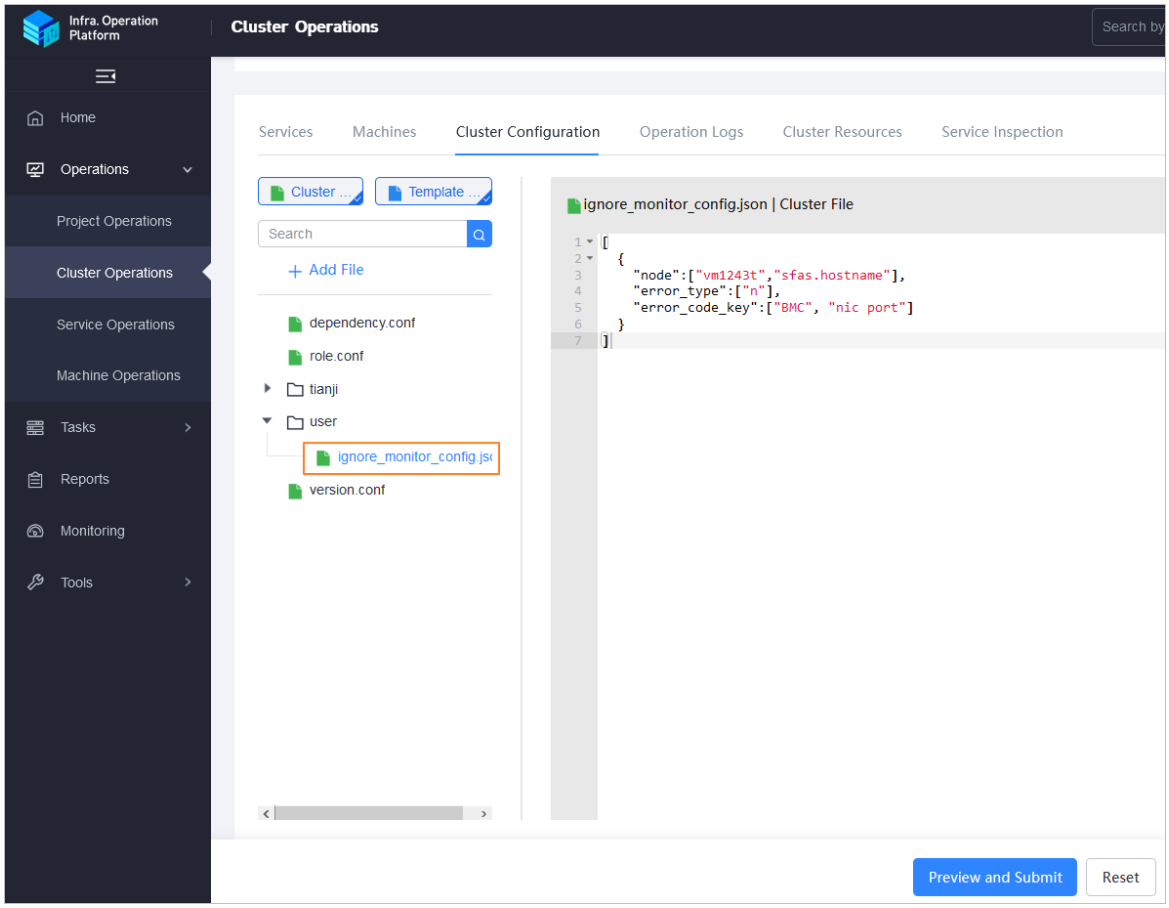

The following table describes parameters in the configuration file.

Parameter	Description	Remarks
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	Remarks
error_code_key	The keyword 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**.

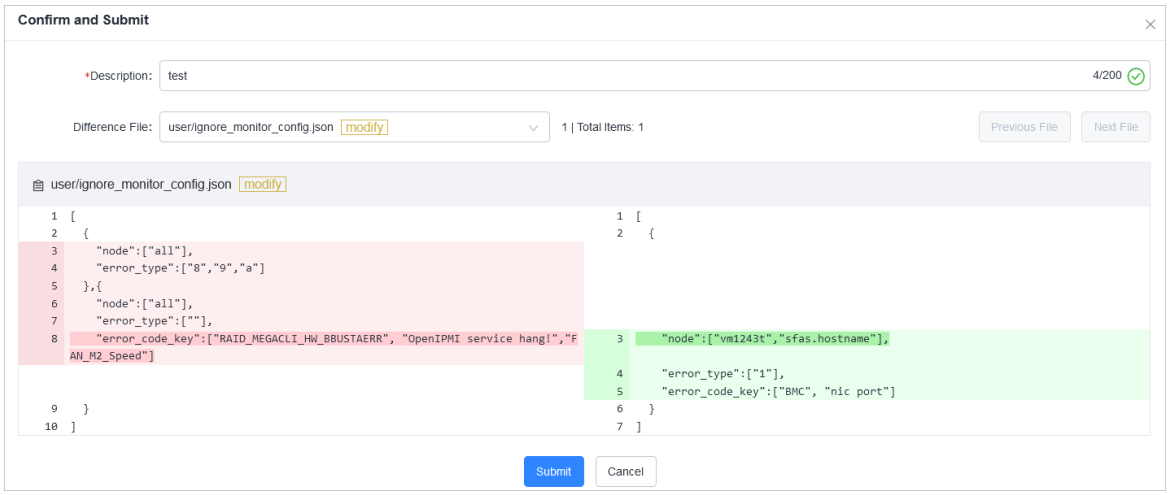

10. Click the **Operation Logs** tab to view related records.

1.6.4.2.4.4. 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.

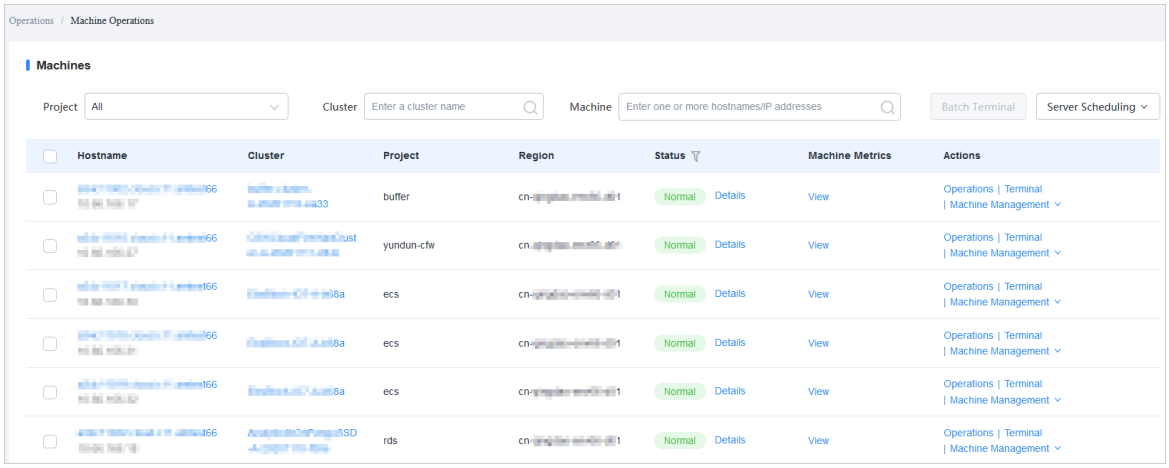

Parameter	Description
Hostname	The hostname of the machine. Click a hostname to go to the Machine Details page.
Cluster	The cluster where the machine is deployed. Click a cluster name to go to the Cluster Details page.
Status	The status of the machine. Click the icon to filter machines. Click **Details**. 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.

1.6.4.2.5. 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.

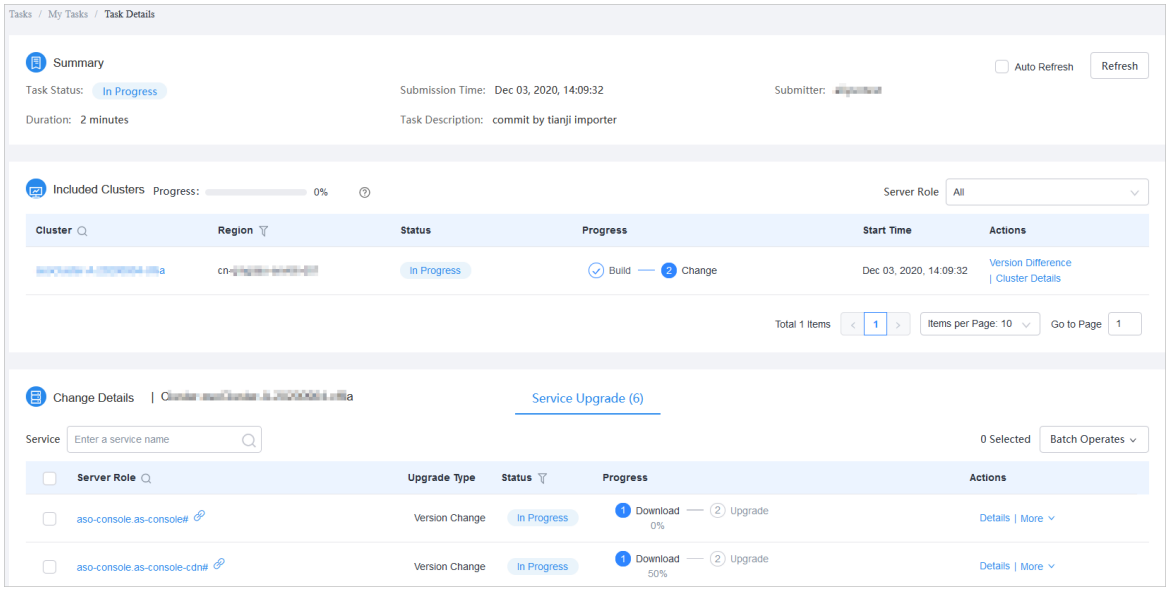

1.6.4.2.6. Reports

1.6.4.2.6.1. 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.

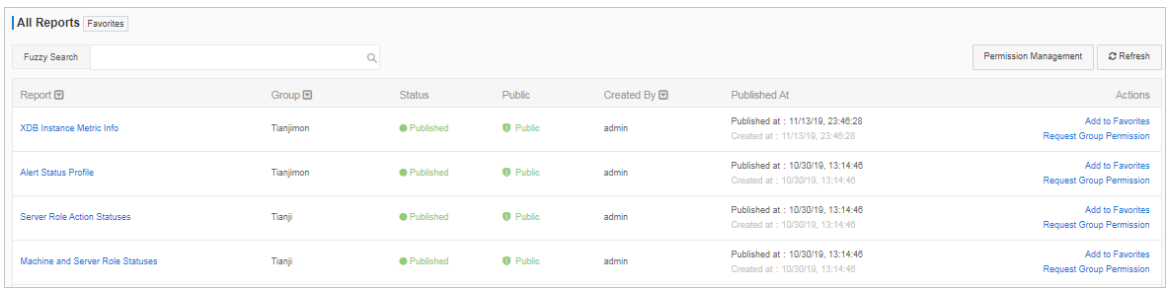

The following table describes information about reports.

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.

1.6.4.2.6.2. 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**.

1.6.4.2.7. Monitoring center

You can view the alert status, alert rules, and alert history in the monitoring center.

1.6.4.2.7.1. 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.

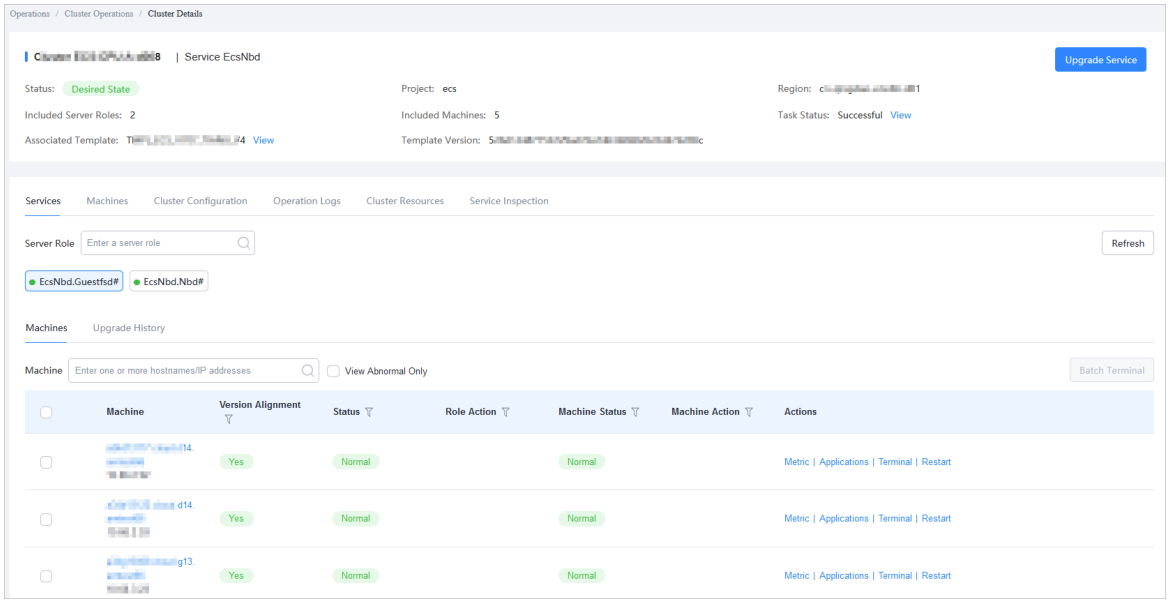

1.6.4.2.7.2. View the alert status

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**.

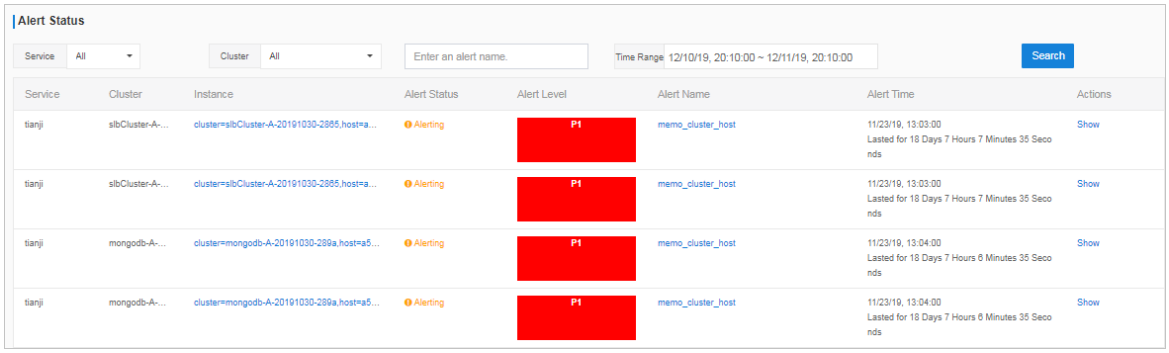

4. (Optional)Search for an alert by service name, cluster name, alert time range, or alert name.
5. View alert details on the **Alert Status** page. The following table describes the related parameters.

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.

1.6.4.2.7.3. 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**.

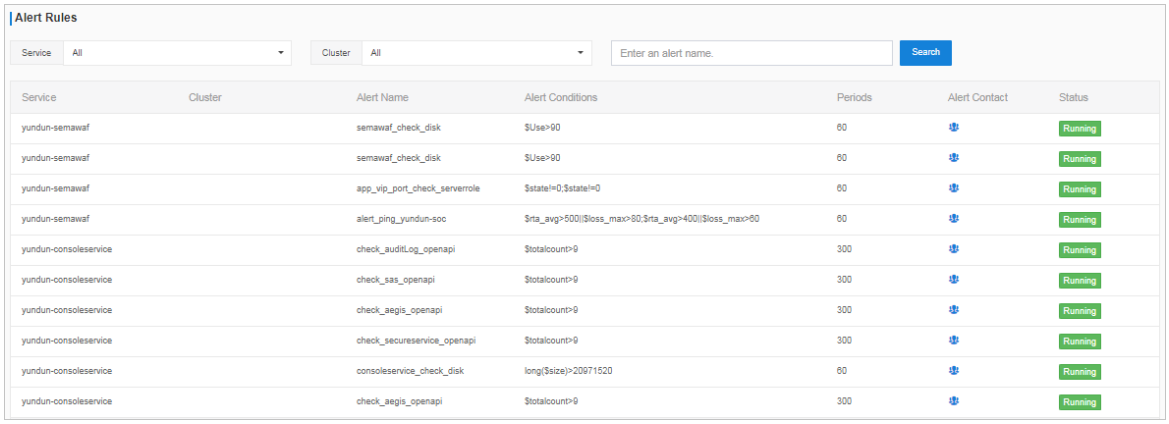

4. (Optional)Search for alert rules by service name, cluster name, or alert name.
5. View alert rules on the **Alert Rules** page. The following table describes the related parameters.

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.

1.6.4.2.7.4. 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**.

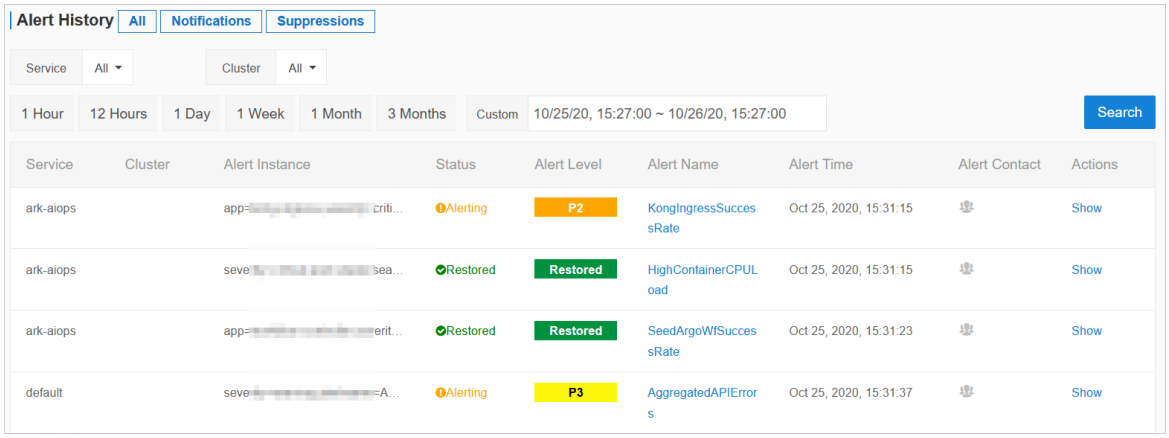

4. (Optional)Search for an alert by service name, cluster name, alert cycle, or alert time range.
5. View the alert history on the **Alert History** page. The following table describes the related parameters.

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.

1.6.4.2.8. Tools

1.6.4.2.8.1. 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**.

1.6.4.2.8.2. 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.

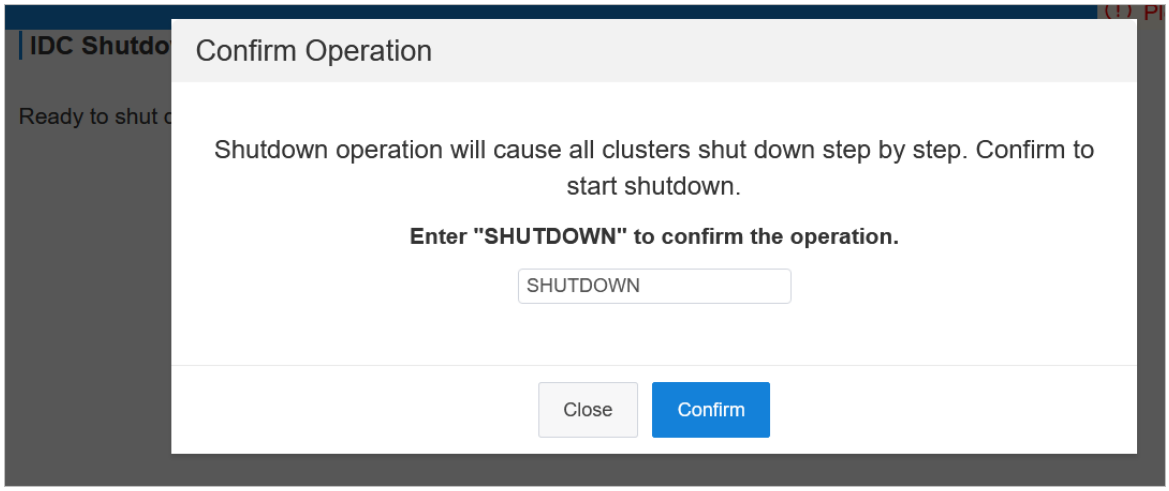

5. View the data center shutdown progress and the statuses of clusters, machines, and server roles.

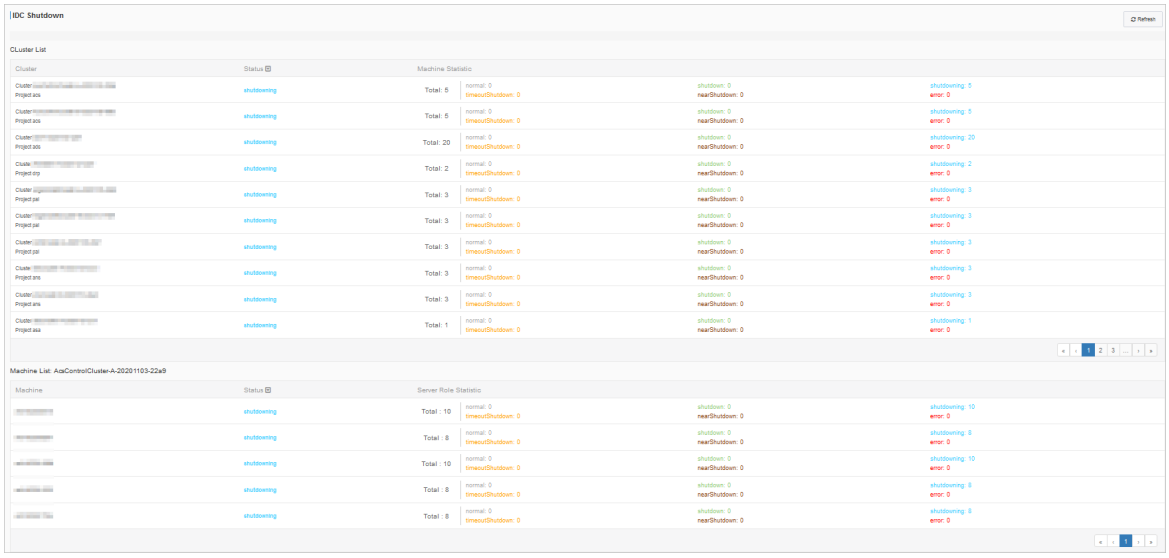

It takes a long time to shut down all clusters and machines within an environment. You can view the shutdown progress on the **IDC Shutdown** page. The following statuses are available for clusters, machines, and server roles:

- **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.

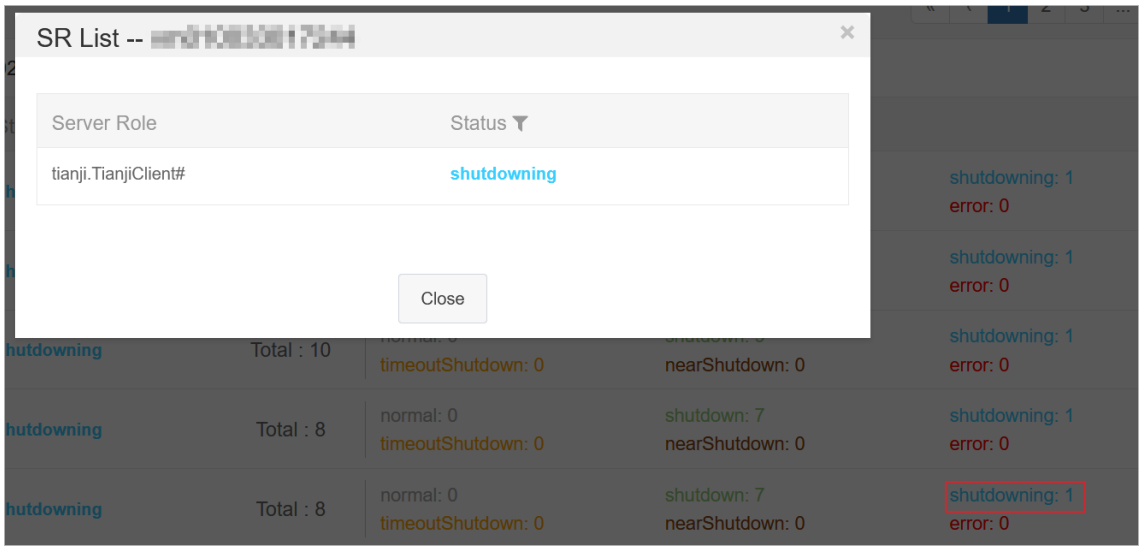

> **Note**
>
> In the left-side navigation pane, click **Go**. On the **All Reports** page, enter the entire or part of **Machine Power On or Off Statuses of Clusters** in the **Fuzzy Search** search box. In the search results, click **Machine Power On or Off Statuses of Clusters** to view the status 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.

1.6.4.2.8.3. 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**.

1.6.4.2.9. Metadata operations

In this version, you can use only command lines to perform metadata operations.

1.6.4.2.9.1. Common parameters

Common parameters consist of the common request parameters and the common response parameters.

Common request parameters

Common request parameters are request parameters that you must use when you call each API.

Parameter descriptions

Name	Type	Required	Description
Action	String	Yes	The API name. For more information about the valid values, see APIs on the control side and APIs on the deployment side.

Common response parameters

Each time you send a request to call an API, the system returns a unique identifier, regardless of whether the call is successful.

Parameter descriptions

Name	Type	Required	Description
RequestID	String	Yes	The request ID. The request ID is returned, regardless of whether the API call is successful.
Code	String	No	The error code.
Message	String	No	The reason of failure, which appears when the API call fails.
Result	The type varies with the request, which is subject to the returned result of the specific API.	No	The request result, which appears when the API call is successful.

Note

- If the API call is successful, RequestID is returned and the HTTP return code is 200.
- If the API call fails, RequestID, Code, and Message are returned and the HTTP return code is 4xx or 5xx.

Instance types

```
{
  "rds.mys2.small":{
    "cpu":2,
    "memory":4096,
    "disk":51200,
    "max_connections":60
  },
  "rds.mys2.mid":{
    "cpu":4,
    "memory":4096,
    "disk":51200,
    "max_connections":150
  },
  "rds.mys2.standard":{
    "cpu":6,
    "memory":4096,
    "disk":51200,
    "max_connections":300
  },
  "rds.mys2.large":{
    "cpu":8,
    "memory":7200,
    "disk":102400,
    "max_connections":600
  },
  "rds.mys2.xlarge":{
    "cpu":9,
    "memory":12000,
    "disk":204800,
    "max_connections":1500
  },
  "rds.mys2.2xlarge":{
    "cpu":10,
    "memory":20000,
    "disk":512000,
    "max_connections":2000
  }
}
```


1.6.4.2.9.2. Make API requests

This topic describes how to make API requests from the control and deployment sides.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the left-side navigation pane, choose **Operations > Machine Operations**.
3. Select a project from the drop-down list or enter a cluster or machine name to search for the machine.
4. Make API requests.

- Make API requests from the control side
  a. Find the machine and click **Terminal** in the **Actions** column to log on to the machine.
  b. On the command line, enter the following command and press the Enter key to obtain the value of intranet-domain:

```
grep 'intranet-domain' /cloud/app/tianji/TianjiClient#/service_manager/current/conf.global/kv.json
```

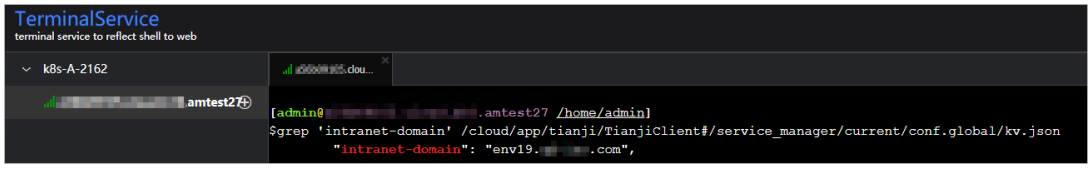

  c. Use one of the following methods to make API requests from the control side. ListInstance is used in the example.
     - GET request

```
curl 'xdb-master.xdb.{intranet-domain}:15678?Action=ListInstance'
```

     - POST request

```
curl 'xdb-master.xdb.{intranet-domain}:15678' -X POST -d '{"Action":"ListInstance"}'
```

- Make API request from the deployment side
  a. Find the machine and record the IP address in the Hostname column.
  b. Use one of the following methods to make API requests from the deployment side. CheckState is used in the example.

     Assume that the IP address of the machine is 127.0.XX.XX.

     - GET request

```
curl '127.0.XX.XX:18765?Action=CheckState&Port=3606'
```

     - POST request

```
curl '127.0.XX.XX:18765' -X POST -d '{"Action":"CheckState","Port":3606}'
```

1.6.4.2.9.3. APIs on the control side

DescribeInstance

You can call this operation to query an instance.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeInstance.

Parameter	Type	Required	Description
InstanceName	String	Yes	The name of the instance.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
InstanceID	Integer	Yes	The ID of the instance.
InstanceName	String	Yes	The name of the instance.
Domain	String	Yes	The domain name.
Port	Integer	Yes	The port of the instance.
PaxosPort	Integer	Yes	The communication port between instance nodes.
InstanceDir	String	Yes	The directory of the instance.
Level	String	Yes	The instance type.
User	String	Yes	The username.
Password	String	Yes	The password.
Config	String	No	The custom my.cnf configuration of the instance in the JSON format.
LeaderIP	String	No	The IP address of the leader node.
ActionName	String	Yes	The name of the operation.
ActionStatus	String	Yes	The status of the operation.
Description	String	Yes	The description of the instance.

Parameter	Type	Required	Description
IsDeleted	Integer	No	Indicates whether the instance was deleted. Valid values: 0 and 1. 0 indicates no, and 1 indicates yes.
NodeList	[]NodeInfo	Yes	The information of the instance node.

The following table describes the parameters of NodeInfo.

