ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Agility SE

Operations and Maintenance Guide

Product Version: 2006, Internal: V3.3.0

Document Version: 20201106

Alibaba Cloud

Legal disclaimer

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

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

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

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

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

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

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

Document conventions

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

Table of Contents

1.ASO overview

Apsara Stack Operations (ASO) is an operations management system developed for the Apsara Stack operations management personnel, such as field operations engineers, operations engineers on the user side, and operations management engineers, operations security personnel, and audit personnel of the cloud platform. ASO provides operations engineers with information about running conditions of the system in a timely manner and allows them to perform O&M operations.

ASO provides the following features:

- Monitoring and alerting

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

- Resource management

  The Resource Management module monitors and manages hardware devices in the data center. You can monitor and manage the overall status information, monitoring metrics, alert delivery status, and port traffic of physical servers, physical switches, and network security devices.

- Inventory management

  The Inventory Management module allows you to view the resource usage and inventory of various services and manage system resources effectively.

- Products

  The Products module allows you to access the O&M services of other products on the cloud platform and to configure ISV access configurations.

- NOC

  The Network Operation Center (NOC) module provides operations capabilities such as the visualization of end-to-end monitoring, automated implementation, automated fault location, and network traffic analysis to enhance the efficiency of network operations engineers, reduce the operations risk, and improves the quality of Apsara Stack services.

- Storage operations center

  The Storage Operation Center module contains Apsara Distributed File System and miniOSS.

- Task management

  The Task Management module allows you to perform O&M operations without using command lines.

- System management

  The System Management module provides features such as user management, two-factor authentication, role management, department management, logon policy management, application whitelist, server password management, operation logs, and authorization. As the module for centralized management of accounts, roles, and permissions, System Management supports the Single Sign-On (SSO). After you log on to the ASO console, you can perform O&M operations on all components of the cloud platform or be redirected to the O&M page without providing the username or password.

2.Preparations before operations

2.1. Prepare an operations account

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

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

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

> **Note** For a more fine-grained division of the operations role, the administrator can create a basic role based on Operation Administrator Manager (OAM) in the Appendix, grant permissions to the role, and then grant the role to the corresponding operations account. >

2.2. Log on to the ASO console

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

Prerequisites

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

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

Procedure

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

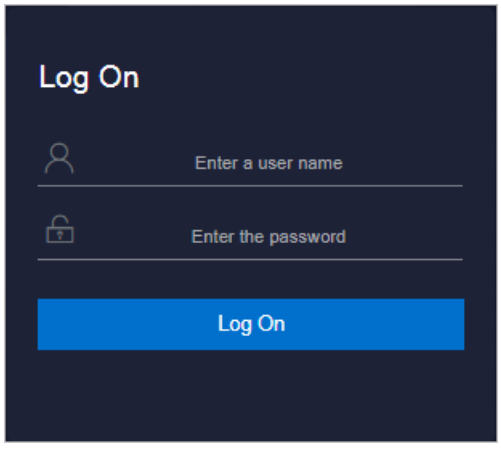

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

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

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

2.3. ASO console overview

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

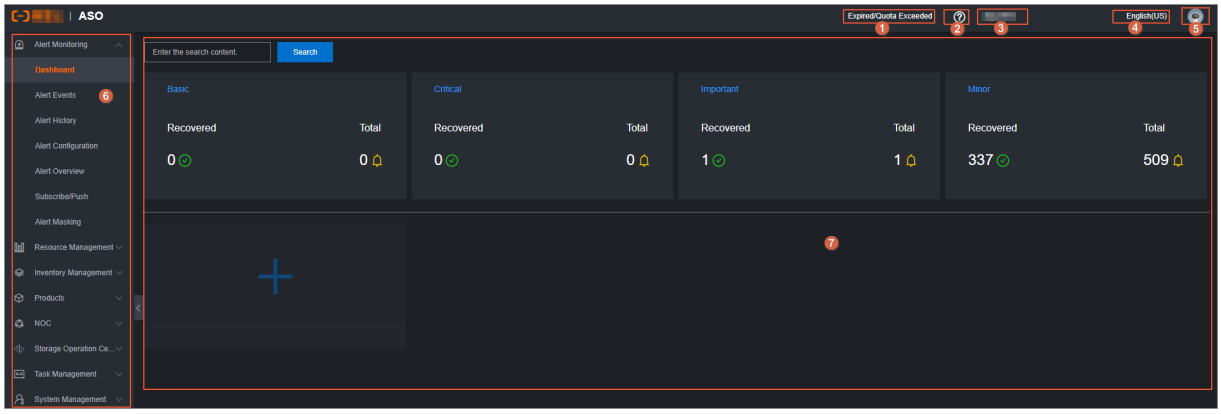

The following table describes the modules.

Section		Description
1	Authorization	Click this area to go to the **Authorization** page and then view the authorization conditions of services.
2	Help center	In the help center, you can view the alert knowledge base and upload other HTML documents that are related to O&M.
3	Current logon user	The current logon user name.
4	Language	The language of the current environment.

Section		Description
5	Information of the current logon user	Move your pointer over the profile picture in the upper-right corner of the page. In the drop-down list that appears, you can select Personal Information to view the personal information of the current user and modify the password, click Logon Settings to configure logon parameters, or click Log Out to log off from the ASO console.
6	Left-side navigation pane	Select an O&M operation.
7	Operation area	The information display and operation area.

3.System settings

3.1. Default operations roles

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

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

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

3.2. System Management

System Management centrally manages the departments, roles, and users involved in Apsara Stack Operations (ASO), making it easy to grant different resource access permissions to different users. As the core module for centralized permission management, the user center integrates the functions such as department management, role management, logon policy management, and user management.

3.2.1. Department management

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

Context

By default, after ASO is deployed, a root department is created. You can create departments under the root department. Departments are displayed in a hierarchy and you can create sub-departments under each level of departments.

A department that is created under the root department is a level-1 department and a department that is created under a level-1 department is a level-2 department. In ASO, sub-departments of a department refer to all levels of departments under the department. Departments reflect the tree structure of an enterprise or organization. Each user can belong to only one department.

Procedure

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

2. On the **Department Management** page, perform the following operations:
   - Add a department

     Click **Add Department** in the upper-left corner of the page. In the **Add Department** dialog box, specify **Department Name**, and then click **OK**. Then, you can view the created department under your selected catalog.
   - Modify a department

     Select the department to be modified in the catalog tree and click **Modify Department** in the upper part of the page. In the **Modify Department** dialog box, modify **Department Name**, and then click **OK**.
   - Delete a department

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

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

3.2.2. Role management

You can add custom roles in the Apsara Stack operations (ASO) console to more efficiently grant permissions to users.

Context

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

Procedure

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

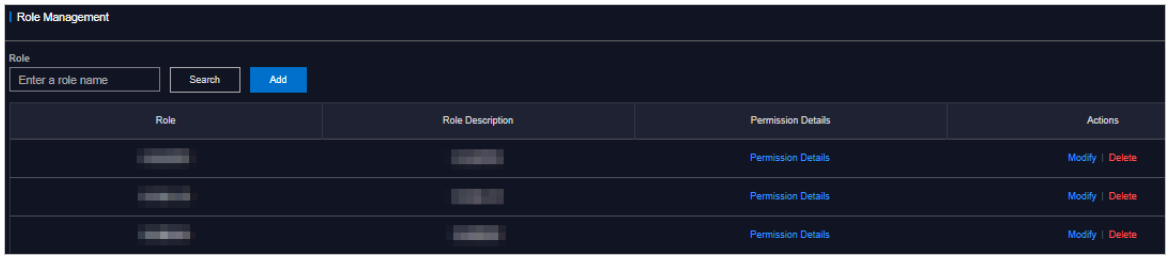

2. On the **Role Management** page, perform the following operations:
   - Query roles

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

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

- Add a role

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

  Click **Add** in the upper part of the page. In the **Add** dialog box, specify **Role Name**, **Role Description**, and **Basic Role**, and then click **OK**.

- Modify a role

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

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

- Delete a role

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

3.2.3. Logon policies

The administrator can configure logon policies to control the logon time and IP addresses of users.

Context

The system has a default policy as the initial configuration. You can configure logon policies to better control the read and write permissions of users and improve the system security.

Procedure

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

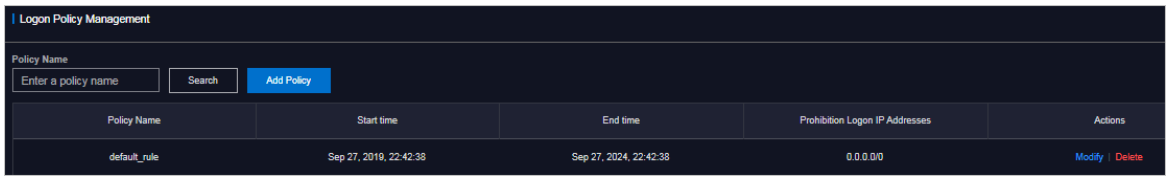

2. On the **Logon Policy Management** page, perform the following operations:

   - Query policies

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

   - Add a policy

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

   - Modify a policy

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

   - Delete a policy

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

3.2.4. User management

You can create users as an administrator and assign different user roles 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 Role management.

Context

User management provides different permissions for different users.

Procedure

1. In the left-side navigation pane, choose **System Management > Users**.The **Users** tab appears.
2. Perform the following operations:
   - Query users

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

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

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

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

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

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

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

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

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

- Bind a logon policy

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

- Query personal information of the current user

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

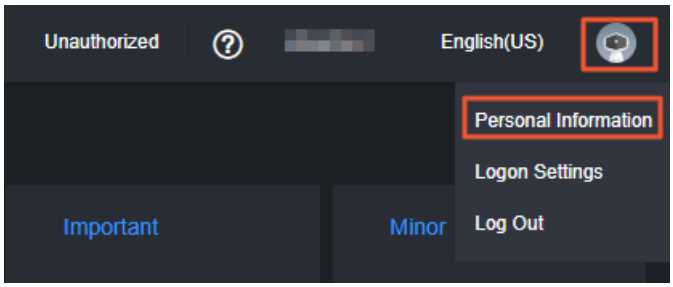

- Logon settings

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

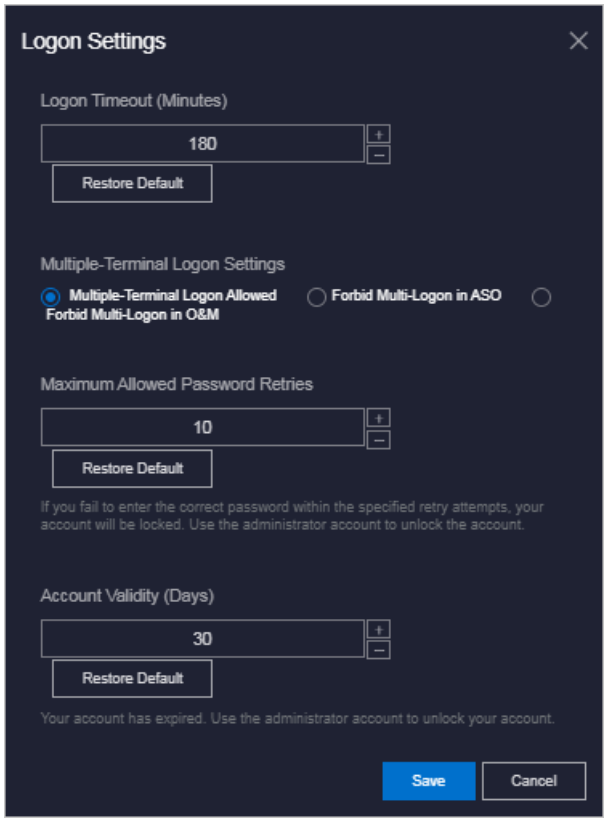

3.2.5. Two-factor authentication

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

Context

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

- Google two-factor authentication

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

- USB key authentication

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

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

- PKI authentication

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

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

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

Procedure

1. In the left-side navigation pane, choose **System Management > Two Factor Authentication**.
2. On the Two Factor Authentication page, you can perform the following operations:
   - Google two-factor authentication
     a. Set **Current Authentication Method** to **Google Two-Factor Authentication**.
     b. Click **Add User** in the upper-right corner of the page. In the Add User dialog box, enter a username and click OK. The added user is displayed in the user list.
     c. Find the user for whom you want to enable Google two-factor authentication, and then click **Create Key** in the **Actions** column. After the **Added** message appears, **Show Key** is displayed in the **Actions** column. Click **Show Key**, and the key is displayed in plain text.

d. Enter the key in the Google Authenticator app on your mobile phone. The app dynamically generates a verification code for logon based on the time and key. With two-factor authentication enabled, you are required to enter the verification code on your app when you log on to the system.

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

e. To disable two-factor authentication, click **Delete Key** in the **Actions** column.

- USB key authentication
  a. Set **Current Authentication Method** to **USB Key Authentication**.
  b. In the upper-right corner of the **Authentication Server Configuration** section, click **Add Server**. In the Add User dialog box, specify the **IP Address** and **Port** parameters for the server, and then click OK. The added server is displayed in the server list. Click **Test** to test the connectivity of the authentication server.
  c. In the upper-righter corner of the **User List** section, click **Add User**. The added user is displayed in the user list.
  d. Find the user for whom you are about to enable the USB key authentication, and then click **Bind Serial Number** in the **Actions** column. In the dialog box, enter the serial number to bind the user account with this serial number.

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

  e. Then, click **Enable Authentication** in the **Actions** column.
- PKI authentication
  a. Set **Current Authentication Method** to **PKI Authentication**.
  b. In the upper-right corner of the **Authentication Server Configuration** section, click **Add Server**. In the Add Server dialog box, specify the **IP Address** and **Port** parameters for the server. The added server is displayed in the server list. Click **Test** to test the connectivity of the authentication server.
  c. In the **User List** section, click **Add User**. In the Add User dialog box, specify **Username**, **Full Name**, and **ID Card Number**, and then click OK. The added user is displayed in the user list.
  d. (Optional)Find the user for whom you want to enable the PKI authentication, and then click **Bind** in the **Actions** column. Enter the full name and ID card number of the user to bind the user account with the name and ID card number.
  e. Then, click **Enable Authentication** in the **Actions** column.
- No authentication

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

3.2.6. Application whitelists

The system administrator can add, modify, or delete the application whitelist.

Context

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

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

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

Procedure

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

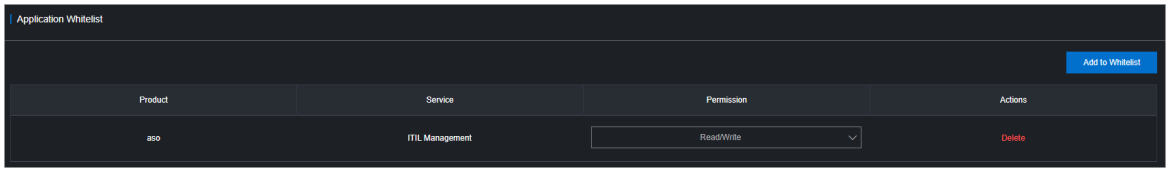

2. On the **Application Whitelist** page, perform the following operations:
   - Add a whitelist

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

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

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

3.2.7. Server password management

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

Context

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

- The system automatically collects information of all the servers in the Apsara Stack environment.
- The server password is automatically updated periodically.
- You can configure the password expiration period and password length.

- You can manually update the password of one or more servers at a time.
- The system records the history of server password updates.
- You can search for server passwords by product, hostname, or IP address.

Procedure

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

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

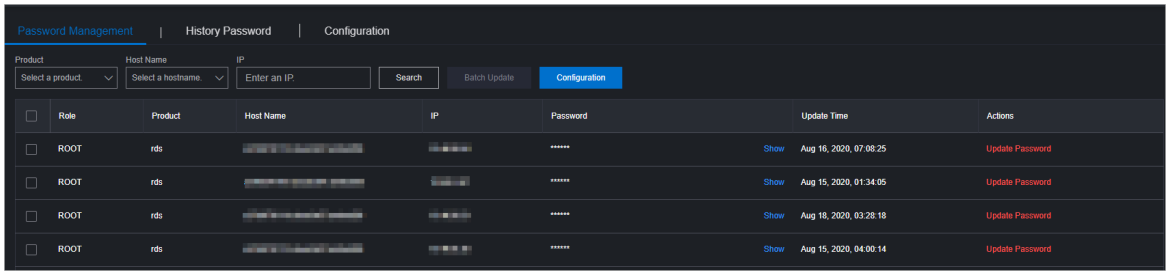

2. Perform the following operations:
   - Search for servers

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

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

        The passwords of the selected servers are updated.
   - Configure the password expiration period
     a. On the **Password Management** tab, select one or more servers.
     b. Click **Configuration** in the upper part of the tab.

c. In the **Configuration Item** dialog box, specify **Password Expiration Period** and **Unit**. Click **OK**.

Server passwords are updated immediately after the configuration and will be updated again after an expiration period.

- View the history of server password updates

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

- Show historical passwords of servers

  a. On the **History Password** tab, find a server.

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

- View and modify the password configuration policy

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

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

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

3.2.8. Operations logs

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

Context

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

Procedure

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

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

   - Delete logs

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

- Export logs

  Click the icon to export the displayed logs.

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

3.2.9. View authorization information

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

Prerequisites

Make sure that the current logon user has administrator permissions. Only a user with administrator permissions can view the trial authorization information or enter the authorization code to view the formal authorization information on the **Authorization Details** tab.

When a non-administrator user accesses this page, a message indicating that the user has insufficient permissions is displayed.

Procedure

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

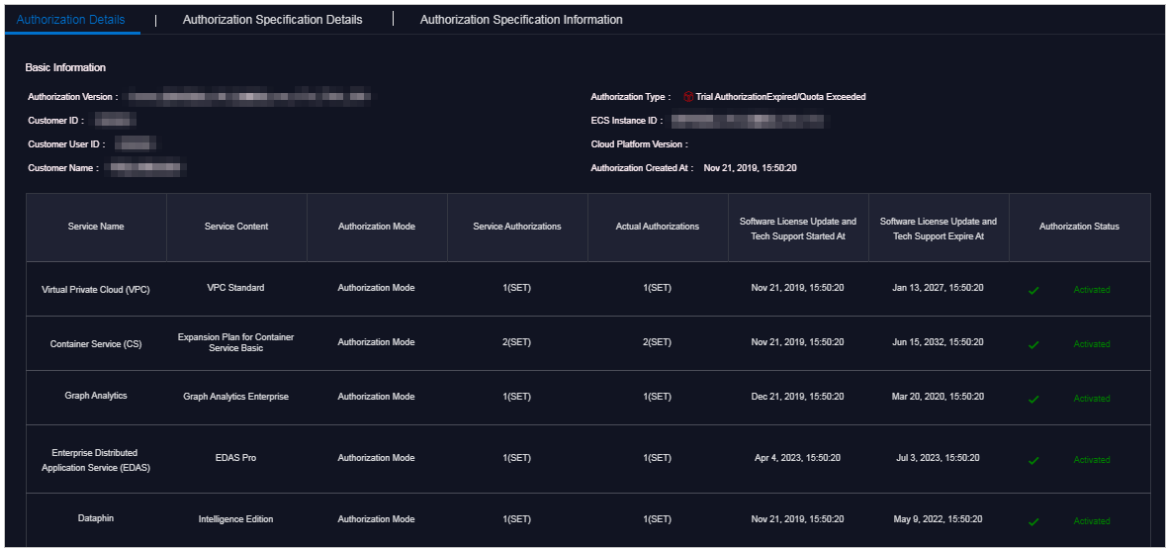

2. Perform the following operations to view the authorization information.

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

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

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

The following table describes the detailed authorization information.

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

- Click the **Authorization Specification Details** tab to view the authorization specification

information of a service.

The following table describes the authorization specification information and the corresponding description.

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.

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

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

3.2.10. Menu settings

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

3.2.10.1. Add a level-1 menu

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

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **System Management > Menu Settings**.
3. Click **Add**.
4. On the displayed page, complete the configurations for the level-1 menu you are about to add.

   For more information about the configurations, see the following table.