Parameter	Type	Required	Description
InstanceID	Integer	Yes	The ID of the instance.
InstanceName	String	Yes	The name of the instance.
IP	String	Yes	The IP address of the instance node.
NodeID	Integer	Yes	The ID of the instance node.
ActionName	String	Yes	The name of the operation.
ActionStatus	String	Yes	The status of the operation.
Description	String	Yes	The description of the instance.
IsDeleted	Integer	No	Indicates whether the instance was deleted. Valid values: 0 and 1. 0 indicates no, and 1 indicates yes.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DescribeInstance&InstanceName=xdb-meta'
```

Sample success responses

```
{
  "Result": {
    "ActionName": "",
    "Level": "rds.mys2.standard",
    "InstanceID": 1,
    "LeaderIP": "10.39.XX.XX",
    "Config": "{}",
    "Description": "",
    "ActionStatus": "",
    "Domain": "xdb-meta.xdb.env8c-inc.com",
    "PaxosPort": 11606,
    "InstanceName": "xdb-meta",
    "User": "xdb",
    "Password": "xdb",
    "Port": 3606,
    "IsDeleted": 0,
    "InstanceDir": "/apsarapangu/disk1/xdb/xdb_instance_3606",
    "NodeList": [
      {
        "ActionStatus": "",
        "ActionName": "",
        "Description": "",
        "InstanceID": 1,
        "IP": "10.38.XX.XX",
        "InstanceName": "xdb-meta",
        "NodeID": 1,
        "IsDeleted": 0
      },
      {
        "ActionStatus": "",
        "ActionName": "",
        "Description": "",
        "InstanceID": 1,
        "IP": "10.39.XX.XX",
        "InstanceName": "xdb-meta",
        "NodeID": 2,
        "IsDeleted": 0
      },
      {
        "ActionStatus": "",
        "ActionName": "",
        "Description": "",
        "InstanceID": 1,
        "IP": "10.39.XX.XX",
        "InstanceName": "xdb-meta",
        "NodeID": 3,
        "IsDeleted": 0
      }
    ]
  },
  "RequestID": "3CFCBA07-3D87-4A99-B8C1-E861A7D1A573"
}
```

ListInstance

You can call this operation to list the basic information of instances.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ListInstance.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
InstanceNames	String	Yes	The list of one or more instance names.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=ListInstance'
```

Sample success responses

```
{
  "Result": {
    "InstanceNames": [
      "xdb-meta",
      "xdb-instance-1",
      "xdb-instance-2",
      "xdb-instance-3"
    ]
  },
  "RequestID": "A921B8C7-C833-417C-B46A-E0CE129EBE48"
}
```

CreateInstance

You can call this operation to create an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CreateInstance.
InstanceName	String	Yes	The name of the instance.
User	String	Yes	The username.
Password	String	Yes	The password.
Level	String	Yes	Instance types
Config	String	No	The custom my.cnf configuration of the instance in the JSON format. The key must be the same as the value of the field in my.cnf. The value must be of the String type.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=CreateInstance&InstanceName=xdb-instance-1&User=admin&password=xdb&Level=rds.mys2.small'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "8BCB3B39-6140-459F-B283-F83C03ADC3CA"
}
```

DeleteInstance

You can call this operation to delete an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DeleteInstance.
InstanceName	String	Yes	The name of the instance.

Response parameters

Common response parameters

Examples

Sample requests

```
curl '127.0.XX.XX:15678? Action=DeleteInstance&InstanceName=xdb-instance-1'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "9C40CCB3-4FAB-4242-9B87-792E8154E5CD"
}
```

RestartInstance

You can call this operation to restart an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to RestartInstance.
InstanceName	String	Yes	The name of the instance.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=RestartInstance&InstanceName=xdb-instance-2'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "47277A23-5FFE-4A46-B65F-E6F2569F44E5"
}
```

UpgradeInstance

You can call this operation to upgrade the minor version of an instance. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to UpgradeInstance.
InstanceName	String	Yes	The name of the instance.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=UpgradeInstance&InstanceName=xdb-instance-2'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "95E8B098-B04A-4BCA-BEBE-DA1D11BBAD4A"
}
```

DescribeTaskProgress

You can call this operation to query the task progress.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeTaskProgress.
RequestID	String	Yes	The ID of the request.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
Progress	String	Yes	The task progress of the instance. Valid values: pending, doing, done, and failed.
Description	String	Yes	The description of the task progress.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DescribeTaskProgress&RequestID=47277A23-5FFE-4A46-B65F-E6F2569F44E5'
```

Sample success responses

```
{
  "Result": {
    "Progress": "done",
    "Description": "Success"
  },
  "RequestID": "AC535130-F40E-4D45-BC05-0F45C8473346"
}
```

ChangeLeaderTo

You can call this operation to change the leader role of an instance to another node.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ChangeLeaderTo.
InstanceName	String	Yes	The name of the instance.
IP	String	Yes	The IP address of the machine where the new leader node is deployed.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=ChangeLeaderTo&InstanceName=xdb-instance-1&IP=10.39.XX.XX'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "37638DE5-14C1-4D2E-984F-FEA1F29C9F84"
}
```

ModifyInstanceLevel

You can call this operation to modify the instance type. This operation is an asynchronous task. You can query the task result by calling the DescribeTaskProgress operation based on the request ID in the responses.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ModifyInstanceLevel.
InstanceName	String	Yes	The name of the instance.
Level	String	Yes	The new instance type.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=ModifyInstanceLevel&InstanceName=xdb-instance-1&Level=rds.mys2.mid'
```

Sample success responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "21B91211-BB09-4665-835D-9471A6F07F24"
}
```

DescribeLeader

You can call this operation to query the leader node information of an instance.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeLeader.
InstanceName	String	Yes	The name of the instance.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
LeaderIP	String	Yes	The IP address of the leader node.
Port	Integer	Yes	The port of the instance.
User	String	Yes	The username.
Password	String	Yes	The password.

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DescribeLeader&InstanceName=xdb-meta'
```

Sample success responses

```
{
  "Result": {
    "LeaderIP": "10.27.XX.XX",
    "Password": "xdb",
    "Port": 3606,
    "User": "xdb"
  },
  "RequestID": "2F05EE81-DC47-478E-9CA9-9AE8CA809151"
}
```

RecreateNode

Recreates an instance node.

Description

Uses other available nodes to recreate an instance node by backup and recovery. This is an asynchronous task. You can view the task result by calling the DescribeTaskProgress API based on the RequestID in the responses.

Request parameters

For more information about the common request parameters, see Common request parameters.

Name	Type	Required	Description
Action	String	Yes	The parameter specified by the system. Value: RecreateNode
InstanceName	String	Yes	The instance name.
IP	String	Yes	The IP address of the instance node to be recreated.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=RecreateNode&InstanceName=xdb-instance-1&IP=10.39.XX.XX'
```

Sample responses

```
{
  "Result": "Task has created, you can use api(DescribeTaskProgress) to get task progress.",
  "RequestID": "7F079E11-1DE9-4148-A9FA-683E4C58F9C2"
}
```

CreateDatabase

You can call this operation to create a database and a user.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CreateDatabase.
InstanceName	String	Yes	The name of the instance.
DBName	String	Yes	The name of the database.
User	String	Yes	The username.
Password	String	Yes	The password.

Response parameters

Common response parameters

Examples

Sample requests

```
curl '*db-master.*db.env8c-inc.com:15678? Action=CreateDatabase&InstanceName=*db-instance-1&DBName=***&User=***&Password=***_password'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "A2BEF74F-5C3A-4CEF-A2B8-C14C71E36569"
}
```

DeleteDatabase

You can call this operation to delete a database.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DeleteDatabase.
InstanceName	String	Yes	The name of the instance.
DBName	String	Yes	The name of the database to be deleted.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DeleteDatabase&InstanceName=xdb-instance-1&DBName=xdb'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "23F75A0A-B1D6-4341-BD5B-1A5F3FD45848"
}
```

DeleteUser

You can call this operation to delete a user.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DeleteUser.
InstanceName	String	Yes	The name of the instance.
User	String	Yes	The username.

Parameter	Type	Required	Description
Host	String	No	The range of IP addresses of hosts to which the user logs on. If you do not specify this parameter, the user is deleted from all the IP addresses.

Response parameters

Common response parameters

Examples

Sample requests

```
curl 'xdb-master.xdb.env8c-inc.com:15678? Action=DeleteUser&InstanceName=xdb-instance-1&User=admin&Host=10.39.XX.XX'
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "6A82AFF6-2B4D-48EF-868D-BBA54667D846"
}
```

1.6.4.2.9.4. APIs on the deployment side

CheckHealth

You can call this operation to check whether an instance node assumes a leader role and is able to read and write data.

Description

You can call this operation to check whether an instance node assumes a leader role. An instance node is considered to be healthy only when it assumes a leader role and is able to read and write data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CheckHealth.
Port	Integer	Yes	The port of the instance node.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
health	Boolean	Yes	The health status.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=CheckHealth&Port=3606'
```

Sample success responses

```
{
  "Result": {
    "health": true
  },
  "RequestID": "304B69CE-1566-4E87-B618-233F40238FFF"
}
```

```
{
  "Message":"{\"health\": false}",
  "Code":"NodeNotHealth",
  "RequestID":"E939DB9B-4337-4B1C-8680-F62BEDD645DC"
}
```

CheckState

You can call this operation to check whether the status of an instance node is normal.

Description

You can call this operation to check whether the status of an instance node is normal. Two scenarios are available:

- The node assumes a leader role and is able to read and write data.
- The node assumes a follower role and is able to read data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to CheckState.
Port	Integer	Yes	The port of the instance node.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
IP	String	Yes	The IP address of the instance node.
Port	Integer	Yes	The port of the instance.
Role	String	Yes	The role of the instance node.
Writeable	String	Yes	Indicates whether the instance node can write data.
Readable	String	Yes	Indicates whether the instance node can read data.
State	String	Yes	The status of the instance node. If the instance node is normal, *GOOD* is returned. Otherwise, *ERROR* is returned.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=CheckState&Port=3606'
```

Sample success responses

```
{
  "Result": {
    "Readable": true,
    "State": "GOOD",
    "Role": "Follower",
    "Port": 3606,
    "IP": "10.39.XX.XX"
  },
  "RequestID": "45A59426-46D3-4709-8DD6-CD9F243336E0"
}
```

DescribeNodeStatus

You can call this operation to query the status of an instance node.

Description

You can call this operation to query the status of an instance node. A leader node must be able to read and write data. A follower node must be able to read data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to DescribeNodeStatus.
Port	Integer	Yes	The port of the instance node. This parameter must be specified if the instance is in single_machine mode.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
IP	String	Yes	The IP address of the instance node.
Port	Integer	Yes	The port of the instance node.
Role	String	Yes	The role of the instance node.
Writeable	String	Yes	Indicates whether the node can write data.
Readable	String	Yes	Indicates whether the node can read data.
ConnectionCount	Integer/String	Yes	The number of connections. If this parameter value cannot be obtained, unknown is returned.
MaxConnectionCount	Integer/String	Yes	The maximum number of connections. If this parameter value cannot be obtained, unknown is returned.

Parameter	Type	Required	Description
ConnectionPercent	Integer/String	Yes	The usage of connections. If this parameter value cannot be obtained, unknown is returned.
QPS	Integer/String	Yes	The queries per second (QPS). If this parameter value cannot be obtained, unknown is returned.
CpuPercent	Integer/String	Yes	The CPU utilization. If this parameter value cannot be obtained, unknown is returned.
MemoryPercent	Integer/String	Yes	The memory usage. If this parameter value cannot be obtained, unknown is returned.
DiskPercent	Integer/String	Yes	The disk usage. If this parameter value cannot be obtained, unknown is returned.
State	String	Yes	The status of the instance node. If the instance node is normal, GOOD is returned. Otherwise, ERROR is returned.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=DescribeNodeStatus&Port=3606'
```

Sample success responses

```
{
  "Result": {
    "CpuPercent": 2.74,
    "IP": "10.39.XX.XX",
    "Readable": true,
    "MemoryPercent": 56.13,
    "State": "GOOD",
    "Role": "Follower",
    "MaxConnectionCount": 500,
    "ActiveThreadCount": 34,
    "Writeable": false,
    "ConnectionCount": 37,
    "DiskPercent": 3.0,
    "ConnectionPercent": 7.4,
    "QPS": 15.95,
    "Port": 3606
  },
  "RequestID": "D18328B1-78A9-4F3E-BB2E-B27AB7683C19"
}
```


ListNode

You can call this operation to query the basic information of instance nodes.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to ListNode.

Response parameters

The following table describes response parameters other than Common response parameters.

Parameter	Type	Required	Description
Nodes	String	Yes	The list of one or more instance node names.

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=ListNode'
```

Sample success responses

```
{
  "Result": {
    "Nodes": [
      "xdb-instance-1",
      "xdb-instance-2",
      "xdb-instance-3",
      "xdb-meta"
    ]
  },
  "RequestID": "3F7BB536-FA3F-4597-A3DF-E5830F5A3A21"
}
```

BackupNode

You can call this operation to back up the data of an instance node and transfer the data to a specified location. An nc command is used to specify a port for receiving the data.

Request parameters

The following table describes request parameters other than Common request parameters.

Parameter	Type	Required	Description
Action	String	Yes	The operation that you want to perform. Set the value to BackupNode.
Port	Integer	Yes	The port of the instance.
TargetIP	String	Yes	The IP address of the location to which you want to transfer data.
TargetPort	Integer	Yes	The port of the location to which you want to transfer data.

Response parameters

Common response parameters

Examples

Sample requests

```
curl '127.0.XX.XX:18765? Action=BackupNode&Port=3606&TargetIP=10.39.XX.XX&TargetPort='
```

Sample success responses

```
{
  "Result": "Success",
  "RequestID": "6A82AFF6-2B4D-48EF-868D-BBA54667D846"
}
```


1.6.4.2.10. Appendix

1.6.4.2.10.1. Project component info report

This report displays the name and status for each type of project components, including services, server roles, and machines.

Item	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.
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.

1.6.4.2.10.2. IP list

This report displays the IP addresses of physical machines and Docker applications.

IP List of Physical Machines

Item	Description
Project	The project name.
Cluster	The cluster name.

Item	Description
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.

1.6.4.2.10.3. Machine info report

This report displays the statuses of machines and server roles on the machines.

Machine Status

Displays all the machines currently managed by Apsara Infrastructure Management Framework and their corresponding statuses. In the **Global Filter** section at the top of the page, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists, and then click **Filter** on the right to filter the data.

Item	Description
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The machine status.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status.
Status Description	The description about the machine status.

Expected Server Role List

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The name of the expected server role on the machine.

Abnormal Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

Server Role Version and Status on Machine

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.
Server Role Status	The status of the server role.
Target Version	The expected version of the server role on the machine.
Current Version	The current version of the server role on the machine.
Status Description	The description about the status.
Error Message	The exception message of the server role.

Monitoring Status

Select a row in the **Machine Status** section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The server role name.

Item	Description
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.

1.6.4.2.10.4. 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.

Item	Description
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.
Action Name	The ongoing action of the server role.
Action Status	The status of the action.

1.6.4.2.10.5. Machine RMA approval pending list

Some Apsara Infrastructure Management Framework actions (such as restart) on machines and server roles can be triggered by users, but this type of actions must be reviewed and approved. This report is used to process the actions that must be reviewed and approved.

Machine

Displays the basic information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
State	The running status of the machine.
Action Name	The action on the machine.
Action Status	The status of the action on the machine.
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.

1.6.4.2.10.6. Registration vars of services

This report displays values of all service registration variables.

Item	Description
Service	The service name.
Service Registration	The service registration variable.
Cluster	The cluster name.
Update Time	The updated time.

1.6.4.2.10.7. Virtual machine mappings

Use the global filter to display the virtual machines of a specific cluster.

Displays the information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.

Item	Description
Project	The project name.
Cluster	The cluster name.
VM	The hostname of the virtual machine.
Currently Deployed On	The hostname of the physical machine on which the virtual machine is currently deployed.

Item	Description
Target Deployed On	The hostname of the physical machine on which the virtual machine is expected to be deployed.

1.6.4.2.10.8. Service inspector report

Use the global filter to display the service inspection reports of a specific cluster.

Service Inspector: Data is available only for services with inspection configured.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Description	The contents of the inspection report.
Level	The level of the inspection report.

1.6.4.2.10.9. Resource application report

In the **Global Filter** section, select the project, cluster, and machine from the **project**, **cluster**, and **machine** drop-down lists and then click **Filter** on the right to display the corresponding resource application data.

Change Mappings

Item	Description
Project	The project name.
Cluster	The cluster name.
Version	The version where the change occurs.
Resource Process Status	The resource application status in the version.
Msg	The exception message.
Begintime	The start time of the change analysis.
Endtime	The end time of the change analysis.

Changed Resource List

Item	Description
Res	The resource ID.
Type	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.
Parameters	The resource parameters.
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
APP	The application of the server role.
Name	The resource name.
Type	The resource type.
Status	The resource application status.
Parameters	The resource parameters.
Result	The resource application result.
Res	The resource ID.
Reprocess Status	The status of the interaction with Business Foundation System during the VIP resource application.
Reprocess Msg	The error message of the interaction with Business Foundation System during the VIP resource application.
Reprocess Result	The result of the interaction with Business Foundation System during the VIP resource application.
Refer Version List	The version that uses the resource.

Item	Description
Error Msg	The exception message.

1.6.4.2.10.10. 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.
Description	The description of the alert.
Updated At	The update time of the alert.

1.6.4.2.10.11. Relationship of service dependency

This report displays the dependencies among server roles. Use the global filter to display the data of a specific cluster in the list.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.

Item	Description
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.

1.6.4.2.10.12. Check report of network topology

This report checks if network devices and machines have wirecheck alerts.

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Network Instance	The name of the network device.
Level	The alert level.
Description	The description about the alert information.

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Machine Name	The server (machine) name.
Level	The alert level.
Description	The description about the alert information.

1.6.4.2.10.13. Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Status	The running status of the machine.
Clone Progress	The progress of the current clone process.

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
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.

1.6.4.2.10.14. Auto healing/install approval pending report

The list structure is the same as the machine RMA approval pending list, whereas this view is used for the approval during the installation. For more information, see Machine RMA approval pending list.

1.6.4.2.10.15. Machine power on or off statuses of clusters

After a cluster starts or shuts down machines, you can view the related information in this report.

Cluster Running Statuses

If a cluster is starting or shutting down machines, the corresponding data is available in this list. No data indicates that no cluster has machines shut down.

Item	Description
Project	The project name.
Cluster	The cluster name.
Action Name	The startup or shutdown action that is being performed by the cluster.
Action Status	The status of the action.

Server Role Power On or Off Statuses

Displays the power on or off statuses of server roles in the cluster selected in the **Cluster Running Statuses** section.

Select a row in the **Cluster Running Statuses** section to display the information of the corresponding cluster in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
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.

1.6.5. Log O&M

1.6.5.1. Overview of the Kibana log O&M platform

Kibana is an open source analytics and visualization platform. Logs for Apsara Stack Agility services such as ApsaraDB RDS, Xnet2, ASAPI, and POP are accessible to Elasticsearch, Logstash, and Kibana (ELK). You can use Kibana to view and retrieve related logs.

For more information about how to use Kibana 7.2, see Kibana Guide.

1.6.5.2. Log on to the Kibana log O&M platform

This topic describes how to log on to the Kibana log O&M platform.

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 URL of the Apsara Uni-manager Operations Console is in the following format: *region-id*.*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 (*region-id*.*ops*.asconsole.*intranet-domain-id*.com). Then, press the Enter key.

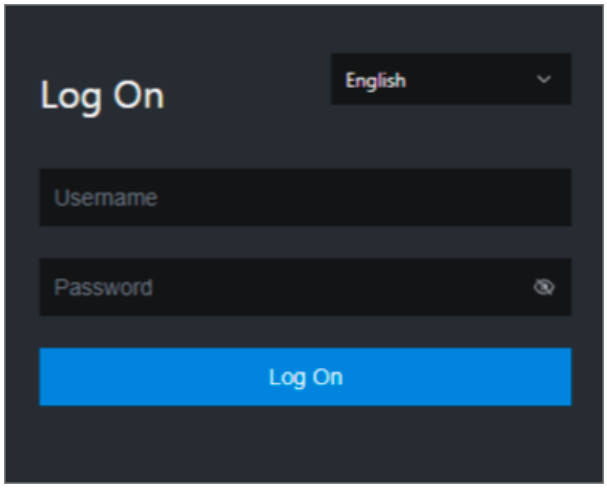

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password used to log on to the Apsara Uni-manager Operations Console from 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**.
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 > Basic O&M** section, click **Kibana Log O&M**.
8. In the dialog box that appears, enter the username and password for the Kibana log O&M platform and click **Log in**.

   > **Note** When you log on to the Kibana log O&M platform for the first time, you must enter the username and password.

1.7. Analysis

1.7.1. 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.

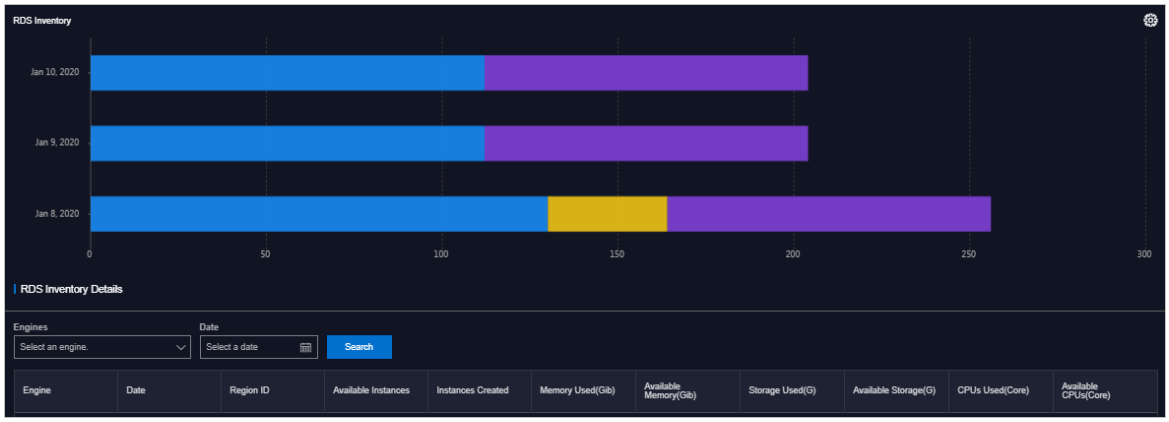

4. View the RDS inventory.
   - The **RDS Inventory** section shows the inventories of different types of RDS instances within the last five days. Different colors indicate different types of RDS instances.
   - The **RDS Inventory Details** section shows the RDS inventory details on multiple pages by **Engines** and **Date**.

1.7.2. 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**.

2.PaaS operations and maintenance

2.1. PaaS console

2.1.1. PaaS console overview

The PaaS console is designed based on the platform and products. The console is mainly used to view, manage, and upgrade the products deployed in the PaaS console. The PaaS console also provides task management capabilities to support orchestration, O&M, and custom extension.

2.1.2. Log on to the PaaS Operations Console

This topic describes how to log on to the Apsara Agility PaaS 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 URL of the Apsara Uni-manager Operations Console is in the following format: *region-id*.*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 (*region-id*.*ops*.asconsole.*intranet-domain-id*.com). Then, press the Enter key.

   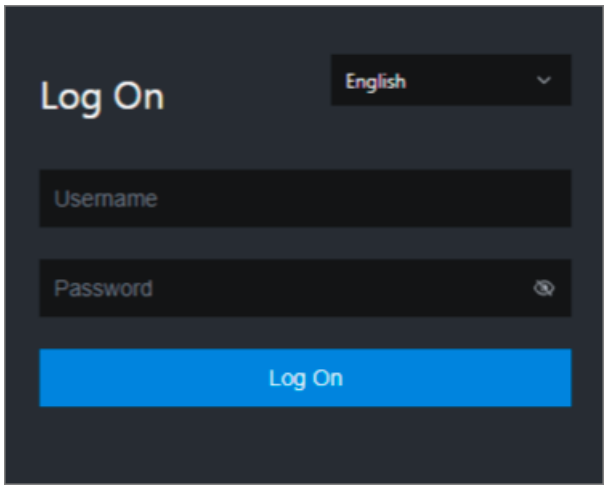

   Note You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

> **Note** Obtain the username and password used to log on to the Apsara Uni-manager Operations Console from 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**.
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 > Basic O&M** section, click **PaaS Console**.

2.1.3. Overview

The Overview module provides you with brief information about the health status of Apsara Stack Agility PaaS OM Platform.

2.1.3.1. Health Panorama

The Health Panorama module provides the overall system health status, including cluster health, product health, and release link health. O&M engineers can identify issues by analyzing the system health status.

2.1.3.1.1. View cluster health

The Cluster Health tab provides the health status of nodes in a cluster. You can go to the homepage of the Grafana service from this tab.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Overview > Health Panorama**.

   The **Cluster Health** tab is displayed.
3. On the **Cluster Health** tab, view details of cluster health, node health, and cluster events.
   - Cluster health details

     You can view cluster health details such as the total nodes, health score, and health status.

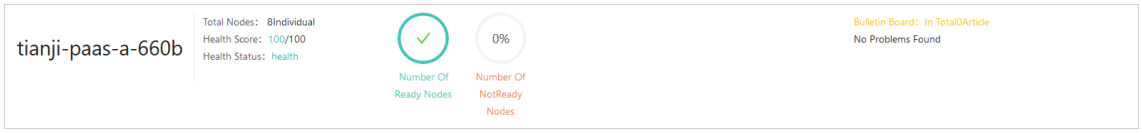

- Node health details

  You can view node health details such as the health status of each node and the cluster topology.

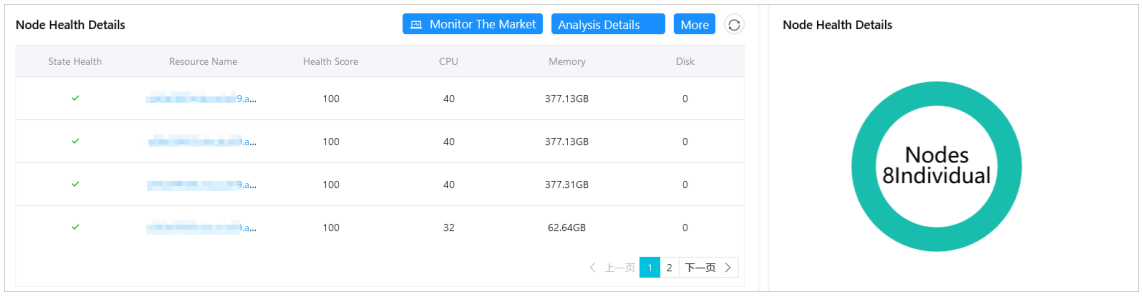

  - Click More to go to the **Nodes** page and view the details of nodes.
  - Click the icon to view the latest node data.
  - Click Analysis Details to view the cluster topology and analysis results of the health status. You can also provide feedback about the accuracy of the analysis results.

    > **Note**
    > - If the health score is lower than 100, the **Analysis Results** section displays the exception information of each unhealthy primary node and its child nodes.
    > - When you move the pointer over a node of the **Clusters Topology**, the exception information of the node is displayed.

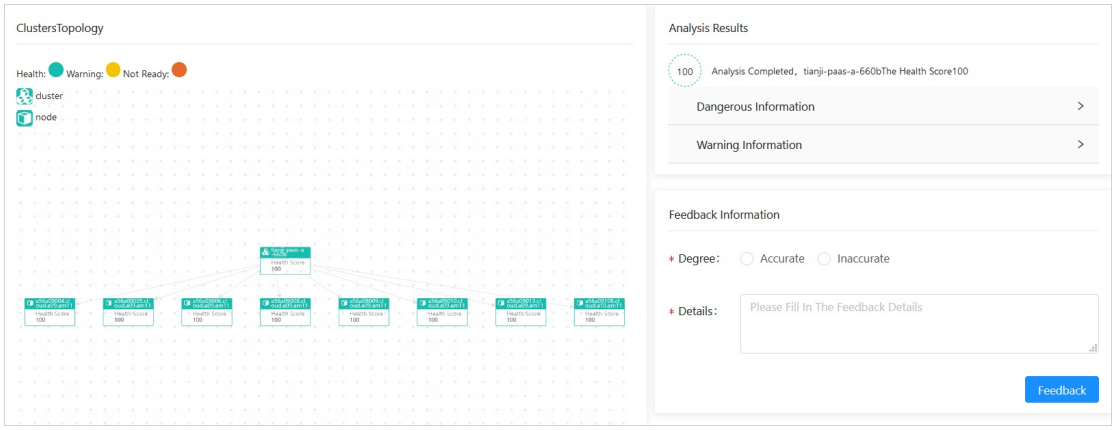

- Cluster events

  You can view all event parsing logs of a cluster.

  - Click **Original Alarm Information 10 Article** to view original alerts that are triggered by cluster events.
  - In the upper-right corner of the **Cluster event** section, click **More** to go to the **Cluster event** page and view all event information of a cluster.

2.1.3.1.2. View product health

The Product Health tab provides the health status of products that are deployed in a cluster. You can go to the homepage of the Grafana service from this tab.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Overview > Health Panorama**.

   The **Cluster Health** tab is displayed.
3. Click the **Product Health** tab to view details of product health.
   - Overall status of product health

     You can view the number, ready status, and risks of products in a cluster.

     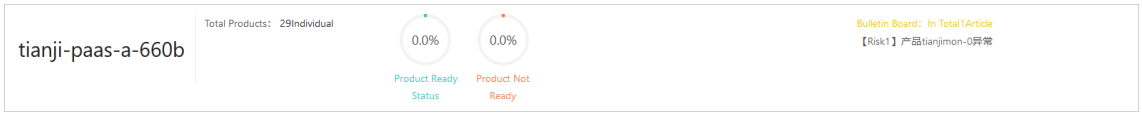

   - Product health details

     You can view the health status of a product.

     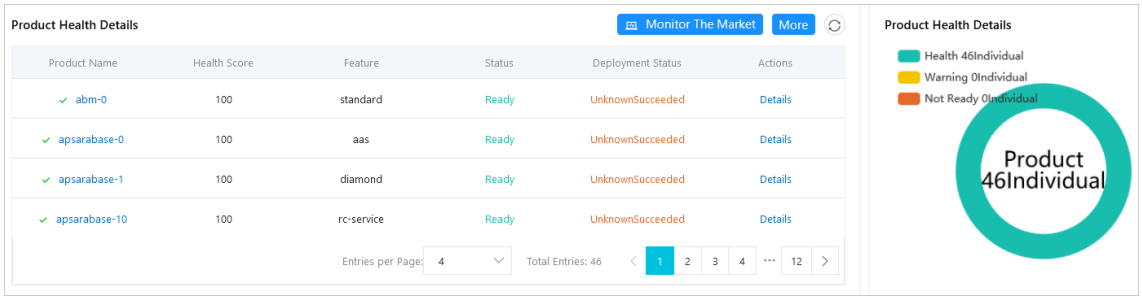

     - Click **More** to go to the **Products** page and view the details of products.
     - Click the icon to view the latest product data.

- Find the product that you want to view, and click **Details** in the **Actions** column. On the page that appears, view the **product topology** and **analysis results** of the health status. You can also provide feedback about the accuracy of the analysis results.

  > **Note**
  >
  > - If the health score is lower than 100, the **Analysis Results** section displays the exception information of each unhealthy primary node and its child nodes.
  > - When you move the pointer over a node of the **Product Topology**, the exception information of the node is displayed.

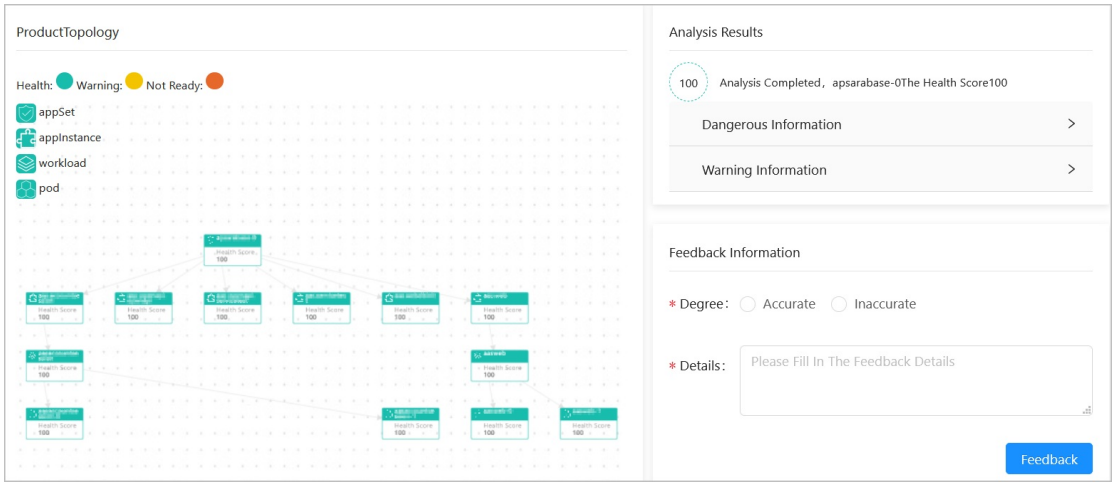

- Cluster events

  You can view all event parsing logs of a cluster.

  - Click **Original Alarm Information 10 Article** to view original alerts that are triggered by cluster events.
  - In the upper-right corner of the **Cluster event** section, click **More** to go to the **Cluster event** page and view all event information of a cluster.

2.1.3.1.3. View release link health

The Release Link Health tab provides the health status of release link components. This helps you better understand the overall health status of a cluster. You can go to the homepage of the Grafana service from this tab.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Overview > Health Panorama**.

   The **Cluster Health** tab is displayed.
3. Click the **Release Link Health** tab to view the health status of the release link.
   - Overall status of release link health

You can view the number of components, health score, and health status of the cluster release link.

- Details of release link health

  You can view the health scores of release link components, donut chart of component distribution, and release link topology.