Configuration	Description

Configuration	Description
Menu Icon	Select the icon of the level-1 menu to be added from the drop-down list.
Menu Name	Enter the name of the level-1 menu to be added in Simplified Chinese, Traditional Chinese, and English.
Menu Order	The order, from top to bottom, of this menu in the level-1 menus.
Show/Hide	Whether to hide this level-1 menu. Turn on or off the switch to hide or show the menu. By default, the menu is not hidden.
Deletable	Whether this level-1 menu can be deleted after being added. Turn on or off the switch to configure whether the menu can be deleted. By default, the menu can be deleted. The setting cannot be modified after being configured.

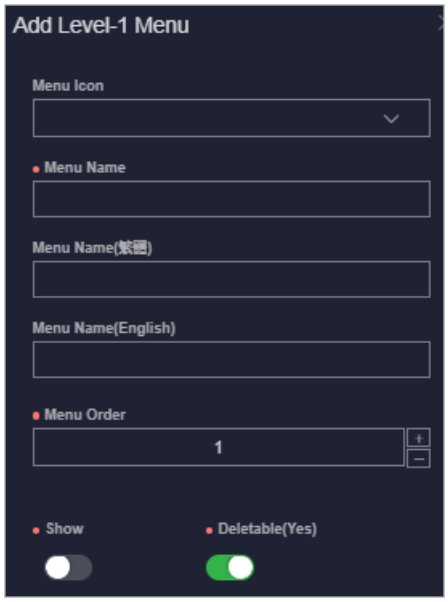

5. Click **OK**.

Result

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

3.2.10.2. Add a submenu

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

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **System Management > Menu Settings**.
3. Add a level-2 menu
   i. Find the level-1 menu to which you are about to add a level-2 menu, and then click **Add** in the

Actions column.

ii. On the displayed page, complete the configurations for the submenu you are about to add.

For more information about the configurations, see the following table.

Configuration	Description
Menu Name	Enter the name of the level-2 menu to be added in Simplified Chinese, Traditional Chinese, and English.
Menu Order	The order, from top to bottom, of this menu in the level-2 menus.
Show/Hide	Whether to hide this level-2 menu. Turn on or off the switch to hide or show the menu. By default, the menu is not hidden.
Deletable	Whether this level-2 menu can be deleted after being added. Turn on or off the switch to configure whether the menu can be deleted. By default, the menu can be deleted. The setting cannot be modified after being configured.
Link Address	Enter the menu path in the format of module name/path name. For example, /Dashboard/#/dashboardView.
Parent Menu	The parent menu of this menu.

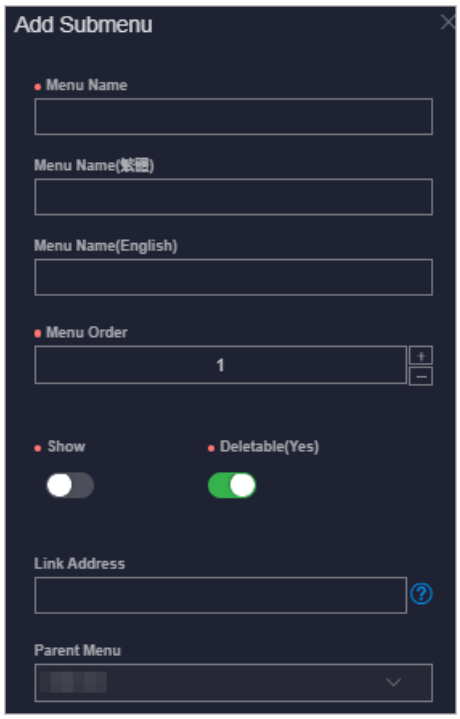

iii. Click **OK**.

Then, you can view the added level-2 menu under the corresponding level-1 menu in the menu list and the left-side navigation pane.

4. Click the button at the left of the level-1 menu to expand the level-2 menus. Add a level-3 menu. For more information, see the preceding step.

> **Note** The system only supports expanding menus of three levels. Therefore, you cannot add submenus for a level-3 menu.

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

3.2.10.3. Hide a menu

This topic describes how to hide a level-1, level-2, or level-3 menu.

Prerequisites

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

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **System Management > Menu Settings**.
3. Then, you can:
   - Hide a level-1 menu

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

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

3.2.10.4. Modify a menu

You can modify the icon, name, and order of an added menu.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **System Management > Menu Settings**.
3. In the menu list, find the level-1, level-2, or level-3 menu you are about to modify and then click **Modify** in the **Actions** column.
4. On the displayed page, modify the icon, name, and order of a level-1 menu, and modify the name, order, and link address of a level-2 or level-3 menu.

3.2.10.5. Delete a menu

You can delete a menu that is no longer in use based on business needs.

Prerequisites

> **Notice** You can only delete menus with **Deletable(Yes)** configured when being added.

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **System Management > Menu Settings**.
3. In the menu list, find the level-1, level-2, or level-3 menu you are about to delete and then click **Delete** in the **Actions** column.
4. In the displayed dialog box, click **OK**.

4.Monitoring

4.1. Alert Monitoring

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

4.1.1. Dashboard

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

Context

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

Procedure

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

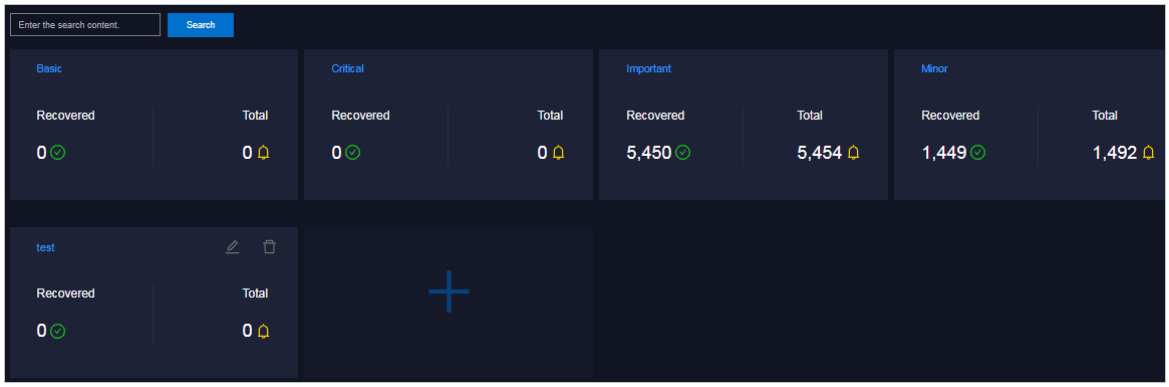

2. Perform the following operations:
   - View the total number of alerts and the number of recovered alerts in the basic, critical, important, and minor monitoring metrics, as well as custom filters.

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

   - Search for alerts

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

   - Add a custom filter

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

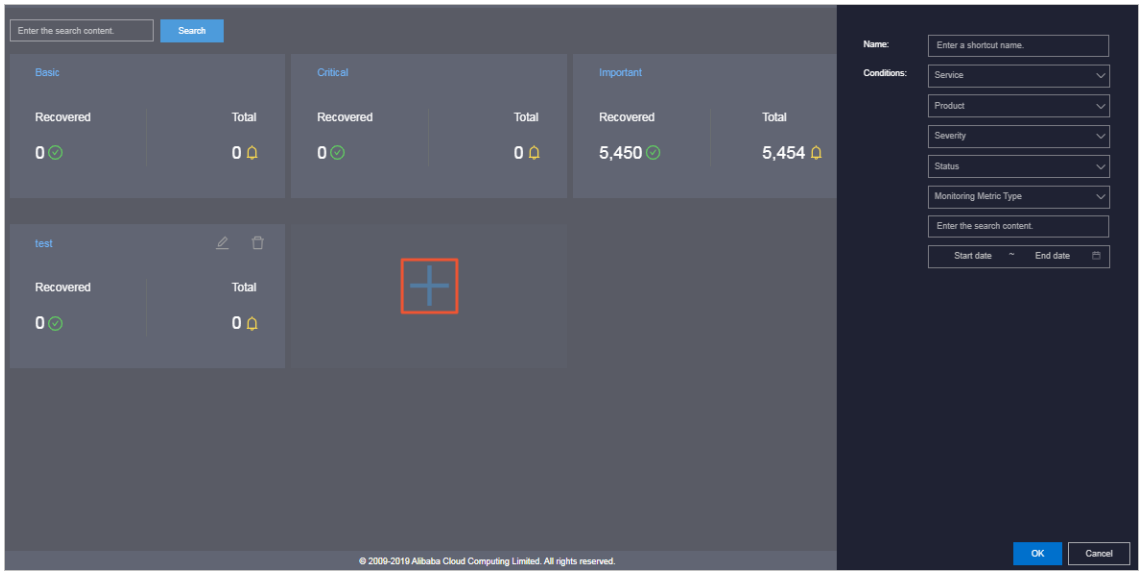

The following table describes the parameters for adding a filter.

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

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

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

- Modify a custom filter

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

- Delete a custom filter

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

4.1.2. Alert events

The Alert Events module displays the information of all alerts generated by the system on different tabs. The alert information is aggregated by monitoring item or product name. You can search for alerts based on filter conditions such as monitoring metric type, product, service, severity, status, and time range when the alert is triggered, and then perform O&M operations on the alerts.

Procedure

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

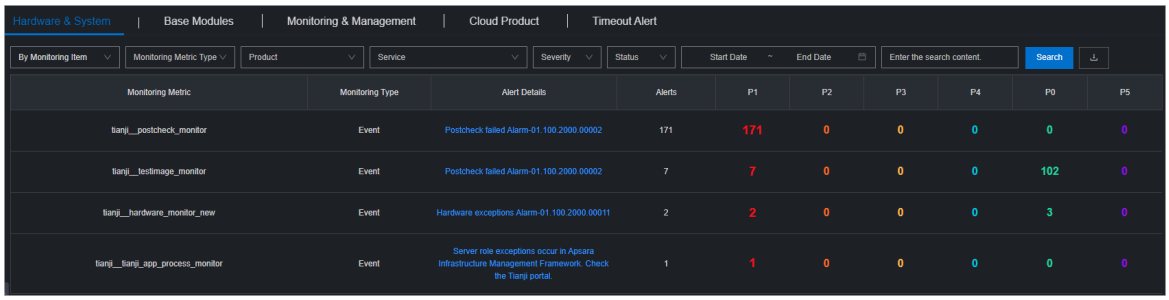

2. You can click the **Hardware & System**, **Base Modules**, **Monitoring & Management**, **Cloud Product**, or **Timeout Alert** tab, and perform the following operations:
   - Search for an alert

     In the upper part of the tab, you can search for an alert by specifying **Monitoring Metric Type**, **Product**, **Service**, **Severity**, **Status**, **Start Date**, **End Date**, or search content.
   - View alert sources
     a. If the alert information is aggregated by **Product Name** on this tab, click + on the left side of the product name to show the monitoring metrics. If the alert information is aggregated by **Monitoring Item** on this tab, skip this step.
     b. Find the monitoring metric and severity of the target alert, and then click the number in the specific severity column.
     c. Move the pointer over the alert source information in blue in the **Alert Source** column to view the alert source details.
   - View the details of a metric
     a. If the alert information is aggregated by **Product Name** on this tab, click + on the left side of the product name to show the monitoring metrics. If the alert information is aggregated by **Monitoring Item** on this tab, skip this step.
     b. Find the monitoring metric and severity of the target alert, and then click the number in the specific severity column.
     c. Click the alert details in blue in the **Alert Details** column. On the **Alert Details** page, you can view the alert information such as the alert description, reference, impact scope, and resolution.
   - View the original alert information of an alert

a. If the alert information is aggregated by **Product Name** on this tab, click + on the left side of the product name to show the monitoring metrics. If the alert information is aggregated by **Monitoring Item**, skip this step.

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

c. Click the number in blue in the **Alerts** column. The **Alerts** pane appears.

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

- Process alerts

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

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

  - If an alert is being processed by operations engineers, choose **Actions > Process** in the **Actions** column to set to **In Process**.
  - If the alert has been processed, choose **Actions > Processed** in the **Actions** column to set the alert status to **Processed**.
  - To view the whole processing flow of an alert, choose **Actions > Alert Tracing** in the **Actions** column.
  - View the recent monitoring data

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

- Export reports

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

4.1.3. Alert history

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

Procedure

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

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

   - Export the alert list

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

   - View alert sources

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

- View the details of a metric

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

- View the original alert information

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

- View the alert duration

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

4.1.4. Alert configuration

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

4.1.4.1. Alert contacts

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

Procedure

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

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

   - Add an alert contact

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

   - Modify an alert contact

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

   - Delete an alert contact

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

4.1.4.2. Alert contact groups

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

Procedure

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

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

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

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

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

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

4.1.4.3. Configure static parameters

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

Context

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

Procedure

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

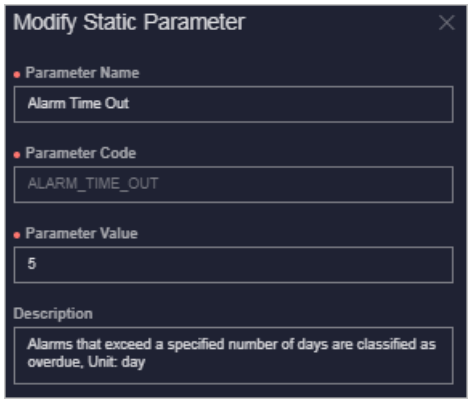

Parameter	Description
Parameter Name	Enter a parameter name related to the configuration.
Parameter Value	Enter the parameter value. The default value is 5, indicating five days. After you complete the configuration, you can choose **Alert Monitoring > Alert Events** and then click the **Timeout Alert**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

6. Click **OK**.

4.1.5. Alert overview

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

Procedure

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

   The **Alert Overview** page appears.

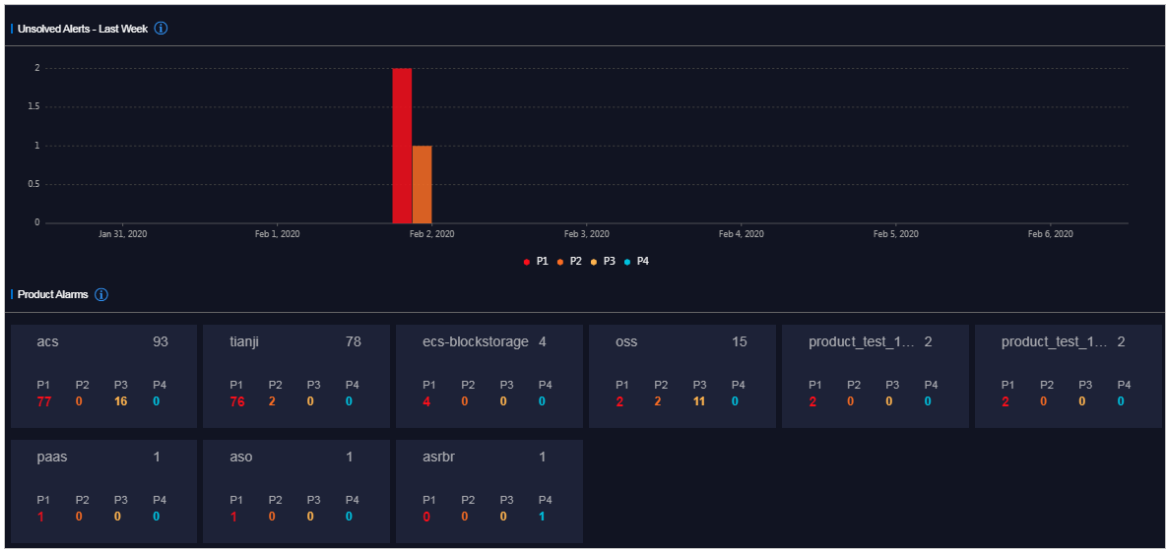

- The column chart in the upper part of the page shows the number of unresolved alerts for the last seven days.
- The section in the lower part of the page shows the alert statistics in the current system by service.

4.1.6. Alert subscription and push

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

Procedure

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

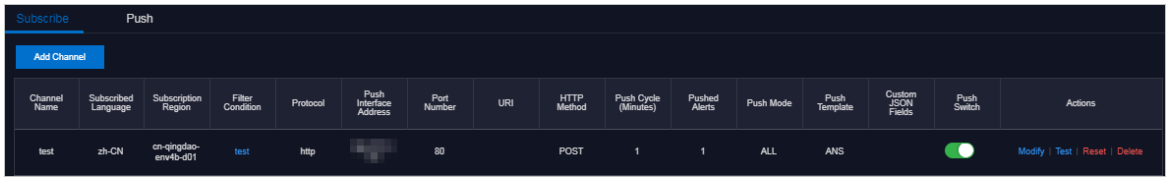

2. On the **Subscribe** tab, click **Add Channel**.
3. In the **Add Subscription** pane, configure the following parameters.

Parameter	Description
Channel Name	The name of the subscription channel.
Subscribed Language	The subscription language. Valid values: Chinese and English.
Subscription Region	The region where the subscription is located.

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

Parameter	Description
Push Switch	Specifies whether to push alerts. If the switch in this pane is not turned on, after you configure the subscription channel, you can enable the push feature in the **Push Switch** column.

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

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

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

4.1.7. Alert masking

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

4.1.7.1. Add masking rules

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

Procedure

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

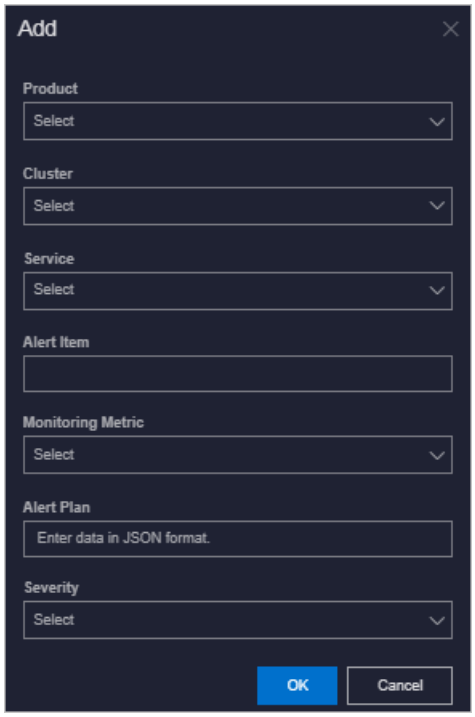

Parameter	Description
Product	Optional. The product to which the alerts to be masked belong.
Cluster	Optional. The cluster to which the alerts to be masked belong.
Service	Optional. The service to which the alerts to be masked belong.
Alert Item	Optional. The alert name to be masked. **Note** When you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

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

4. Click **OK**.

Result

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

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

4.1.7.2. Remove the masking

You can remove the masking for masked alerts.