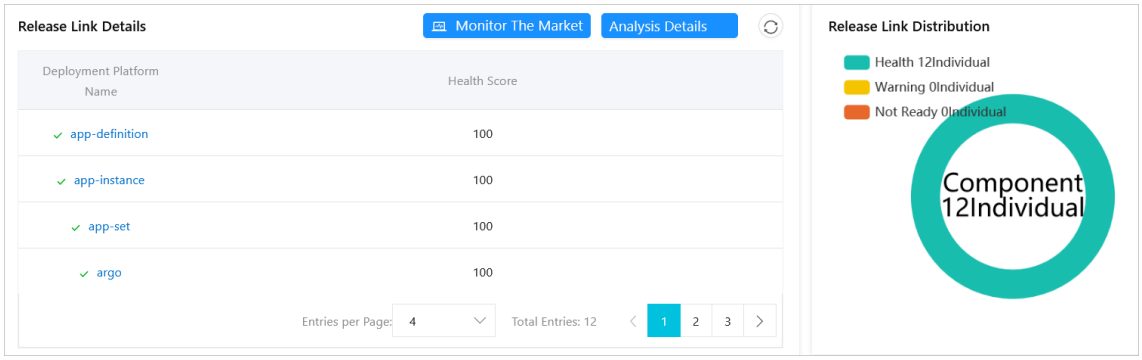

  - In the upper-right corner of the Release Link Details section, click the icon to view the latest component data.
  - Click **Analysis Details** to view the **release link topology** and **analysis results** of the health status. You can also provide feedback about the accuracy of the analysis results.

    > **Note**
    > - If the health score is lower than 100, the **Analysis Results** section displays the exception information of each unhealthy primary node and its child nodes.
    > - When you move the pointer over a node of the **Release Link Topology**, the exception information of the node is displayed.

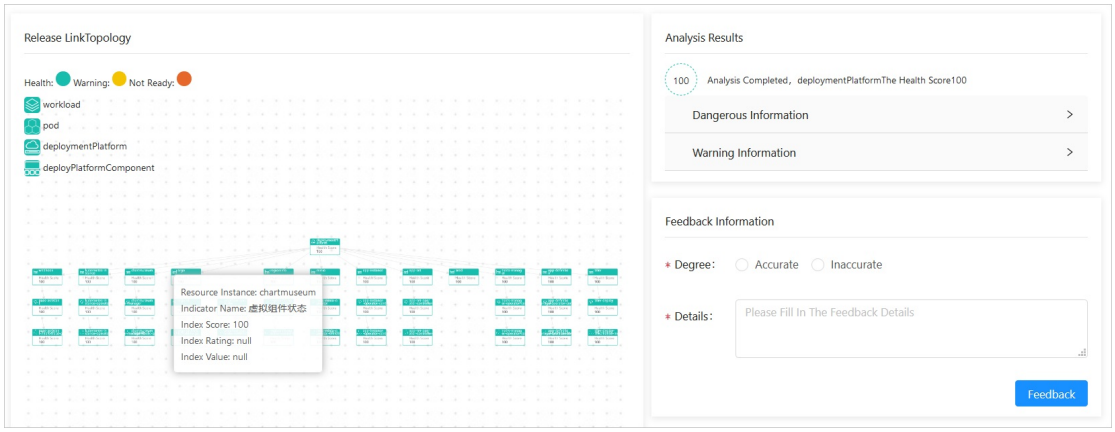

- Cluster events

  You can view all event parsing logs of a cluster.

  - Click **Original Alarm Information 10 Article** to view original alerts that are triggered by cluster events.

- In the upper-right corner of the **Cluster event** section, click **More** to go to the **Cluster event** page and view all event information of a cluster.

2.1.3.2. Alert events

The Alert Events page displays all alert events and all aggregated alert events by alert or product name.

2.1.3.2.1. View aggregated alert events by alert name

Apsara Agility PaaS Operations Console allows you to classify alert events based on various aggregation methods. You can quickly find the desired alert event to in the aggregated view of alert events by **alert name**.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Events**.
   By default, the **Alert Aggregation** tab is displayed.
3. In the upper part of the page, select the cluster that you want to view from the drop-down list.

   By default, all alert events that are aggregated by alert name are displayed.
4. In the alert name view, the Alert Aggregation tab displays all aggregated alert events by alert name. The aggregated alert event list includes the following columns: Alert Name, Details, Total Alerts, Severity, and Actions.

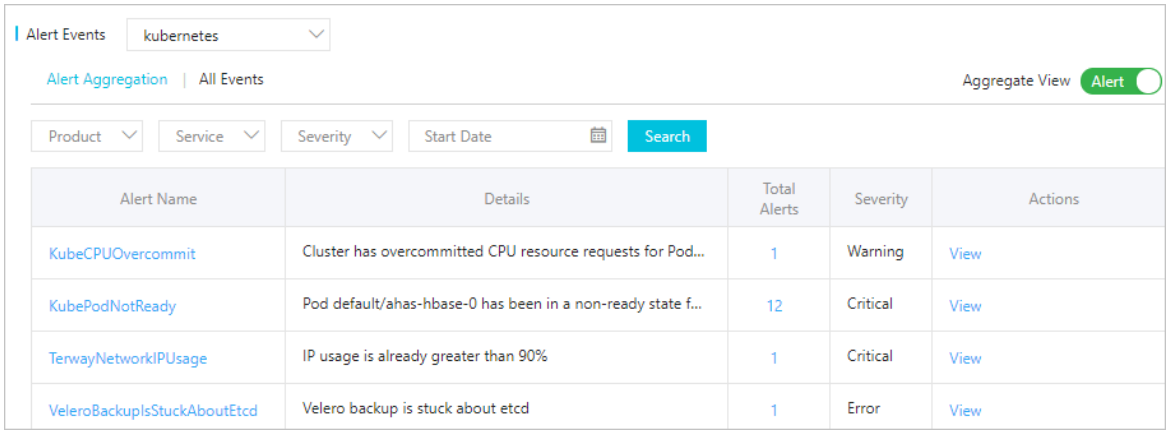

5. (Optional)In the search box at the top of the tab, set Product, Service, Severity, and Start Date, and then click **Search** to query aggregated alert events that meet the conditions.
6. Find the aggregated alert event that you want to view. Click the name in the **Alert Name** column, the number in the **Total Alerts** column, or **View** in the **Actions** column to view details of individual alert events within the aggregated alert events.

   The alert details include the following columns: Status, Start Time, End Time, Update Time, and Label.

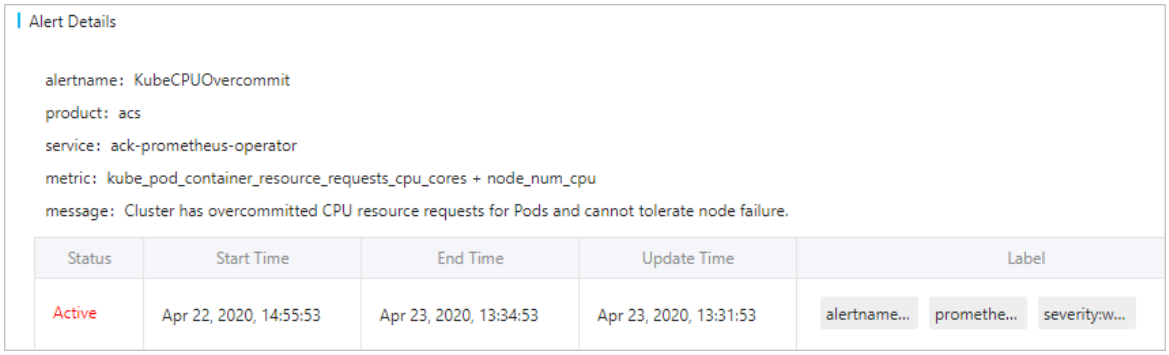

2.1.3.2.2. View alert events aggregated by product name

On the Alert Aggregation tab, you can view alert events aggregated by product name.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Overview > Alert Events**.
   The **Alert Aggregation** tab is displayed by default.
3. In the upper part of the page, select the cluster that you want to view from the drop-down list.

   By default, all alert events that are aggregated by alert name are displayed.
4. Turn off the **Aggregate View** to switch to the product name view.

   In the product name view, the **Alert Aggregation** tab displays all alert events aggregated by product name.

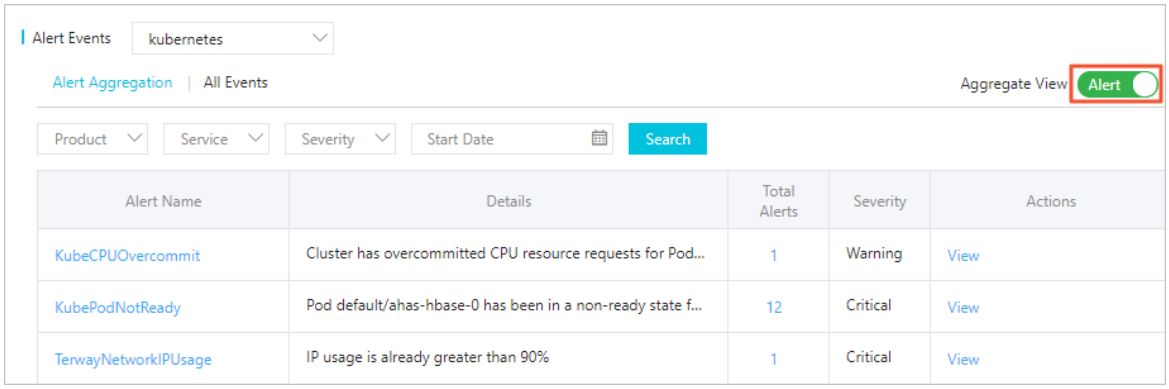

5. (Optional)In the upper part of the page, set Product, Service, Severity, and Start Date, and then click **Search** to query aggregated alert events that meet the specified conditions.
6. Find the aggregated alert that you want to view. Click the name in the **Alert Name** column, the number in the **Total Alerts** column, or **View** in the **Actions** column to view details of individual alert events within the aggregated alert.

   The alert event list includes the following columns: Status, Start Time, End Time, Update Time, and Label.

2.1.3.2.3. View all alert events

On the All Events tab, you can view all alert events that are generated in the PaaS Operations Console.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Overview > Alert Events**.
3. Click the **All Events** tab.
4. All alert events are displayed on the tab. The alert event list includes the following columns: Alert Name, Start Time, End Time, Update Time, Status, Details, Severity, and Label.

2.1.3.3. Environment model

The environment model displays the logical relationships among the region, zone, environment, and data center subsets in the Apsara Agility PaaS Operations Console.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Overview > Environment model**.
3. On the **Environment model** page, view the logical relationships among the region, zone, environment, and data center subsets in the PaaS Operations Console.

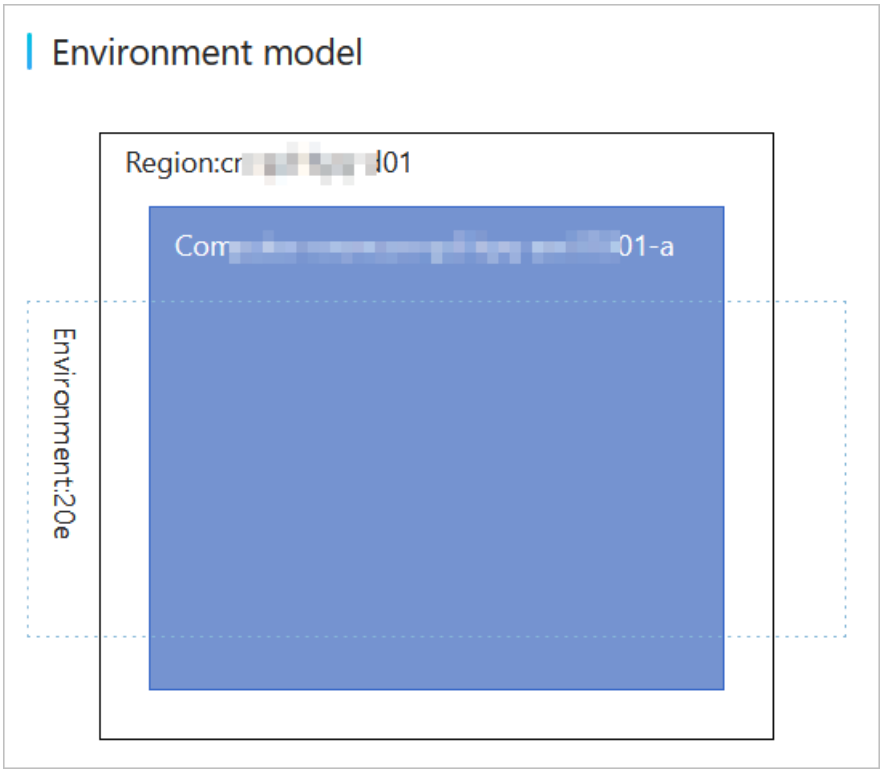

2.1.4. Clusters

2.1.4.1. View the cluster list

On the Clusters page, you can view the status, CPU usage statistics, memory usage statistics, IP address usage, registration time, and kubeconfig connection information of the clusters that are managed by the PaaS Operations Console.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Clusters > Clusters**.
3. On the **Clusters** page, view the clusters that are managed by the PaaS Operations Console.

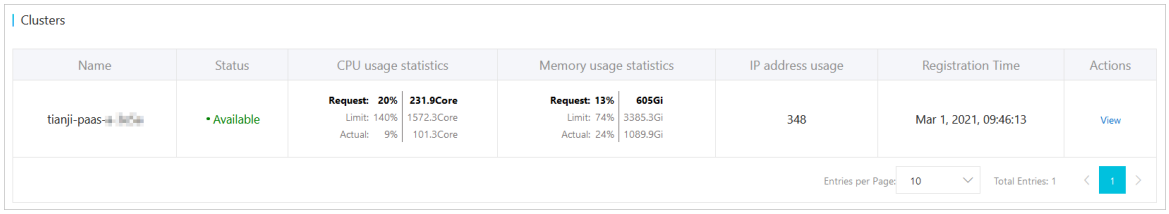

4. Find the cluster that you want to view, and click **View** in the **Actions** column to view the kubeconfig connection information of the cluster.

2.1.4.2. Node management

You can add node tags or taints for clusters to manage scheduling policies.

2.1.4.2.1. View node details

You can view the details of each node deployed in the Apsara Agility PaaS Operations Console, including the node name, status, IP address, role, operating system, version, and resource usage.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters > Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster that you want to view from the drop-down list.
4. On the **Nodes** page, view the name, IP address, and resource usage of each node.

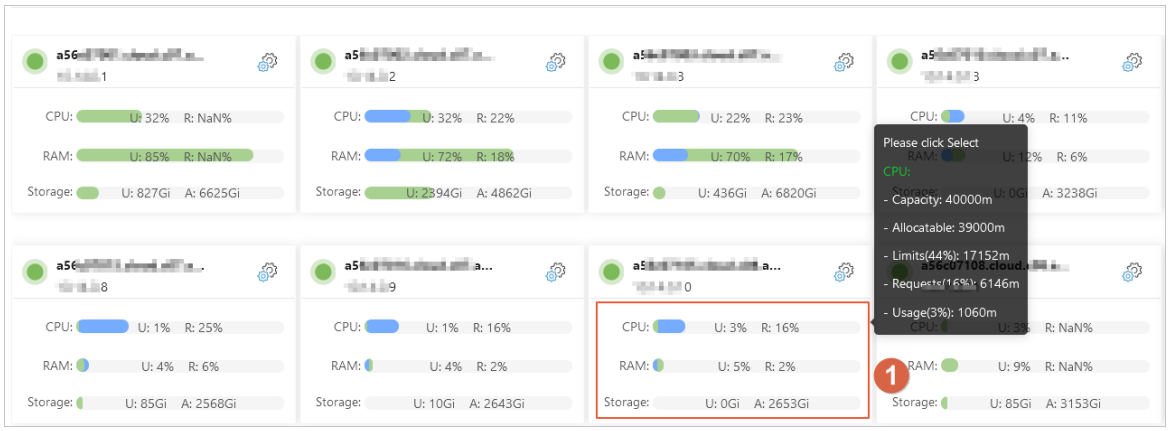

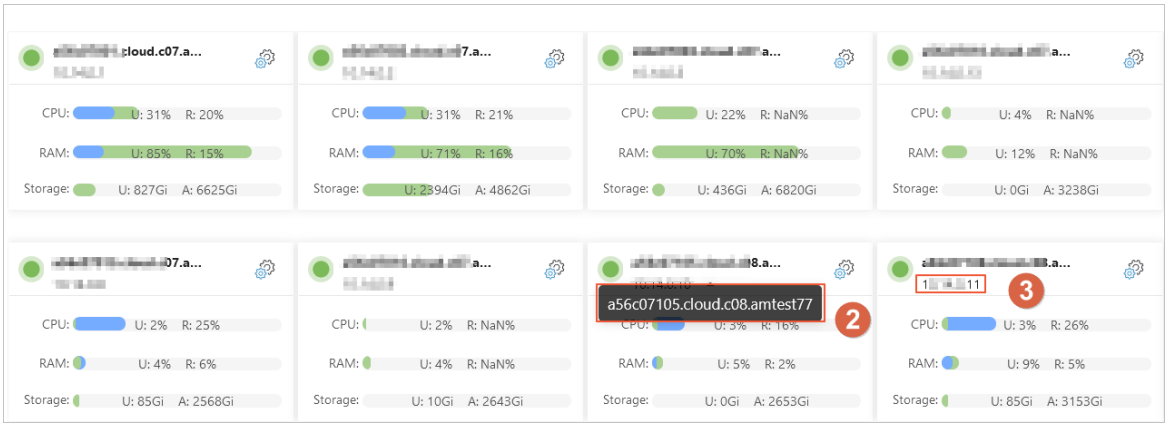

No.	Item	Description
1	Node resource usage	The node resources include the following items: ○ CPU ○ RAM ○ Storage Move the pointer over a resource item to view the usage details.
2	Node name	Move the pointer over a node name to view the complete node name.
3	IP	The IP address of the node.

5. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node that you want to view is displayed.
6. In the upper-right corner of the node, click the icon. On the Node management page, view node information including details, labels, taints, and disks.

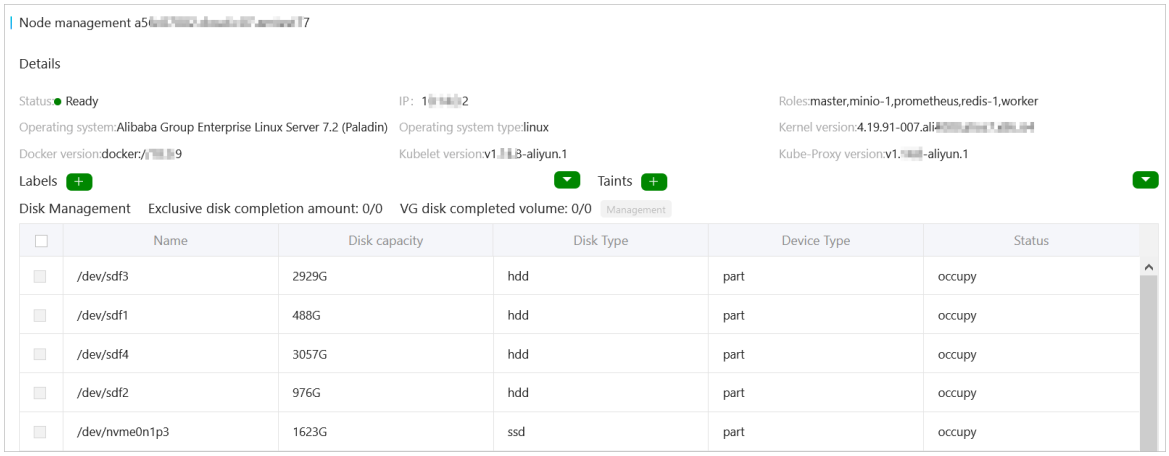

2.1.4.2.2. Add a tag

You can add tags to nodes for subsequent cluster scheduling, configuration, and behavior customization. You can add tags to nodes on the **Nodes** or **Node management** page. On the **Node management** page, you can add tags to multiple nodes.

Add a tag on the Nodes page

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster that you want to manage from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. On the **Nodes** page, select one or more nodes to which you want to add tags, and click **Add Label**.

   > **Note**
   >
   > The border color of the selected node changes to blue.

6. In the **Add Label to Node** dialog box, perform the following operations:

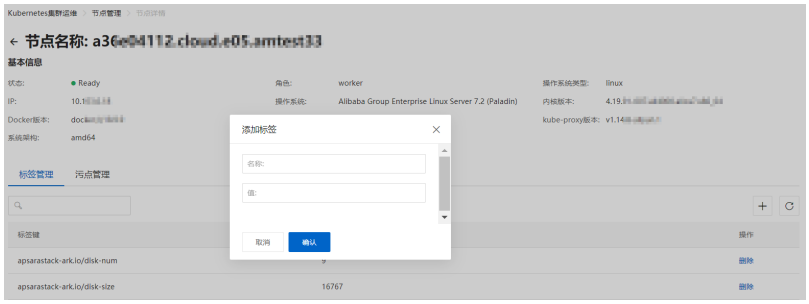

   - Add a built-in tag

     In the Add Label to Node dialog box, click a tag in the Built-in Labels field. The tag name is automatically added into the Key field. Set Value and click **OK**.

     The following table describes the parameters.

Parameter	Description
Built-in Labels	Existing tags in the system. Valid values: ■ **Hypervisor failure-domain**: During virtualization output, virtual machines are distributed across different physical machines. This tag can be used to distribute pods to different physical machines. ■ **Zone failure-domain**: distributes Kubernetes nodes to different zones. ■ **Region failure-domain**: distributes Kubernetes nodes to different regions.
Key	After you click a tag in the Built-in Labels field, the tag name is automatically added into the **Key** field. You can also set **Key** to specify a custom tag.
Value	The custom tag value.

- Add a custom tag

  In the Add Label to Node dialog box, set Key and Value, and then click **OK**.

Add a tag on the Node management page

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster that you want to manage from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. In the upper-right corner of the node to which you want to add tags, click the  icon.
6. On the **Node management** page, click the  icon next to **Labels**.

7. In the **Add Label to Node** dialog box, perform the following operations:

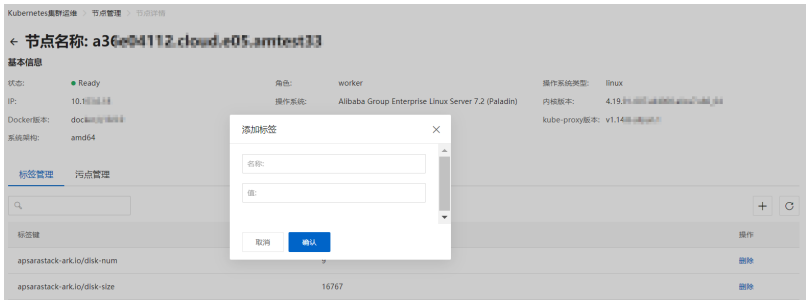

   - Add a built-in tag

     Click a tag in the Built-in Labels field. The tag name is automatically filled into the Key field. Set Value and then click **OK**.

   - Add a custom tag

     Set Key and Value, and then click **OK**.

2.1.4.2.3. Add a taint

You can add taints to nodes for subsequent pod scheduling. You can add taints to nodes on the **Nodes** or **Node management** page. On the **Node management** page, you can add taints to multiple nodes.

Add a taint on the Nodes page

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster that you want to manage from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. On the **Nodes** page, select one or more nodes to which you want to add taints, and click **Add Label**.

> **Note**
>
> The border color of the selected node changes to blue.

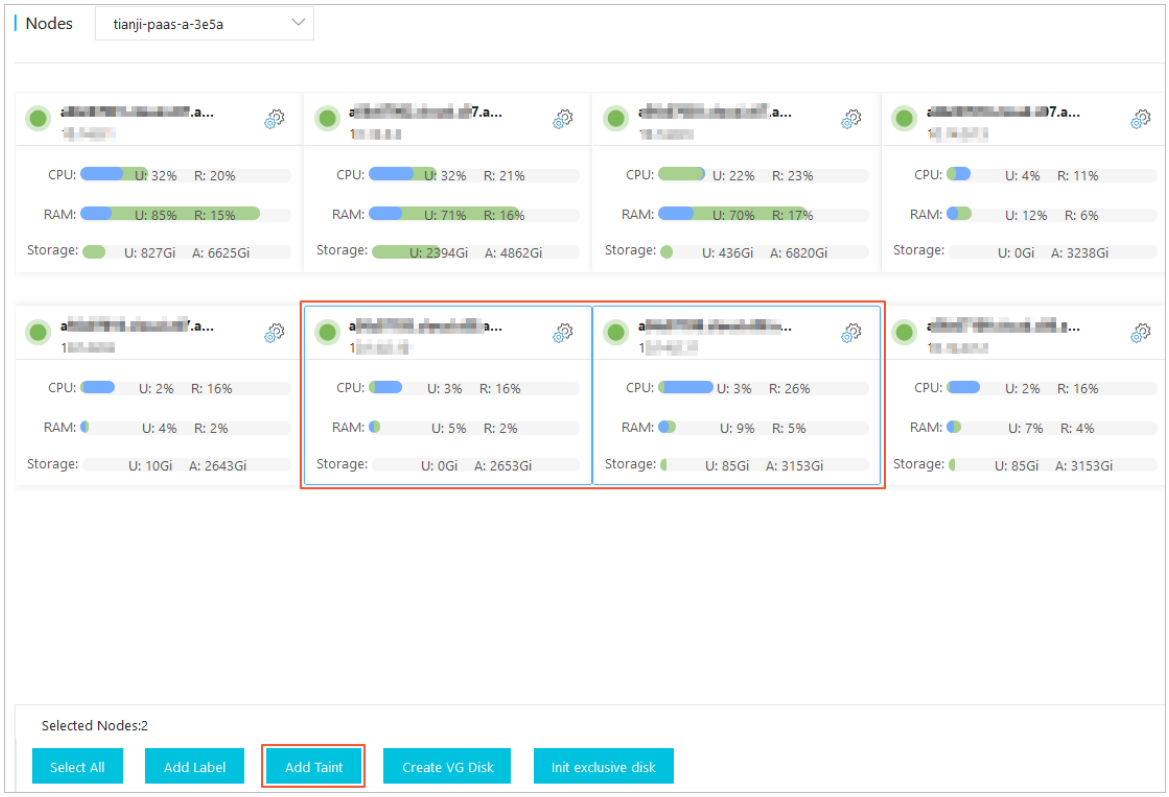

6. In the **Add Taint** dialog box, set **Key**, **Value**, and **Effect**.

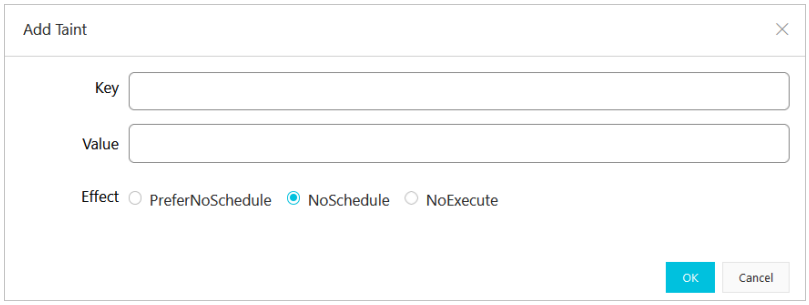

The following table describes the parameters.

Parameter	Description
Key	The custom taint key.
Value	The custom taint value.

Parameter	Description
Effect	○ **PreferNoSchedule**: Kubernetes avoids scheduling pods that do not tolerate the taint onto the node. ○ **NoSchedule**: Pods that do not tolerate the taint are not scheduled on the node. ○ **NoExecute**: Pods are evicted from the node if they are already running on the node, and are not scheduled onto the node if they are not running on the node.

7. Click **OK**.

Add a taint on the Node management page

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster that you want to manage from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. In the upper-right corner of the node to which you want to add taints, click the icon.
6. On the **Node management** page, click the icon next to **Taints**.

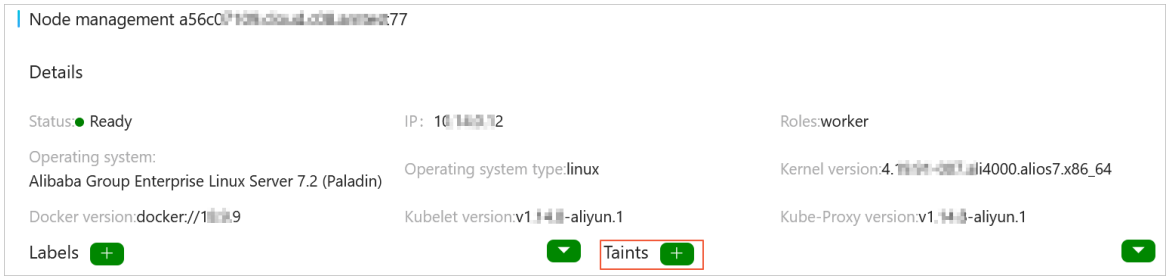

7. In the **Add Taint** dialog box, set **Key**, **Value**, and **Effect**.

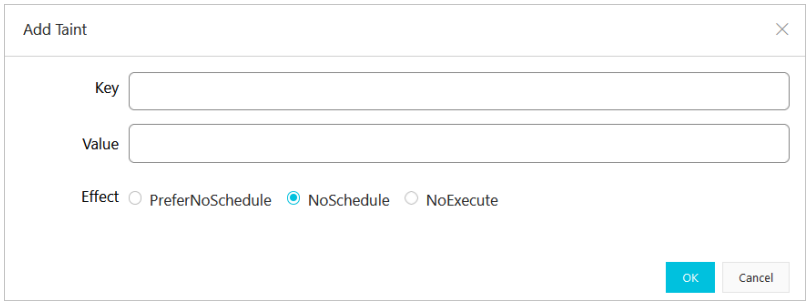

8. Click **OK**.

2.1.4.2.4. Delete a tag or taint

You can delete built-in or custom tags or taints from nodes. Kubernetes-defined tags of nodes cannot be deleted.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Clusters > Nodes**.
3. (Optional)In the upper-left corner of the Nodes page, select the cluster from the drop-down list.
4. (Optional)In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node that you want to manage is displayed.
5. In the upper-right corner of the node, click the icon.
6. Click the icon next to **Labels** or **Taints** to show the tag list or the taint list.

7. Move the pointer over the label or taint that you want to delete, and click **Delete**.

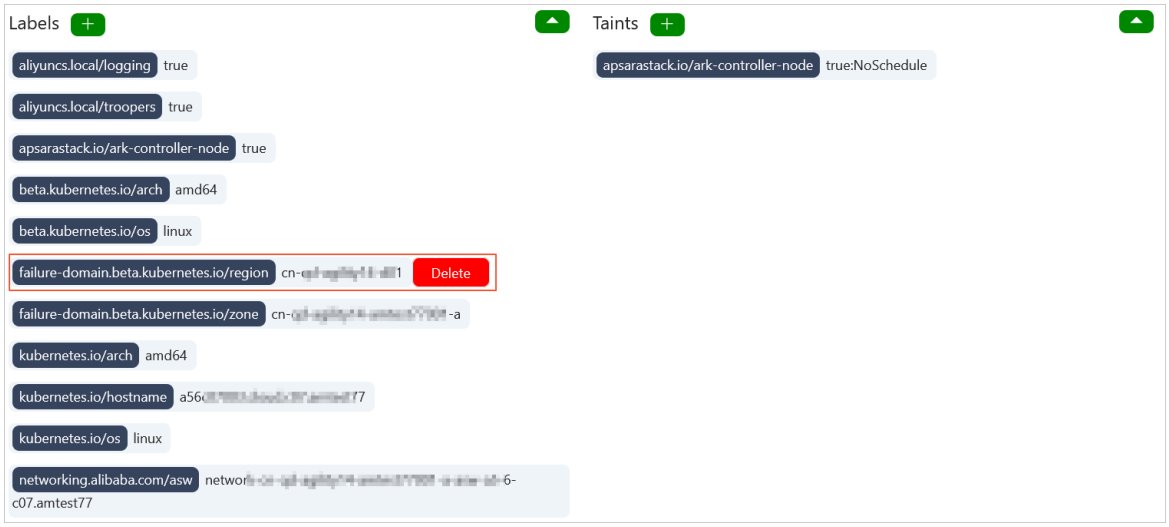

8. In the **Confirm** message, click **OK**.

2.1.4.2.5. Create a VG Disk or initiate a dedicated disk for a single node

Apsara Agility PaaS Operations Console allows you to create a VG disk or initiate a dedicated disk for a single node by using graphical interfaces. This improves efficiency because you can configure nodes on the spot and decouple from the strong dependencies in Deployment Planner.