Procedure

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

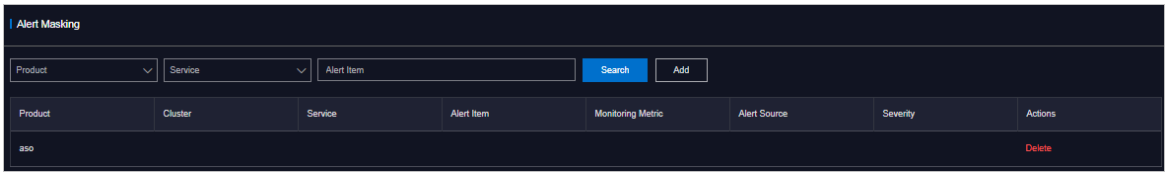

4. In the message that appears, click **OK**.

Result

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

4.2. Physical servers

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

4.2.1. View the physical server information

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

Product tab

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

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

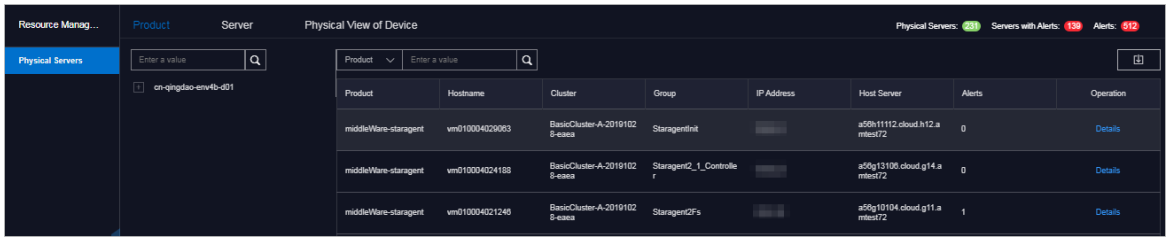

2. On the Product tab, perform the following operations to view the physical server information:
   - Expand the left-side navigation tree by selecting a region, product, and cluster in sequence to view the list of physical servers where a cluster of a service is located.
   - In the left-side search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

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

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

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

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

Server tab

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

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

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

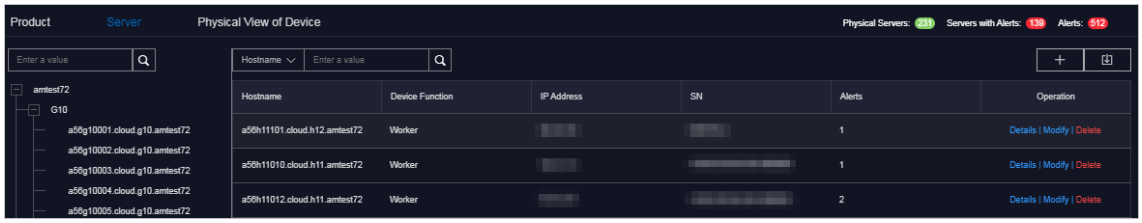

4. To view the details of a physical server, enter the hostname, IP address, device function, or serial number (SN) in the right-side search box and press the Enter key.
5. Find the physical server whose details you are about to view and then click **Details** in the **Actions** column. On the **Physical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

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

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

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

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

The Physical View of Device tab

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

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

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

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

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

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

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

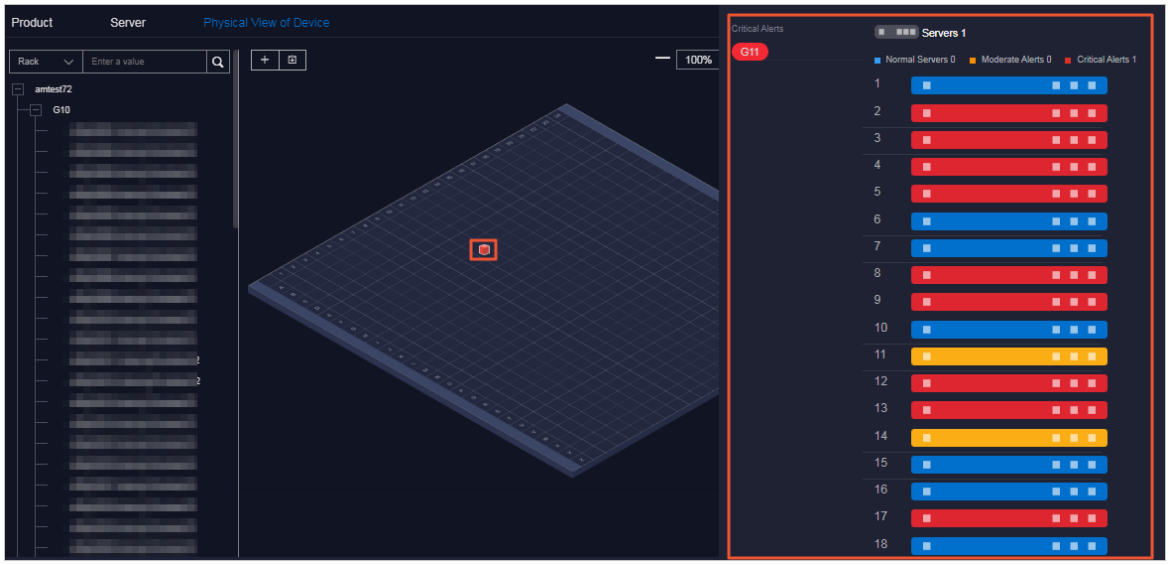

4. To view the details of a physical server, perform the following operations:

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

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

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

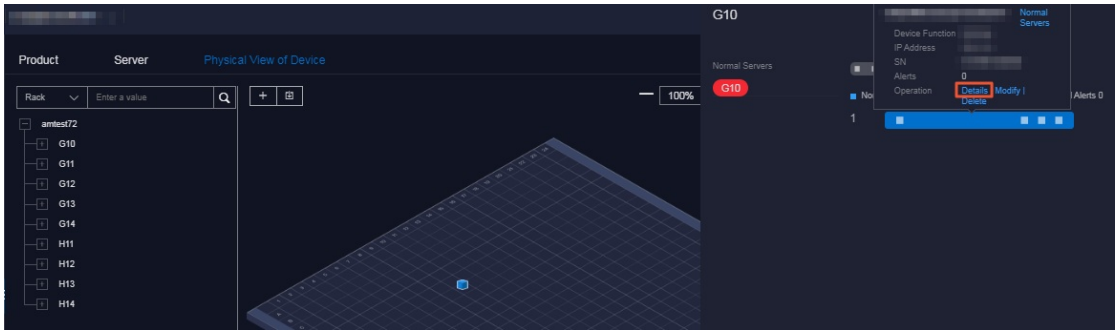

iv. On the **Physical Server Details** page, view the basic information, monitoring details, and alert information of the physical server.

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

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

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

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

4.2.2. Add physical servers

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

Procedure

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

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

   The following table describes the parameters.

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

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

5. Click **OK**.

4.2.3. Modify a physical server

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

Server tab

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

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

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

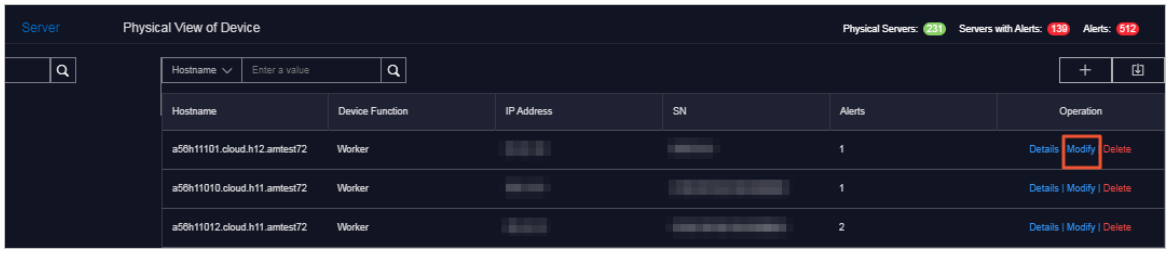

5. In the **Modify Physical Server** pane, modify the physical server information. You can modify the following physical server information: zone, data center, rack, room, physical server name, memory size, disk size, CPU cores, rack group, server type, server role, serial number, operating system template,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

Physical View of Device tab

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

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

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

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

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

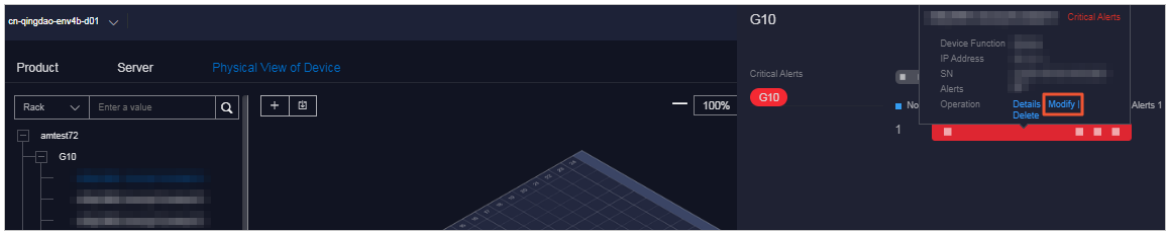

6. In the **Modify Physical Server** pane, modify the physical server information. You can modify the following physical server information: zone, data center, rack, room, physical server name, memory size, disk size, CPU cores, rack group, server type, server role, serial number, operating system template,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

4.2.4. Export server information

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

Product tab

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

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

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

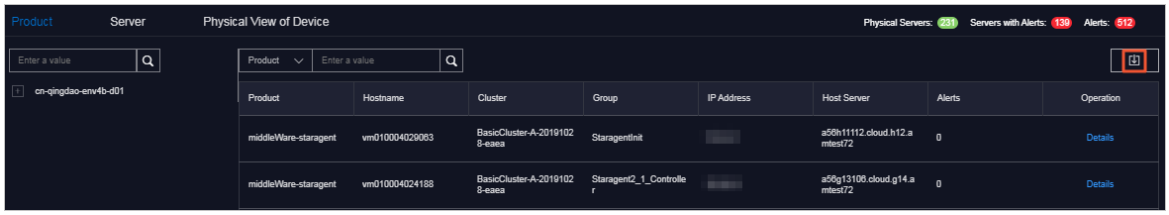

2. In the upper-right corner of the tab, click the icon to export the information of all the physical servers of all services to your local computer.

Server or Physical View of Device tab

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

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

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

4.2.5. Delete a physical server

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

Server tab

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

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

Physical View of Device tab

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

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

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

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

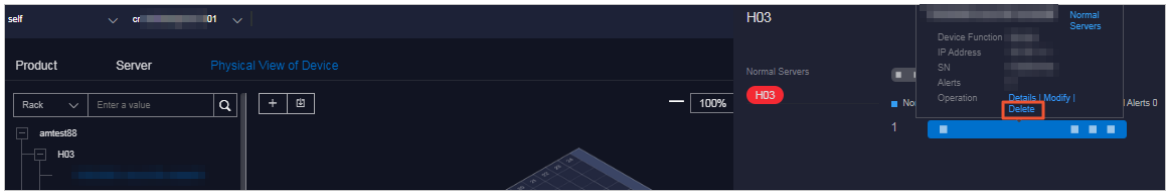

6. In the message that appears, click **OK**.

4.3. Inventory Management

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

4.3.1. View the RDS inventory

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

Procedure

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

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

2. View the RDS inventory.

   On this page:

   - The **RDS Inventory** section shows the inventory of different RDS services for the last five days. Different RDS services are displayed in different colors.
   - You can query the RDS inventory by pages by specifying **Engines** or **Date** in the **RDS Inventory Details** section.

4.3.2. View the OSS inventory

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

Procedure

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

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

2. View the OSS inventory.

   On this page:

   - The **Inventory Availability History (TB)** section shows the available OSS inventory for the last five days.
   - The **Current Inventory Usage (TB)** section shows the amount and percentage of OSS inventory currently being used.
   - You can query the OSS inventory details by pages by specifying **Date** in the **OSS Bucket Inventory Details** section.

4.4. Storage operation center

The Storage Operation Center module contains Apsara Distributed File System and miniOSS.

4.4.1. Apsara Distributed File System

The Apsara Distributed File System module shows the overview information, cluster information, node information, and the statuses of clusters.

4.4.1.1. Overview

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

Procedure

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

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

   - **Overview**

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

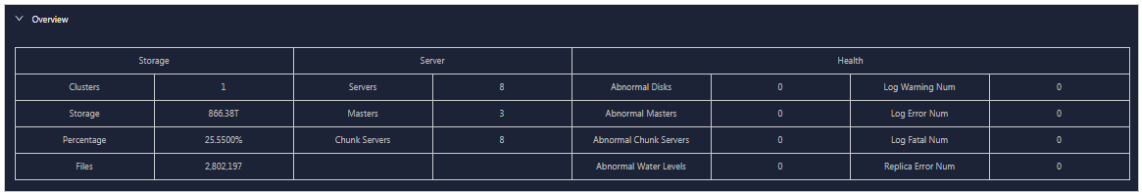

- **Heatmap of Health**

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

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

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

- **Data of Top 5 Services**

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

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

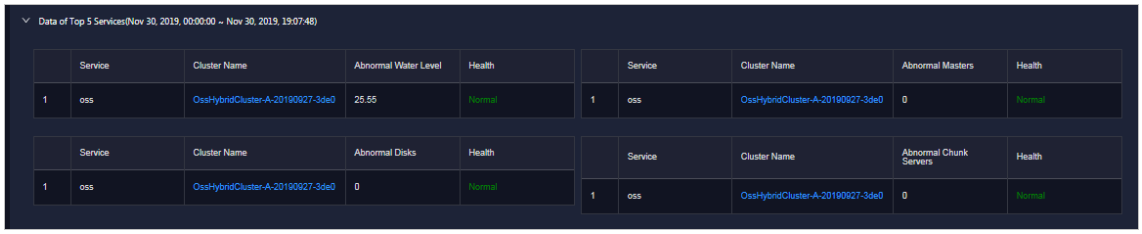

4.4.1.2. Cluster information

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

Procedure

1. Log on to the ASO console.
2. In the left-side navigation pane, choose **Storage Operation Center > Pangu > Cluster Information**.

   By default, data of the first cluster in the **Cluster Name** drop-down list is displayed.
3. Select the cluster that you are about to view from the **Cluster Name** drop-down list and then view the following information.

> **Note** All the accessed clusters that are not in the closed status in the current environment are available for you to select from the **Cluster Name** drop-down list.

- **Overview**

  Displays the storage space, server information, and health information of the selected cluster. Values of **Abnormal Water Levels**, **Abnormal Masters**, **Abnormal Chunk Servers**, and **Abnormal Disks** in the **Health** section are displayed in red if they are larger than zero.

- **Alarm Monitor**

  Displays the alert information of the selected cluster. You can perform a fuzzy search based on a keyword.

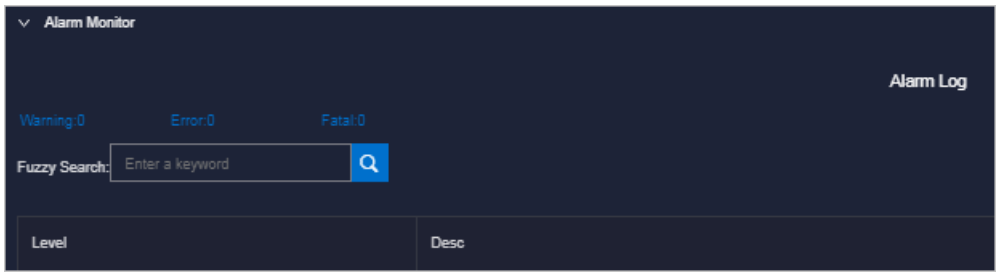

- **Replica**

  Displays the replica information of the selected cluster.

- **Run Chart of Clusters**

  Displays the charts of historical water levels, predicted water levels, number of files, number of chunk servers, and number of disks for the selected cluster.

  Predicted Water Levels predicts the run chart of the next seven days.

> **Note** Predicted Water Levels has values only if Historical Water Levels has a certain amount of data. Therefore, some clusters may only have historical water levels, without predicted water levels.

- **Rack Information**

  Contains Servers in Rack and Storage. Where,

  - **Servers in Rack** displays the number of servers in each rack of the selected cluster.

    

  - **Storage** displays the total storage and used storage in each rack of the selected cluster.

    

4.4.1.3. Node information

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

Procedure

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

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

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

following information is displayed:

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

- **Master Info**

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

- **Chunk Server Info**

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

4.4.1.4. Operations and maintenance

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

Procedure

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

- Dark red indicates that the cluster has a fatal error.
- Grey indicates that the cluster is disabled.

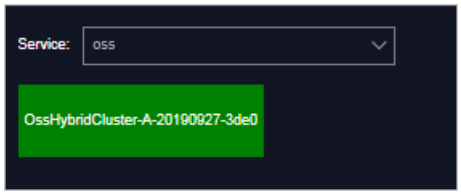

3. Move the pointer over a cluster name to view the service name, server name, and IP address to which the cluster belongs.

4.4.1.5. Product configuration

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

Procedure

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

   The following table describes the parameters.

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

Section		Description
	Fatal Error Threshold	When the storage usage of the cluster is greater than or equal to this value, a fatal-error alert is triggered and the health heatmap of the cluster is displayed in dark red. Value range: (0,100]. If this parameter is not specified, a fatal-error alert is triggered by default when the water value of the cluster is greater than or equal to 92%.
Chunk Server	**Warn Threshold (Abnormal Chunk Server Quantity)**	When the number of abnormal chunk servers is greater than or equal to this value, a warning alert is triggered and the health heatmap of the cluster is displayed in yellow. If this parameter is not specified, a warning alert is triggered by default when the number of abnormal chunk servers is greater than or equal to 1.
	Error Threshold (Chunk Server Ratio)	If the ratio of abnormal chunk servers to all the chunk servers is greater than this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. If this parameter is not specified, an error alert is triggered by default when the ratio of abnormal chunk servers to all the chunk servers is greater than or equal to 10%.
Disk	**Warn Threshold (Abnormal Disk Quantity)**	When the number of abnormal disks is greater than or equal to this value, a warning alert is triggered and the health heatmap of the cluster is displayed in yellow. If this parameter is not specified, a warning alert is triggered by default when the number of abnormal disks is greater than or equal to 1.
	Error Threshold (Abnormal Disk Ratio)	When the ratio of abnormal disks to all the disks is greater than this value, an error alert is triggered and the health heatmap of the cluster is displayed in red. If this parameter is not specified, an error alert is triggered by default when the ratio of abnormal disks to all the disks is greater than or equal to 10%.

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

4. Click **Save**.

4.4.2. miniOSS

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

4.4.2.1. Monitoring dashboard

The Monitoring Dashboard module allows you to view the overview, bucket watermark heatmap, user quota watermark heatmap, watermark trend, and network traffic trend of miniOSS in the system, and download logs to your local computer.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > miniOSS > Monitoring Dashboard**.
2. On the page that appears, view the following information:
   - **Overview**

     This section displays 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 value is displayed in red.

     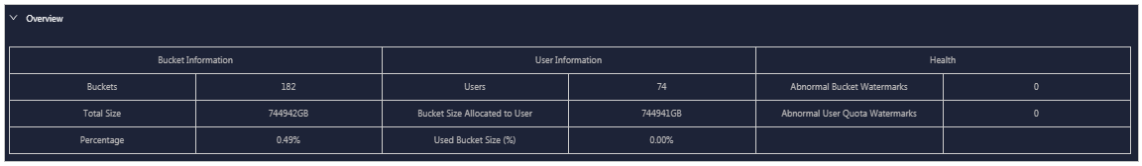

   - **Bucket Watermark Heatmap**

     This section displays the bucket capacity usage.

     The number of sections in **Bucket Watermark Heatmap** is the same as the value of Buckets in **Overview**. Buckets in different statuses are displayed in different colors:

     - Green indicates that the bucket works 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.

     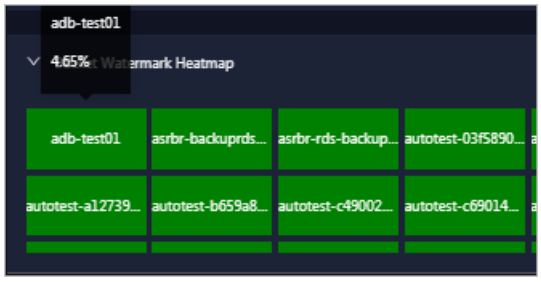

   - **User Quota Watermark Heatmap**

This section displays the user quota watermark information.

User quota watermark = Used capacity of all buckets of the user/Total capacity of all buckets of the user. Buckets of different watermark values are displayed in different colors:

- Green indicates that the bucket works 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 displays the historical water levels and predicted water levels of a user or bucket. Watermark represents the disk utilization, and watermark of a user indicates the disk usage of buckets.

  Data in the watermark trend comes from 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 watermark trend.

  > **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 watermarks is displayed on the right. Click a bucket name in the top 10 data to view the watermark trend of the bucket on the left.

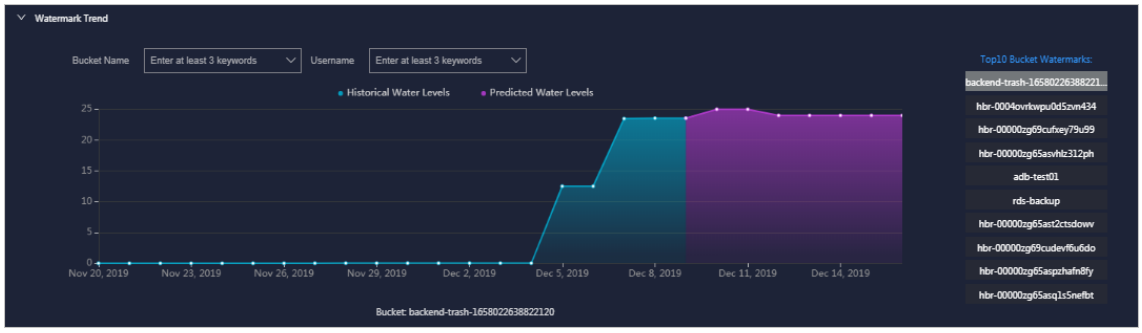

- **Network Traffic Trend**

  This section displays 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.