Prerequisites

- Base services are deployed and are at the desired state.

- Each node for which you create a VG disk has at least one idle disk.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster that you want to manage from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. Click the  icon in the upper-left corner of the node.
6. In the disk list, select the disk that is in the **Idle state**. The state of a disk is displayed in the Status column.

   You can select an idle disk based on various factors such as the disk capacity and the disk type.
7. Click **Management**. In the **Disk Management** dialog box, configure the parameters.

Option	Description
Create VG Disk	Perform the following operations to create a VG disk: i. Select **Create VG Disk**. ii. Click **OK**. The "**All configurations added successfully**" message appears. Wait until the node state changes to **Ready**.
Init exclusive disk	Perform the following operations to initiate a dedicated disk: i. Select **Init exclusive disk**. ii. Select the corresponding product and component. iii. Click **OK**. The "**All configurations added successfully**" message appears. Wait until the node state changes to **Ready**.

2.1.4.2.6. Create VG disks for multiple nodes

Apsara Agility PaaS Operations Console allows you to create VG disks for multiple nodes by using graphical interfaces.

Prerequisites

- Base services are deployed and are at the desired state.
- Each node for which you create a VG disk has at least one idle disk.
- The node that has disk partitions is individually processed because mutiple VG disks cannot be created for such type of nodes at the same time.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. Select the nodes for which you want to create VG disks, and click **Create VG Disk**.

   > **Note**
   >
   > - The border color of the selected node changes to blue.
   > - If the node you select has disk partitions, the "**There are part types of disks in the following nodes, please operate separately!**" message appears.

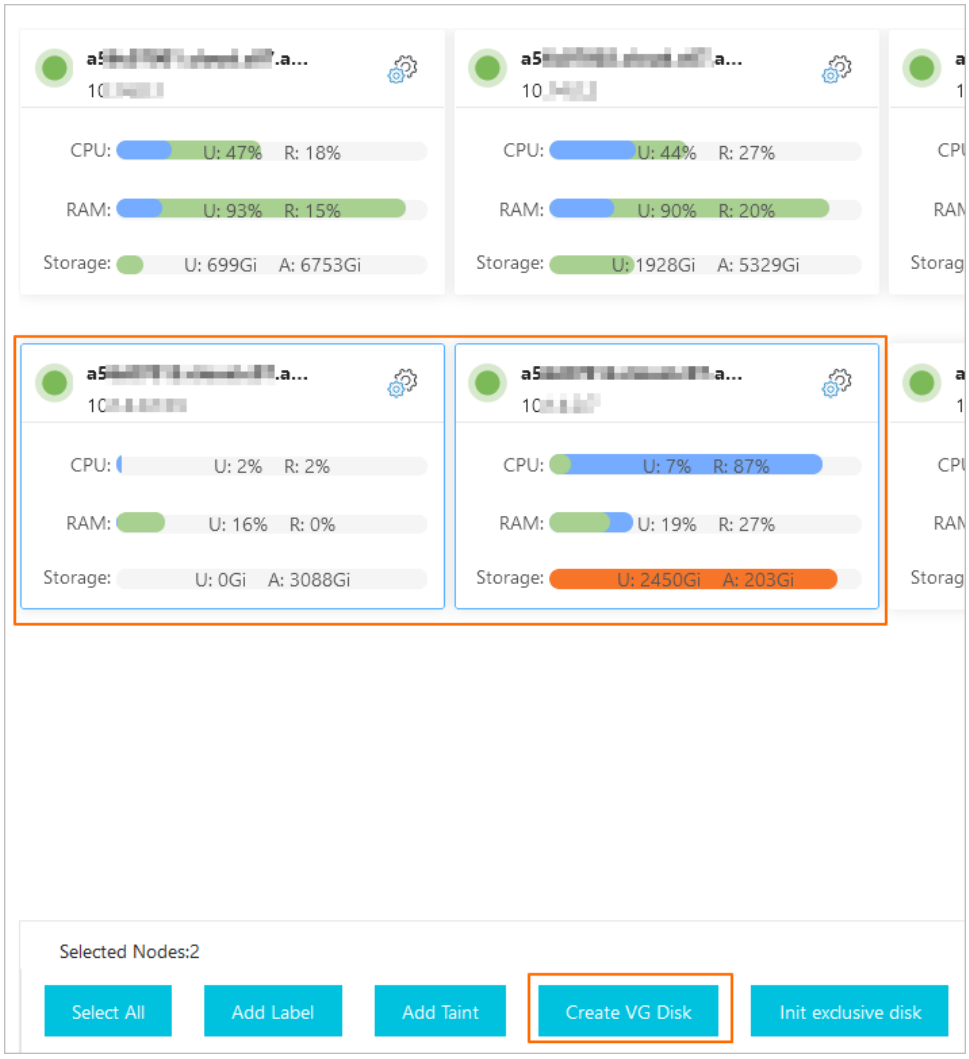

6. In the **Create VG Disk** dialog box, configure the parameters.

Parameter	Description
SSD	Specifies whether the type of the required physical disk is SSD. If you turn on SSD, the type of the required physical disk is SSD. This parameter applies to each disk.
Size	The storage capacity of the required physical disk. We recommend that you specify the minimum storage capacity for the required physical disk. This parameter applies to each disk.
Number	The number of the required physical disks. This parameter applies to each node.

7. Click **OK**.

Wait until each node is in the **Ready** state.

2.1.4.2.7. Initialize dedicated disks for multiple nodes

Apsara Agility PaaS Operations Console allows you to initialize dedicated disks for multiple nodes by using graphical interfaces.

Prerequisites

- Base services are deployed and are at the desired state.
- Each node for which you initialize the dedicated disk has at least one idle disk.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Clusters** > **Nodes**.
3. (Optional) In the upper-left corner of the Nodes page, select the cluster from the drop-down list.
4. (Optional) In the upper-right corner of the Nodes page, enter the node name in the search box and click the search icon. The node is displayed.
5. Select the nodes for which you want to initialize dedicated disks, and click **Init exclusive disk**.

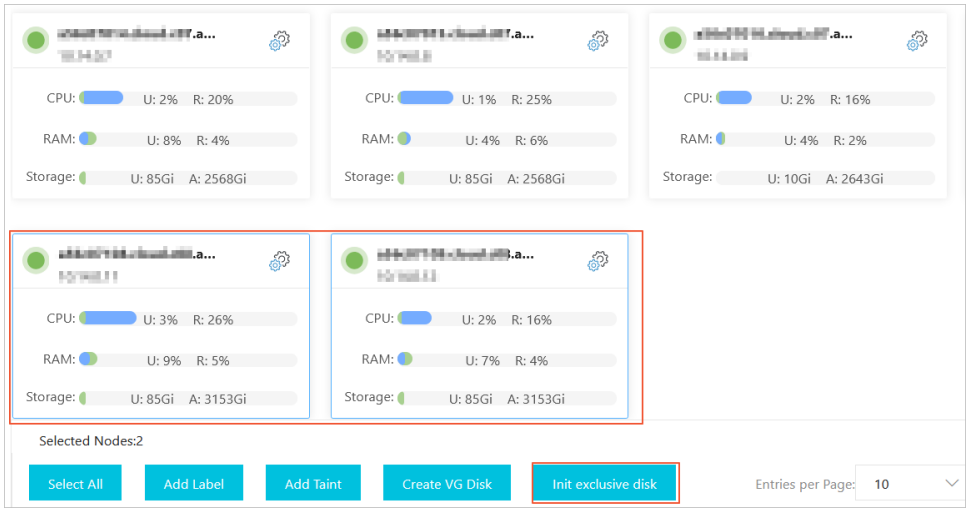

6. Click **Batch selection**. In the Batch selection dialog box, configure the parameters.

Parameter	Description
Type	The type of the disk.
Size	The storage capacity of the required physical disk. We recommend that you specify the minimum storage capacity for the required physical disk. This parameter applies to each disk.

Parameter	Description
Number	The number of the required physical disks. This parameter applies to each node.

7. Click **OK**.

   The details of the dedicated disk to be initialized are displayed, including the **name**, **owning node**, **disk capacity**, **disk type**, **device type**, and **status**.

8. Click **Initialize**. In the **Disk Management** dialog box, select the product and the component.
9. Click **OK**.

   Wait until each node is in the **Ready** state.

2.1.4.3. Query event details

On the Cluster event page, you can view all event parsing logs of clusters that are deployed in the Apsara Stack Agility PaaS console.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Clusters > Clusters**.
3. In the upper-left corner of the page, filter events by cluster name, namespace, and log type to view details of event logs.

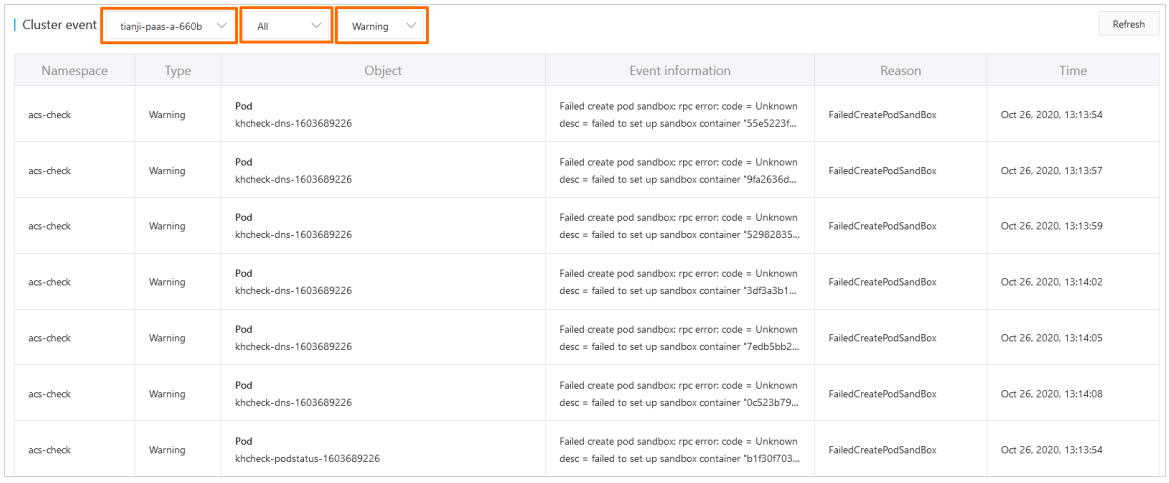

> **Note** In the upper-right corner of the page, you can view the latest information of event logs by clicking **Refresh**.

The following table describes relevant fields in the event list.

Field	Description
Namespace	The namespace that is associated with the event.

Field	Description
The type of the service.	The type of the event.
Object	The Kubernetes object that corresponds to the event.
Reason	The reason why the event was triggered.
Time	The time when the event was triggered.

2.1.5. Intelligent O&M

Inspection is the key to preventing faults and is essential for routine O&M. Apsara Agility PaaS Operations Console performs inspections on base services and cloud services by using the Intelligent O&M feature. This way, faults can be detected before your bussiness is affected.

2.1.5.1. View details of an inspection case

On the **Inspection Case Set** page, you can view inspection cases and their details.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Intelligence Operations** > **Inspection Framework**.

   On the **Inspection Case Set** page, you can view inspection cases.
3.
4. Find the desired inspection case and view its details.

> Note
>
> On the **Inspection Case Set** page, click the icon in the column that you want to filter.
>
> Enter a filter condition in the search box and click **Search**. Inspection cases that meet the filter condition are displayed.

Field	Description
Product	The product or component that is inspected by the inspection case.
Use Case ID	The ID of the inspection case. In the **Use Case ID** column, you can click an inspection case ID to view the custom resource (CR) definition of the inspection case.

Field	Description
Use Case Name	The name of the inspection case.
Use Case Description	The inspection description of the inspection case.
Cycle(second)	The execution cycle of the inspection case that is in the enabled state. Unit: seconds.
Overtime Time(second)	The execution timeout period of the inspection case that is in the enabled state. Unit: seconds.
Last State	The state of the latest execution of the inspection case. ○ **Succeeded**: The inspection case is successfully executed. ○ **Running**: The inspection case is being executed. ○ **Failed**: The inspection case fails to be executed.
Enabled State	Specifies whether to enable periodic inspection. ○ Enabled: The periodic inspection feature is enabled. ○ Deactivate: The periodic inspection feature is disabled.
Last Result	To view the result details of the latest execution of the inspection case, click **Click To View Results**.

Field	Description
Operating	○ To view the parameter details, click **See Details**. **Note** To copy a line, click the icon next to the line.  ○ To view the inspection history of the inspection case, click **View Inspection History**.

2.1.5.2. Enable periodic inspection

Apsara Agility PaaS Operations Console allows you to enable periodic inspection by using graphical interfaces.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Intelligence Operations** > **Inspection Framework**.

   On the **Inspection Case Set** page, you can view inspection cases.
3. On the **Inspection Case Set** page, find the inspection case for which you want to enable periodic inspection.

   Note

   On the **Inspection Case Set** page, click the icon in the column that you want to filter.

   Enter a filter condition in the search box and click **Search**. Inspection cases that meet the filter condition are displayed.
4. Click Deactivate in the **Enabled State** column to enable periodic inspection.

   When the icon in the **Enabled State** column changes to Enabled, periodic inspection is enabled.

2.1.5.3. Manually trigger inspections

Apsara Agility PaaS Operations Console allows you to manually trigger inspections by using graphical interfaces.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Intelligence Operations** > **Inspection Framework**.

   On the **Inspection Case Set** page, you can view inspection cases.
3. On the **Inspection Case Set** page, find the inspection cases that you want to trigger.

   > **Note**
   >
   > On the **Inspection Case Set** page, click the  icon in the column that you want to filter.
   >
   > Enter a filter condition in the search box and click **Search**. Inspection cases that meet the filter condition are displayed.
4. Select the inspection cases that you want to trigger.

   > **Note**
   >
   > You can select one or more inspection cases.
5. Click **Manually trigger inspection**. In the message that appears, click **Determine**.

   To view the inspection history, click **View Inspection History** in the **Operating** column.

2.1.6. Product center

2.1.6.1. Product list

The product list displays the information about all products deployed in the PaaS console, including their names and versions. In the product list, you can perform O&M operations and view product resources or register variables. You can also remove products that are no longer needed.

2.1.6.1.1. View product details

You can view the details of products that are deployed in the Apsara Agility PaaS Operations Console, including their names, versions, and components.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center** > **Products**.
3. In the product list, find the product that you want to view and click **Details** in the **Actions** column.

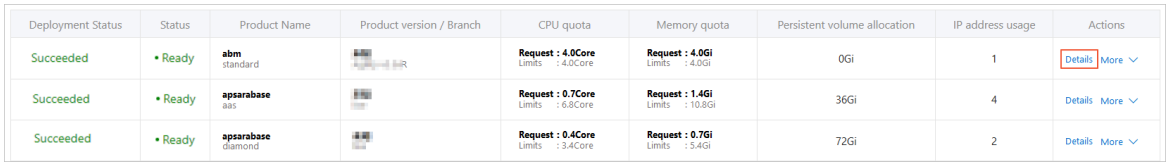

4. On the **Overview** page, view the name, version, components, release status, and resource usage of the product.

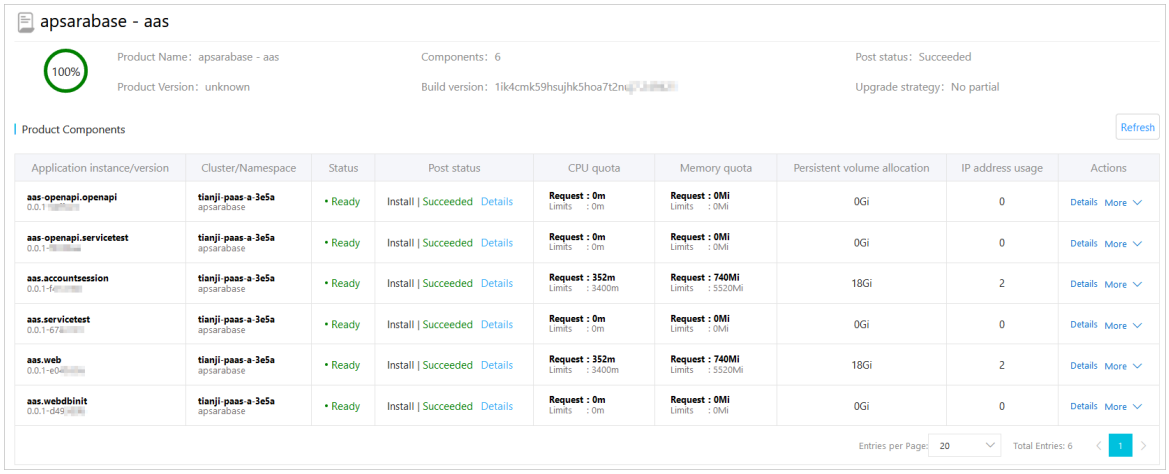

2.1.6.1.2. View product versions

You can view versions of products that are deployed in the Apsara Agility PaaS Operations Console.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. On the **Products** page, find the product that you want to view and choose **More > Version Information** in the **Actions** column.
4. In the dialog box that appears, view the version information of the product.

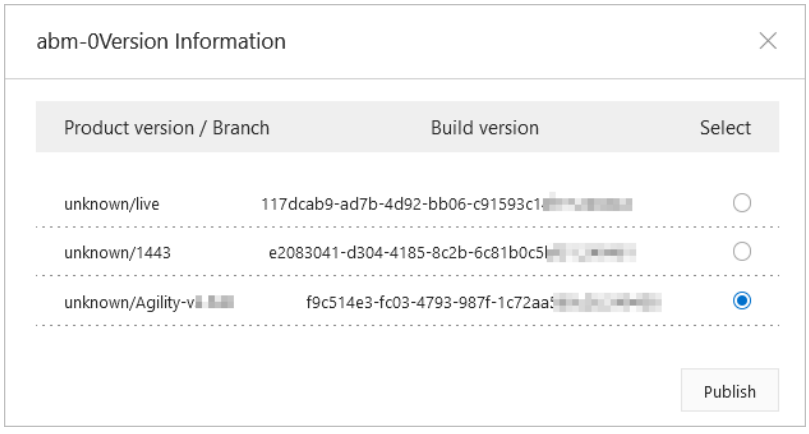

Note In the lower-right corner of the dialog box, you can click **Publish** to go to the Deploy&Upgrade page. For more information about how to deploy and upgrade products, see Deployment and upgrade.

2.1.6.1.3. View component information

You can view the component details in the Product Components section of the Overview page of a product.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. In the product list, find the product and click **Details** in the **Actions** column.
   In the **Product Components** section of the Overview page, view the deployment information of components, such as the component status, release status, cluster, namespace, component name, component version, and resource usage.
4. Find the component that you want to view and click **Details** in the **Actions** column to view details of the component.
   On the **Component Details** page, you can view the resource information of the deployed component.

2.1.6.1.4. View the release status of a product component

You can view release status details of a product component.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. On the **Products** page, find the product that you want to view and click **Details** in the **Actions** column.
4. In the Product Components section of the **Overview** page, find the component that you want to view and click **Details** in the **Post status** column.
5. In the **Details** dialog box, move the pointer over the component to view tasks. Move the pointer over each task to view subtasks. Move the pointer over each subtask to view details of the subtask.

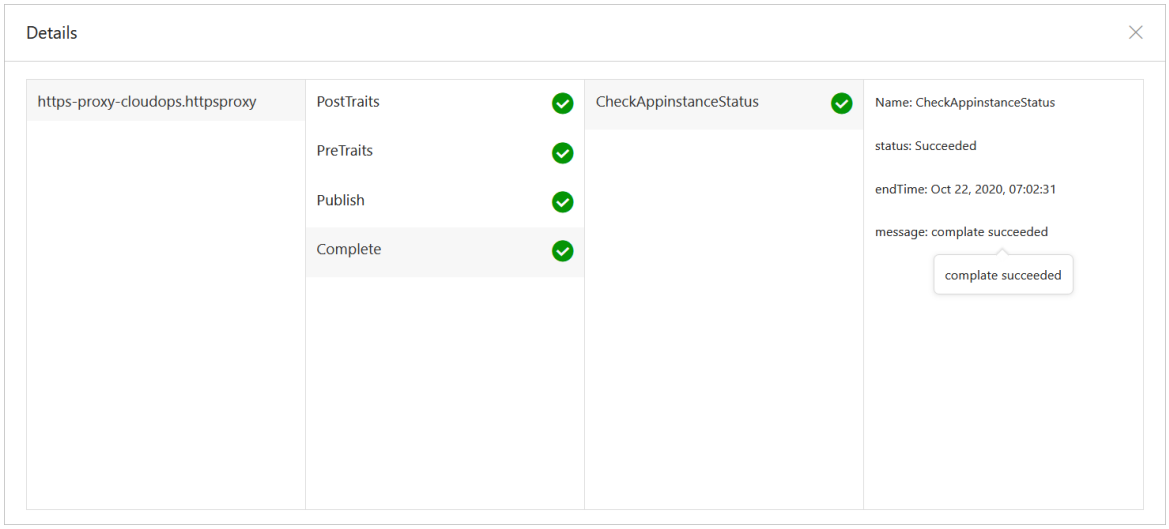

2.1.6.1.5. View the deployment progress of product components

You can view the deployment progress of product components.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. In the product list, find the product that you want to view and click **Details** in the **Actions** column.
4. In the **Product Components** section, find the component that you want to view and choose **More > Deployment Progress** in the **Actions** column.

5. On the **Product Component Deployment Progress** page, click the deployment nodes in sequence to view the deployment progress and logs of the current component.

> **Note** You can click **LOGS** in the lower-left corner of the SUMMARY tab to view the deployment logs.

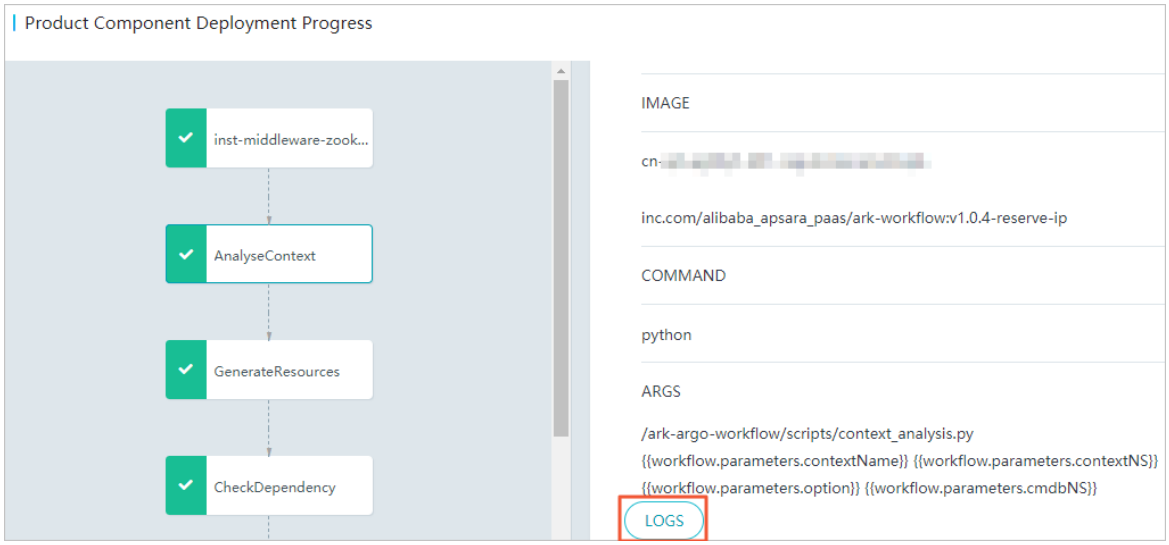

2.1.6.1.6. Log on to a web terminal

The StatefulSets and Deployments tabs of the Component Details page list available terminals. Browser-based terminals are used for O&M management and troubleshooting.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. In the product list, find the product and click **Details** in the **Actions** column.
   In the **Product Components** section of the Overview page, view the deployment information of components, such as the component status, release status, cluster, namespace, component name, component version, and resource usage.
4. In the **Product Components** section of the Overview page, find the desired component and click **Details** in the **Actions** column.
5. In the **StatefulSets** or **Services** section, find the component to which you want to log on, and click **Start Terminal** in the **Actions** column.

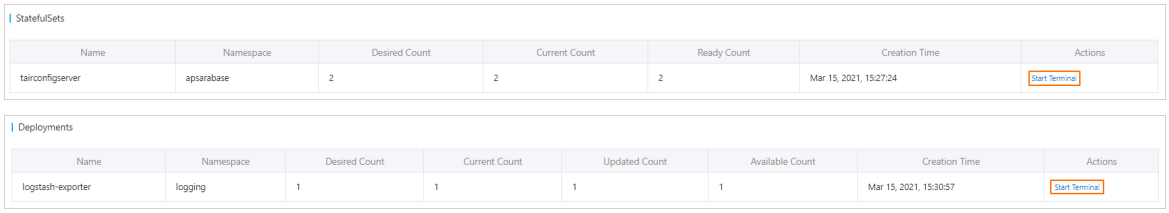

   In the panel that appears, the pods to be logged on are displayed based on the number of the component replicas.
6. Select the pod to which you want to log on, and click **OK** to start the terminal process.

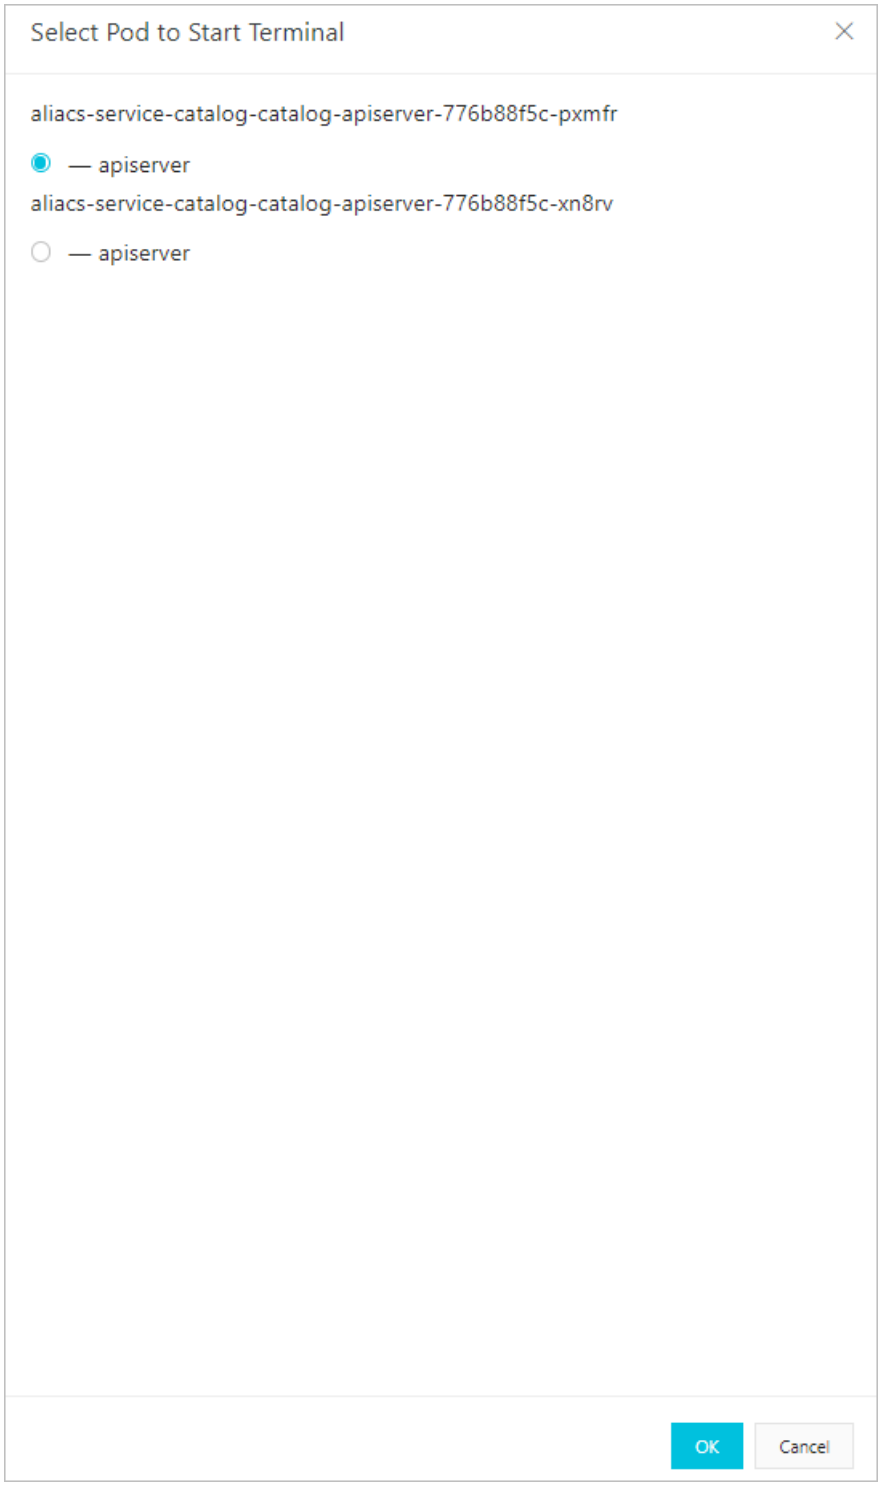

2.1.6.1.7. Perform O&M operations

The O&M Actions page displays the O&M operations that are available to a product. You can also perform O&M operations on this page.

Procedure

1. Log on to the PaaS console.

2. In the left-side navigation pane, choose **Product Center > Products**.
3. On the Products page, find the product for which you want to perform O&M operations and click **Details** in the **Actions** column.
4. In the left-side navigation pane, click **O&M Actions**.
5. Perform O&M operations that are available to the product.

2.1.6.1.8. View a resource report

The Resource Report page displays the information of all resources that a product has requested from the PaaS console. The resource type can be cni (ip), db, vip, dns, and accesskey.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. In the product list, find the product that you want to view and click **Details** in the **Actions** column.
4. In the left-side navigation pane, click **Resource Report**.
5. View the information of resources.

   By default, all resources are displayed. You can click the up and down arrows next to **Resource Owner** to sort resources. You can also click the icon next to **Type** to filter resources.

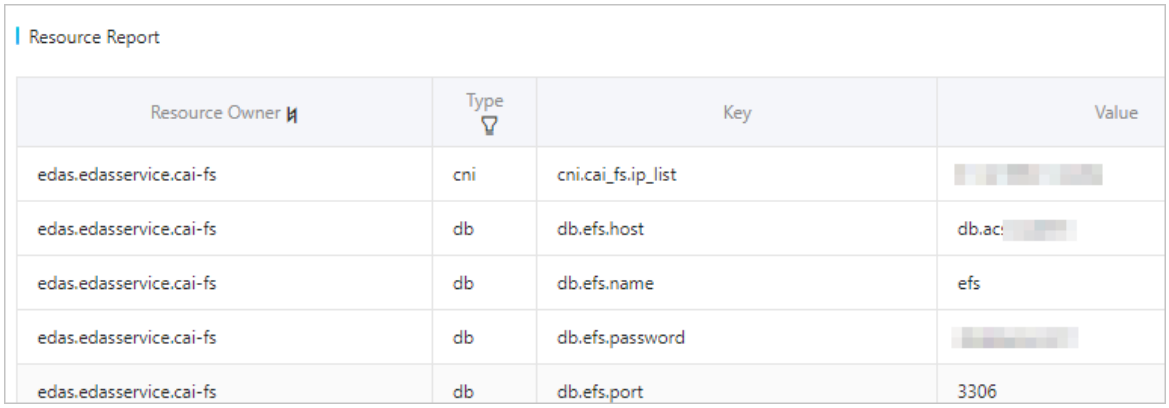

   The following table describes the fields in the resource report.

Field	Description
Resource Owner	The name of the component to which the resource belongs.
Type	The type of the resource.
Key	The attribute name of the resource.
Value	The attribute value of the resource.

2.1.6.1.9. View service registration variables

The Service Registration Variables page displays the values of all service registration variables. You can view the service registration variables of a product. The service registration variables report for a product lists the variables that the product can deliver to other products or components.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Products**.
3. In the product list, find the product that you want to view and click **Details** in the **Actions** column.
4. In the left-side navigation pane, click **Service Registration Variables**.
5. View the information of service registration variables.

   By default, all service registration variables are displayed. You can click the up and down arrows next to **Resource Owner** to sort service registration variables. You can also click the icon next to **Resource Owner** to filter service registration variables.

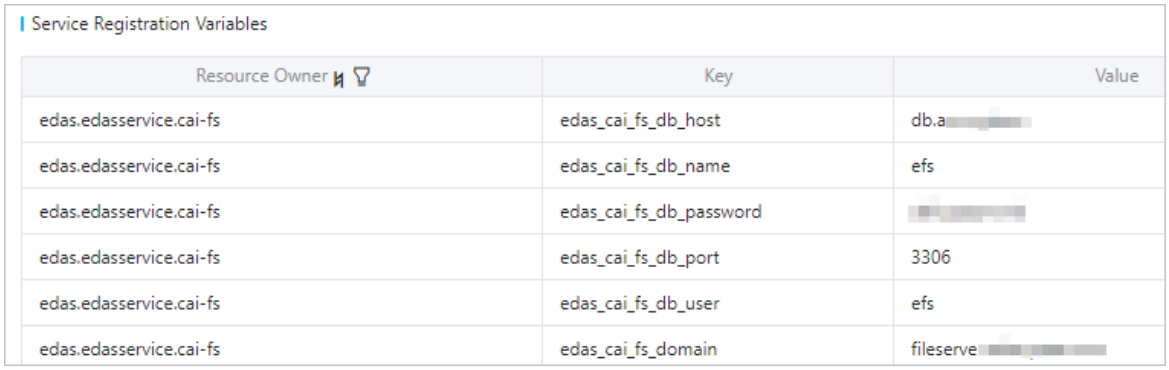
Service Registration Variables

Resource Owner	Key	Value
edas.edasservice.cai-fs	edas_cai_fs_db_host	db.a
edas.edasservice.cai-fs	edas_cai_fs_db_name	efs
edas.edasservice.cai-fs	edas_cai_fs_db_password	
edas.edasservice.cai-fs	edas_cai_fs_db_port	3306
edas.edasservice.cai-fs	edas_cai_fs_db_user	efs
edas.edasservice.cai-fs	edas_cai_fs_domain	fileserve

   The following table describes the fields for service registration variables.

Field	Description
Resource Owner	The name of the component to which the resource belongs.
Key	The variable name that is registered on the configuration management database (CMDB) and can be used by this product or other product components.
Value	The variable value that is registered on the CMDB.

2.1.6.2. Deployment and upgrade

This topic describes how to perform batch upgrade and incremental deployment. You can customize product features when you deploy a product. If the product supports custom configuration, the system goes to the custom configuration page.

Prerequisites

The deployment and upgrade packages are imported to the Apsara Agility PaaS Operations Console. For more information, see Import deployment and upgrade packages to the PaaS Operations Console.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Product Center > Deploy&Upgrade**.

   The **System Packages** page displays the deployment packages that have been imported to the PaaS Operations Console.
3. On the **System Packages** page, find the deployment package that you want to upgrade. Click **Publish** in the **Actions** column to start the deployment or upgrade process.

   > **Note**
   > - If multiple deployment and upgrade packages exist, you can enter a system ID in the search box to search for deployment packages.
   > - In the **Select Products** step, click the number in the **Components** column to view the components and versions of the current product.

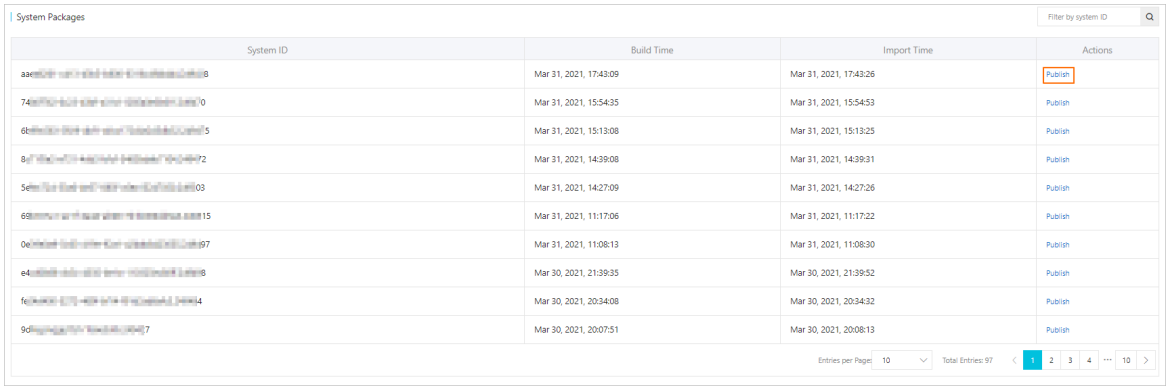

4. Select the required features and click **Next**.

   > **Note**
   >
   > The system can automatically parse dependencies among products. When the Automatic Dependency Processing check box is selected, the system checks whether dependencies exist between the deployed products and the products that you want to deploy. Then, the system selects the products that have dependencies with the products that you want to deploy.
   > - If you want to manually select the products to be deployed, you can clear the **Automatic Dependency Processing** check box.
   > - If you select products that have been deployed, the system upgrades these products.
   > - If you select products that have not been deployed, the system incrementally deploys these products.

   If the custom configuration feature is enabled for a selected product, the **Customize Configurations** step is displayed. Otherwise, the **resource planning** step is displayed.
5. In the **Customize Configurations** step, configure the parameters and click **Save**. Then, click **Next**.
   If the custom configuration feature is enabled for a selected product, the **resource planning** step is displayed. Otherwise, the **Node Planning** step is displayed.
6. In the **resource planning** step, configure the parameters and click **Save**. Then, click **Next**.
7. In the **Node Planning** step, verify that the node planning is correct and click **Next**.

> **Note** If you want to modify the node planning, click the **Reselect** button.

8. In the **Preview** step, check the information of the products to be deployed.

   A type of icon to the left of each item in the **Product&Feature** column indicates a type of deployment state of the product:

   - : the product is to be deployed.
   - : the product has been deployed and does not need to be upgraded.
   - : the product is to be upgraded. You can click the icon to check the differences.

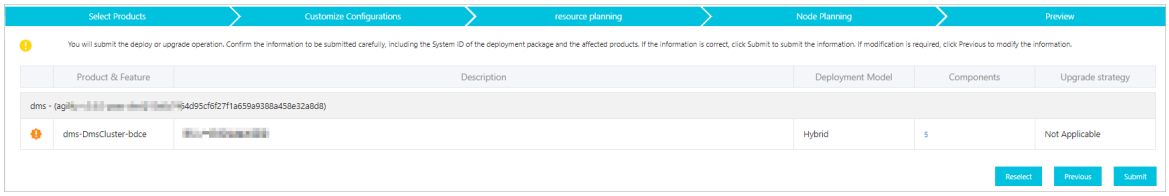

9. Click **Submit** to start the deployment or upgrade process.

   After the deployment or upgrade process starts, you can view the progress on the Task Instances page. To view the progress, choose **Task Center > Task Instances**.

2.1.7. Task center

The Task Center module provides general task management capabilities. You can view and run task templates, and view, suspend, resume, terminate, and delete tasks.

2.1.7.1. Task templates

The Task Templates page lists all task templates, both imported and preset.

2.1.7.1.1. View a task template

You can view information of all task templates on the Task Templates page, such as the name, description, parameters, and workflow definition.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Task Center > Task Templates**.
3. Find the task template that you want to view. Click **View** in the **Actions** column.

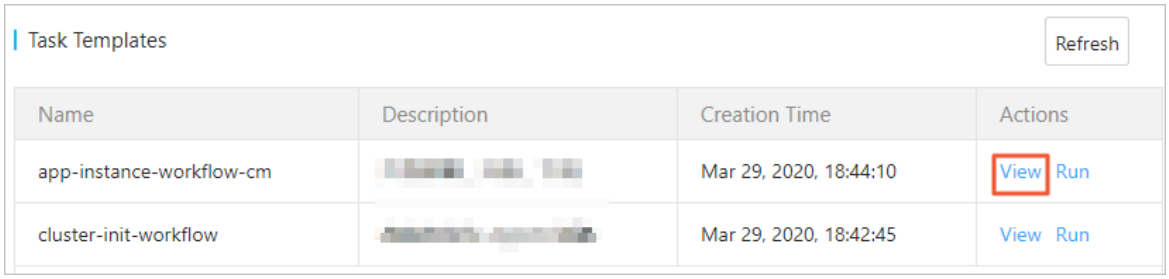

4. In the pane that appears, view the name, description, parameters, and workflow definition of the task template.

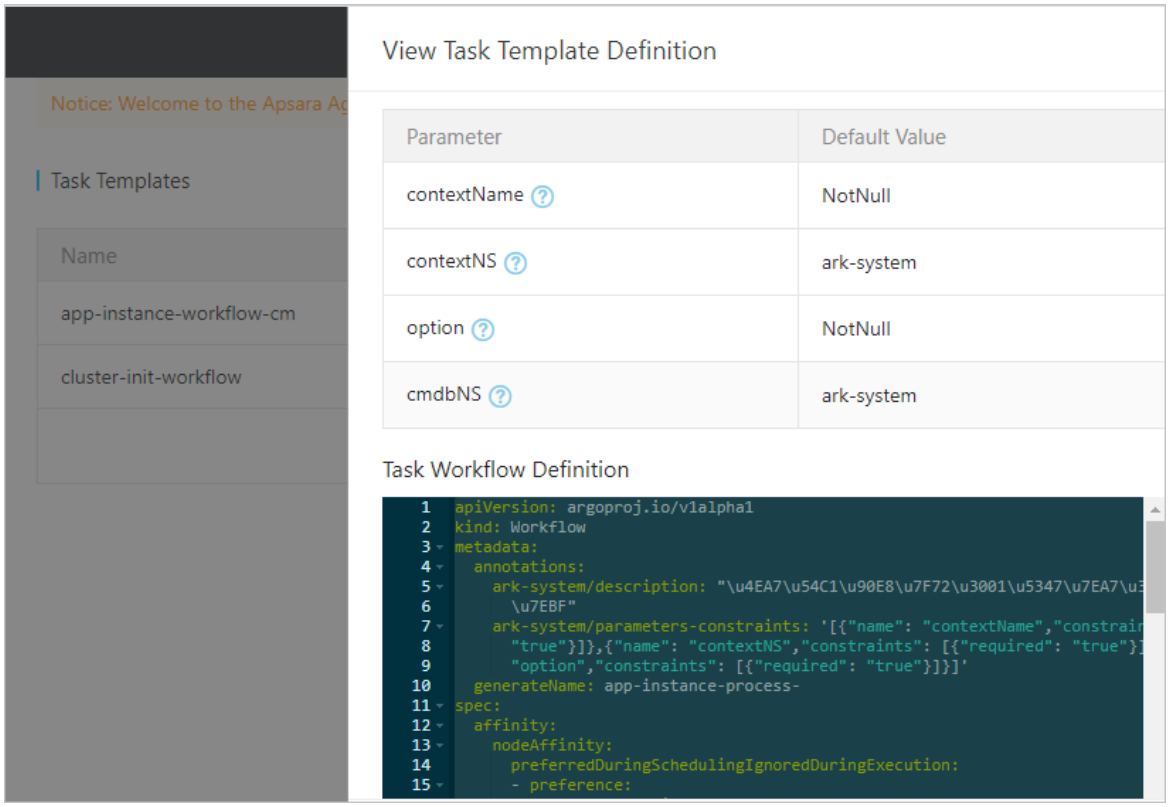

2.1.7.1.2. Run a task

You can run a task on the Task Templates page.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Task Center > Task Templates**.
3. Find the target task template. Click **Run** in the **Actions** column.
4. In the pane that appears, set Task Instance Name and Action Parameters.

> **Note** If the task instance name is not specified, the system automatically generates a task instance name. We recommend that you enter a recognizable name for easy query.

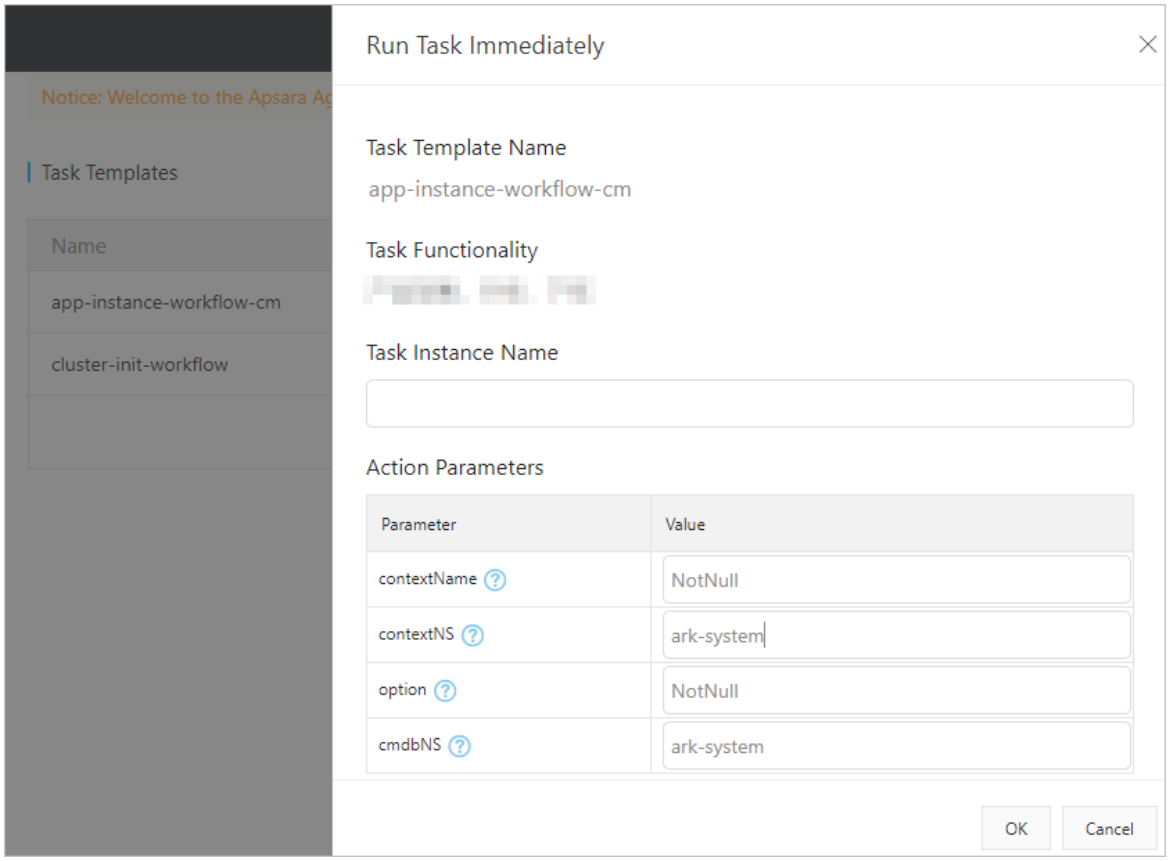

5. Click **OK**.

2.1.7.2. Task instances

The Task Instances page displays information of all tasks. On this page, you can view, suspend, resume, terminate, retry, and delete tasks.

2.1.7.2.1. View task details

After you run a task, you can view the progress, logs, and parameters of the task on the Task Instances page.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Task Center > Task Instances**.
3. In the task instance list, view the state of each task.

   The following states are available:

   - **Succeeded**: The task is successfully executed.
   - **Running**: The task is being executed.
   - **Running (Suspended)**: The task is suspended.
   - **Failed**: The task failed.
   - **Failed (Terminated)**: The task is terminated.

4. Find the task that you want to view and click **View** in the **Actions** column. Then, you are redirected to the **Task Instance Details** page.
5. On the Task Instance Details page, click the task nodes in sequence to view the information and logs of the current task.

> **Note** To view task logs, you can click **LOGS** in the lower-left corner of the panel that appears.

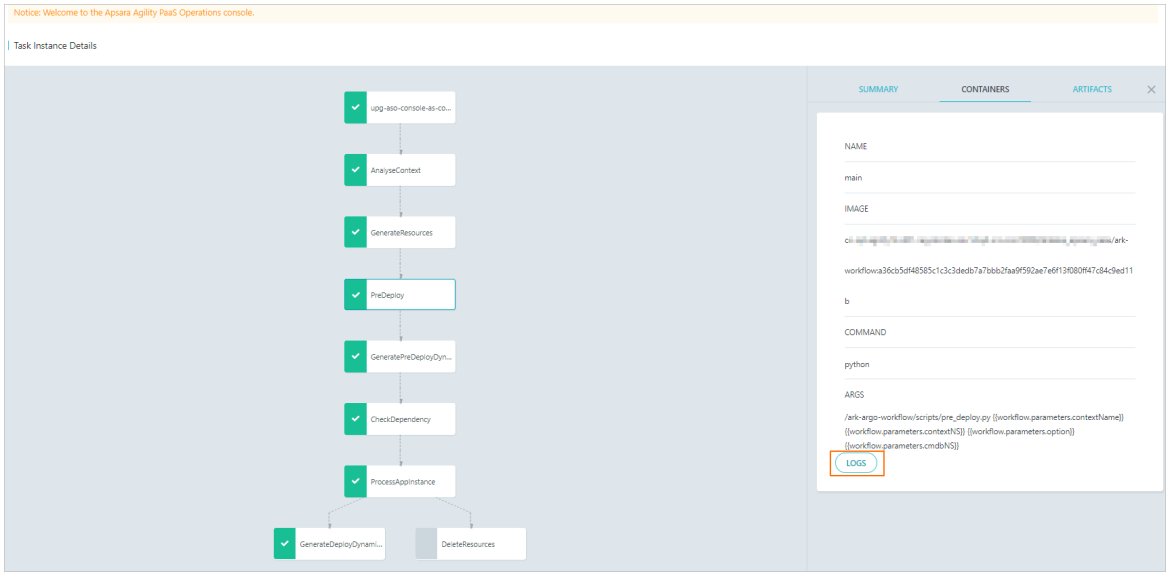

2.1.7.2.2. Suspend a task

You can suspend a task in the Running state. Then, the task status becomes Running (Suspended).

Prerequisites

The task is in the **Running** state.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Task Center > Task Instances**.
3. In the task instance list, find the task in the **Running** state that you want to suspend. Choose **More > Suspend** in the **Actions** column.

   After a successful operation, the task status changes from **Running** to **Running (Suspended)** in the **Status** column.

2.1.7.2.3. Resume a task

After a task is suspended, the task is in the Running (Suspended) state. Then, you can click Resume in the Actions column to resume the task.

Procedure

1. Log on to the PaaS console.

2. In the left-side navigation pane, choose **Task Center > Task Instances**.
3. In the task instance list, find the task in the **Running (Suspended)** state that you want to resume. Choose **More > Resume** in the **Actions** column.

   After the task is resumed, the task state changes from **Running (Suspended)** to **Running** in the **Status** column.

2.1.7.2.4. Delete a task

You can delete a task in any state. If a task is in the Running state, this operation enables the system to immediately terminate the task and delete the task record. If a task is in a state other than Running, this operation enables the system to immediately delete the task record.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Task Center > Task Instances**.
3. In the task instance list, find the task that you want to delete. Click **Delete** in the **Actions** column.

2.1.8. Platform management

Apsara Agility PaaS Operations Console allows you to configure virtual IP addresses (VIPs) for your SLB instance. This way, you can access cloud services that are deployed in the PaaS Operations Console.

2.1.8.1. Modify VIP addresses in the configuration file

Apsara Agility PaaS Operations Console allows you to import external virtual IP (VIP) addresses. This topic describes how to modify VIPs in a configuration file.

Prerequisites

Base services are at the desired state.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Platforms** > **VIP configuration**.
3. Export the configuration file that you want to modify.

   In the upper-right corner of the page, click **Export configuration** to save the configuration file to your on-premises machine.
4. Modify VIP addresses in the configuration file
5. Import the configuration file.

   In the upper-right corner of the page, click **Import configuration** to upload the modified configuration file to the PaaS Operations Console.

2.1.8.2. Enable a VIP configuration

This topic describes how to enable a virtual IP address (VIP) configuration.

Prerequisites

The configuration file that stores the modified VIP addresses is imported to the PaaS Operations Console.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Platforms** > **VIP configuration**.
3. On the **Load balancing service list** page, find the service for which you want to enable a VIP configuration.
4. Click **Enable** in the **Actions** column.

2.1.9. Platform diagnostics

The PaaS console provides platform-level diagnostics. This module collects information about the console and products deployed in the console, presents summary diagnostic results, and allows you to download detailed diagnostic results. The module aims to improve user experience of diagnostics.

2.1.9.1. Diagnostic items

The Diagnostic Items page displays all diagnostic items in the PaaS console. On this page, you can view, execute, and delete diagnostic items.

2.1.9.1.1. View a diagnostic item

You can view details about the current diagnostic item, such as the name, type, description, start time, deletion protection, and definition.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics** > **Diagnostic Items**.
3. Find the diagnostic item that you want to view. Click **View** in the **Actions** column.
4. In the panel that appears, view details of the diagnostic item.

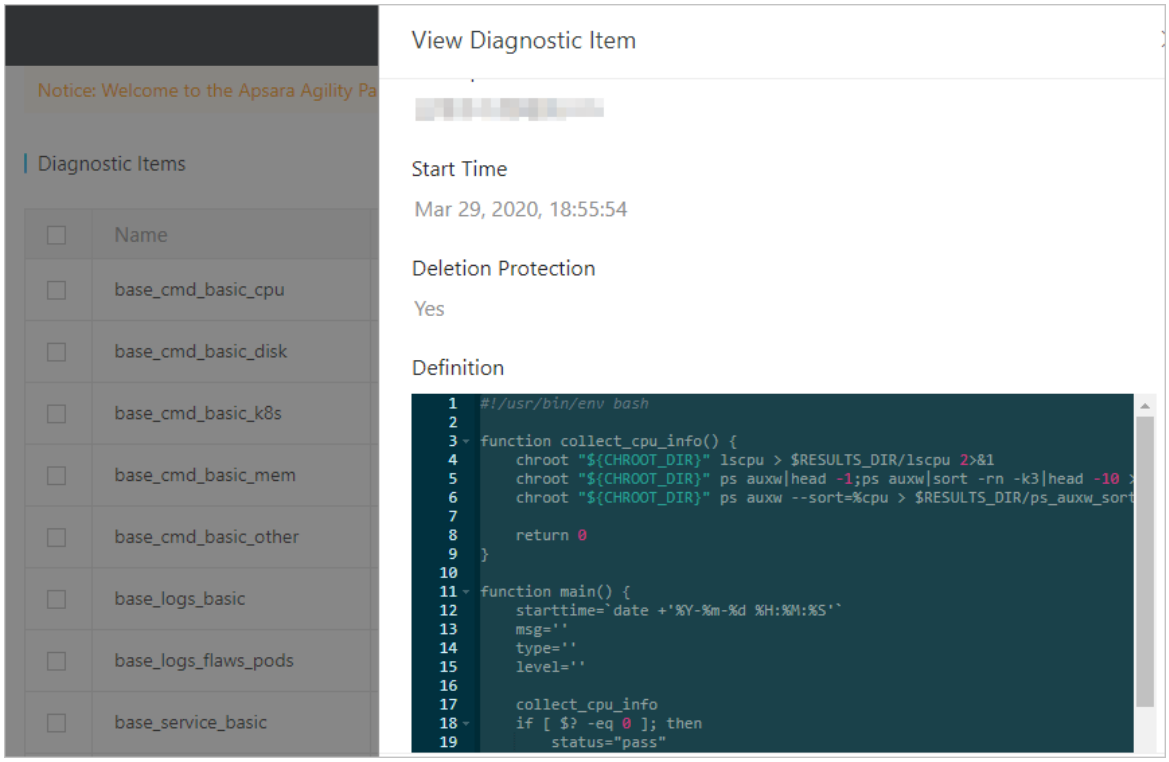

2.1.9.1.2. Execute diagnostic items

On the Diagnostic Items page, you can execute diagnostic items.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Items**.
3. Select one or more diagnostic items and click **Submit Diagnosis**.

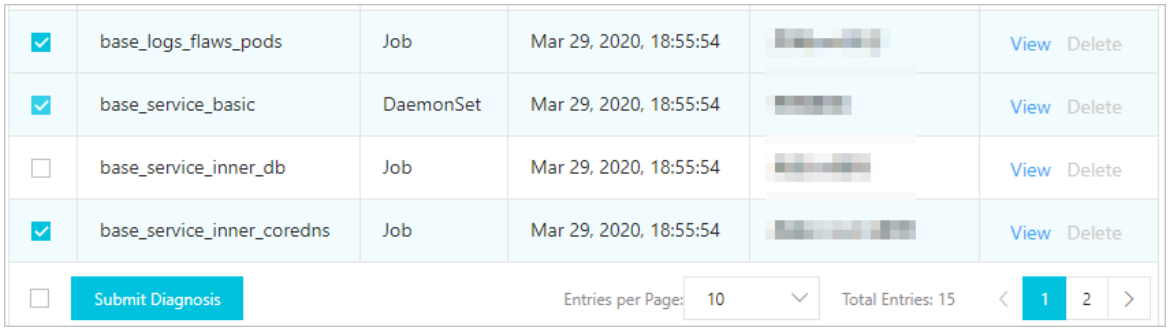

4. In the message that appears, click **OK**.

2.1.9.1.3. Delete a diagnostic item

You can delete a diagnostic item. You can only delete imported diagnostic items, but not the diagnostic items preset by the system.

Procedure

1. Log on to the PaaS console.

2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Items**.
3. Find the diagnostic item that you want to delete. Click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

2.1.9.2. Diagnostic tasks

The Diagnostic Tasks page displays all diagnostic tasks. On this page, you can view diagnostic progress, view diagnostic reports, download diagnostic reports, terminate diagnostic tasks, and delete diagnostic tasks.

2.1.9.2.1. View diagnostic progress

After you start a diagnostic task, you can view its diagnostic progress on the Diagnostic Tasks page.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Tasks**.
3. Find the diagnostic task that you want to view. Click **Diagnostic Progress** in the **Actions** column.
4. On the **Diagnostic Progress** page, click the task nodes in sequence to view the diagnostic progress and logs of the current diagnostic task.

2.1.9.2.2. View a diagnostic report

After a diagnostic task is complete, you can view its diagnostic report.

Prerequisites

You can view the diagnostic report only for a diagnostic task in the **Succeeded** state.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Tasks**.
3. Find the diagnostic task that you want to view. Click **View Report** in the **Actions** column.
4. In the pane that appears, view the diagnostic results, including the name, status, and details.

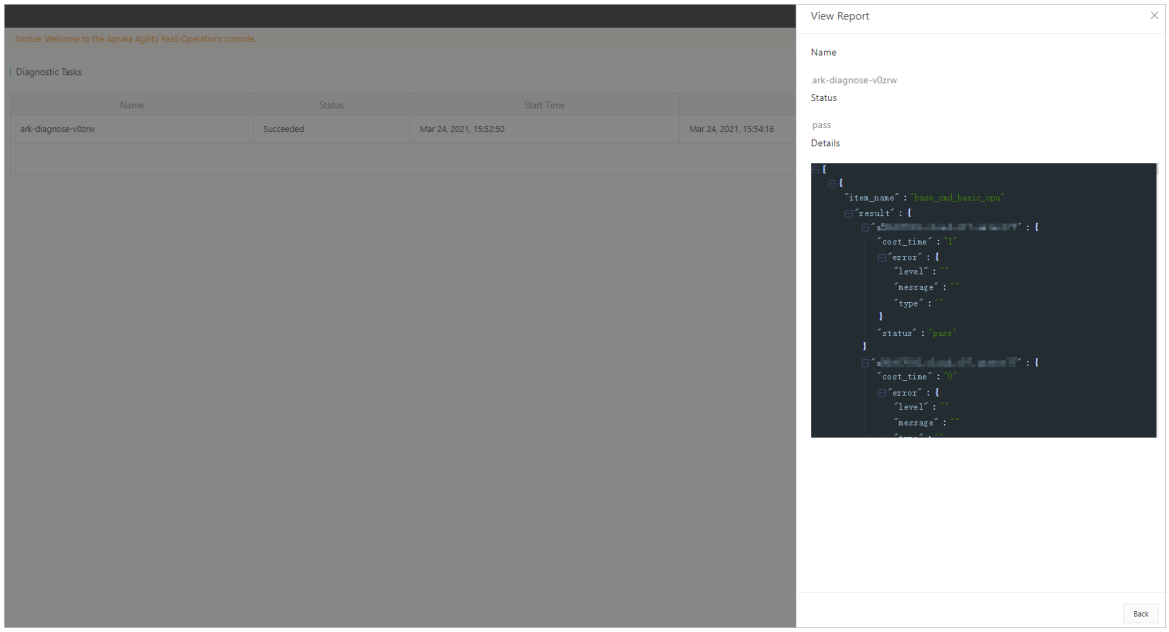

2.1.9.2.3. Download a diagnostic report

After a diagnostic task is complete, you can download its diagnostic report to your on-premises machine for offline query and analysis.

Prerequisites

You can download the diagnostic report only for a diagnostic task in the **Succeeded** state.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Tasks**.
3. Find the diagnostic task that you want to manage. Click **Download** in the **Actions** column.

2.1.9.2.4. Terminate a diagnostic task

You can terminate a diagnostic task in the Running state.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Tasks**.
3. Find the diagnostic task that you want to terminate. Choose **More > Stop** in the **Actions** column.
4. In the message that appears, click **OK**.

2.1.9.2.5. Delete a diagnostic task

You can delete a diagnostic task that is no longer needed.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Platform Diagnostics > Diagnostic Tasks**.
3. Find the diagnostic task that you want to delete. Choose **More > Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

2.1.10. Alerts

The **Alerts** module implements unified management of alerts in the PaaS console. You can view alert rules, notification channels, and alert events. You can also configure alert rules and notification channels in the **Alerts** module.

2.1.10.1. Alert rule groups

An alert rule must belong to an alert rule group. You can create alert rule groups and add alert rules to alert rule groups.

2.1.10.1.1. Create an alert rule group

You can create an alert rule group. When you create an alert rule group, you must add an alert rule to the group.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Groups**.
3. (Optional)In the upper part of the Rule Groups page, select the cluster that you want to manage from the drop-down list.
4. In the upper-right corner of the page, click **Create Rule Group**.
5. In the **Create Rule Group** dialog box, configure the parameters.

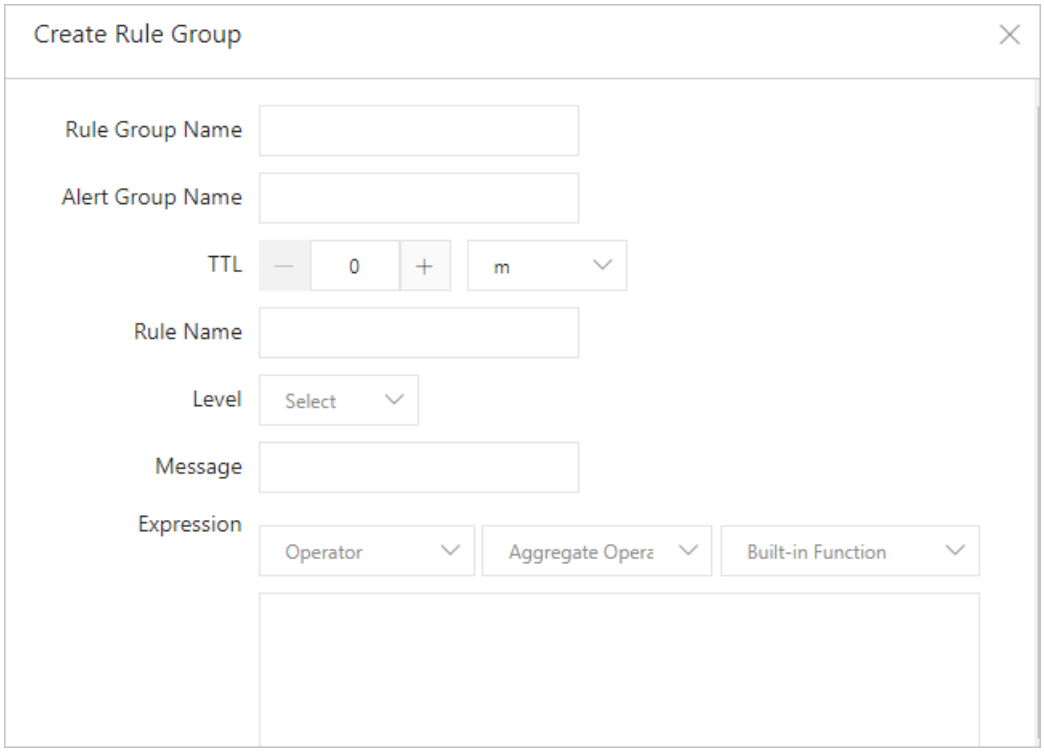

Parameter	Description
Rule Group Name	The globally unique name of the alert rule group.
Alert Group Name	The globally unique name of the alert group. An alert rule group must have an alert group.
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: hours. ○ **m**: minutes. ○ **s**: seconds.
Rule Name	The globally unique name of the alert rule.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

6. Click **Submit**.

2.1.10.1.2. Create an alert rule

After you create an alert rule group, you can add an alert rule to the group.