  In the network traffic trend:

  - 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.

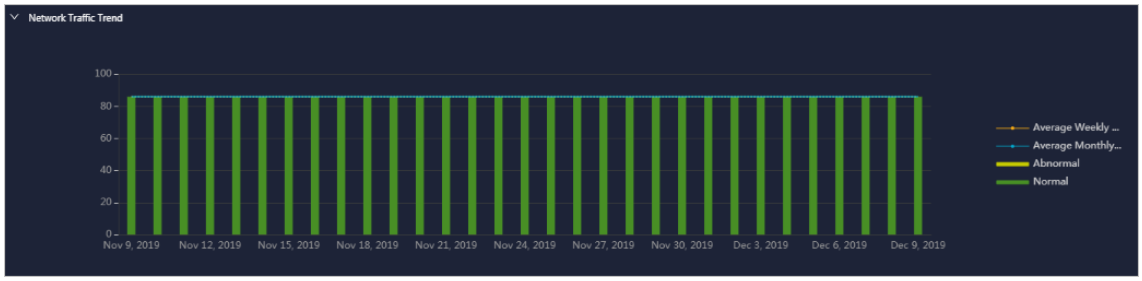

3. (Optional)In the **Download Log** section, click **Download Log Package**, and then use the download URL to download logs to your local computer for subsequent review and analysis.

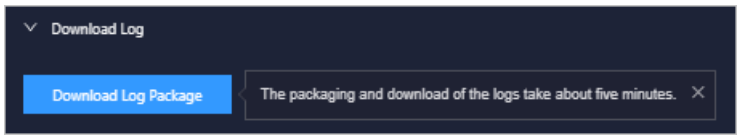

4.4.2.2. User management

The User Management module consists of User List, Bucket List of User, and Network Traffic Control. You can use this module to view user information, network traffic bandwidth, and the list of user buckets.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > miniOSS > User Management**.
2. On the page that appears, perform the following operations:
   - View the 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 SecretKey of the user.

     In the **User List** section, enter three characters of the username, such as def, and then 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 user that meets the search condition.

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

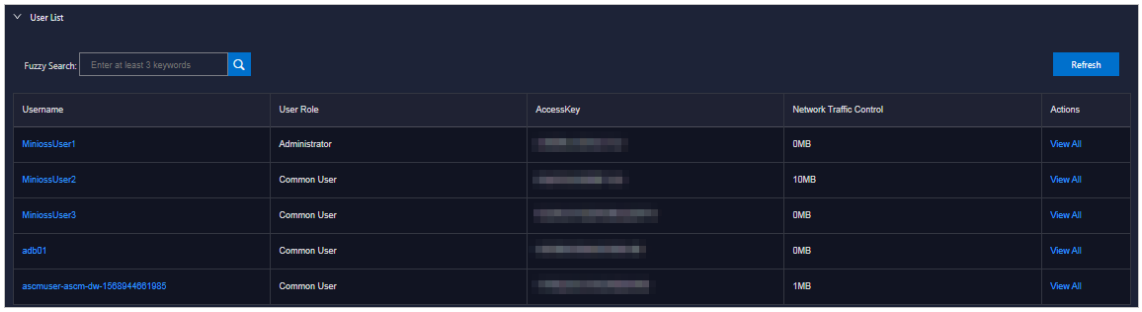

   - View the bucket information of a user

Click a username in the **User List** section. View the 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.

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

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

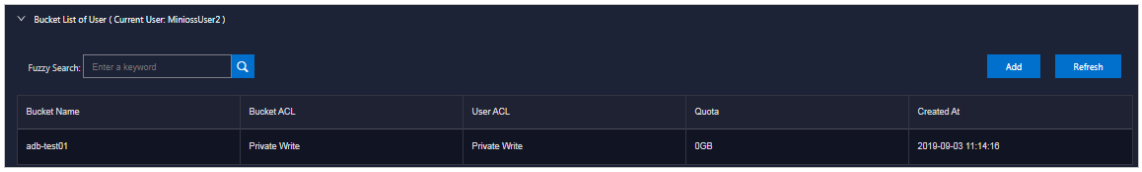

- Add a bucket for a common user

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

Find the common user for whom you are about to add a bucket in the **User List** section, and then click the username. In the **Bucket List of User** section, click **Add**. In the dialog box that appears, select the bucket and User ACL, enter the quota, and then click **OK**.

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

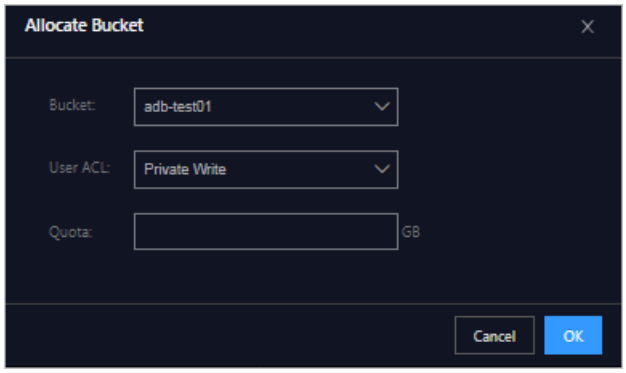

- View the network traffic bandwidth of a user

  Find the user whose network traffic bandwidth you are about to view in the **User List** section, and then click the username. View the network traffic bandwidth of the user in the **Network Traffic Control** section.

- Modify the network traffic bandwidth of a user

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

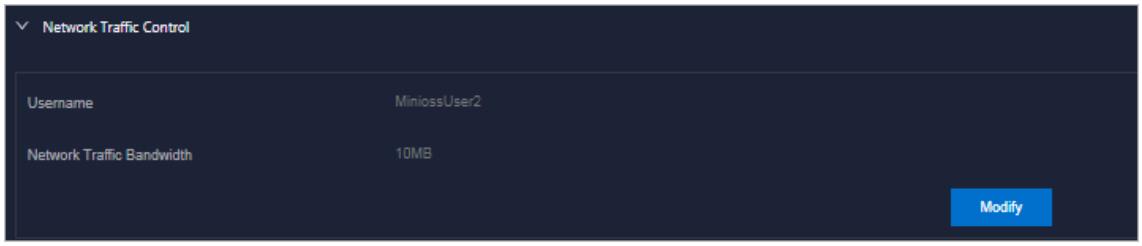

4.4.2.3. Permission and quota management

The Permission/Quota Management module allows you to view the bucket list and user list of bucket, and add, modify, and delete a bucket.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > miniOSS > Permission/Quota Management**.
2. On the page that appears, perform the following operations:
   - View the bucket information

     By default, all the buckets are displayed in the bucket list. In the **Bucket List** section, you can view the basic information, including the bucket name, bucket ACL, quota, and network traffic bandwidth, of all the buckets.

     Enter a keyword of the bucket name in the search box in the upper-left corner, and then press the Enter key or click the search icon. The information of the bucket that meets the search condition is displayed.

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

   - Add a bucket

     In the **Bucket List** section, click **Add**. In the dialog box that appears, enter the bucket name, and then 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 then click **OK**.

   - Delete a bucket

     In the **Bucket List** section, click **Delete** in the **Actions** column corresponding to a bucket. In the dialog box 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.

     Enter a keyword of the username in the search box in the upper part of the page, and then press the Enter key or click the search icon. The information of the user that meets the search condition is displayed.

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

4.4.2.4. Array monitoring

The Array Monitoring module allows you to view the running status of each device.

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > miniOSS > Array Monitoring**.
2. View the running status of each device.

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

   The 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.

   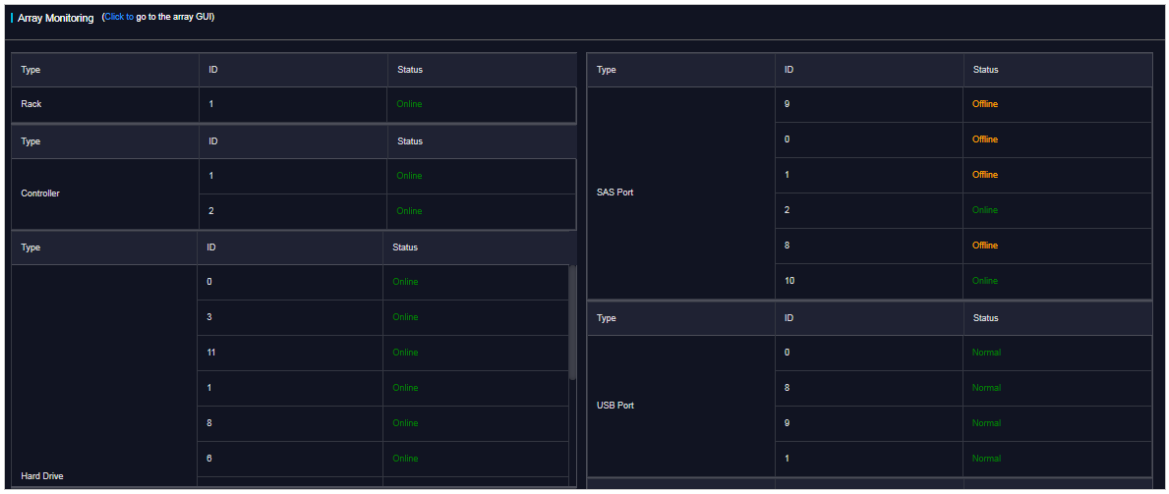

   Values in different colors indicate different states:

   - 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.

4.4.2.5. System management

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

Procedure

1. In the left-side navigation pane, choose **Storage Operation Center > miniOSS > System Management**.
2. On the page that appears, perform the following operations:
   - Modify the bucket watermark threshold

a. In the **Bucket Watermark Threshold** section, click **Modify**.

b. Enter a positive number less than or equal to 100 as the warning value, error value, and fatal error value. Make sure that the warning value is less than the error value which is less than the fatal error value.

c. Click **Save**.

- Modify the user watermark threshold

a. In the **User Watermark Threshold** section, click **Modify**.

b. Enter a positive number less than or equal to 100 as the warning value, error value, and fatal error value. Make sure that the warning value is less than the error value which is less than the fatal error value.

c. Click **Save**.

5.Operations tools

5.1. NOC

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

5.1.1. Network topology

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

Procedure

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

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

   > **Note**
   >
   > The colors of connections between network devices indicate the connectivity between the network devices:
   >
   > - Green: The link works properly.
   > - Red: The link has an error.
   > - Grey: The link is inactive.

   By default, if **Topology Type** is set to **Standard Topology**, the **Refresh Alert** switch is turned on. You can turn off **Refresh Alert**, and then devices or link status in the topology are not updated after new alerts are triggered.

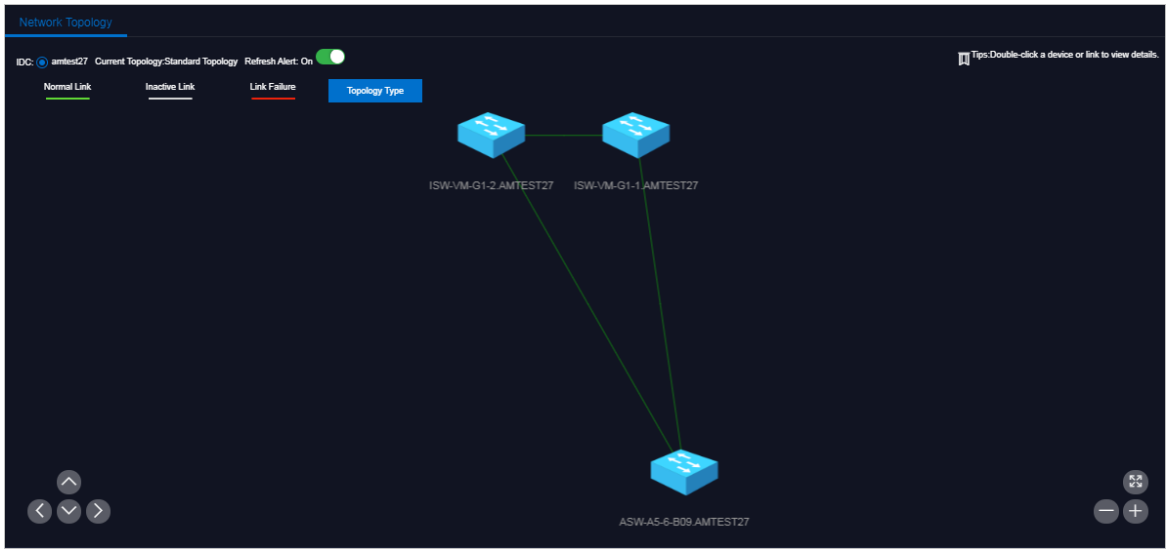

3. In the topology, double-click a connection between two devices to view the links and alerts between the two devices.
4. In the topology, double-click a physical network device to view the basic information and node alerts of the device on the right.

5.1.2. Resource management

The Resource Management module is used to manage network-related resources, including the information of physical network element devices, virtual network products, and IP addresses.

5.1.2.1. Device management

The **Device Management** page displays the basic information, running status, traffic monitoring, and logs of physical network element devices, and allows you to configure the collection settings of network devices.

5.1.2.1.1. View the network monitoring information

The Network Monitoring tab allows you to view the basic information, running status, and traffic monitoring of Apsara Stack physical network devices and check the health status of network devices in a timely manner.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. In the **Device Management** page, click the **Network Monitoring** tab.

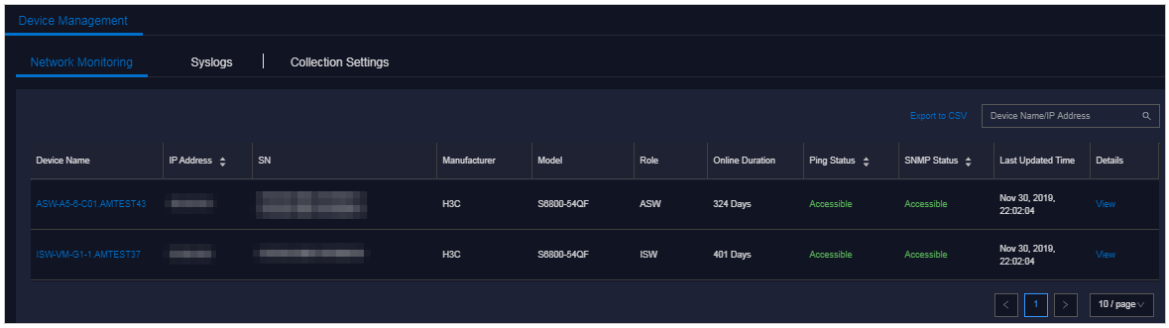

3. Perform the following operations:
   - View the basic information, ping status, and SNMP status of physical network devices in Apsara Stack.

     > **Note** You can also click **Export to CSV** to export network device information to your local computer.

     If a device has a business connectivity or gateway connectivity problem, the value in the Ping Status column or SNMP Status column turns from green to red. The operations personnel are required to troubleshoot the problem.
   - In the search box in the upper-right corner, enter the device name or IP address to search for the monitoring information of a specific device.
   - View the port information and alert information of a device.
     a. Click a device name, or click **View** in the **Details** column corresponding to a device.
     b. View the port list, port working status, and other link information of the device in the **Port** column.

c. View the alert information of the device in the **Alert Info** column.

During routine O&M, pay attention to the alert list of the device. Typically, if no data is displayed in the **Alert Info** column, it indicates that the device is operating normally.

If alert events occur, unrecovered alert events are displayed in the list. You must handle these exceptions in time. After you handle exceptions, the alert events are automatically cleared from the list.

- View the traffic information of a device for a specified port and time range.
  a. Click a device name, or click **View** in the **Details** column corresponding to a device.
  b. Search for the port that you are about to view by using the search box in the upper-right corner of the **Port** section. Click **View** in the **Details** column corresponding to the port.

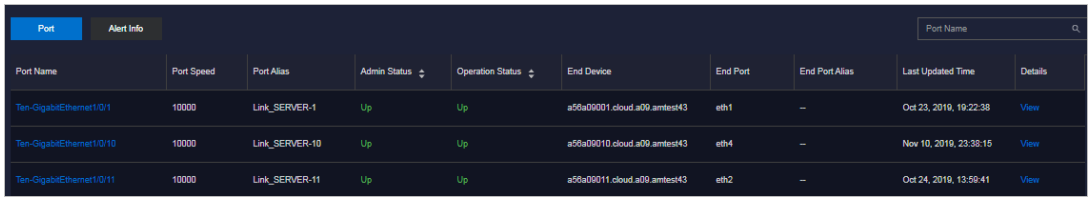

c. Select a time range on the right, and then click **Search** to view the traffic in the selected time range.

You can select 5MIN, 30MIN, 1H, or 6H in the **Quick Query** section to view the traffic within 5 minutes, 30 minutes, 1 hour, or 6 hours.

5.1.2.1.2. View logs

The Syslogs tab allows you to view logs of physical network element devices, providing necessary data for fault location and diagnosis information collection if a fault occurs.

Context

During the daily inspection, you can search for logs generated by a specific network device during a specific time range on the **Syslogs** tab.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. On the **Device Management** page, click the **Syslogs** tab.
3. In the upper-right corner of the tab, select a device name from the drop-down list, select a time range, and then click **Search** to check whether the device has generated system logs in the specified time range.

   If the device has a configuration exception or does not have any generated logs for the specified time range, no search results will be returned.

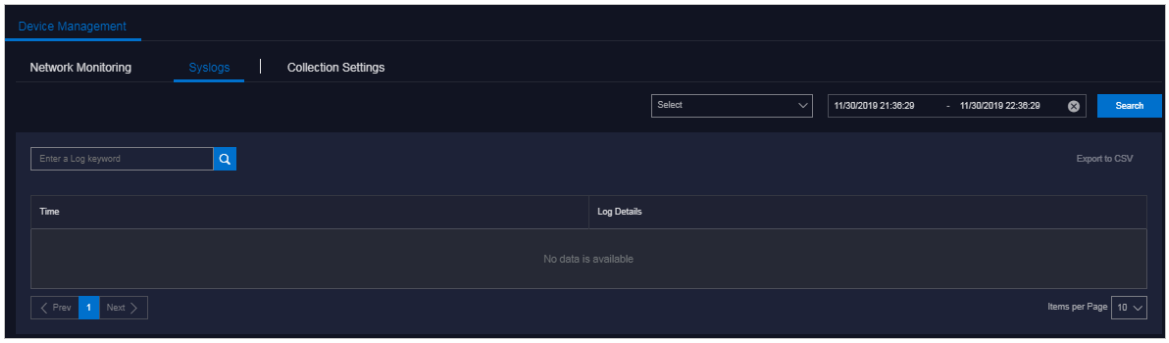

4. (Optional)You can filter the search results based on the log keyword.
5. (Optional)Click **Export to CSV** in the upper-right corner o to export the search results to your local computer.

5.1.2.1.3. Collection settings

The **Collection Settings** tab allows you to configure the collection interval of physical network element devices and manage OOB network segments.

5.1.2.1.3.1. Modify the collection interval

You can modify the collection interval to adjust the time interval of collection.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. On the **Device Management** page, click the **Collection Settings** tab.
3. In the **Collection Interval Settings** section, modify the values.

   > **Note** To cancel your modification before submission, click **Reset** in the upper-right corner to reset the collection interval to the previous version.

4. Click **Submit**. One minute later, the modified collection interval of the network device information takes effect.

5.1.2.1.3.2. Add an OOB network segment

If this is the first time you are using the Network Elements feature of Network Operation Center (NOC), you must add the device loopback network segment planned by the current Apsara Stack network device, which is typically the network segment of the netdev.loopback field in the IP address planning list.

Context

The OOB Network Segments section is used to configure the management scope of a physical network element device. Typically, operations engineers are required to add the loopback network segment where the network device to be managed resides.

In the Apsara Stack scenario, a loopback network segment is used to configure the management scope of a physical network element device. To expand the network and the loopback network segment, you must add the network segment involved in the expansion to the management scope. The procedure to add an expanded network segment is the same as that used to add the loopback network segment for the first time. Then, you can search for the network segment of the managed device on this page.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. On the **Device Management** page, click the **Collection Settings** tab.
3. In the lower part of the **OOB Network Segments** section, click **Add Network Segment**.
4. In the Add Network Segment dialog box, enter the network segment that contains the mask information and subnet mask, and select an IDC.