Prerequisites

An alert rule group is created. For more information about how to create an alert rule group, see Create an alert rule group.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Groups**.
3. (Optional)In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. Find rule group for which you want to add alert rules. Click **Create Rule Group** in the **Actions** column.
   The **Rules** page appears. You can view all alert rules in the alert rule group.

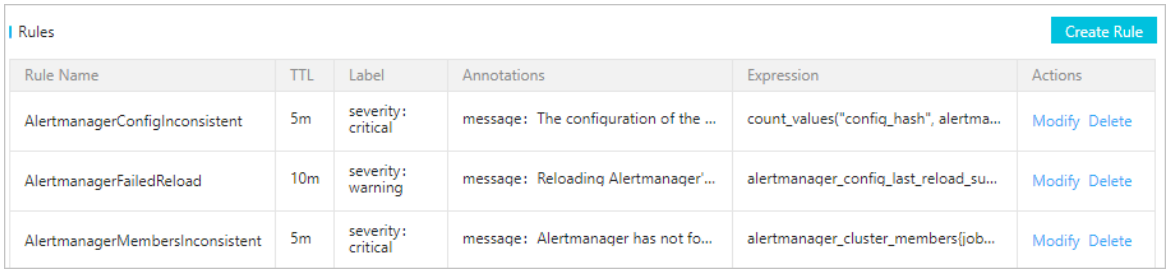

5. In the upper-right corner of the page, click **Create Rule**.
6. In the **Create Rule** dialog box, configure the parameters.

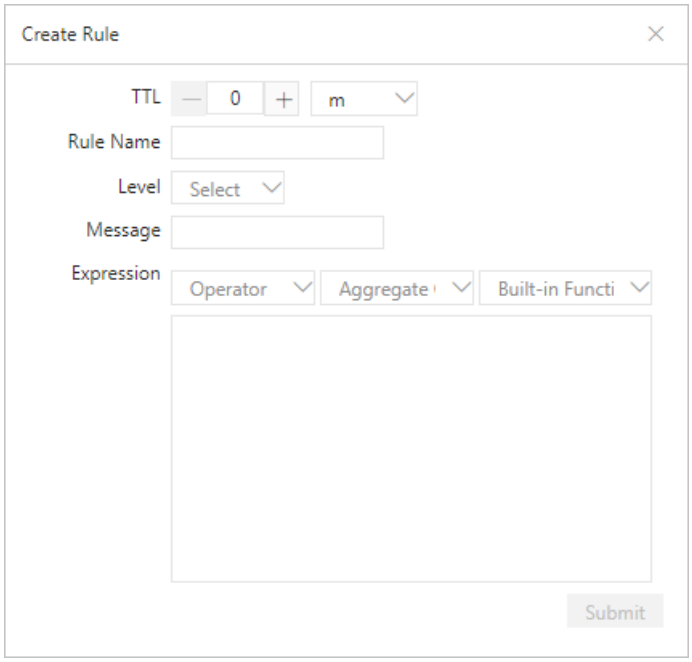

Parameter	Description
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: hours. ○ **m**: minutes. ○ **s**: seconds.
Rule Name	The globally unique name of the alert rule.
Level	The severity of the alert. Valid values: ○ **Warning**: warning alert. ○ **Critical**: critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

7. Click **Submit**.

2.1.10.1.3. Modify an alert rule

You can modify an alert rule.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Groups**.
3. (Optional)In the upper part of the Rule Groups page, select the cluster that you want to manage from the drop-down list.
4. On the **Rule Groups** page, view all alert rule groups defined in the system.
5. Find the rule group that the rule belongs. Click **Modify Rule** in the **Actions** column.
6. On the **Rules** page, view all alert rules in the rule group.
7. Find the rule that you want to modify. Click **Modify** in the **Actions** column.
8. Modify the TTL, Level, Message, and Expression parameter settings of the alert rule.

Parameter	Description

Parameter	Description
TTL	The length of time that an error persists before an alert is sent. ○ **h**: hours. ○ **m**: minutes. ○ **s**: seconds.
Level	The severity of the alert. Valid values: ○ **Warning** ○ **Critical**
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

9. Click **Submit**.

2.1.10.1.4. Delete an alert rule

You can delete an alert rule that is no longer needed from an alert rule group.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Groups**.
3. (Optional)In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. On the **Rule Groups** page, view all alert rule groups defined in the system.
5. Find the rule group for the target rule. Click **Modify Rule** in the **Actions** column.
6. On the **Rules** page, view all alert rules in the rule group.
7. Find the target rule. Click **Delete** in the **Actions** column.
8. In the Disable Alert Rule message, click **Confirm**.

2.1.10.1.5. Delete an alert rule group

You can delete an alert rule group that is no longer needed.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Alert Groups**.

The **Rule Groups** page appears.

3. (Optional)In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. Find the rule group that you want to delete. Click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

2.1.10.2. Notification channels

You can view and modify notification channel settings on the Notification Channels page.

2.1.10.2.1. View notification channel settings

You can view the current notification channel settings.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Notification Channels**.
3. In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. In the **Global Settings**, **Routing**, and **Receiver** sections, view the relevant information.

2.1.10.2.2. Modify notification channel settings

You can modify notification channel settings such as global settings, routing, and receivers.

2.1.10.2.2.1. Modify global settings

You can modify global settings, such as the resolve_timeout, smtp_info, and notifications settings.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Notification Channels**.
3. In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. In the upper-right corner of the page, click **Edit**.
5. In the **Global Settings** section, modify the resolve_timeout, smtp_info, and notifications settings.

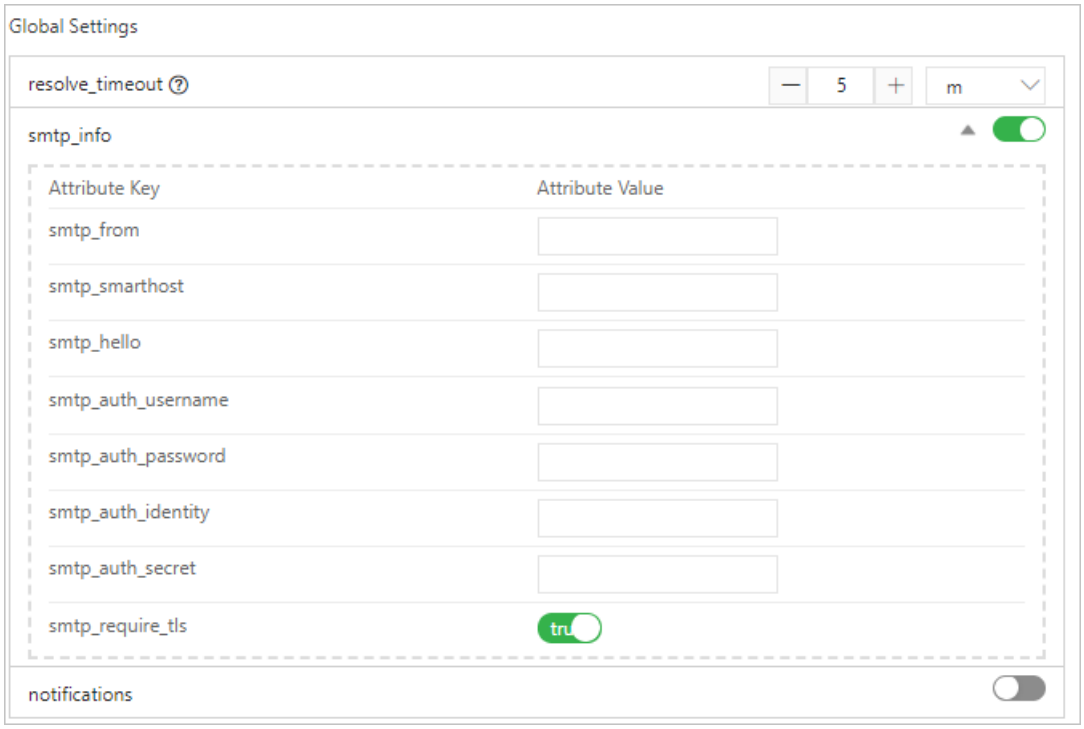

Item	Description
resolve_timeout	Specifies the time period before an alert is marked as resolved if the Alertmanager does not receive further notifications of the alert.
smtp_info	Specifies global SMTP information. To modify the settings, turn on the switch on the right and then click the Show icon. You can configure the following parameters: ○ **smtp_from**: the source email address that is used to send alerts. ○ **smtp_smarthost**: the SMTP server endpoint and port number for the source email address used to send alerts. Example: smtp_smarthost:smtp.example.com:465. ○ **smtp_hello**: the default hostname that identifies the SMTP server. ○ **smtp_auth_username**, **smtp_auth_password**: the username and password for the source email address that is used to send alerts. ○ **smtp_auth_identity**: specifies the PLAIN SMTP authentication method. ○ **smtp_auth_secret**: specifies the CRAM-MD5 SMTP authentication method. ○ **smtp_require_tls**: the default SMTP TLS configuration. By default, this parameter is set to **true**. However, starttls errors may occur if the parameter is set to **true**. We recommend that you set the parameter to **false**.

Item	Description
notifications	The alert channels. The PaaS Operations Console allows you to configure alert channels including Slack, VictorOps, PagerDuty, Opsgenie, HipChat, and WeChat. To modify the settings, turn on the switch on the right and then click the Show icon. You can configure the following parameters: ○ **slack_api_url**: the API URL for Slack notifications. ○ **victorops_api_key**: the VictorOps API key. ○ **victorops_api_url**: the VictorOps API URL. ○ **pagerduty_url**: the destination URL for API requests. ○ **opsgenie_api_key**: the Opsgenie API key. ○ **opsgenie_api_url**: the destination URL for Opsgenie API requests. ○ **hipchat_api_url**: the source URL for API requests. ○ **hipchat_auth_token**: the authentication token. ○ **wechat_api_url**: the WeChat API URL. ○ **wechat_api_secret**: the WeChat API key. ○ **wechat_api_corp_id**: the WeChat API corporate ID.

6. In the upper-right corner of the page, click **Save**.

2.1.10.2.2.2. Modify routing settings

You can modify global routing settings, and create or delete sub-routes.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Notification Channels**.
3. In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. In the upper-right corner of the page, click **Edit**.
5. In the **Routing** section, perform the following operations:
   - Modify routing settings

     You can modify the default route or subroutes.

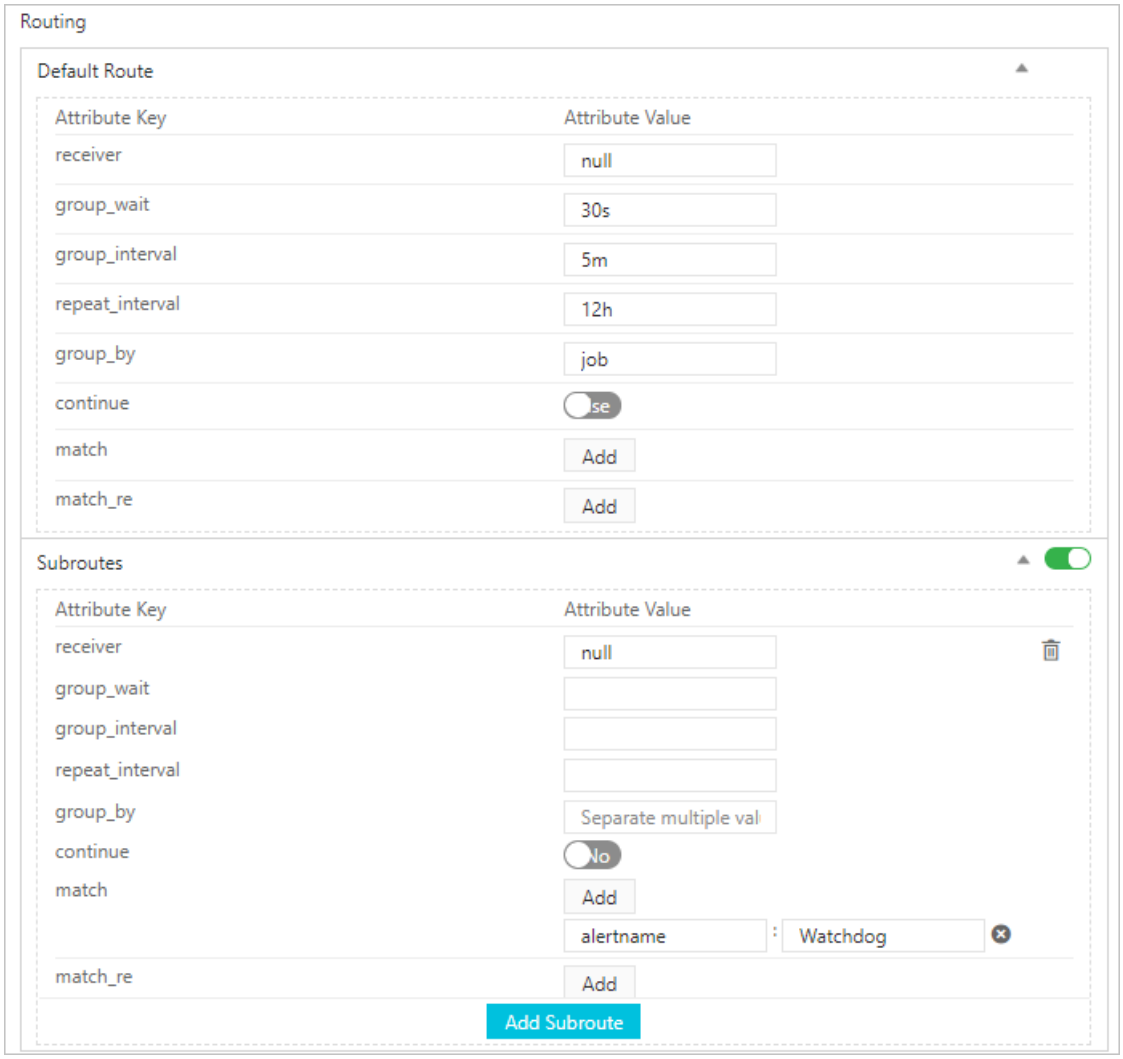

Item	Description

Item	Description
Default Route	The global route. You can configure the route information based on the actual environment. ■ **receiver**: the name of the alert receiver. ■ **group_wait**: specifies the waiting time to initialize a message when a new alert group is created. This method ensures that the system can have enough time to obtain multiple alerts for the same alert group, and then trigger an alert message. ■ **group_interval**: specifies the waiting time to send a new alert message. ■ **repeat_interval**: specifies the waiting time to resend an alert message. ■ **group_by**: the tag list. It is the regrouping tag list after alert messages are received. For example, all received alert messages that contain the `cluster=A` and `alertname=Latncy High` tags are aggregated into a group. ■ **continue**: specifies whether an alert matches subsequent nodes. ■ **match**: Click **Add** and specify a receiver for matched alerts. ■ **match_re**: Click **Add**. Enter a regular expression and specify a receiver for alerts that match the regular expression.
Subroutes	Configure subroutes in a similar way to the global route, so that you can export an alert type to another location. ■ **receiver**: the name of the alert receiver. ■ **group_wait**: specifies the waiting time to initialize a message when a new alert group is created. This method ensures that the system can have enough time to obtain multiple alerts for the same alert group, and then trigger an alert message. ■ **group_interval**: specifies the waiting time to send a new alert message. ■ **repeat_interval**: specifies the waiting time to resend an alert message. ■ **group_by**: the tag list. It is the regrouping tag list after alert messages are received. For example, all received alert messages that contain the `cluster=A` and `alertname=Latncy High` tags are aggregated into a group. ■ **continue**: specifies whether an alert matches subsequent nodes. ■ **match**: click **Add**. Enter the key and value of a tag and specify a receiver for alerts that match the tag. ■ **match_re**: click **Add**. Enter a regular expression based on the key and value of a tag and specify a receiver for alerts that match the regular expression.

- Create a subroute

  To export an alert type to another location, you can click **Add Subroute** in the lower part of the **Routing** section to configure a new subroute.

- Delete a subroute

  In the **Routing** section, find a subroute that is no longer needed and click the Delete icon to delete the subroute.

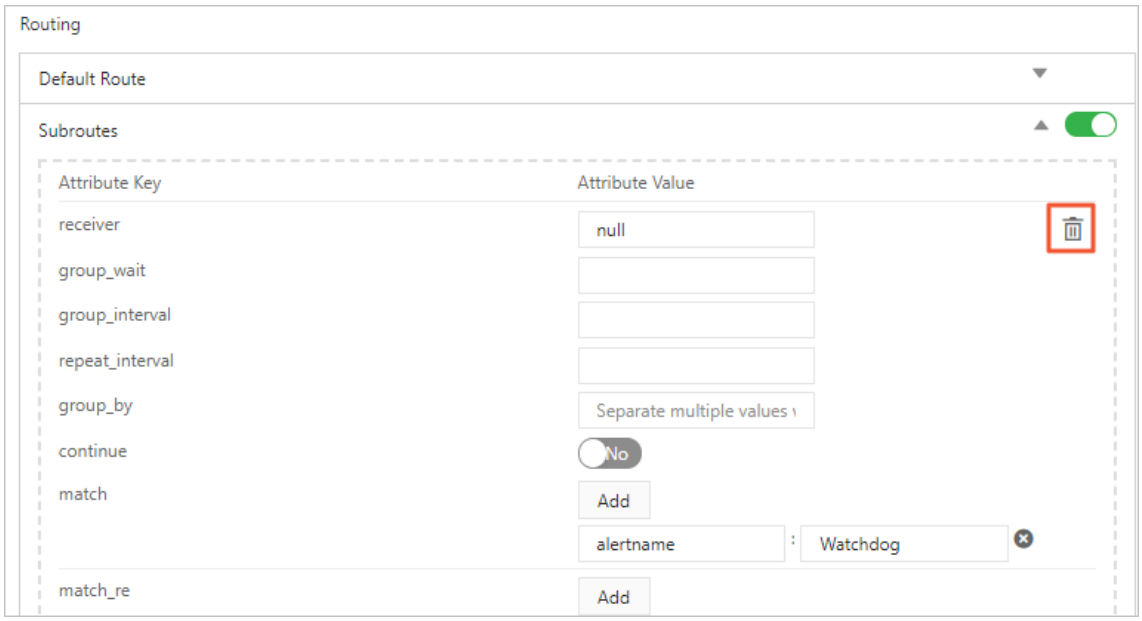

6. In the upper-right corner of the page, click **Save**.

2.1.10.2.2.3. Modify receiver settings

You can create, modify, or delete alert receiver settings.

Procedure

1. Log on to the PaaS console.
2. In the left-side navigation pane, choose **Alerts > Notification Channels**.
3. In the upper part of the page, select the cluster that you want to manage from the drop-down list.
4. In the upper-right corner of the page, click **Edit**.
5. In the **Receivers** section, perform the following operations:
   - Modify receiver settings

     Modify the name and type of a receiver.

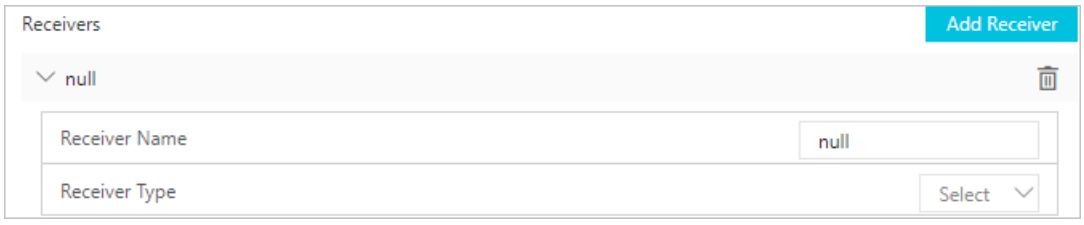

Item	Description
Receiver Name	The name of the alert receiver.
Receiver Type	Valid values for **Receiver Type**: **webhook** and **email**. If Receiver Type is set to **webhook**, you must configure the following parameters: ■ **url**: the URL of the alert receiver. ■ **send_resolved**: specifies whether to send messages for resolved alerts. Default value: **No**. If Receiver Type is set to **email**, you must configure the following parameters: ■ **send_resolved**: specifies whether to send messages for resolved alerts. Default value: **No**. ■ **to**: the destination email address for alerts. ■ **from**: the source email address that is used to send alerts. ■ **smarthost**: the server address and port number for the source email address that is used to send alerts. ■ **hello**: the default hostname that identifies the email server. ■ **auth_username**: the username for the source email address that is used to send alerts. ■ **auth_password**: the password for the source email address that is used to send alerts. ■ **auth_secret**: specifies the CRAM-MD5 authentication method. ■ **auth_identity**: specifies the PLAIN authentication method. ■ **require_tls**: the default TLS configuration. The default value is **Yes**. However, errors may occur if the parameter is set to **Yes**. We recommend that you set the parameter to **No**.

- Add a receiver

  In the upper-right corner of the **Receivers** section, click **Add Receiver**. Configure the parameters.

- Delete a receiver

  In the **Receivers** section, find the receiver that you want to delete and click the Delete icon to delete a receiver that is no longer needed.

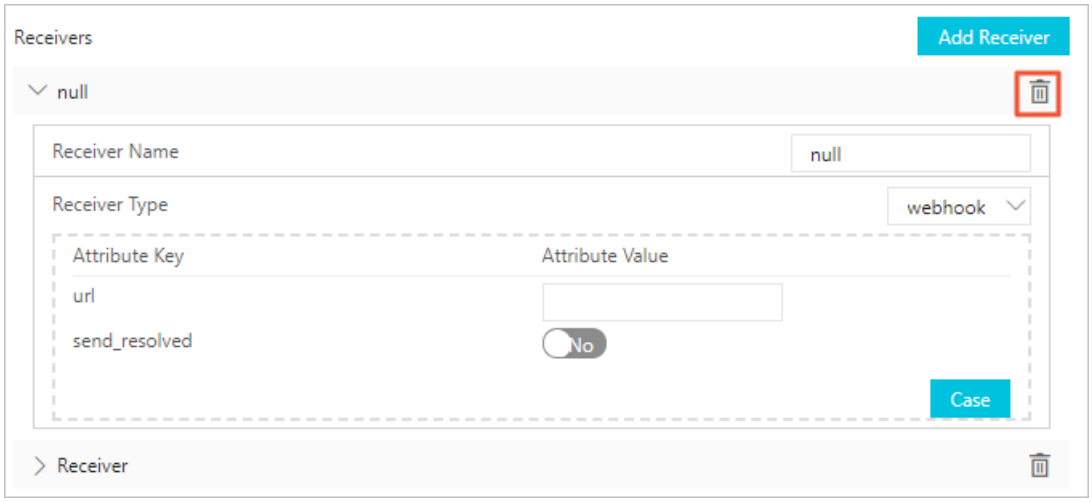

6. In the upper-right corner of the page, click **Save**.

2.1.11. Query history events

You can query the operation records that are generated in the Apsara Agility PaaS Operations Console. The read and write records that are generated in the PaaS Operations Console are stored in Elasticsearch. The records include the operation type, operator, source IP address, event name, operation object, operation time, request parameters, and interface endpoint.

Procedure

1. Log on to the PaaS Operations Console.
2. In the left-side navigation pane, choose **Operation audit** > **Historical event query**.
3. (Optional) In the **Search** section, query the records of the history events.
   i. Click **Advanced Search** to expand the **search** area.
   ii. Set the filter conditions including the operation type, start time, end time, operator, operation name, and operation object.
   iii. Click **Search**.

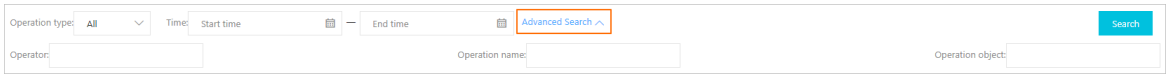

4. In the history event list, find the history event that you want to view.

   View the **operation type**, **operator**, **source IP address**, **event name**, **operation object**, and **operation time** of the history event.
5. Click the + icon to the left of the history event to view the request parameters and interface endpoint.

2.1.12. Appendix

2.1.12.1. Import deployment and upgrade packages to the PaaS Operations Console

Before you deploy or upgrade products, you must import the deployment and upgrade packages to the Apsara Agility PaaS Operations Console.

1. Upload the installation disk used for deployment and upgrade to the bootstrap node in the onsite environment.
2. Log on to the bootstrap node by using SSH.
3. Run the following command to import deployment packages and generate a deployment package list:

```
sh upgrade.sh [$Packages_Path].iso
```

> Note
>
> Replace {packages -path}.iso with the actual storage path of the iso file on the installation disk.

2.1.12.2. Exception troubleshooting for inspection cases

When an inspection case fails to be executed, you can find solutions by analyzing the possible causes and impact scope.

2.1.12.2.1. check-k8s-dns-hostnet

This topic describes how to troubleshoot exceptions that occur when the check-k8s-dns-hostnet inspection case is executed.

Inspection item CURL_DNS_SERVER_NAME

The following table describes the information related to the CURL_DNS_SERVER_NAME inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the server can be accessed by using the domain name that DNS provides.
Possible cause	• The Kubernetes service is abnormal. • The Kubernetes network is abnormal.
Impact scope	DNS is unavailable for pods on some nodes.

Item	Description
Solution	1. Check whether the pods that use DNS run normally. 2. Check whether the network is connected.

Inspection item CURL_DNS_SERVER_IP

The following table describes the information related to the CURL_DNS_SERVER_IP inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the server can be accessed by using the IP address that DNS provides.
Possible cause	• The Kubernetes service is abnormal. • The Kubernetes network is abnormal.
Impact scope	DNS is unavailable for pods on some nodes.
Solution	1. Check whether the pods that use DNS run normally. 2. Check whether the network is connected.

Inspection item CURL_DNS_ENDPOINT

The following table describes the information related to the CURL_DNS_ENDPOINT inspection item.

Item	Description
Product	K8s
Inspection level	Critical

Item	Description
Inspection description	This inspection item is used to check whether the server can be accessed by using the backend IP address that DNS provides.
Possible cause	The Kubernetes network is abnormal.
Impact scope	DNS is unavailable for pods on some nodes.
Solution	Contact Alibaba Cloud technical support.

2.1.12.2.2. check-docker-overlay-mount

This topic describes how to troubleshoot exceptions that occur when the check-docker-overlay-mount inspection case is executed.

Inspection item CreateContainerError_Check

The following table describes the information related to the CreateContainerError_Check inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether a container can be created.
Possible cause	The overlay file system that Docker uses cannot be mounted.
Impact scope	Pods on some nodes cannot be created.

Item	Description

Solution	1. Log on to the corresponding master node by using SSH.
	2. Run the following command to stop Docker:
	`sudo systemctl stop docker`
	3. Run the following command to delete the files that are stored in the following path:
	`rm -rf /var/lib/docker`
	4. Use the `vim` command to edit the `/etc/docker/daemon.json` file.
	`{ "storage-driver": "overlay" }`
	5. Run the following command to start Docker:
	`sudo systemctl start docker`
	Warning
	This may introduce security risks. Contact Alibaba Cloud technical support.

2.1.12.2.3. check-k8s-apiserver-crash

This topic describes how to troubleshoot exceptions that occur when the check-k8s-apiserver-crash inspection case is executed.

Inspection item K8S_APISERVER_CHECK

The following table describes the information related to the K8S_APISERVER_CHECK inspection item.

Item	Description
Product	K8s
Inspection level	Critical

Item	Description
Inspection description	This inspection item is used to check whether the Kubernetes API servers run normally.
Possible cause	Some API server processes may be out of the control of Kubernetes.
Impact scope	Some API server has exceptions, which results in kubectl being unavailable.

Solution	1. Identify the node on which the pod of the abnormal API servers is located.
	2. Log on to the node by using SSH. Run the following command to check API servers for any exceptions.
	`docker ps -a |grep kube-apiserver`
	○ If API servers have exceptions, run the following commands to clear the abnormal API servers:
	`docker stop <$ApiServer>`
	`docker rm <$ApiServer>`
	○ If API servers do not have exceptions, run the following command to check whether the API servers that are not managed by Docker exist:
	`ps aux |grep kube-apiserver`
	If the API servers that are not managed by Docker exist, stop the abnormal API servers.
	3. If the problem persists after you perform the preceding steps, contact Alibaba Cloud technical support.

2.1.12.2.4. check-k8s-cs

This topic describes how to troubleshoot exceptions that occur when the check-k8s-cs inspection case is executed.

Inspection item K8S_CS_CHECK

The following table describes the information related to the K8S_CS_CHECK inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	Check whether the Kubernetes components are normal.
Possible cause	• Etcd is abnormal. • The Kubernetes API server is abnormal. • The Kubernetes scheduler is abnormal. • The Kubernetes controller manager is abnormal.
Impact scope	Kubernetes is unavailable, which results in kubectl being unavailable.
Solution	• Analyze exceptions of Etcd by viewing the state and logs. • Analyze exceptions of the Kubernetes API server by viewing the state and logs. • Analyze exceptions of the Kubernetes scheduler by viewing the state and logs. • Analyze exceptions of the Kubernetes controller manager by viewing the state and logs.

Inspection item MASTER_POD_CHECK

The following table describes the information related to the MASTER_POD_CHECK inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	Check whether Kubernetes master pods are normal.

Item	Description
Possible cause	• Etcd is abnormal. • The Kubernetes API server is abnormal. • The Kubernetes scheduler is abnormal. • The Kubernetes controller manager is abnormal.
Impact scope	Kubernetes is unavailable, which results in kubectl being unavailable.
Solution	• Analyze exceptions of Etcd by viewing the state and logs. • Analyze exceptions of the Kubernetes API server by viewing the state and logs. • Analyze exceptions of the Kubernetes scheduler by viewing the state and logs. • Analyze exceptions of the Kubernetes controller manager by viewing the state and logs.

2.1.12.2.5. check-kube-proxy-pod

This topic describes how to troubleshoot exceptions that occur when the check-kube-proxy-pod inspection case is executed.

Inspection item KUBE_PROXY_CHECK

The following table describes the information related to the KUBE_PROXY_CHECK inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the kube-proxy pod runs normally.
Possible cause	The forwarding policies of iptables are incorrectly configured.

Item	Description
Impact scope	• The Kubernetes services cannot be used. • Applications within a Kubernetes cluster cannot be interconnected. • Kubernetes clusters cannot provide external services.
Solution	1. Log on to the corresponding node and run the following command to check Proxy pods for any exceptions: `kubectl -n kube-system get pod | grep kube-proxy` 2. Run the `kubectl logs` command to query the node logs. 3. If the problem persists after you perform the preceding steps, contact Alibaba Cloud technical support.

2.1.12.2.6. check-network-control-plane

This topic describes how to troubleshoot exceptions that occur when the check-network-control-plane inspection case is executed.

Inspection item NETWORK_CONTROL_PLANE_CHECK

The following table describes the information related to the NETWORK_CONTROL_PLANE_CHECK inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the network control plane run normally.
Possible cause	The network configuration has errors in the IaaS layer.