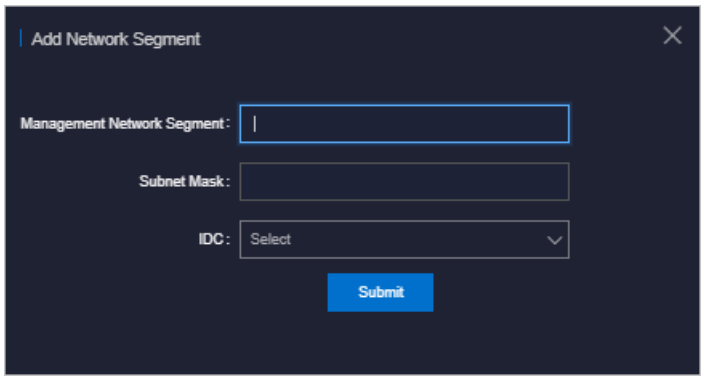

5. Click **Submit**. The initial data entry is completed.

   To modify or delete an OOB network segment, find it in the list, and then click **Edit** or **Delete** in the **Actions** column.

5.1.2.1.3.3. View the OOB network segment information

You can search for and view the network segment information of your managed device.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Network Elements**.
2. On the **Device Management** page, click the **Collection Settings** tab.
3. In the **OOB Network Segments** section, click **Refresh** in the upper-right corner of the section.

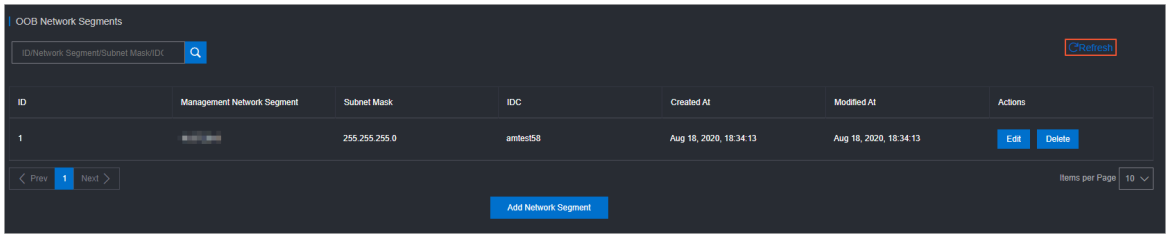

4. In the list, view the network segment information of your managed device.

> **Note** You can search for the information of a specific network segment by entering a keyword in the search box.

5.1.2.2. View the instance monitoring information

The Instance Monitoring tab allows you to view the basic information and water level of an instance, including the bps and pps.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Server Load Balancers**.
2. Click the **Instance Monitoring** tab.
3. Select the cluster where the target instance resides from the cluster drop-down list. Enter the VIP address that you are about to search for in the field, and then click **Search**.
4. View the water level data of the VIP address. Select a time range, and then click **Search**. Alternatively, select 5MIN, 30MIN, 1H, or 6H in the Quick Query section to view the operating water level graph of the VIP address in a specific time range.

5.1.3. Alert management

The Alert Management module provides you with the real-time alert dashboard, history alert dashboard, and the alert settings function.

5.1.3.1. View and process current alerts

You can view and process current alerts on the Current Alerts tab.

Procedure

1. In the left-side navigation pane, choose **NOC > Alert Management > Alert Dashboard**.
2. Click the **Current Alerts** tab.
3. Enter a keyword in the search box in the upper-right corner, and then click **Search**.Alerts that meet the search conditions are displayed.
4. (Optional)You can filter the search results by device name, device IP address, or alert name.
5. Click **Details** in the **Details** column corresponding to an alert to view detailed information about the alert.
6. Find the reason why the alert is triggered and then process the alert.
   - If the alert does not affect the normal operation of the system, you can click **Ignore** in the **Actions** column to ignore the alert.
   - If the alert is no longer significant, you can click **Delete** in the **Actions** column to delete the alert.

     After the alert is processed, you can search for it on the **History Alerts** tab.
7. (Optional)Click **Export to SCV** to export the alert information to your local computer.

5.1.3.2. View historical alerts

You can view historical alerts on the History Alerts tab.

Procedure

1. In the left-side navigation pane, choose **NOC > Alert Management > Alert Dashboard**.
2. Click the **History Alerts** tab.
3. Select Alert Source, Alerting IP Address, Alerting Device, Alert Name, Alert Item, or Alerting Instance from the drop-down list, and then enter a keyword in the field. Select a time range, and then click **Search**.Alerts that meet the search conditions are displayed.
4. Click **Details** in the **Details** column corresponding to an alert to view detailed information about the alert.
5. (Optional)Click **Export to SCV** to export the alert information to your local computer.

5.1.3.3. Add a trap

If the initially configured trap subscription cannot meet the monitoring requirements, you can add a trap as needed for monitoring match.

Context

The trap in this topic is the Simple Network Management Protocol (SNMP) trap. SNMP trap is a part of SNMP and a mechanism that devices being managed (here refers to network devices such as switches and routers) send SNMP messages to the NOC monitoring server. If an exception occurs on the side being monitored, namely the switch monitoring metrics have an exception, the SNMP agent running in a switch sends an alert event to the NOC monitoring server.

Procedure

1. In the left-side navigation pane, choose **NOC > Resource Management > Alert Configuration**.
2. On the **Alert Settings** page, click **Configure Trap**.
3. In the **Configure Trap** dialog box, configure the parameters.

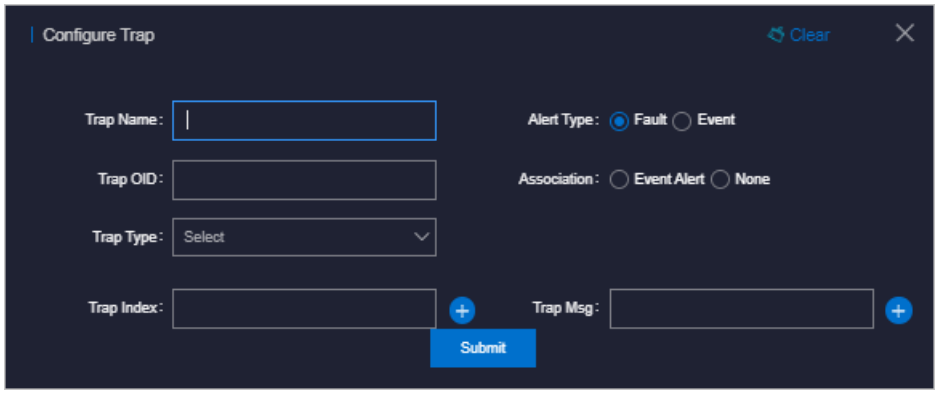

The following table describes the parameters.

Parameter	Description	Example
Trap Name	The name of the alert event.	linkdown or BGPneighbor down. You can customize the value.

Parameter	Description	Example
Trap OID	The OID of the alert event.	.1.3.6.1.4.1.25506.8.35.12.1.12 Configure the value based on the device document. You cannot customize the value.
Trap Type	The type of the alert event. Select a value from the drop-down list.	None
Trap Index	The index ID of the alert item. This value is the KV information in the trap message, which is used to identify the alert object. Typically, this value can be an API name, protocol ID, or index ID. Configure the value based on the device document. You cannot customize the value.	None
Trap Msg	The message of the alert item. This value is the KV information in the trap message, which is used to identify the alert data. Typically, this value can be the additional information of the alert item, such as a system message or a message indicating the location of the state machine or the current status. Configure the value based on the device document. You cannot customize the value.	None
Alert Type	Specifies whether the alert is of the fault type or the event type.	None
Association	Specifies whether the alert has an event alert. If **Alert Type** is set to **Fault** and the alert has an associated alert, set Association to Event Alert and then add the trap of the associated alert.	None

4. Click **Submit**. After the configuration is submitted, the system checks whether the Trap OID and Trap Name are the same as the existing ones. If not, the configuration of the trap is complete.

After the trap is added, the alert events of the configured Trap OID will be monitored and displayed on the **Current Alerts** and **History Alerts** tabs of **Alert Management**.

5.1.3.4. View traps

You can view traps configured in the current system.

Procedure

1. In the left-side navigation pane, choose **NOC > Alert Management > Alert Configuration**.
2. Enter a keyword in the search box in the upper-right corner, and then click **Search**.

> **Note** After the search results are displayed, you can click **Export to CSV** in the upper-right corner to export the trap information to your local computer.

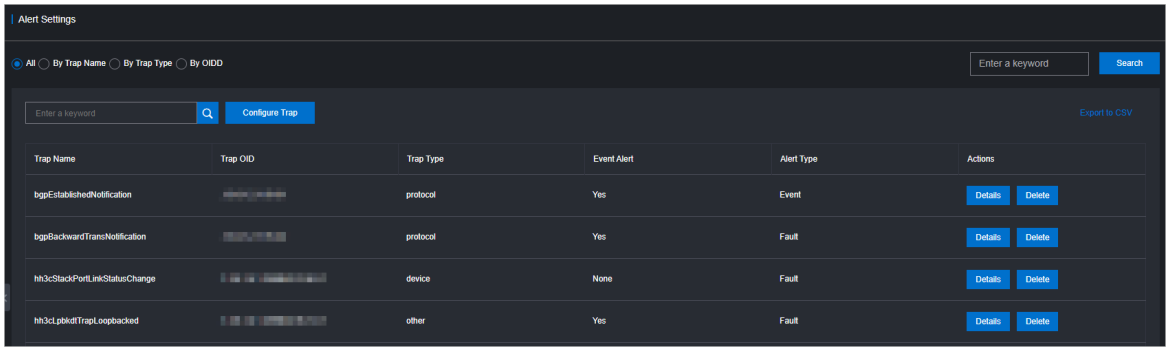

3. (Optional)You can filter the search results by trap name, trap type, or OID.
4. Move the pointer over **Details** in the **Actions** column corresponding to a trap to view detailed information about the trap.

> **Note** If a trap is no longer needed, you can click **Delete** in the **Actions** column corresponding to the trap.

5.2. Task Management

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

5.2.1. Overview

The Task Management module has the following functions:

- Supports viewing task overview and creating tasks quickly.
- Supports the following four methods to run tasks: manual execution, scheduled execution, regular execution, and advanced mode.
- Supports the breakpoint function, which allows a task to stop between its two scripts and wait for manual intervention.
- Supports searching for tasks by name, status, and created time.
- Supports uploading the .tar package as the script.

5.2.2. View the task overview

The Task Overview page shows the overall running conditions of tasks in the system. You can also create a task on this page.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Overview**.

   The **Task Overview** page appears.

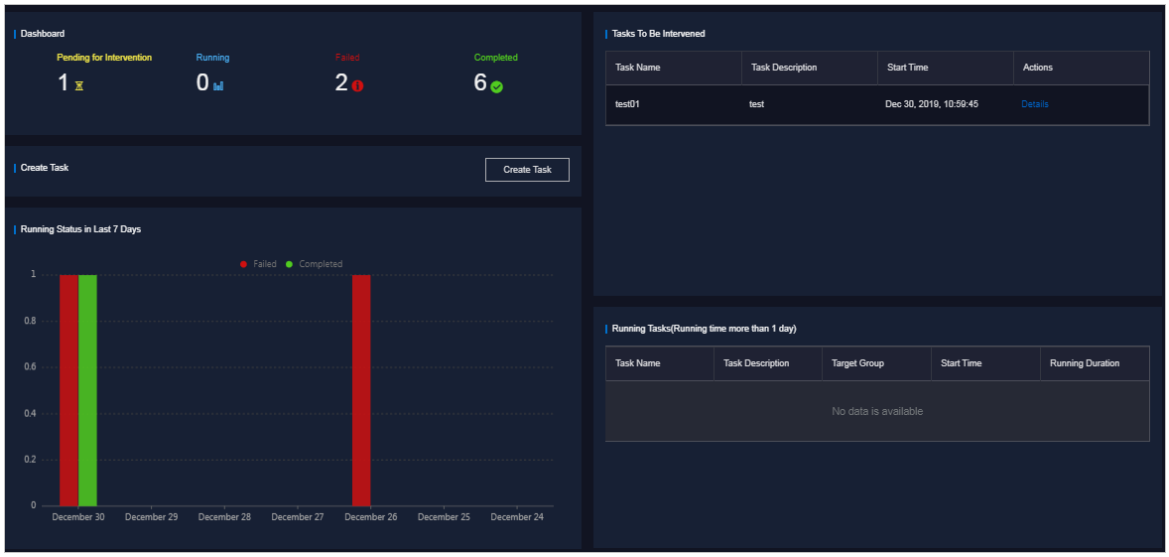

2. You can perform the following operations:
   - In the **Dashboard** section, view the number of tasks that are in the **Pending for Intervention**, **Running**, **Failed**, or **Completed** state in the system.

     Click a state or number to view the task list of the corresponding state.
   - In the **Create Task** section, click **Create Task** to create an operations task.

     For more information about how to create a task, see Create a task.
   - If a task has a breakpoint and reaches the breakpoint, the task stops and waits for manual confirmation. You can view and process tasks that require manual intervention in the **Tasks To Be Intervened** section.
   - In the **Running Status in Last 7 Days** section, view the running trend of tasks and whether tasks are successful within the last seven days.
   - In the **Running Tasks** section, view tasks running within the last 24 hours.

5.2.3. Create a task

You can make daily changes as tasks to run in the ASO console.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. Click **Create**.
3. In the Create Task dialog box, configure the parameters.

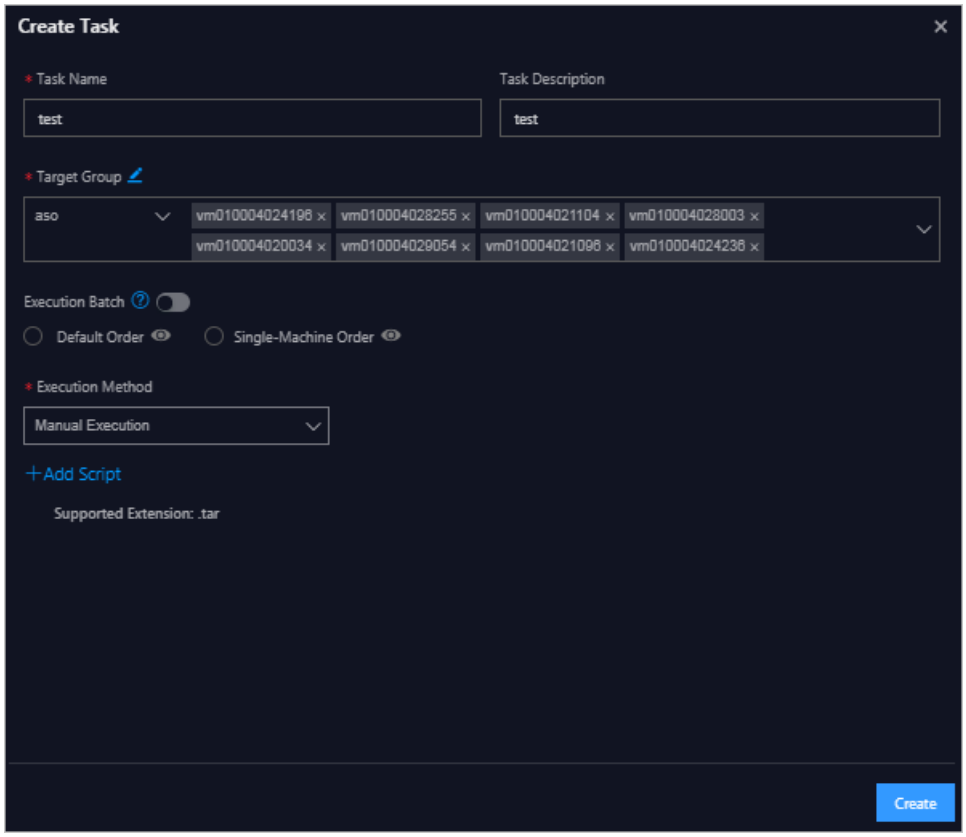

Parameter	Description
Task Name	The name of the operations task.
Task Description	The description of the operations task.
Target Group	The task target. You can use one of the following methods to configure the target group: ○ Select from the drop-down list by selecting a **product**, **cluster**, **service**, **server role**, and **virtual machine (VM) or physical machine** in sequence. ○ Select a product. Enter the VM or physical machine in the field and then press the Enter key. You can enter multiple VMs or physical machines in sequence. ○ Click the icon next to **Target Group**. In the dialog box that appears, enter the target group, with one VM or physical machine in one line. Click **OK**.

Parameter	Description
Execution Batch	Optional. This option appears after you specify the target group. If **Execution Batch** is not specified, **Target Group** is displayed in the **Target Group** column in the task list, which can be viewed by choosing **Task Management > Task Management**. If you specify **Execution Batch**, **Batch Execution Policy** is displayed in the **Target Group** column. You can set **Execution Batch** to one of the following 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 other machines in batch 5. You can adjust the batch for machines. By default, if the number of machines is greater than 10, the machines are allocated to different batches, with one machine in batch 1, three machines in batch 2, five machines in batch 3, N/3-1 (an integer) machines in batch 4, N/3-1 (an integer) machines in batch 5, until all of the machines are allocated. N is the total number of servers in the cluster. You can adjust the batch for machines. ○ **Single-Machine Order**: By default, each batch has one machine. You can adjust the batch for machines.
Execution Method	If **Execution Batch** is specified, **Execution Method** can only be set to **Manual Execution**. If **Execution Batch** is disabled, you can select one of the following execution methods: ○ **Manual Execution**: You must manually start the task. With **Manual Execution** specified, you must click **Start** in the **Actions** column to run the task after the task is created. ○ **Scheduled Execution**: Select the execution time. The task automatically runs when the execution time is reached. ○ **Regular Execution**: Select the time interval and times to run the task. The task runs again if the execution condition is met. ○ **Advanced**: Configure the command to run the task periodically.

Parameter	Description
Add Script	Click **Add Script**. Select one or more .tar packages to upload the script file. After the upload, you can delete and re-upload the script. After you upload the script, if **Execution Method** is set to **Manual Execution**, you must specify whether to enable **Intervention Required**. If manual intervention is enabled, the task will stop and wait for manual intervention after you run the script.

4. Click **Create**.

Result

The created task is displayed in the task list.

5.2.4. View the execution status of a task

After a task starts, you can view the execution status of the task.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Query** to search for tasks.
3. Find the task that you want to view, and then click **Target Group** or **Batch Execution Policy** in the **Target Group** column.

> **Note** If **Execution Batch** is not selected when you create a task, **Target Group** is displayed in the **Target Group** column. If you select **Execution Batch** when you create a task, **Batch Execution Policy** is displayed in the **Target Group** column.

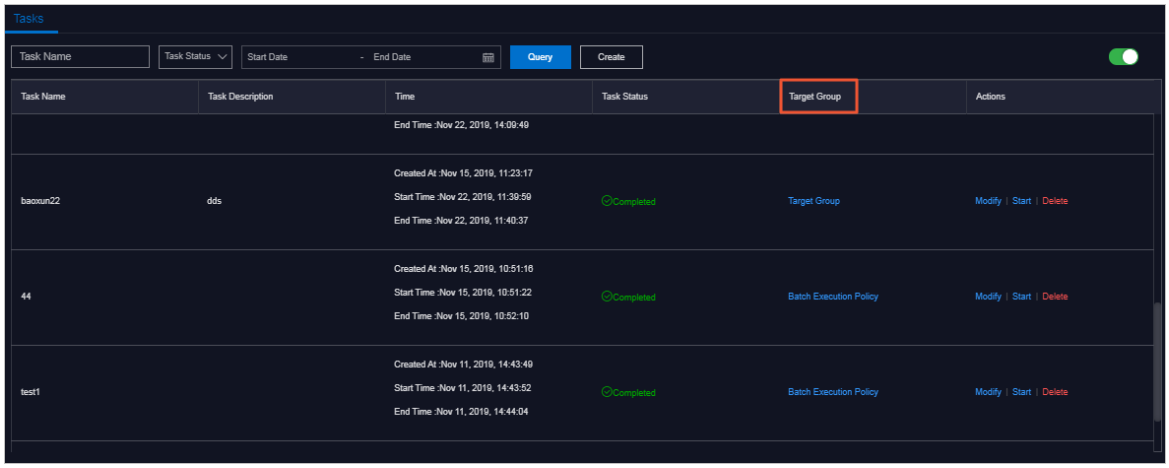

4. In the dialog box that appears, view the task execution status based on the machine color. Click a machine to view the execution results of the task.

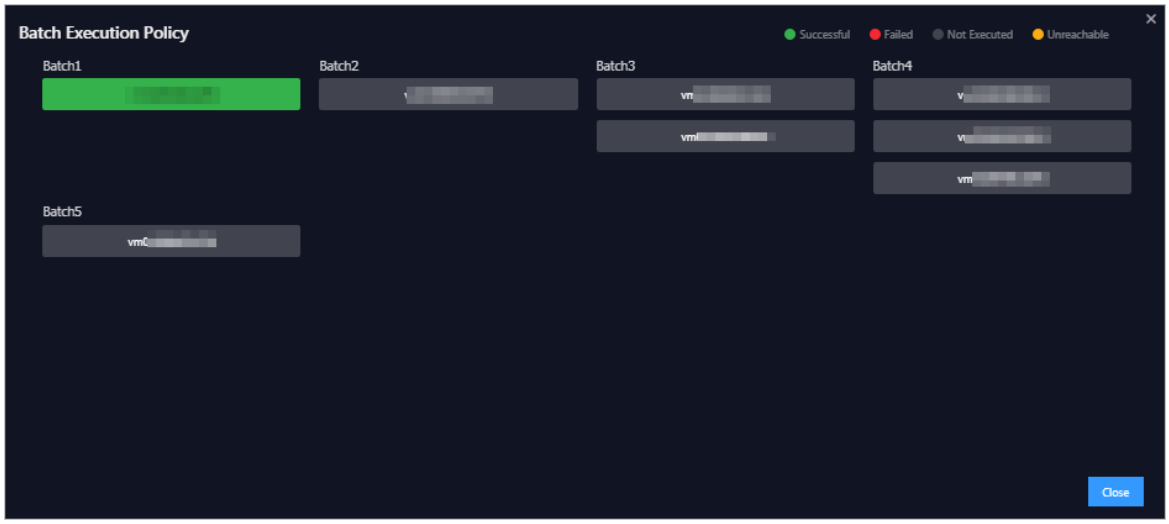

5.2.5. Start a task

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

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Query** to search for tasks.
3. Find the task that you are about to start, and then click **Start** in the **Actions** column.
4. In the dialog box that appears, select the batches to start, and then click **Start**.

   For a new task, after you click **Start** for the first time, the system will indicate that the task is started. The virtual machines (VMs) or physical machines in batch 1 start to run the task. Click **Start** again and you can select VMs or physical machines in one or more batches to run the task.

   If the task has enabled Intervention Required, you must intervene the script after you click **Start**. The **Task Status** turns to **Pending for Intervention**, and you can continue to run the task only by clicking **Continue** in the **Actions** column.

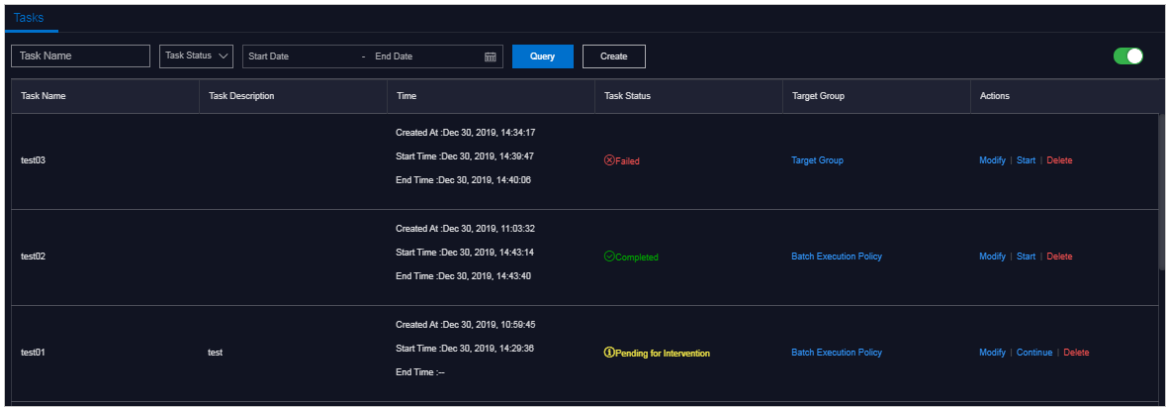

5.2.6. Delete a task

You can delete tasks that are no longer needed.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Management**.
2. (Optional)Enter the task name, select the task status, start date, and end date, and then click **Query** to search for the task.
3. Find the task to be deleted, and then click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

5.2.7. Process tasks to be intervened

If a task reaches a breakpoint, the task will stop and wait for manual confirmation. The task will continue only after receiving manual confirmation.

Procedure

1. In the left-side navigation pane, choose **Task Management > Task Overview**.
2. In the **Tasks To Be Intervened** section, find the task to be intervened, and then click **Details** in the **Actions** column.

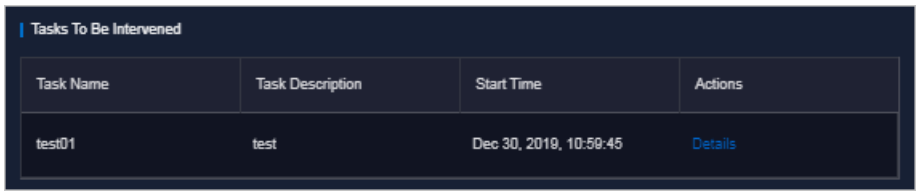

3. On the **Task Details** tab, check the information and then click **Continue** to continue to run the task.

5.2.8. Configure the XDB backup task

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

Procedure

1. In the left-side navigation pane, choose **Task Management > Commonn Tasks > XDB Backup**.
2. On the **XDB Backup** page, configure the XDB backup task information.

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 from the drop-down list by selecting a **product**, **cluster**, **service**, **server role**, and **virtual machine (VM) or physical machine**. ○ 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 other machines in batch 5. You can adjust the batch for machines as needed. By default, if the number of machines is greater than 10, the machines are allocated to different batches, with one machine in batch 1, three machines in batch 2, five machines in batch 3, N/3-1 (an integer) machines in batch 4, N/3-1 (an integer) machines in batch 5, until all of the machines are allocated. N is the total number of servers in the cluster. You can adjust the batch for machines as needed. ○ **Single-Machine Order**: By default, each batch has one machine. You can adjust the batch for machines as needed. If **Execution Batch** is not specified, **Execution Batch** will be disabled by default. **Target Group** is displayed in the **Target Group** column in the task list, which can be viewed by choosing **Task Management > Task Management**. If **Execution Batch** is specified and saved, **Execution Batch** will be enabled automatically, and **Batch Execution Policy** is displayed in the **Target Group** column.

Parameter	Description
Execution Method	If **Execution Batch** is specified, **Execution Method** can only be set 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. With **Manual Execution** 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 runs 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 is run again. ○ **Advanced**: 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.

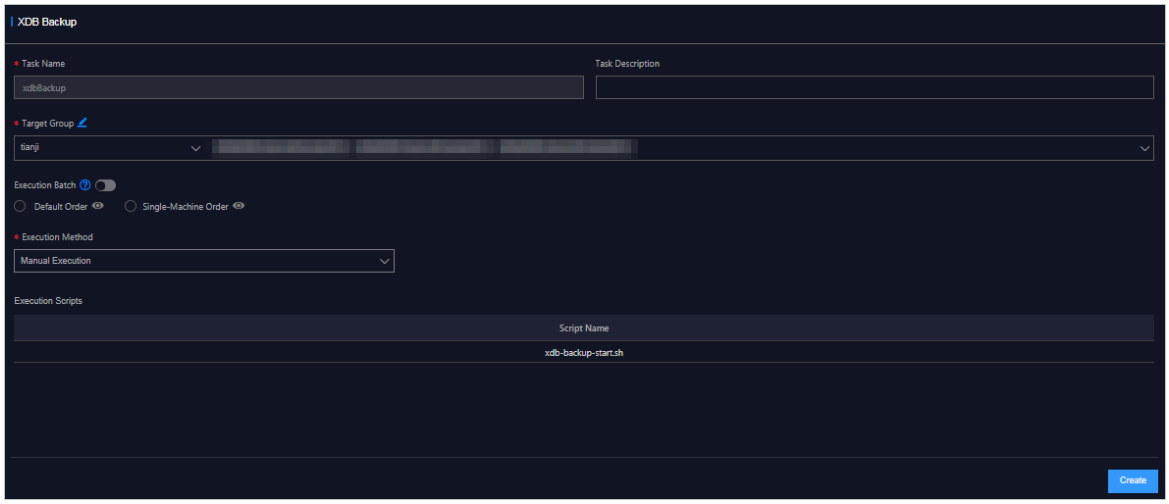

3. Click **Create**.

   You can view the created XDB task in the task list by choosing **Task Management > Task Management**. The system automatically runs the XDB backup task when the task execution condition is met. If **Execution Method** of the XDB backup task is specified as **Manual Execution**, 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, to modify the information of the backup task, you can click **Modify** in the lower part of the **XDB Backup** page.

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

5.3. Apsara Infrastructure Management Framework

5.3.1. Old version

5.3.1.1. What is Apsara Infrastructure Management Framework?

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

5.3.1.1.1. Overview

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

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

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

Core functions

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

5.3.1.1.2. Basic concepts

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

project

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

cluster

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

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

service

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

service instance

A service that is deployed on a cluster.

server role

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

server role instance

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

application

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

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

rolling

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

service configuration template

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

associated service template

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

final status

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

dependency

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

upgrade

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

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

5.3.1.2. Log on to Apsara Infrastructure Management Framework

This topic describes how to log on to Apsara Infrastructure Management Framework.

Prerequisites

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

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

Procedure

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

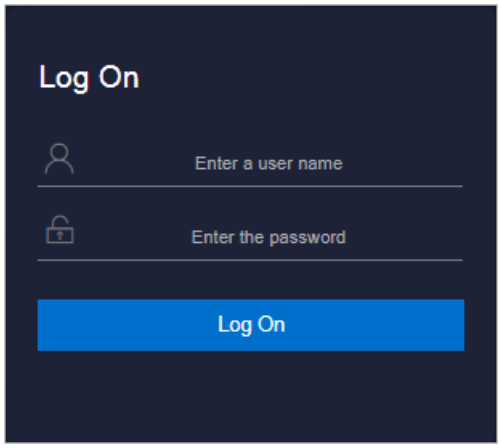

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, select **Products**.
6. In the product list, select **Apsara Infrastructure Management Framework**.

5.3.1.3. Web page introduction

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

5.3.1.3.1. Introduction on the home page

After you log on to Apsara Infrastructure Management Framework, the home page appears. This topic allows you to get a general understanding of the basic operations and functions of Apsara Infrastructure Management Framework.

Log on to Apsara Infrastructure Management Framework. The home page appears, as shown in Home page of Apsara Infrastructure Management Framework.

Home page of Apsara Infrastructure Management Framework

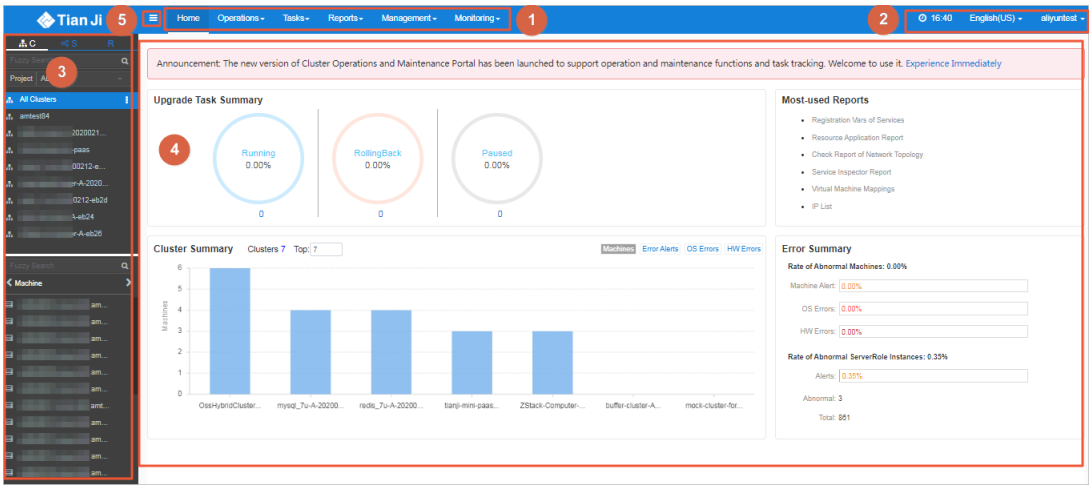

For more information about the descriptions of functional areas on the home page, see Descriptions of functional areas.

Descriptions of functional areas

Area		Description

1	Top navigation bar	• **Operations**: the quick entrance of Operations & Maintenance (O&M) operations, which allows operations engineers to quickly find the corresponding operations and operation objects. This menu consists of the following sections:
		  ○ **Cluster Operations**: performs O&M operations on and manages clusters with the project permissions, such as viewing the cluster status.
		  ○ **Service Operations**: manages services with the service permissions, such as viewing the service list information.
		  ○ **Machine Operations**: maintains and manages all the machines in Apsara Infrastructure Management Framework, such as viewing the machine status.
		• **Tasks**: A rolling task is generated after you modify the configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects.
		• **Reports**: displays the monitoring data in tables and provides the function of searching for different reports.
		• **Monitoring**: effectively monitors metrics in the process of system operation and sends alert notifications for abnormal conditions. This menu includes the functions of displaying alert status, modifying alert rules, and searching for the alert history.

Area		Description

2	Function buttons in the upper-right corner	•  :
		  ◦ **TJDB Synchronization Time**: the generated time of the data that is displayed on the current page.
		  ◦ **Final Status Computing Time**: the computing time of the final-status data that is displayed on the current page.
		  After data is generated, the system processes the data at maximum speed. As an asynchronous system, Apsara Infrastructure Management Framework has some latency. The time helps explain why the current data results are generated and determine whether the current system has a problem.
		• 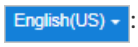 : In the English environment, click this drop-down list to switch to another language.
		•  : The logon account information. Click this drop-down list and select **Logout** to log out of Apsara Infrastructure Management Framework.

3	Left-side navigation pane	In the left-side navigation pane, you can directly view the logical structure of the Apsara Infrastructure Management Framework model. You can view the corresponding detailed data analysis and operations by selecting different levels of nodes in the left-side navigation pane. For more information, see Introduction on the left-side navigation pane.
4	Home page	Displays the summary of related tasks or information as follows: • **Upgrade Task Summary**: the numbers and proportions of running, rolling back, and paused upgrade tasks. • **Cluster Summary**: the numbers of machines, error alerts, operating system errors, and hardware errors for different clusters. • **Error Summary**: the metrics for the rate of abnormal machines and the rate of abnormal server role instances. • **Most-used Reports**: links of the most commonly used statistics reports, which facilitates you to view the report information.
5	Button used to collapse/expand the left-side navigation pane	If you are not required to use the left-side navigation pane when performing O&M operations, click this button to collapse the left-side navigation pane and increase the space of the content area.

5.3.1.3.2. Introduction on the left-side navigation pane

The left-side navigation pane has three common tabs: **C** (cluster), **S** (service), and **R** (report). With some operations, you can view the related information quickly.

Cluster

Fuzzy search is supported to search for the clusters in a project, and you can view the cluster status, cluster operations information, service final status, and logs.

In the left-side navigation pane, click the **C** tab. Then, you can:

- Enter the cluster name in the search box to search for the cluster quickly. Fuzzy search is supported.
- Select a project from the **Project** drop-down list to display all the clusters in the project.
- Move the pointer over  at the right of a cluster and then perform operations on the cluster as instructed.
- Click a cluster and all the machines and services in this cluster are displayed in the lower-left corner. Move the pointer over  at the right of a machine or service and then perform operations on the machine or service as instructed.
- Click the **Machine** tab in the lower-left corner. Double-click a machine to view all the server roles in the machine. Double-click a server role to view the applications and then double-click an application to view the log files.
- Click the **Service** tab in the lower-left corner. Double-click a service to view all the server roles in the service. Double-click a server role to view the machines, double-click a machine to view the applications, and double-click an application to view the log files.
- Double-click a log file. Move the pointer over  at the right of the log file and then select **Download** to download the log file.

  Move the pointer over a log file and then click **View** at the right of the log file to view the log details based on time. On the **Log Viewer** page, enter the keyword to search for logs.

Service

Fuzzy search is supported to search for services and you can view services and service instances.

In the left-side navigation pane, click the **S** tab. Then, you can:

- Enter the service name in the search box to search for the service quickly. Fuzzy search is supported.
- Move the pointer over  at the right of a service and then perform operations on the service as instructed.
- Click a service and all the service instances in this service are displayed in the lower-left corner. Move the pointer over  at the right of a service instance and then perform operations on the service instance as instructed.

Report

Fuzzy search is supported to search for reports and you can view the report details.

In the left-side navigation pane, click the **R** tab. Then, you can:

- Enter the report name in the search box to search for the report quickly. Fuzzy search is supported.
- Click **All Reports** or **Favorites** to display groups of different categories in the lower-left corner.

Double-click a group to view all the reports in this group. Double-click a report to view the report details on the right pane.

5.3.1.4. Cluster operations

This topic describes the actions about cluster operations.

5.3.1.4.1. View cluster configurations

By viewing the cluster configurations, you can view the basic information, deployment plan, and configurations 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 displays the following information:

   - **Cluster**

     The cluster name. Click the cluster name to go to the Cluster Dashboard page.
   - **Scale-Out/Scale-In**

     The number of machines or server roles that are scaled out or in. Click the link to go to the Cluster Operation and Maintenance Center page.
   - **Abnormal Machine Count**

     The statistics of machines whose status is not Good in the cluster. Click the link to go to the Cluster Operation and Maintenance Center page.
   - **Final Status of Normal Machines**

     Displays whether the cluster reaches the final status. Select **Clusters Not Final** to display clusters that do not reach the final status. Click the link to go to the Service Final Status Query page.
   - **Rolling**

     Displays whether the cluster has a running rolling task. Select **Rolling Tasks** to display clusters that have rolling tasks. Click the link to go to the Rolling Task page.
3. (Optional)Select a project from the **Project** drop-down list and/or enter the cluster name in the **Cluster** field to search for clusters.
4. Find the cluster whose configurations you are about to view and then click **Cluster Configuration** in the **Actions** column. The **Cluster Configuration** page appears.

   For more information about the **Cluster Configuration** page, see Cluster configurations.

Cluster configurations

Category	Item	Description
	Cluster	The cluster name.
	Project	The project to which the cluster belongs.

Category	Item	Description
Basic Information	Clone Switch	○ **Mock Clone**: The system is not cloned when a machine is added to the cluster. ○ **Real Clone**: The system is cloned when a machine is added to the cluster.
	Machines	The number of machines in the cluster. Click **View Clustering Machines** to view the machine list.
	Security Verification	The access control among processes. Generally, the non-production environment uses the default configurations and does not perform the verification. In other cases, customize the configurations based on actual requirements to enable or disable the verification.
	Cluster Type	○ RDS ○ NETFRAME ○ T4: a special type that is required by the mixed deployment of e-commerce. ○ Default: other conditions.
Deployment Plan	Service	The service deployed in the cluster.
	Dependency Service	The service that the current service depends on.
Service Information	Service Information	Select a service from the **Service Information** drop-down list and then the configurations of this service are displayed.
	Service Template	The template used by the service.
	Monitoring Template	The monitoring template used by the service.
	Machine Mappings	The machines included in the server role of the service.
	Software Version	The software version of the server role in the service.
	Availability Configuration	The availability configuration percentage of the server role in the service.
	Deployment Plan	The deployment plan of the server role in the service.

Category	Item	Description
	Configuration Information	The configuration file used in the service.
	Role Attribute	Server roles and the corresponding parameters.