Item	Description
Impact scope	• The network of the Kubernetes pod is unavailable. • Kubernetes applications cannot communicate with each other.
Solution	1. Check the pod logs of the unhealthy network component based on the output of the inspection task. 2. Indentify the issue of network configuration in the IaaS layer based on the logs. 3. If the problem persists after you perform the preceding steps, contact Alibaba Cloud technical support.

2.1.12.2.7. check-node-network

This topic describes how to troubleshoot exceptions that occur when the check-node-network inspection case is executed.

Inspection item NODE_NETWORK_CHECK

The following table describes the information related to the NODE_NETWORK_CHECK inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	Check whether the network of the node is normal.
Possible cause	Network problems exist in the IaaS layer.
Impact scope	• Kubernetes nodes cannot communicate with each other. • The Kubernetes cluster runs abnormally.

Item	Description
Solution	1. Use the ping command to test the connectivity of the faulty node. 2. Check whether the basic network configuration is normal. 3. If the problem persists after you perform the preceding steps, contact Alibaba Cloud technical support.

2.1.12.2.8. check-pod-network

This topic describes how to troubleshoot exceptions that occur when the check-pod-network inspection case is executed.

Inspection item CHECK_POD_TO_NODE

The following table describes the information related to the CHECK_POD_TO_NODE inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether pods and nodes are interconnected.
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods.
Impact scope	• Container Service for Kubernetes (ACK) pods cannot communicate with each other. Pods cannot communicate with nodes. • The ACK clusters run abnormally.

Item	Description
Solution	1. Check whether the pods for which the networking plug-in is installed are normal. 2. Check the basic networking configurations. 3. If the problem persists after you perform the preceding steps, contact Alibaba Cloud technical support.

Inspection item CHECK_POD_TO_POD

The following table describes the information related to the CHECK_POD_TO_POD inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the network between pods can be connected.
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods.
Impact scope	The network between the ACK pods cannot be connected.
Solution	1. Check whether the pods for which the networking plug-in is installed are normal. 2. Check the basic networking configurations.

Inspection item CHECK_POD_TO_CLUSTER_IP

The following table describes the information related to the CHECK_POD_TO_CLUSTER_IP inspection item.

Item	Description
Product	K8s

Item	Description
Inspection level	Critical
Inspection description	This inspection item is used to check whether the network between the pod and the cluster of the service can be connected.
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods. • DNS has exceptions.
Impact scope	The network between pods cannot be connected. DNS is unavailable.
Solution	1. Check whether the pods for which the networking plug-in is installed are normal. 2. Check the basic networking configurations. 3. Check whether the kube-proxy node is normal.

Inspection item CHECK_POD_TO_NODEPORT

The following table describes the information related to the CHECK_POD_TO_NODEPORT inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the network between the pod and the NodePort-type service can be connected.
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods. • DNS has exceptions.

Item	Description
Impact scope	The network between pods cannot be connected. DNS is unavailable.
Solution	1. Check whether the pods for which the networking plug-in is installed are normal. 2. Check the basic networking configurations. 3. Check whether the kube-proxy node is normal. 4. Check whether the node has a firewall.

Inspection item CHECK_POD_TO_DNS_ENDPOINT

The following table describes the information related to the CHECK_POD_TO_DNS_ENDPOINT inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the pod can be addressed by using the backend IP address that is provided by DNS.
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods. • DNS has exceptions.
Impact scope	The network between pods cannot be connected. DNS is unavailable.
Solution	1. Check whether the pod that uses DNS is normal. 2. Check the basic networking configurations. 3. Check whether the kube-proxy node is normal.

Inspection item CHECK_POD_TO_DNS_CLUSTERIP

The following table describes the information related to the CHECK_POD_TO_DNS_CLUSTERIP inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the pod can be accessed by using the IP address that DNS provides.
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods. • DNS has exceptions.
Impact scope	The network between pods cannot be connected. DNS is unavailable.
Solution	1. Check whether the pod that uses DNS is normal. 2. Check whether the network is connected.

Inspection item CHECK_POD_TO_DNS_NAME

The following table describes the information related to the CHECK_POD_TO_DNS_NAME inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the pod can be accessed by using the domain name that DNS provides.

Item	Description
Possible cause	• Network problems exist in the IaaS layer. • Network problems exist in pods. • DNS has exceptions.
Impact scope	The network between pods cannot be connected. DNS is unavailable.
Solution	1. Check whether the pod that uses DNS is normal. 2. Check whether the network is connected.

2.1.12.2.9. check-k8s-namespace

This topic describes how to troubleshoot exceptions that occur when the check-pod-network inspection case is executed.

CInspection item CHECK_K8S_NAMESPACE

The following table describes the information related to the CHECK_K8S_NAMESPACE inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	Check whether the namespace is normal.
Possible cause	An namespace in the Terminating state exists.
Impact scope	The namespace is unavailable.

Item	Description
Solution	1. Log on to the following node and run the following command to check webhook services for any exceptions. `kubectl api-resources --verbs=list --namespaced -o name | xargs -n 1 kubectl get --ignore-not-found -n default` 2. Troubleshoot the webhook services if necessary.

2.1.12.2.10. check-k8s-node

This topic describes how to troubleshoot exceptions that occur when the check-k8s-node inspection case is executed.

Inspection item CHECK_K8S_NODE

The following table describes the information related to the CHECK_K8S_NODE inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether the node is in the normal state.
Possible cause	• The Kubelet node is abnormal. • The Docker node is abnormal. • The disk usage is high. • The CPU utilization is high. • An accidental deletion occurs.
Impact scope	• The node cannot be used normally. • The existing applications that are deployed on the node are evicted to other nodes.

Item	Description
Solution	1. Check whether the CPU, memory, and disk resources of the node are exhausted. 2. Restart the kubelet node. **Notice** This may introduce security risks to the instance. We do not recommend that you perform this operation.

2.1.12.2.11. check-k8s-pod

This topic describes how to troubleshoot exceptions that occur when the check-k8s-pod inspection case is executed.

Inspection item CHECK_K8S_POD

The following table describes the information related to the CHECK_K8S_POD inspection item.

Item	Description
Product	K8s
Inspection level	Critical
Inspection description	This inspection item is used to check whether a pod is abnormal.
Possible cause	• Something is wrong with the service itself. • The node is abnormal.
Impact scope	The pod runs abnormally.

Item	Description
Solution	1. Log on to the abnormal node. 2. Run the `kubectl describe pod` command to query pod details. 3. Run the `kubectl logs <$Pod_Name>` command to query the logs of a pod. 4. Identify the cause of an exception based on the logs.

2.1.12.2.12. check-arkwebhook-svc

This topic describes how to troubleshoot exceptions of the check-arkwebhook-svc inspection case.

Inspection item ARKWEBHOOK_SERVICE

The following table describes the information related to the ARKWEBHOOK_SERVICE item.

Item	Description
Product	Ark
Inspection level	Critical
Inspection description	Check whether the ark-webhook service is normal.
Possible cause	• The pod of the ark-webhook is abnormal. • The network is abnormal.
Impact scope	All pods fail to be created.

Item	Description

Solution	1. Log on to the corresponding master node by using SSH.
	2. Run the following command to query the state of the specified pod:
	`kubectl get pod -n ark-webhook-system`
	3. Run the following command to query the logs of the specified pod:
	`kubectl logs -n ark-webhook-system <$Pod_Name>`
	4. Identify the cause of an exception based on the logs.

2.1.12.2.13. check-bridge-console

This topic describes how to troubleshoot exceptions that occur when the check-bridge-console inspection case is executed.

Inspection item BRIDGE_CONSOLE_API_USABILITY

The following table describes the information related to the BRIDGE_CONSOLE_API_USABILITY inspection item.

Item	Description
Product	Ark
Inspection level	Critical
Inspection description	This inspection item is used to check whether the PaaS Operations Console runs normally.
Possible cause	• Abnormal pods exist in the PaaS Operations Console. • The network is abnormal.
Impact scope	The PaaS Operations Console is unavailable.

Item	Description

Solution	1. Log on to the corresponding master node.
	2. Run the following command to check whether the application in the PaaS Operations Console is in the **Ready** state.
	`kubectl get appinstance -nark-system|grep bridge-`
	3. Log on to the PaaS Operations Console again to check whether the logon is normal.
	4. Run the following command to query the logs of a pod in the PaaS Operations Console:
	`kubectl logs <$Pod_Name>`
	5. Identify the cause of an exception based on the logs.

2.1.12.2.14. check-common-ingress

This topic describes how to troubleshoot exceptions that occur when the check-common-ingress inspection case is executed.

Inspection item COMMON_INGRESS_ACCESS

The following table describes the information related to the COMMON_INGRESS_ACCESS inspection item.

Item	Description
Product	Cluster-Init
Inspection level	Critical
Inspection description	This inspection item is used to check whether the common-ingress component can be accessed.
Possible cause	• The Voyager pod is abnormal. • The network is abnormal.
Impact scope	The access request is abnormal.

Item	Description

Solution	1. Log on to the corresponding master node by using SSH.
	2. Run the following command to query the state of the pod:
	`kubectl get |grep "voyager-ingress-common"`
	3. Run the following command to query the logs of the pod:
	`kubectl logs <$Pod_Name>`
	4. Identify the cause of an exception based on the logs.

2.1.12.2.15. check-seed-status

This topic describes how to troubleshoot exceptions that occur when the check-seed-status inspection case is executed.

Inspection item SEED

The following table describes the information related to the SEED inspection item.

Item	Description
Product	Ark
Inspection level	Critical
Inspection description	This inspection item is used to check the health of the seed base module.
Possible cause	• Multiple application replicas that are managed by the seed base module are abnormal. • Minio is abnormal. • Kubernetes is abnormal.
Impact scope	The deployment and update of the base and cloud services are affected.

Item	Description

Solution	1. Log on to the corresponding master node.
	2. Run the following command to query the status of cm:
	`kubectl get cm ark.cmdb.seed.status -n ark-system -o yaml`
	3. Check the application whose status is not **Ready** based on the message of cm.
	4. Run the following command to query the status of the seed base module:
	`docker ps|grep seed`
	5. Run the following command to query the docker logs of the seed base module:
	`kubectl logs seed`
	6. Identify the cause of an exception based on the logs.

2.1.12.2.16. check-nginx-ingress

This topic describes how to troubleshoot exceptions that occur when the check-nginx-ingress inspection case is executed.

Inspection item NGINX_INGRESS_CHECK

The following table describes the information related to the NGINX_INGRESS_CHECK inspection item.

Item	Description
Product	cluster-init
Inspection level	Critical
Inspection description	This inspection item is used to check whether the NGINX Ingress Controller is available
Possible cause	The nginx-ingress pod is abnormal.
Impact scope	The access request is abnormal.

Item	Description

Solution	1. Log on to the corresponding master node by using SSH.
	2. Run the following command to query the status of the nginx-ingress pod.
	`kubectl get -n acs-system |grep "nginx-ingress-controller"`
	3. Run the following command to query the logs of the nginx-ingress pod.
	`kubectl logs -nacs-system <$Pod_Name>`
	4. Identify the cause of an exception based on the logs.

2.1.12.2.17. check-operator-reconcile-queue-length

This topic describes how to troubleshoot exceptions that occur when the check-operator-reconcile-queue-length inspection case is executed.

Inspection item OPERATOR_RECONCILE_QUEUE_LENGTH_CHECK

The following table describes the information related to the OPERATOR_RECONCILE_QUEUE_LENGTH_CHECK inspection item.

Item	Description
Product	Ark
Inspection level	Warning
Inspection description	This inspection item is used to check the reconcile queue length of the operator.
Possible cause	The number of upgrade tasks exceeds the limit.
Impact scope	The upgrade duration is affected.
Solution	Wait until the Ark operator completes the task.

2.1.12.2.18. check-synchronizer

This topic describes how to troubleshoot exceptions that occur when the check-synchronizer inspection case is executed.

Inspection item SYNCHRONIZER_CHECK

The following table describes the information related to the SYNCHRONIZER_CHECK inspection item.

Item	Description
Product	Tianji-provider
Inspection level	Critical
Inspection description	This inspection item is used to check whether the internal interface of the synchronizer pod is accessible.
Possible cause	• The Apsara Infrastructure Management Framework server fails to start when the synchronizer pod is deployed. • The Apsara Infrastructure Management Framework service is abnormal.
Impact scope	• Only deployment services are affected. Runtime services are not affected. • The synchronizer pod may do not respond when base services start.

Solution	1. Log on to the corresponding master node by using SSH.
	2. Run the following command to query the status of the synchronizer pod:
	`kubectl get -n acs-system |grep "<$Pod_Name>"`
	3. Run the following command to query the logs of the synchronizer pod:
	`kubectl logs -nacs-system <$Pod_Name>`
	4. Identify the cause of an exception based on the logs of the pod.

2.1.12.2.19. check-registory

This topic describes how to troubleshoot exceptions that occur when the check-registory inspection case is executed.

Inspection item REGISTORY_SERVICE

The following table describes the information related to the REGISTORY_SERVICE inspection item.

Item	Description
Product	Basic
Inspection level	Critical
Inspection description	This inspection item is used to check whether the image repository is accessible.
Possible cause	The image repository is abnormal.
Impact scope	The start of a pod is affected.
Solution	1. Log on to the corresponding master node by using SSH. 2. Run the following command to pull an image from the image repository. `docker pull <$Image_Path>` 3. Check whether the network can be connected. 4. If the network cannot be connected and the image cannot be pulled, we recommend that you check whether the image repository itself is operating normally.

2.1.12.2.20. check-app-status

This topic describes how to troubleshoot exceptions that occur when the check-app-status inspection case is executed.

Inspection item APPINSTANCE_STATUS

The following table describes the information related to the APPINSTANCE_STATUS inspection item.

Item	Description
Product	Basic
Inspection level	Critical
Inspection description	This inspection item is used to check whether appinstance is in the normal state.
Possible cause	The backend service is faulty.
Impact scope	The corresponding application feature is abnormal.

Solution	1. Log on to the corresponding master node by using SSH to go to the arkshoot environment.
	2. Run the following command to query the ID of the workflow:
	`argo list -nark-system|grep <$ Application name >`
	3. Run the following command to identify the cause of the deployment failure:
	`argo get -nark-system <$ Workflow ID >`
	4. Handle an exception based on the identified cause.
	5. If you suspect that the deployment fails due to network issues, run the following command to modify the appInstance.spec.OrderId parameter:
	`kubectl edit appinstance -nark-system <$ Application name >`
	The workflow is automatically retried for deployment.

Inspection item APPSET_STATUS

The following table describes the information related to the APPSET_STATUS inspection item.

Item	Description
Product	Basic
Inspection level	Critical
Inspection description	This inspection item is used to check whether appset is in the normal state.
Possible cause	The backend service is faulty.
Impact scope	The corresponding application feature is abnormal.

Item	Description

Solution	1. Find the failed appinstance in the appset, and check the failure logs of the workflow that corresponds to the appinstance.
	◦ If the deployment fails due to network issues, modify the appInstance.spec.OrderId parameter toredeploy the appinstance.
	◦ If code errors occur, contact Alibaba Cloud technical support.
	2. Run the following command to query the appInstance that is in the NotReady state.
	`kubectl get appinstance -nark-system -lapps.mwops.alibaba-inc.com/appset-name=ark`
	3. Log on to the corresponding master node by using SSH to go to the arkshoot environment.
	4. Run the following command to query the ID of the workflow:
	`argo list -nark-system|grep <$ Application name >`
	5. Run the following command to identify the cause of the deployment failure:
	`argo get -nark-system <$ Workflow ID >`
	6. Handle an exception based on the identified cause.
	7. If you suspect that the deployment fails due to network issues, run the following command to modify the appInstance.spec.OrderId parameter:
	`kubectl edit appinstance -nark-system <$ Application name >`
	The workflow is automatically retried for deployment.

2.1.12.2.21. check-prometheus

This topic describes how to troubleshoot exceptions that occur when the check-prometheus inspection case is executed.

Inspection item PrometheusWalCheck

The following table describes the information related to the PrometheusWalCheck inspection item.

Item	Description
Product	tianjimon
Inspection level	Critical
Inspection description	This inspection item is used to check whether the WAL files in Prometheus surges.
Possible cause	The frequency of indicator reporting is high, which results in increasing pressure.
Impact scope	Promethues is unstable for monitoring.

Item	Description

Solution	1. Log on to the corresponding master node by using SSH and go to the command-line terminal.
	2. Save the following code as the run.sh script.
	`ret=1 while [$ret != 0]; do kubectl exec -n monitoring prometheus-tianjimon-prometheus-prome-prometheus-0 -c prometheus -- sh -c 'ls wal >/dev/null' ret=$? sleep 1 done kubectl exec -n monitoring prometheus-tianjimon-prometheus-prome-prometheus-0 -c prometheus -- sh -c 'find wal -type f -mmin +180 -exec rm -rf {} \;' kubectl exec -n monitoring prometheus-tianjimon-prometheus-prome-prometheus-0 -c prometheus -- sh -c 'find *.tmp -type d -mtime +14 -exec rm -rf {} \;'`
	kubectl exec -n monitoring prometheus-tianjimon-prometheus-prome-prometheus-0 -c prometheus -- sh -c 'ls wal >/dev/null' ret=$? sleep 1 done kubectl exec -n monitoring prometheus-tianjimon-prometheus-prome-prometheus-0 -c prometheus -- sh -c 'find wal -type f -mmin +180 -exec rm -rf {} \;' kubectl exec -n monitoring prometheus-tianjimon-prometheus-prome-prometheus-0 -c prometheus -- sh -c 'find *.tmp -type d -mtime +14 -exec rm -rf {} \;'
	3. 执行run.sh脚本。
	`sh run.sh`

3.Operations of basic cloud products

3.1. ApsaraDB RDS

3.1.1. Architecture

3.1.1.1. System architecture

3.1.1.1.1. 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 for MySQL supports both physical and logical backups.

You can flexibly configure the backup start time within off-peak hours. All backup files are retained for seven days.

Temporary backup

You can create temporary backup files when necessary. Temporary backup files are retained for seven days.

Log management

ApsaraDB RDS for MySQL generates binlogs that you can download for local incremental backup.

Instance cloning

A cloned instance is a new instance whose data and settings are the same as those of the primary instance. This feature allows you to restore data of the primary instance or create multiple instances that are the same as the primary instance.

3.1.1.1.2. Monitoring system

ApsaraDB RDS provides multi-dimensional monitoring services across the physical, network, and application layers to ensure business availability.

Performance monitoring

ApsaraDB RDS provides nearly 20 metrics for system performance monitoring, such as disk capacity, input/output operations per second (IOPS), connections, CPU utilization, network traffic, transactions per second (TPS), queries per second (QPS), and cache hit rate. You can obtain the running status information of instances within the past year.

SQL auditing

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 auditing to locate problems and check instance security.

Threshold alerts

ApsaraDB RDS provides alert SMS notifications in the event of exceptions in instance status or performance.

These exceptions include instance locking, disk capacity, IOPS, connections, and CPU. You can configure alert thresholds and up to 50 alert contacts (of which five are effective at a time). When an instance exceeds the threshold, an SMS notification is sent to the alert contacts.

Web operation logs

The system records all modification operations in the ApsaraDB RDS console for administrators to check. These logs are retained for up to 30 days.

3.1.1.1.3. Control system

If a host or instance stops responding, it switches services over within 30 seconds after the high-availability (HA) component detects an exception. This ensures that applications run normally.

3.1.1.1.4. 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.

3.1.2. 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 URL of the Apsara Uni-manager Operations Console is in the following format: *region-id*.*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 (*region-id*.*ops*.asconsole.*intranet-domain-id*.com). Then, press the Enter key.

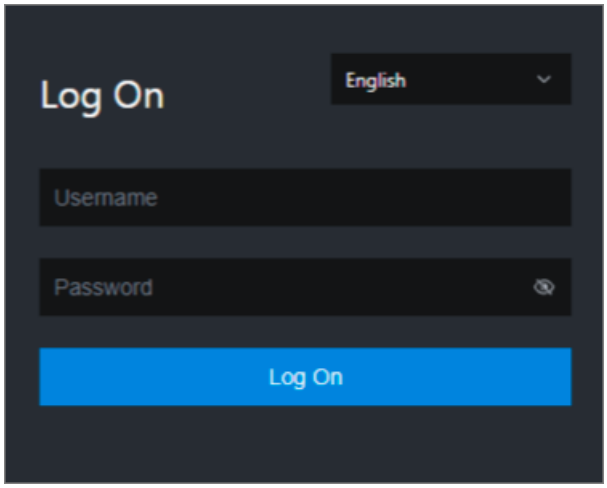

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

3.1.3. Manage instances

This topic describes how to manage ApsaraDB RDS instances. You can view instance details, logs, 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, you can perform the following operations:
   - View instances

     View the instances that belong to your account on the **Instance Management** tab, as shown in Instances.

     Instances

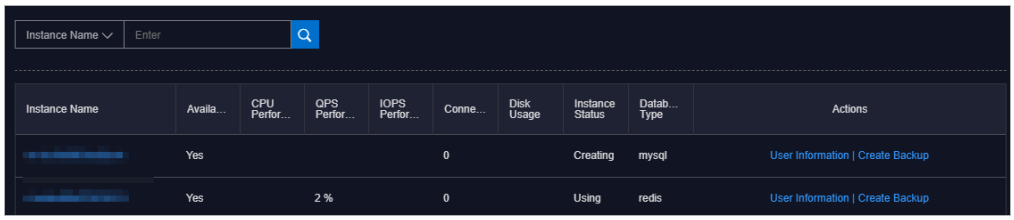

- View instance details

  Click the ID of an instance to view its details, as shown in Instance details. You can switch your service 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 result in data loss if data is not synchronized between the primary and secondary instances.

  Instance details

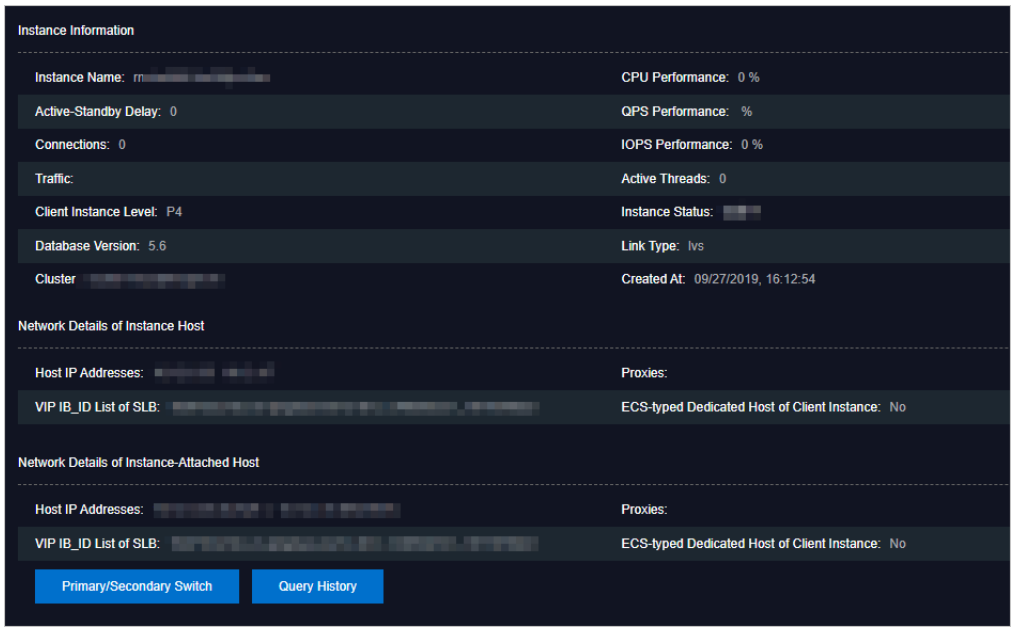

- **View user information**

  Click **User Information** in the **Actions** column corresponding to an instance, as shown in User information.

  User information

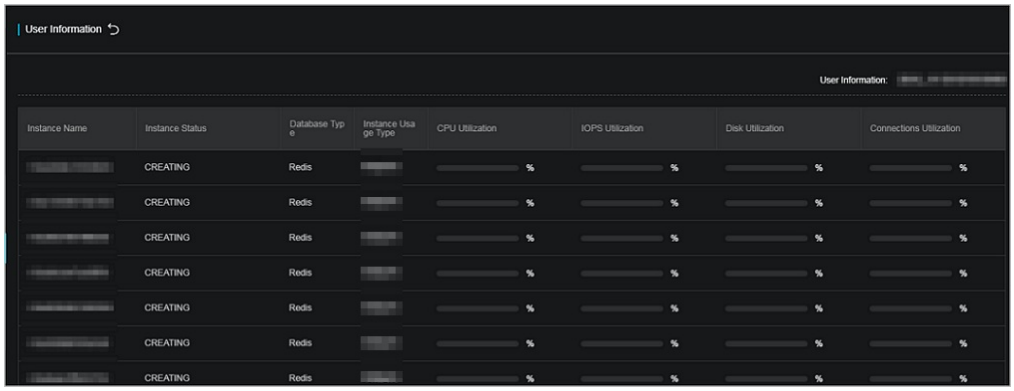

- **Create backups**

  For ApsaraDB RDS for MySQL instances, click **Create Backup** in the **Actions** column to view the backup information, as shown in Backup information. You can also click **Create Single Database Backup** on the Backup Information page to back up a single database.

  Backup information

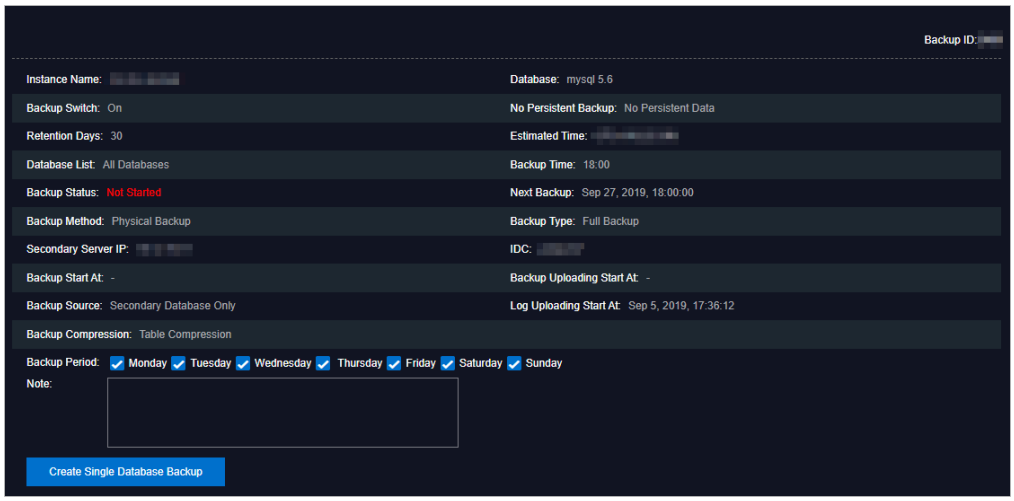

3.1.4. Manage hosts

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.

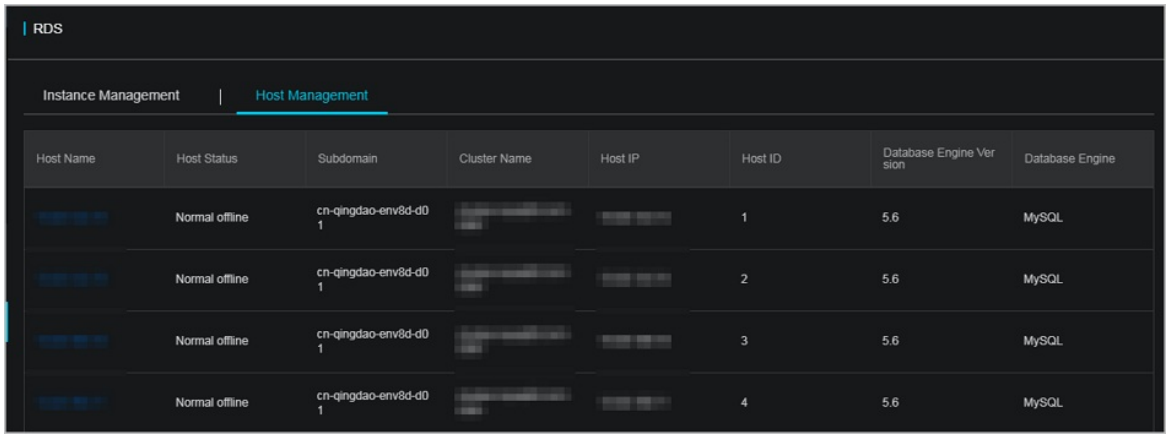

4. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.

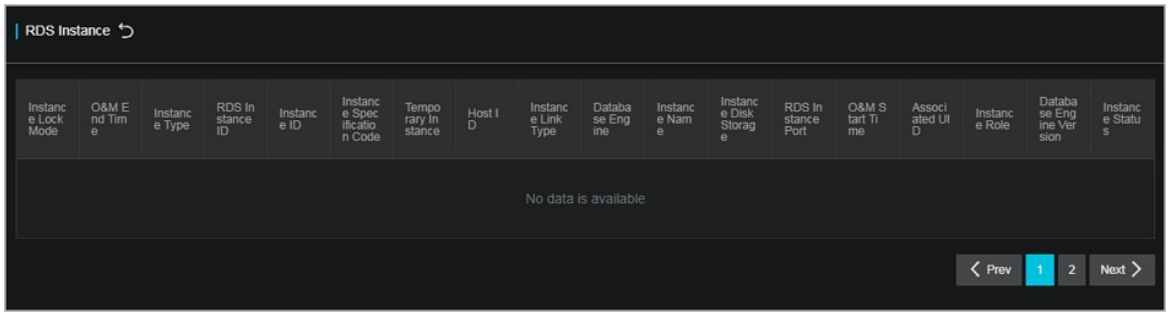

3.1.5. Security maintenance

3.1.5.1. 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.

3.1.5.2. 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.Appendix

4.1. Operation Access Manager (OAM)

4.1.1. 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 the 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

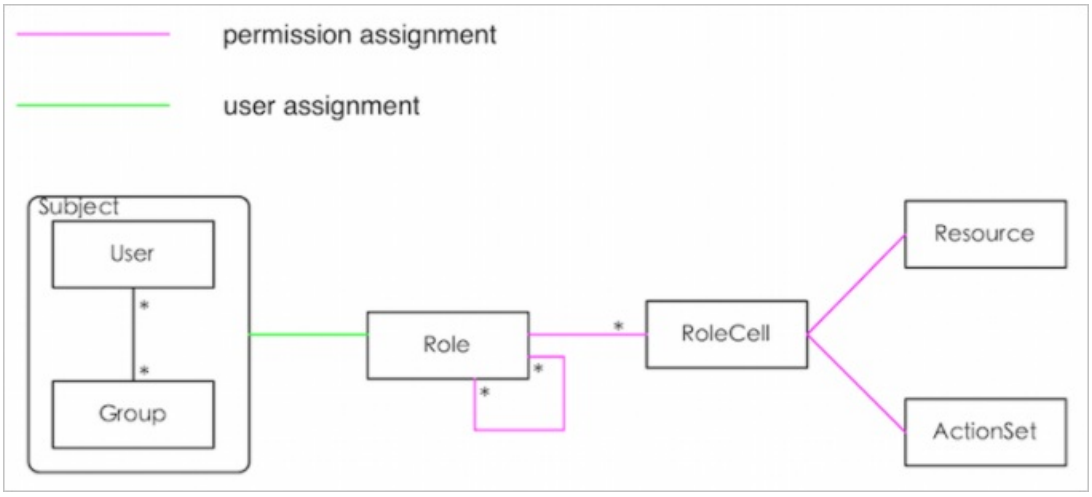

4.1.2. Usage instructions

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.

4.1.3. Quick Start

By completing the steps in this guide, you will learn how to create and assign roles for O&M.

4.1.3.1. 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 URL of the Apsara Uni-manager Operations Console is in the following format: *region-id*.*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 (*region-id*.*ops*.asconsole.*intranet-domain-id*.com). Then, press the Enter key.