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

5.3.1.4.2. View the cluster dashboard

The cluster dashboard allows you to view the basic information and related statistics of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You have two ways to go to the **Cluster Dashboard** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over at the right of a cluster and then select **Dashboard**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click the cluster name.
3. On the **Cluster Dashboard** page, you can view the cluster information, including the basic information, final status information, rolling job information, dependencies, resource information, virtual machines, and monitoring information. For more information about the descriptions, see the following table.

Item	Description

Item	Description
Basic Cluster Information	Displays the basic information of the cluster as follows: ○ **Project Name**: the project name. ○ **Cluster Name**: the cluster name. ○ **IDC**: the data center to which the cluster belongs. ○ **Final Status Version**: the latest version of the cluster. ○ **Cluster in Final Status**: whether the cluster reaches the final status. ○ **Machines Not In Final Status**: the number of machines that do not reach the final status in the cluster when the cluster does not reach the final status. ○ **Real/Pseudo Clone**: whether to clone the system when a machine is added to the cluster. ○ **Expected Machines**: the number of expected machines in the cluster. ○ **Actual Machines**: the number of machines in the current environment. ○ **Machines Not Good**: the number of machines whose status is not Good in the cluster. ○ **Actual Services**: the number of services that are actually deployed in the cluster. ○ **Actual Server Roles**: the number of server roles that are actually deployed in the cluster. ○ **Cluster Status**: whether the cluster is starting or shutting down machines.
Machine Status Overview	The statistical chart of the machine status in the cluster.
Machines in Final Status	The numbers of machines that reach the final status and those that do not reach the final status in each service of the cluster.
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
Service Instances	Displays the service instances deployed in the cluster and the related final status information. ○ **Service Instance**: the service instance deployed in the cluster. ○ **Final Status**: whether the service instance reaches the final status. ○ **Expected Server Roles**: the number of server roles that the service instance expects to deploy. ○ **Server Roles In Final Status**: the number of server roles that reach the final status in the service instance. ○ **Server Roles Going Offline**: the number of server roles that are going offline in 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	Displays the upgrade tasks related to the cluster. ○ **Cluster Name**: the name of the upgrade cluster. ○ **Type**: the type of the upgrade task. The options include app (version upgrade) and config (configuration change). ○ **Git Version**: the change version to which the upgrade task belongs. ○ **Description**: the description about the change. ○ **Rolling Result**: the result of the upgrade task. ○ **Submitted By**: the person who submits the change. ○ **Submitted At**: the time when the change is submitted. ○ **Start Time**: the time to start the rolling. ○ **End Time**: the time to finish the upgrade. ○ **Time Used**: the time used for the upgrade. ○ Actions: Click **Details** to go to the **Rolling Task** page. For more information about the rolling task, see View rolling tasks.
Cluster Resource Request Status	○ **Version**: the resource request version. ○ **Msg**: the exception message. ○ **Begintime**: the start time of the resource request analysis. ○ **Endtime**: the end time of the resource request analysis. ○ **Build Status**: the build status of resources. ○ **Resource Process Status**: the resource request status in the version.

Item	Description
Cluster Resource	○ **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 request status. ○ **Error Msg**: the exception message. ○ **Parameters**: the resource parameters. ○ **Result**: the resource request result. ○ **Res**: the resource ID. ○ **Reprocess Status**: the status of interaction with Business Foundation System during the VIP resource request. ○ **Reprocess Msg**: the exception message of interaction with Business Foundation System during the VIP resource request. ○ **Reprocess Result**: the result of interaction with Business Foundation System during the VIP resource request. ○ **Refer Version List**: the version that uses the resource.
VM Mappings	The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster. ○ **VM**: the hostname of the virtual machine. ○ **Currently Deployed On**: the hostname of the physical machine where the virtual machine is currently deployed. ○ **Target Deployed On**: the hostname of the physical machine where the virtual machine is expected to be deployed.
Service Dependencies	The dependencies of service instances and server roles in the cluster, and the final status information of the dependent service or server role. ○ **Service**: the service name. ○ **Server Role**: the server role name. ○ **Dependent Service**: the service on which the server role depends. ○ **Dependent Server Role**: the server role on which the server role depends. ○ **Dependent Cluster**: the cluster to which the dependent server role belongs. ○ **Dependency in Final Status**: whether the dependent server role reaches the final status.

5.3.1.4.3. View the cluster operation and maintenance center

The cluster operation and maintenance center allows you to view the status or statistics of services or machines in the cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You have three ways to go to the **Cluster Operation and Maintenance Center** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over at the right of a cluster and then select **Cluster Operation and Maintenance Center**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, choose **Monitoring > Cluster Operation and Maintenance Center** in the **Actions** column at the right of a cluster.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click a cluster name. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**.
3. View the information on the **Cluster Operation and Maintenance Center** page.

Item	Description
SR not in Final Status	Displays all the server roles that do not reach the final status in the cluster. Click the number to expand a server role list, and click a server role in the list to display the information of machines included in the server role.
Running Tasks	Displays whether the cluster has running rolling tasks. Click **Rolling** to go to the **Rolling Task** page. For more information about the rolling task, see View rolling tasks.
Head Version Submitted At	The time when the head version is submitted. Click the time to view the submission details.

Item	Description
Head Version Analysis	The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses: ◦ **Preparing**: No new version is available now. ◦ **Waiting**: The latest version is found. The analysis module has not started up yet. ◦ **Doing**: The module is analyzing the application that requires change. ◦ **done**: The head version analysis is successfully completed. ◦ **Failed**: The head version analysis failed. The change contents cannot be parsed. If the status is not **done**, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version. Click the status to view the relevant information.
Service	Select a service deployed in the cluster from the drop-down list.
Server Role	Select a server role of a service in the cluster from the drop-down list. **Note** After you select the service and server role, the information of machines related to the service or server role is displayed in the list.
Total Machines	The total number of machines in the cluster, or the total number of machines included in a specific server role of a specific service.
Scale-in/Scale-out	The number of machines or server roles that are scaled in or out.
Abnormal Machines	The number of abnormal machines that encounter each type of the following faults. ◦ **Ping Failed**: A ping_monitor error is reported, and TianjiMaster cannot successfully ping the machine. ◦ **No Heartbeat**: TianjiClient on the machine does not regularly report data to indicate the status of this machine, which may be caused by the TianjiClient problem or network problem. ◦ **Status Error**: The machine has an error reported by the monitor or a fault of the critical or fatal level. Check the alert information and accordingly solve the issue.

Item	Description
Abnormal Services	The number of machines with abnormal services. To determine if a service reaches the final status, see the following rules: ○ The server role on the machine is in the GOOD status. ○ Each application of the server role on the machine must keep the actual version the same as the head version. ○ Before the Image Builder builds an application of the head version, Apsara Infrastructure Management Framework cannot determine the value of the head version and the service final status is unknown. This process is called the change preparation process. The service final status cannot be determined during the preparation process or upon a preparation failure.
Machines	Displays all the machines in the cluster or the machines included in a specific server role of a specific service. ○ Machine search: Click the search box to enter the machine in the displayed dialog box. Fuzzy or batch search is supported. ○ Click the machine name to view the physical information of the machine in the displayed **Machine Information** dialog box. Click **DashBoard** to go to the **Machine Details** page. For more information about the machine details, see View the machine dashboard. ○ Move the pointer over the blank area in the **Final Status** column or the **Final SR Status** column and then click **Details** to view the machine status, system service information, server role status on the machine, and exception message. ○ If no service or server role is selected from the drop-down list, move the pointer over the blank area in the **Running Status** column and then click **Details** to view the running status information or exception message of the machine. If a service and a server role are selected from the corresponding drop-down lists, move the pointer over the blank area in the **SR Running Status** column and then click **Details** to view the running status information or exception message of the server role on the machine. ○ Click **Error**, **Warning**, or **Good** in the **Monitoring Statistics** column to view the monitored items of machines and monitored items of server roles. ○ Click **Terminal** in the **Actions** column to log on to the machine and perform related operations. ○ Click **Machine Operation** in the **Actions** column to restart, out-of-band restart, or clone the machine again.

5.3.1.4.4. View the service final status

The **Service Final Status Query** page allows you to view if a service in a cluster reaches the final status and the final status information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You have two ways to go to the **Service Final Status Query** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over  at the right of a cluster and then choose **Monitoring > Service Final Status Query**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, choose **Monitoring > Service Final Status Query** in the **Actions** column at the right of a cluster.
3. View the information on the **Service Final Status Query** page.

Item	Description
Project Name	The name of the project to which the cluster belongs.
Cluster Name	The cluster name.
Head Version Submitted At	The time when the head version is submitted.
Head Version Analysis	The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses: ○ **Preparing**: No new version is available now. ○ **Waiting**: The latest version is found. The analysis module has not started up yet. ○ **Doing**: The module is analyzing the application that requires change. ○ **done**: The head version analysis is successfully completed. ○ **Failed**: The head version analysis failed. The change contents cannot be parsed. If the status is not **done**, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.
Cluster Rolling Status	Displays the information of the current rolling task in the cluster, if any. The rolling task may not be of the head version.
Cluster Machine Final Status Statistics	The status of all machines in the cluster. Click **View Details** to go to the **Cluster Operation and Maintenance Center** page and view the detailed information of all machines. For more information about the cluster operation and maintenance center, see View the cluster operation and maintenance center.

Item	Description
Final Status of Cluster SR Version	The final status of cluster service version. **Note** Take statistics of services that do not reach the final status, which is caused by version inconsistency or status exceptions. If services do not reach the final status because of machine problems, go to **Cluster Machine Final Status Statistics** to view the statistics.
Final Status of SR Version	The number of machines that do not reach the final status when a server role has tasks.

5.3.1.4.5. View operation logs

By viewing operation logs, you can obtain the differences between different Git versions.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. You have two ways to go to the **Cluster Operation Logs** page:
   - In the left-side navigation pane, click the **C** tab. Move the pointer over at the right of a cluster and then choose **Monitoring > Operation Logs**.
   - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, choose **Monitoring > Operation Logs** in the **Actions** column at the right of a cluster.
3. On the **Cluster Operation Logs** page, click **Refresh**. View the Git version, description, submitter, submitted time, and task status.
4. (Optional)Complete the following steps to view the differences between versions on the **Cluster Operation Logs** page.
   i. Find the log in the operation log list and then click **View Release Changes** in the **Actions** column.
   ii. On the **Version Difference** page, complete the following configurations:
      - **Select Base Version**: Select a base version.
      - **Configuration Type**: Select **Extended Configuration** or **Cluster Configuration**. **Extended Configuration** displays the configuration differences after the configuration on the cluster is combined with the configuration in the template. **Cluster Configuration** displays the configuration differences on the cluster.
   iii. Click **Obtain Difference**.

      The differential file list is displayed.
   iv. Click each differential file to view the detailed differences.

5.3.1.5. Service operations

This topic describes the actions about service operations.

5.3.1.5.1. 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.
Monitoring Templates	The number of monitoring templates.
Service Schemas	The number of service configuration validation templates.
Actions	Click **Management** to view the service instances, service templates, monitoring templates, monitoring instances, service schemas, and detection scripts.

5.3.1.5.2. View the service instance dashboard

The service instance dashboard allows you to view the basic information and 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 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. View the information on the **Service Instance Information Dashboard** page.

Item	Description
Service Instance Summary	Displays the basic information of the service instance as follows: ○ **Cluster Name**: the name of the cluster to which the service instance belongs. ○ **Service Name**: the name of the service to which the service instance belongs. ○ **Actual Machines**: the number of machines in the current environment. ○ **Expected Machines**: the number of machines that the service instance expects. ○ **Target Total Server Roles**: the number of server roles that the service instance expects. ○ **Actual Server Roles**: the number of server roles in the current environment. ○ **Template Name**: the name of the service template used by the service instance. ○ **Template Version**: the version of the service template used by the service instance. ○ **Schema**: the name of the service schema used by the service instance. ○ **Monitoring System Template**: the name of the monitoring system template used by the service instance.
Server Role Statuses	The statistical chart of the current status of server roles in the service instance.
Machine Statuses for Server Roles	The status statistics of machines where server roles are located.
Service 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 data is updated.
Service Alert Status	○ **Alert Name** ○ **Instance Information** ○ **Alert Start** ○ **Alert End** ○ **Alert Duration** ○ **Severity Level** ○ **Occurrences**: the number of times the alert is triggered.

Item	Description
Server Role List	○ Server Role ○ Current Status ○ Expected Machines ○ Machines In Final Status ○ Machines Going Offline ○ Rolling Task Status ○ Time Used: the time used for running the rolling task. ○ Actions: Click Details to go to the Server Role Dashboard page.
Service Alert History	○ Alert Name ○ Alert Time ○ Instance Information ○ Severity Level ○ Contact Group
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. ○ Server Role: the server role name. ○ Dependent Service: the service on which the server role depends. ○ Dependent Server Role: the server role on which the server role depends. ○ Dependent Cluster: the cluster to which the dependent server role belongs. ○ Dependency in Final Status: whether the dependent server role reaches the final status.

5.3.1.5.3. View the server role dashboard

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

Procedure

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

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

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

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

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

5.3.1.6. Machine operations

This topic describes the actions about machine operations.

5.3.1.6.1. View the machine dashboard

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

Procedure

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

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

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

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

5.3.1.7. Monitoring center

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

5.3.1.7.1. Modify an alert rule

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

Procedure

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

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

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

5.3.1.7.2. View the status of a monitoring instance

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

Procedure

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

5.3.1.7.3. View the alert status

The **Alert Status** page allows you to view the alerts generated in different services and the corresponding alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert Status**.
3. (Optional)You can configure the service name, cluster name, alert name, and/or the time range when the alert is triggered to search for alerts.
4. View the alert details on the **Alert Status** page. See the following table for the alert status descriptions.

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Instance	The name of the service instance being monitored. Click the instance to view the alert history of this instance.

Item	Description
Alert Status	Alerts have two statuses: **Restored** and **Alerting**.
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. ○ P1 ○ P2 ○ P3 ○ P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered and how long the alert has lasted.
Actions	Click **Show** to show the data before and after the alert time.

5.3.1.7.4. View alert rules

The **Alert Rules** page allows you to view the configured alert rules.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert Rules**.
3. (Optional)You can configure the service name, cluster name, and/or alert name to search for alert rules.
4. View the detailed alert rules on the **Alert Rules** page. See the following table for the alert rule descriptions.

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Alert Name	The name of the generated alert.
Alert Conditions	The conditions met when the alert is triggered.
Periods	The frequency (in seconds) with which an alert rule is run.
Alert Contact	The groups and members that are notified when an alert is triggered.

Item	Description
Status	The current status of the alert rule. ○ **Running**: Click to stop this alert rule. ○ **Stopped**: Click to run this alert rule.

5.3.1.7.5. View the alert history

The **Alert History** page allows you to view all the history alerts generated in different services and the corresponding alert details.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Monitoring > Alert History**.
3. (Optional)You can configure the service name, cluster name, time range, and/or period to search for alerts.
4. View the history alerts on the **Alert History** page. See the following table for the history alert descriptions.

Item	Description
Service	The name of the service to which the alert belongs.
Cluster	The name of the cluster where the service is located.
Alert Instance	The name of the resource where the alert is triggered.
Status	Alerts have two statuses: **Restored** and **Alerting**.
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. ○ P1 ○ P2 ○ P3 ○ P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members that are notified when an alert is triggered.
Actions	Click **Show** to show the data before and after the alert time.

5.3.1.8. Tasks and deployment summary

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

5.3.1.8.1. View rolling tasks

You can view running rolling tasks and the corresponding 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 display clusters with rolling tasks.
4. In the search results, click **rolling** in the **Rolling** column.
5. On the displayed **Rolling Task** page, view the information in the **Change Task** list and **Change Details** list.

Change Task list

Item	Description
Change Version	The version that triggers the change of the rolling task.
Description	The description about the change.
Head Version Analysis	The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses: ○ **Preparing**: No new version is available now. ○ **Waiting**: The latest version is found. The analysis module has not started up yet. ○ **Doing**: The module is analyzing the application that requires change. ○ **done**: The head version analysis is successfully completed. ○ **Failed**: The head version analysis failed. The change contents cannot be parsed. If the status is not **done**, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.
Blocked Server Role	Server roles blocked in the rolling task. Generally, server roles are blocked because of dependencies.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.

Item	Description
Actions	Click **View Difference** to go to the **Version Difference** page. For more information, see View operation logs. Click **Stop** to stop the rolling task. Click **Pause** to pause the rolling task.

Change Details list

Item	Description
Service Name	The name of the service where a change occurs.
Status	The current status of the service. The rolling status of the service is an aggregated result, which is calculated based on the rolling status of the server role. ○ **succeeded**: The task is successfully run. ○ **blocked**: The task is blocked. ○ **failed**: The task failed.
Server Role Status	The server role status. Click **>** at the left of the service name to expand and display the rolling task status of each server role in the service. Server roles have the following statuses: ○ **Downloading**: The task is being downloaded. ○ **Rolling**: The rolling task is running. ○ **RollingBack**: The rolling task failed and is rolling back.
Depend On	The services that this service depends on or server roles that this server role depends on.
Actions	Click **Stop** to stop the change of the server role. Click **Pause** to pause the change of the server role.

5.3.1.8.2. View running tasks

By viewing running tasks, you can know the information of all the running tasks.

Procedure

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

3. (Optional)You can configure the cluster name, role name, task status, task submitter, Git version, and/or the start time and end time of the task to search for running tasks.
4. Find the task that you are about to view the details and then click **View Tasks** in the **Rolling Task Status** column. The **Rolling Task** page appears. For more information about the rolling task, see View rolling tasks.

5.3.1.8.3. View history tasks

You can view the historical running conditions of completed tasks.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the top navigation bar, choose **Tasks > History Tasks**.
3. (Optional)You can configure the cluster name, Git version, task submitter, and/or the start time and end time of the task to search for history tasks.
4. Find the task that you are about to view the details and then click **Details** in the **Actions** column. The **Rolling Task** page appears. For more information about the rolling task, see View rolling tasks.

5.3.1.8.4. 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
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.

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

5.3.1.9. Reports

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

5.3.1.9.1. 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

Item	Description
Report	The report name. Move the pointer over next to **Report** to search for reports by report name.
Group	The group to which the report belongs. Move the pointer over next to **Group** to filter reports by group name.
Status	Indicates whether the report is published.
Public	Indicates whether the report is public.
Created By	The person who creates the report.
Published At	The published time and created time of the report.
Actions	Click **Add to Favorites** to add this report to your favorites. Then, you can view the report by choosing **Reports > Favorites** in the top navigation bar or moving the pointer over at the right of **Favorites** on the **R** tab in the left-side navigation pane and then selecting **View**.

3. (Optional)Enter the name of the report that you are about to view in the search box.
4. Click the report name to go to the corresponding report details page. For more information about the reports, see Appendix.

5.3.1.9.2. Add a report to favorites

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

Procedure

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

5.3.1.10. Metadata operations

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

5.3.1.10.1. 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
  }
}
```

5.3.1.10.2. Access APIs

This topic describes how to connect to control-side and deployment-side API operations.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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 target machine.
4. Connect to API operations.
   - Connect to control-side API operations
     a. Find the target 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
```

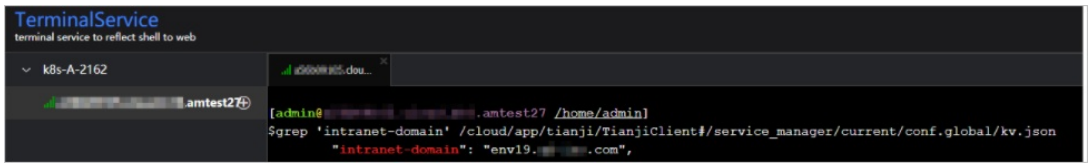

     c. Use one of the following methods to connect to control-side API operations. ListInstance is used in the example.
        - GET request

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

        - POST request

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

   - Connect to deployment-side API operations
     a. Find the target machine and record the IP address in the Hostname column.
     b. Use one of the following methods to connect to deployment-side API operations. CheckState is used in the example.