   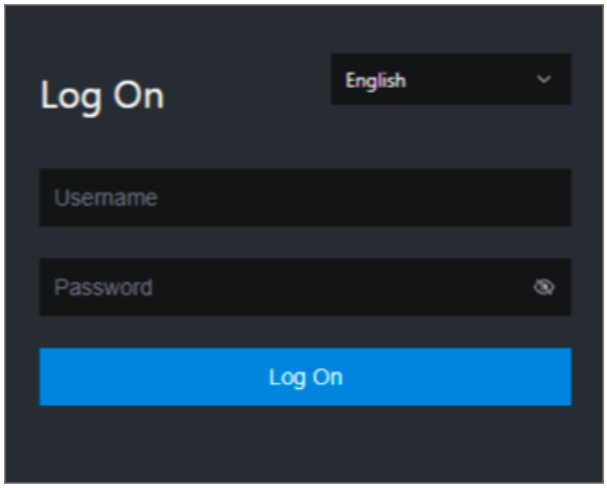

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

4.1.3.2. 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**.

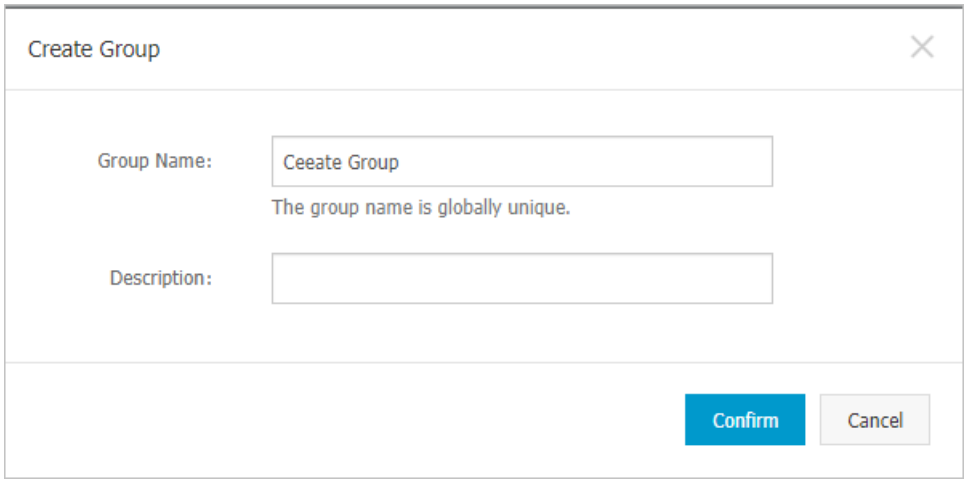

4. Click **Confirm**.

   After the group is created, it is displayed on the **Owned Groups** page.

4.1.3.3. 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.

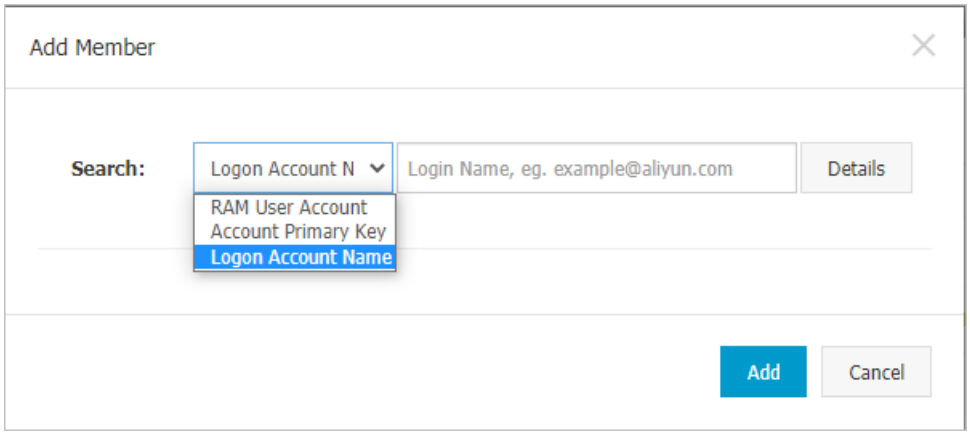

5. Select a search mode, enter the corresponding information, and then click **Details**. Details of the specified user are displayed.

   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.

4.1.3.4. 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**.

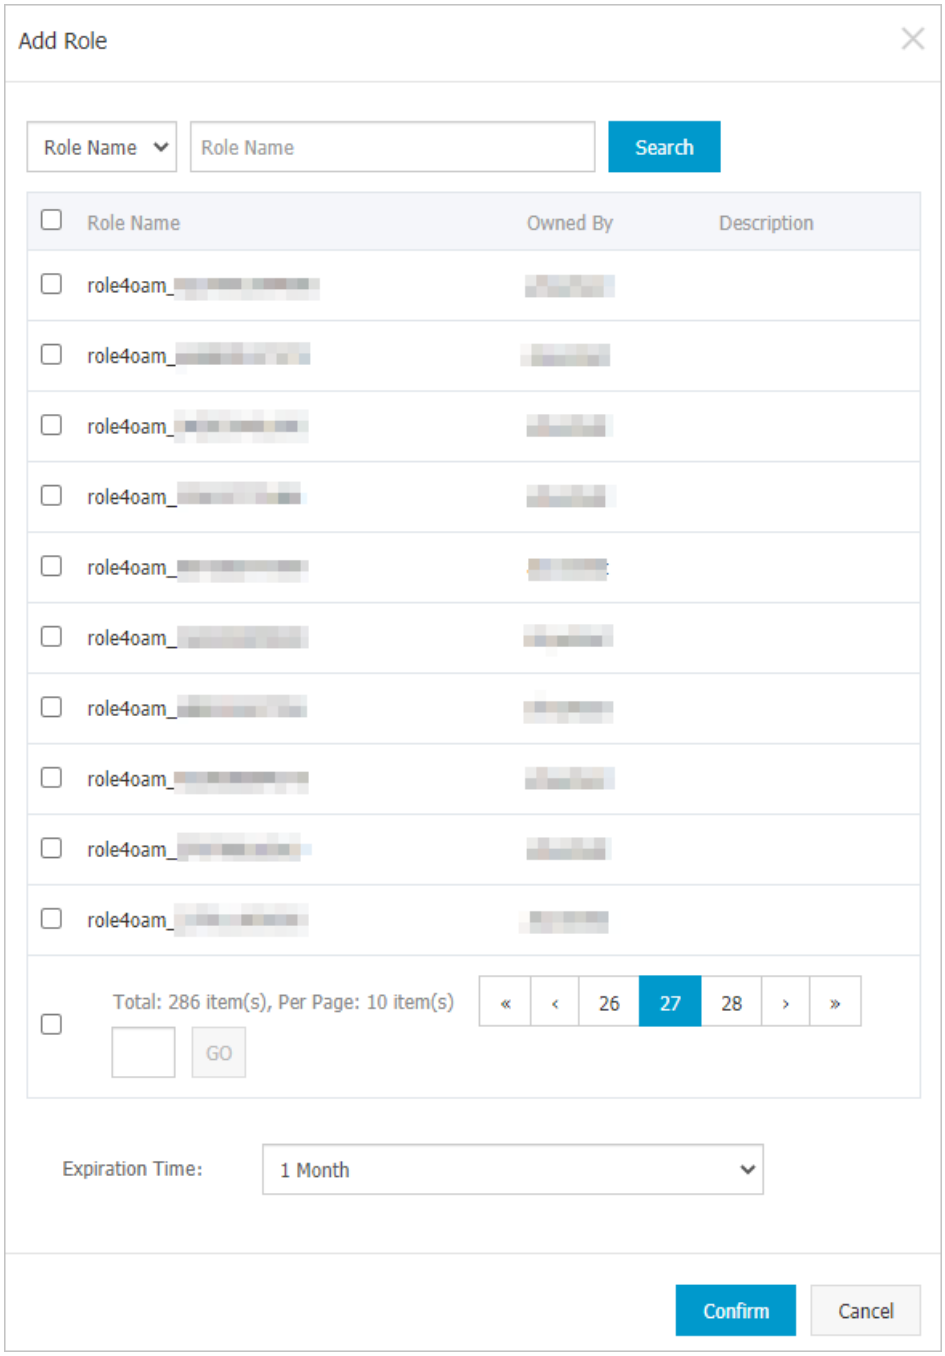

5. Search for roles by **Role Name**. Select one or more roles and set Expiration Time.
6. Click **Confirm**.

   To remove a role from a group, find the role in **Role List** and click **Remove** in the **Actions** column.

4.1.3.5. 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.

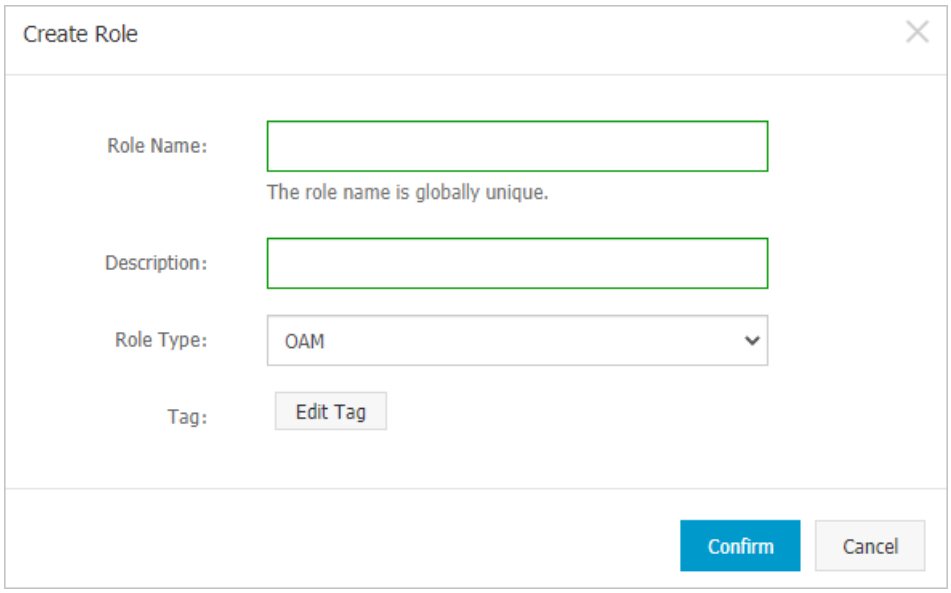

4. In the Create Role dialog box, set **Role Name**, **Description**, and **Role Type**.
5. (Optional)Set tags for the role. Tags can be used to search for roles.
   i. Click **Edit Tag**.

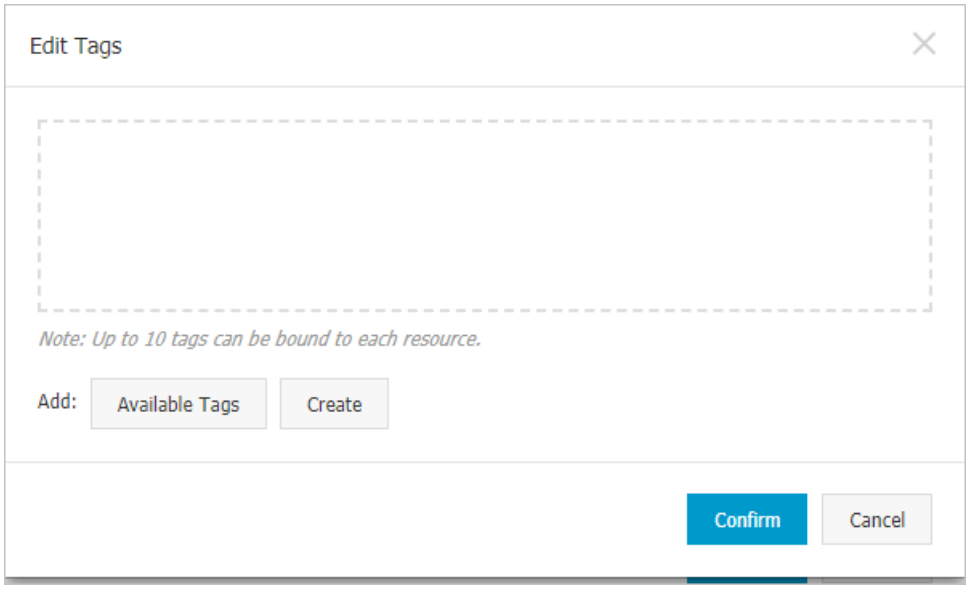

   ii. In the **Edit Tags** dialog box, click **Create**.

iii. Set **Key** and **Value** for the tag and click **Confirm**.

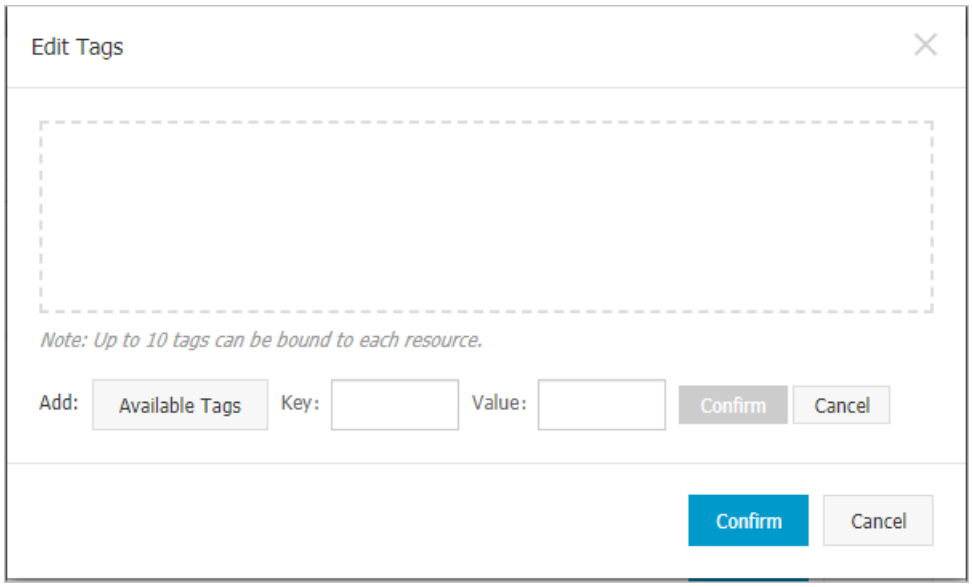

iv. Repeat the preceding step to create more tags.

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.

4.1.3.6. 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.

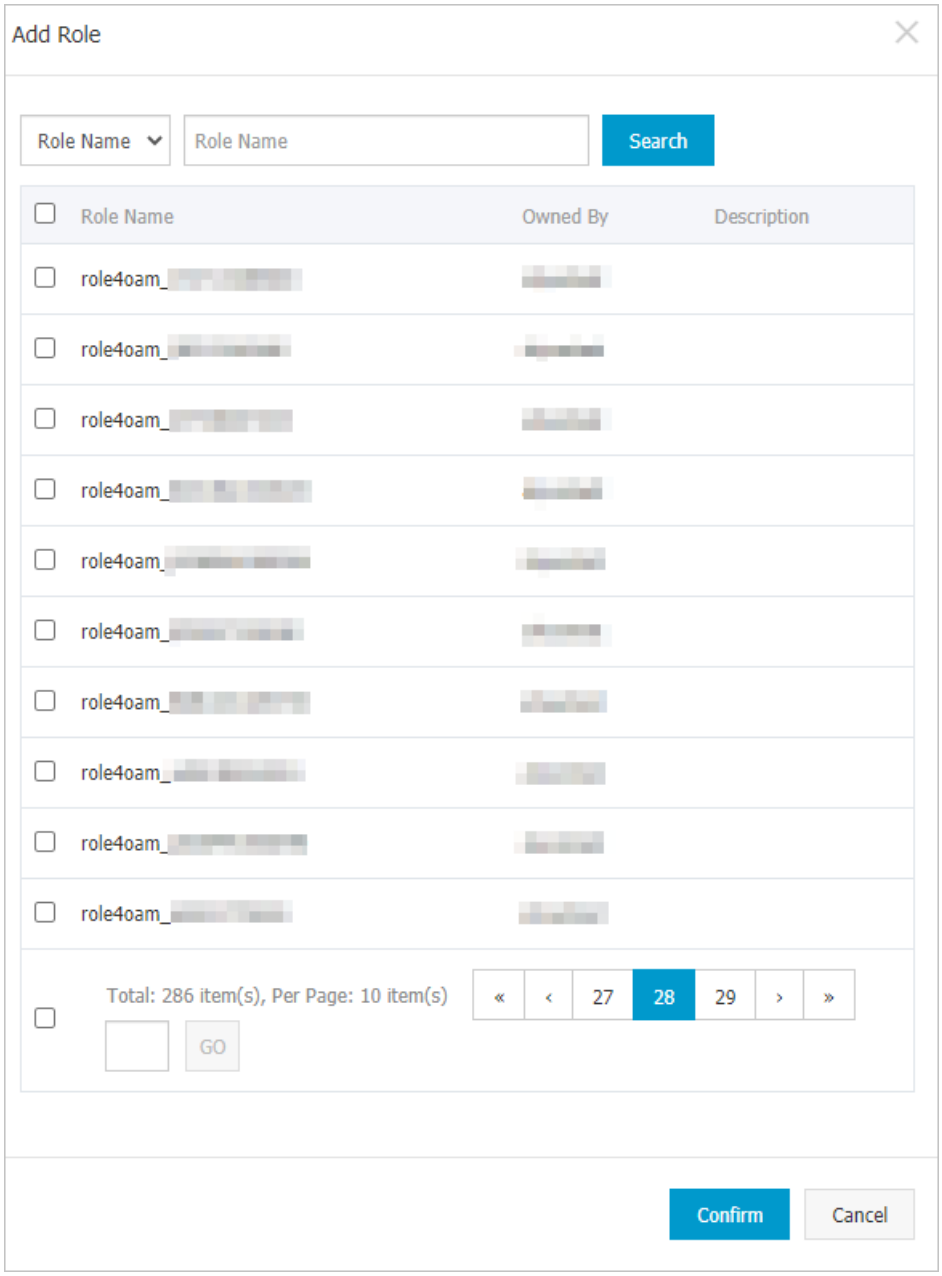

6. Click **Confirm**.

4.1.3.7. Add a resource to a role

You must add resources to a created role.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add a resource and click **Manage** in the **Actions** column.
4. On the Role Information page, click the **Resource List** tab.
5. Click **Add Resource**.

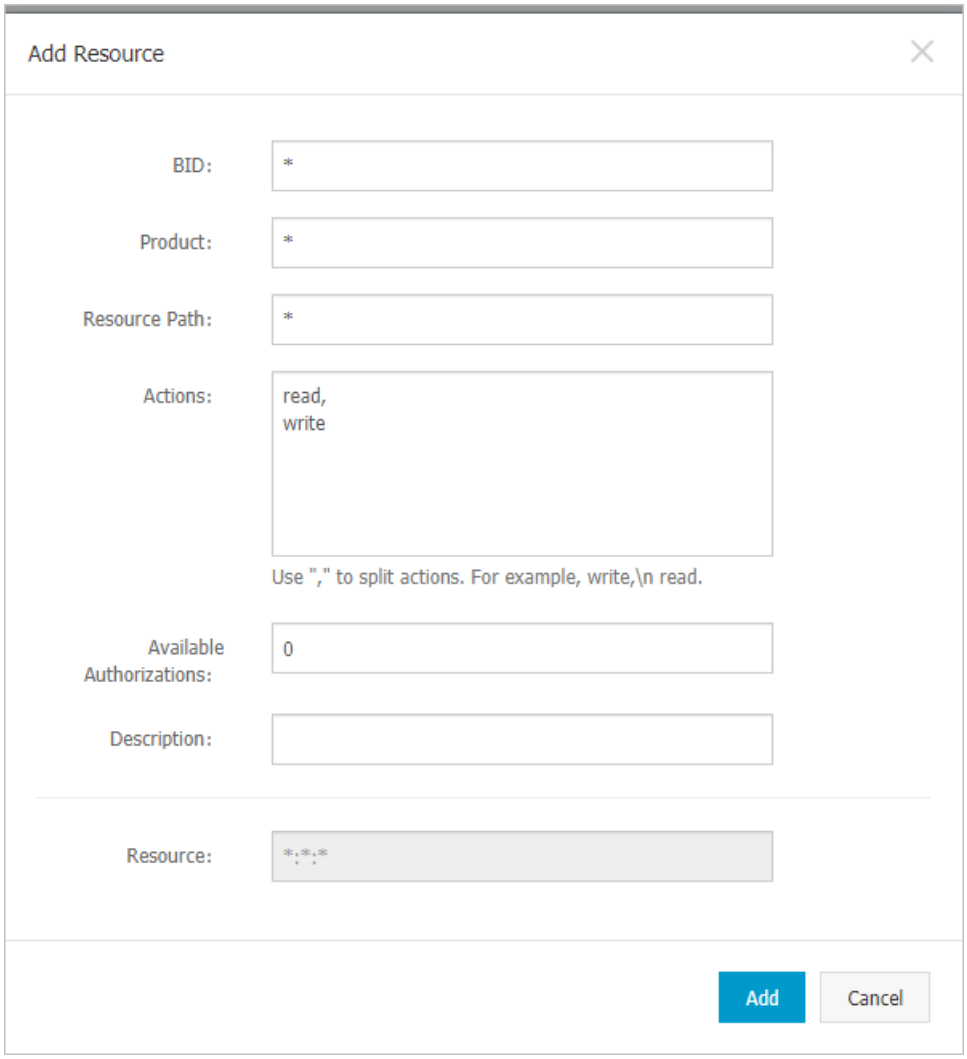

6. In the **Add Resource** dialog box, configure the parameters. For more information, see Parameters.

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**.

4.1.3.8. 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.

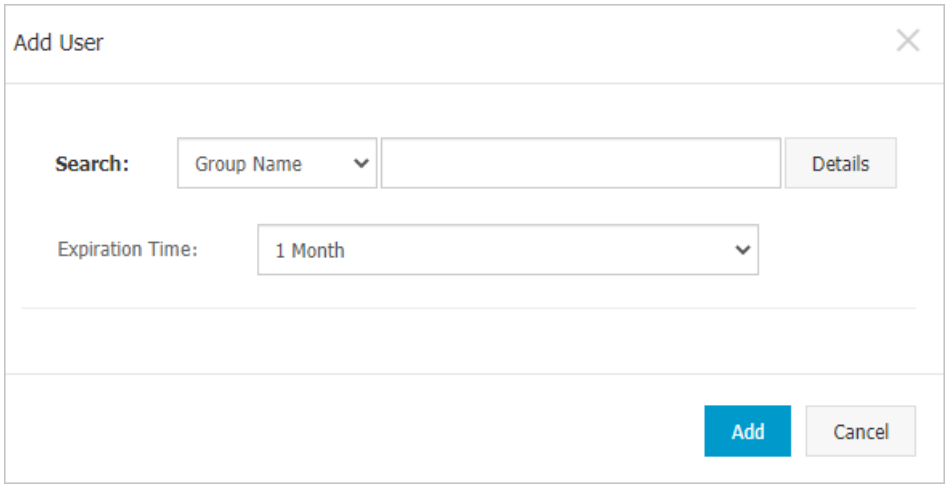

6. Select a search mode and enter corresponding information to search for the user to which you want to assign the role.

   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.

4.1.4. Manage groups

Group management allows you to view, modify, and delete groups.

4.1.4.1. 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**.

4.1.4.2. 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.

4.1.4.3. 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**.

4.1.4.4. 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.

4.1.5. Manage roles

Role management allows you to view, modify, transfer, and delete roles.

4.1.5.1. 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.

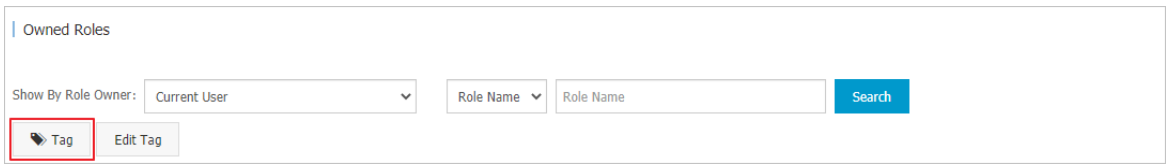

4.1.5.2. Modify role information

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**.

4.1.5.3. 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.

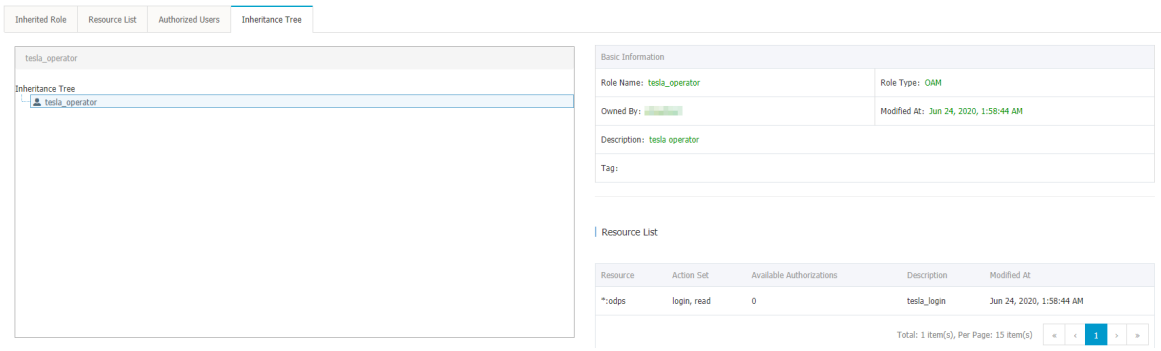

4.1.5.4. Transfer a role

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.

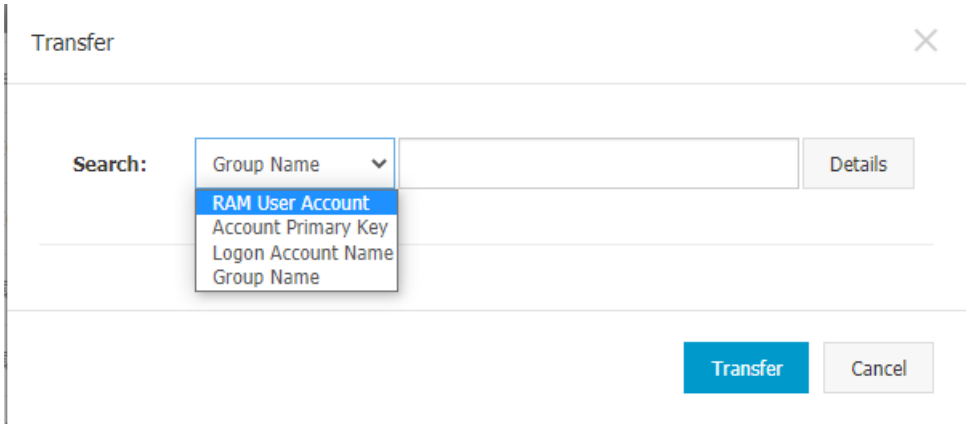

6. Click **Transfer**.

4.1.5.5. 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**.

4.1.5.6. 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, you can view the name, owner, description, modification time, and expiration time of each role assigned to you.

   You can also click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

4.1.5.7. 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.

4.1.6. 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.

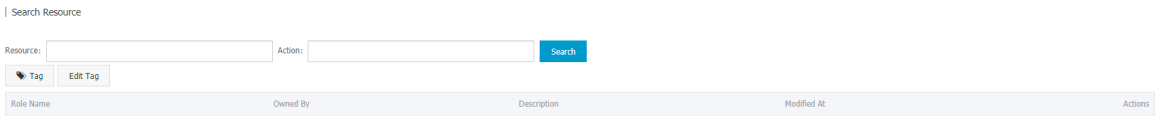

4. Click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

4.1.7. 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.

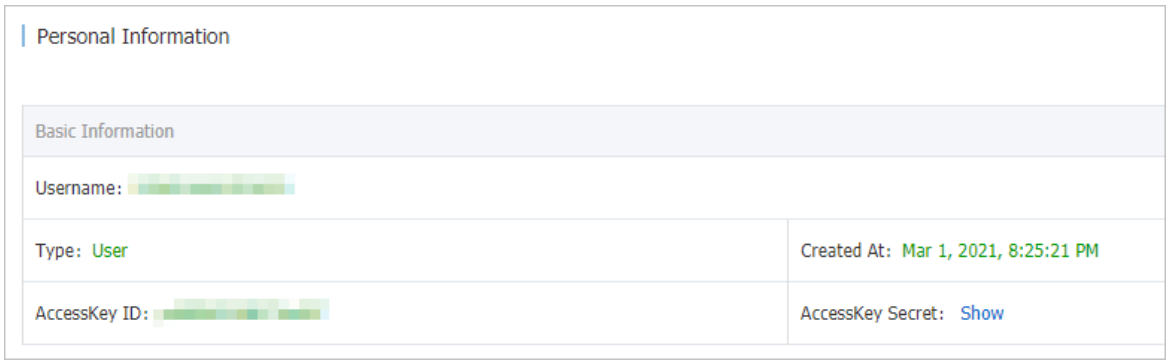

> **Note** You can click **Show** or **Hide** to show or hide the AccessKey secret.

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 (,).

4.1.8. Default roles and permissions

4.1.8.1. Default roles and their functions

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

4.1.8.1.1. Default role of OAM

This topic describes the default role of Operation Access Manager (OAM) and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Super administrator	An administrator with root permissions	*:*	*	10

4.1.8.1.2. Default roles of Apsara Infrastructure Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding 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	Action	Grant option
Tianji_Project read-only	Has the read-only permission on Apsara Infrastructure Management Framework projects, which allows you to view the configurations and statuses of all projects and clusters.	*:tianji:projects	["read"]	0
Tianji_Project administrator	Has all the permissions on Apsara Infrastructure Management Framework projects, which allows you to view and modify the configurations and statuses of all projects and clusters.	*:tianji:projects	["*"]	0

Role	Description	Resource	Action	Grant option
Tianji_Service read-only	Has the read-only permission 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 the 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 the permissions on Apsara Infrastructure Management Framework data centers, which allows you to view and modify the data center information.	*:tianji:idcs	["*"]	0
Tianji administrator	Has all the permissions on Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastructure Management Framework configurations.	*:tianji	["*"]	0

4.1.8.1.3. 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	Action	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

4.1.8.1.4. 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	Action	Grant option
System administrator of the Apsara Uni-manager Operations	Has the permissions to manage platform nodes, physical devices, and virtual resources, back up, restore, and migrate	*:aso:api-adapter:*	["read","write"]	0
		:aso:auth:	["read"]	0
		:aso:backup:	["read","write"]	0
		:aso:cmdb:	["read","write"]	0
		:aso:doc:	["read","write"]	0
		:aso:fullview:	["read","write"]	0
		:aso:init:	["read","write"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read","write"]	0

Console Role	Description	Resource	Action	Grant option
	and migrate product data, and query and back up system logs.	*: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 of the Apsara Uni-manager Operations Console	Has the 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 of the Apsara Uni-manager Operations Console	Has the permissions to audit, track, and analyze the activities of the system administrator and security officer.	*:aso:audit:*	["read","write"]	0
		:aso:auth:	["read"]	0
		:aso:serviceWhitelist:	["read"]	0
Product O&M officer of the Apsara Uni-manager	Has the permissions to perform O&M operations such as data import and export,	*:aso:api-adapter:*	["read"]	0
		:aso:backup:	["read"]	0
		:aso:cmdb:	["read"]	0
		:aso:doc:	["read"]	0
		:aso:fullview:	["read","write"]	0
		:aso:init:	["read"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read"]	0

Role	Description	Resource	Action	Grant option
manager Operations Console	modification, configuration, upgrade, and troubleshooting coordination.			
		: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 of the Apsara Uni-manager Operations Console	Has the 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
Duty observer of the Apsara Uni-manager Operations Console	Has the permissions to view and monitor the dashboard and monitor system alerts.	*:aso:doc:*	["read"]	0
		:aso:fullview:	["read"]	0

4.1.8.1.5. 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	Action	Grant option
PaaS_Operation_Manager	Has all the permissions on the PaaS console.	*:paas-ops:*	["*"]	0

4.1.8.2. Operation permissions on O&M platforms

This topic describes the operation permissions on O&M platforms.

4.1.8.2.1. 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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*: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.
*: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.

Resource	Operation	Description
*:tianji:idcs:[idc]	write	Sets data center properties.
*:tianji:idcs:[idc]	write	Creates a data center.

4.1.8.2.2. 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