        Assume that the IP address of the target 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}'
```

5.3.1.10.3. APIs on the control side

5.3.1.10.4. APIs on the deployment side

5.3.1.11. Appendix

5.3.1.11.1. 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.

5.3.1.11.2. 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.

5.3.1.11.3. 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.

5.3.1.11.4. Rolling info report

This report displays the running and completed rolling tasks and the task-related status.

Choose a rolling action

This list only displays the running rolling tasks. If no rolling task is running, no data is available in the list.

Item	Description
Cluster	The cluster name.
Git Version	The version of change that triggers the rolling task.
Description	The description about the change entered by a user when the user submits the change.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Submitted By	The ID of the user who submits the change.
Rolling Task Status	The current status of the rolling task.
Submitted At	The time when the change is submitted.

Server Role in Job

Select a rolling task in the **Choose a rolling action** section to display the rolling status of server roles related to the selected task. If no rolling task is selected, the server role statuses of all historical rolling tasks are displayed.

Item	Description
Server Role	The server role name.
Server Role Status	The rolling status of the server role.
Error Message	The exception message of the rolling task.

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 that have the rolling task approved by the decider.
Failure Rate	The proportion of machines that have the rolling task failed.
Success Rate	The proportion of machines that have the rolling task succeeded.

Server Role Rolling Build Information

The source version and target version of each application under the server role in the rolling process.

Item	Description
App	The name of the application that requires rolling in the server role.
Server Role	The server role to which the application belongs.
From Build	The version before the upgrade.
To Build	The version after the upgrade.

Server Role Statuses on Machines

Select a server role in the **Server Role in Job** section to display the deployment status of this server role on the machine.

Item	Description
Machine Name	The name of the machine on which the server role is deployed.
Expected Version	The target version of the rolling.
Actual Version	The current version.
State	The status of the server role.
Action Name	The Apsara Infrastructure Management Framework action currently performed by the server role.
Action Status	The action status.

5.3.1.11.5. 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.

5.3.1.11.6. 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.

5.3.1.11.7. 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.

5.3.1.11.8. 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.

5.3.1.11.9. 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.

5.3.1.11.10. Statuses of project components

This report displays the status of all server roles in an abnormal status on machines of the project, and the monitoring information (alert information reported by the server role to Apsara Infrastructure Management Framework monitor) of server roles and machines.

Error State Component Table

Only displays the information of server roles that are not in GOOD status and server roles to be upgraded.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.
Need Upgrade	Whether the current version reaches the final status.
Server Role Status	The current status of the server role.
Machine Status	The current status of the machine.

Server Role Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.

Item	Description
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Machine Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Service Inspector Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

5.3.1.11.11. 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.
Server Role	The server role name.
Dependent Service	The service on which the server role depends.
Dependent Server Role	The server role on which the server role depends.
Dependent Cluster	The cluster to which the dependent server role belongs.
Dependency in Final Status	Whether the dependent server role reaches the final status.

5.3.1.11.12. 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.

5.3.1.11.13. 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.

5.3.1.11.14. 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.

5.3.1.11.15. 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.

Item	Description

Machine Statuses

Displays the running statuses of machines in the selected cluster.

Select a row in the **Statuses on Machines** section to display the information of the corresponding machine in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status of the machine.
Error Message	The exception message.

5.3.2. New version

5.3.2.1. What is Apsara Infrastructure Management Framework?

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

5.3.2.1.1. Introduction

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

Overview

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

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

Core functions

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

5.3.2.1.2. Basic concepts

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

project

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

cluster

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

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

service

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

service instance

A service that is deployed on a cluster.

server role

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

server role instance

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

application

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

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

rolling

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

service configuration template

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

associated service template

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

desired state

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

dependency

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

upgrade

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

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

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

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

Procedure

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

   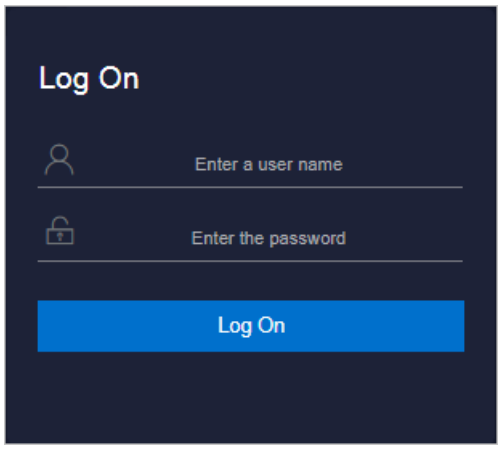

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

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

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

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

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**.
6. In the Apsara Stack O&M section, choose **Basic O&M > Apsara Infrastructure Management Framework**.

5.3.2.3. Homepage introduction

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 the following figure.

Homepage of the Apsara Infrastructure Management Framework console

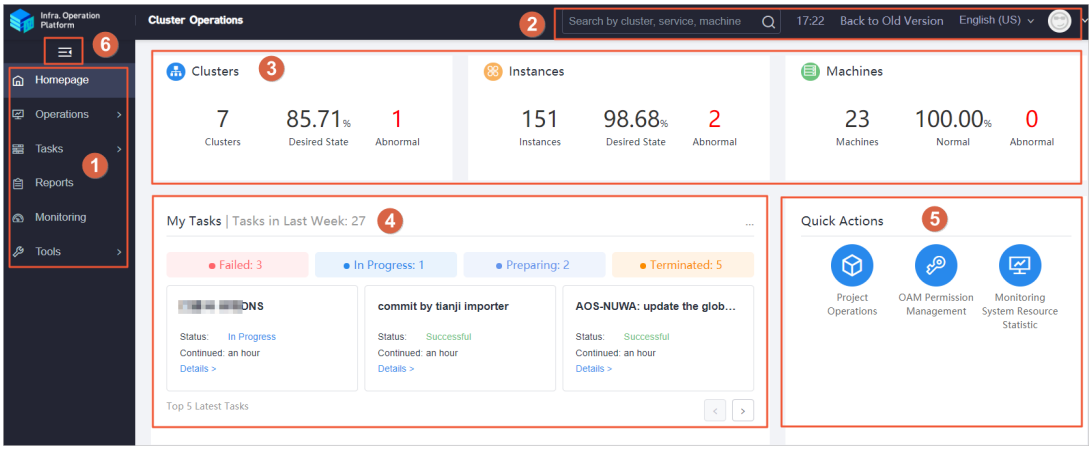

The following table describes the functional sections on the homepage.

Description of functional sections

Section	Description

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 use the project permissions to manage projects.
		  ◦ **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.
		• **Tools**: provides tools such as machine O&M and IDC shutdown.

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 out of your account.

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. • **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 machines in the **Normal** state, and the number of abnormal machines. You can move the pointer over each section and then click Details to go to the Cluster Operations page, Service Operations page, or Machine Operations page.
④	Task status bar	Displays the information of tasks submitted in 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	Displays links of the following common quick actions: • **Project Operations**: allows you to go to the Project Operation page. • **OAM Permission Management**: allows you to go to the Operation Administrator Manager (OAM) console. OAM is a centralized permission management platform in the ASO console. • **Monitoring System Resource Statistic**: 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 Statistic** is displayed only when Monitoring System is deployed in the environment.
⑥	Show/hide button	If you do not need to use the left-side navigation pane, click this button to hide the pane and enlarge the workspace.

5.3.2.4. Project operations

The Project Operations module allows you to search for and view details of a project.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Project Operations**.

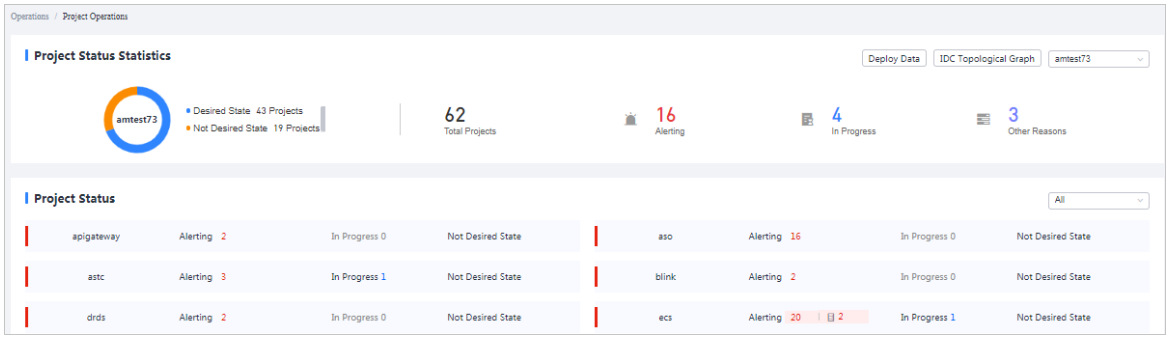

3. On this page, you can:
   - Search for a project

     Click the drop-down list in the upper-right corner of the **Project Status** section. Enter a project name in the search box, and then select the name to search for the project. You can view the numbers of alerts and running tasks for the project and whether the project reaches the desired state.
   - View the details of a project
     - Find the project whose details you are about to view. Click the number at the right of **Alerting**. In the displayed Alert Information dialog box, view the specific monitoring metrics, monitoring types, and alert sources. Click the value in the Alert Source column to view the service details.
     - Find the project whose details you are about to view. Click the number at the right of **In Progress**. In the displayed Tasks dialog box, view the details of Upgrade Service and Machine Change.

5.3.2.5. Cluster operations

This topic describes the actions about cluster operations.

5.3.2.5.1. View the cluster list

This topic describes how to view all clusters and their information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the cluster list:
   - On the **Homepage** page, move the pointer over the **Clusters** section and click Details in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Cluster Operations**.

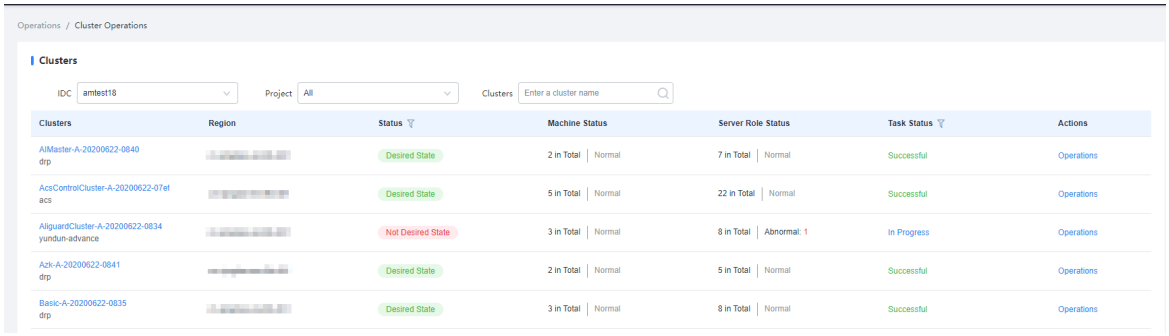

The following table describes the information displayed in the cluster list.

Parameter	Description
Clusters	The name of the cluster. Click the cluster name to view the cluster details.
Region	The region where the cluster resides.
Status	Specifies whether the cluster reaches the desired state. Click the icon to filter states. ○ 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.
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. 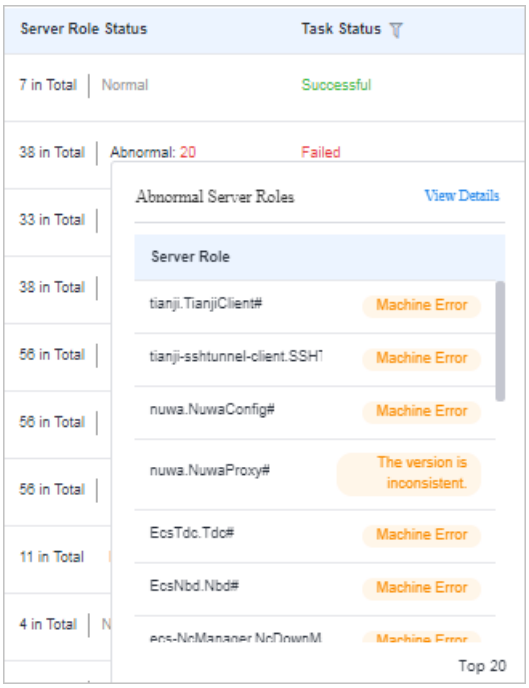
Task Status	The status of the task related to the cluster. Click the icon to filter clusters. Click the task status to view the task details.

5.3.2.5.2. View details of a cluster

This topic describes how to view details of a cluster.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Cluster Operations**.
3. (Optional)Select a project from the 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.

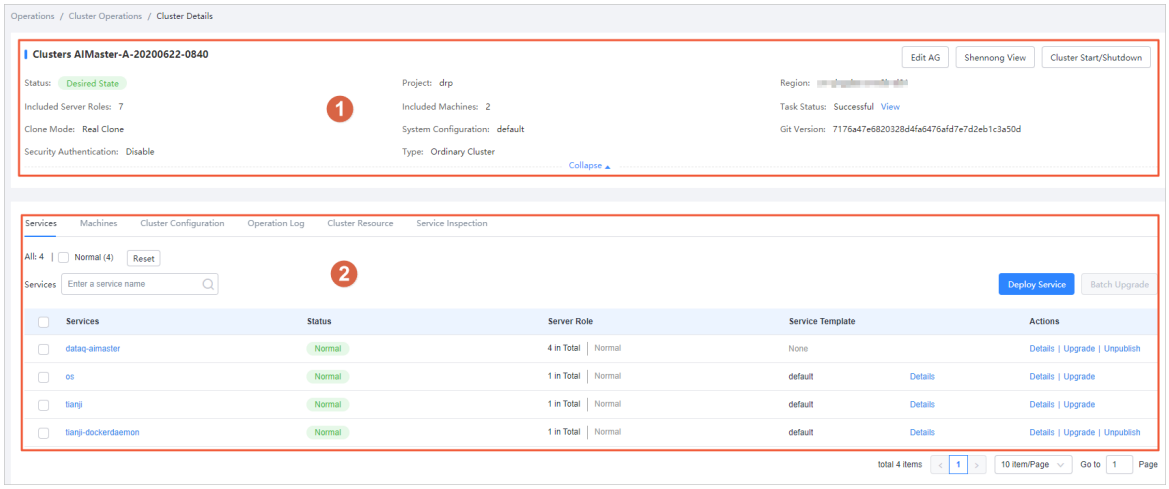

Section	Parameter	Description
①	**Status**	○ **Desired State**: All clusters in a project have reached the desired state. ○ **Not Desired State**: A project has reached the desired state for the first time but a server role has not reached the desired state due to undefined reasons.
	Project	The project to which the cluster belongs.
	Region	The region where the cluster resides.
	Included Server Roles	The number of server roles included in the cluster.
	Included Machines	The number of machines included in the cluster.
	Task Status	The status of the 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.

Section	Parameter	Description
	Git Version	The change version to which the cluster belongs.
	Security Authentication	The access control among processes. By default, security authentication is disabled in non-production environments. You can enable or disable security authentication based on your business requirements.
	Type	◦ **Ordinary Cluster**: an operations unit of machine groups, where 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 certain 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 in the cluster. You can also upgrade or unpublish a service. ◦ **Normal**: The service works normally. ◦ **Not Deployed**: No machine is deployed on the service. ◦ **Changing**: Some server roles in the service are changing. ◦ **Operating**: No server role is changing, but a server role is performing Operations and Maintenance (O&M) operations. ◦ **Abnormal**: No server role is changing or the machines where server roles are deployed are not performing O&M operations. However, the server role status is **not good** or the version that the service runs on the machines is different from the desired state configuration.
	Machines	The running status and monitoring status of each machine in the cluster. You can also view details of server roles that are deployed on each machine.
	Cluster Configuration	The configuration file used in the cluster.
	Operations Log	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 in the cluster.

5.3.2.5.3. View operation logs

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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 Log** tab.
   - In the left-side navigation pane, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, find the target cluster and click **Operations** in the **Actions** column. On the Cluster Details page, click the **Operation Log** tab.

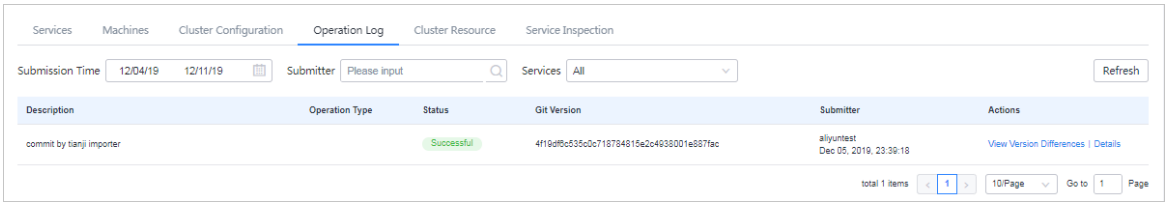

3. View the version differences on the **Operation Log** tab.
   i. Find the target operation log and click **View Version Differences** in the **Actions** column.
   ii. On the **Version Differences** page, select a basic version from the **From** drop-down list. Then, a difference file is displayed in the lower part of the page.
   iii. Select a difference file from the **Different File** drop-down list to view the content of each difference file.

5.3.2.6. Service operations

5.3.2.6.1. View the service list

This topic describes how to view all services and their information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the service list:
   - On the **Homepage** page, move the pointer over the **Instances** section and click Details in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Service Operations**.

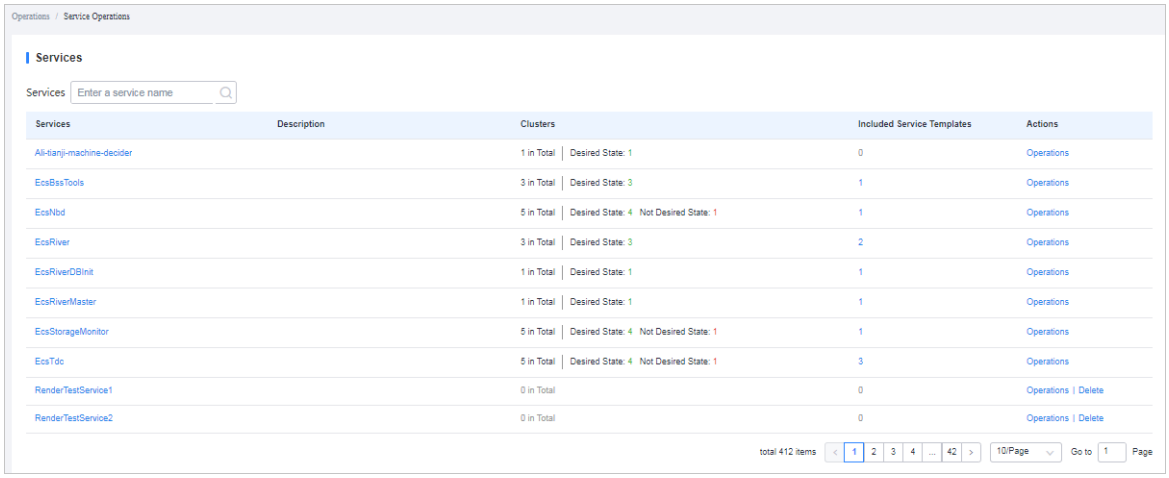

The following table describes the information displayed in the service list.

Parameter	Description
Services	The name of the service. Click the service name to view the service details.
Clusters	The number of clusters where 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.

3. (Optional)Enter a service name in the search box to search for the service.

5.3.2.6.2. View details of a server role

This topic describes how to view details of a server role.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Service Operations**.
3. (Optional)Enter a service name in the search box to search for the service.
4. Click the service name or click **Operations** in the **Actions** column.

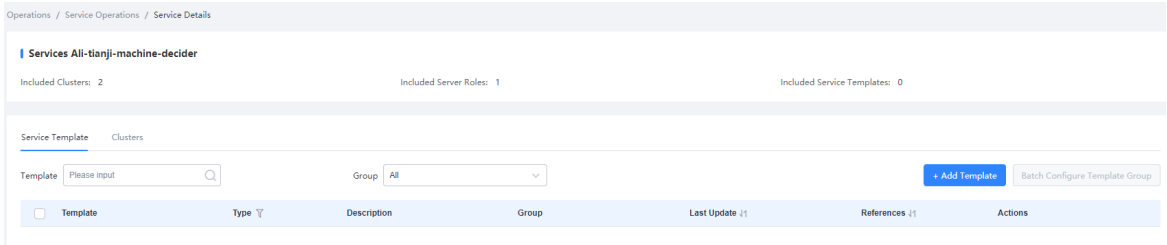

5. On the **Clusters** tab, click a status in the **Server Role Status** column to view the server roles included in a cluster.

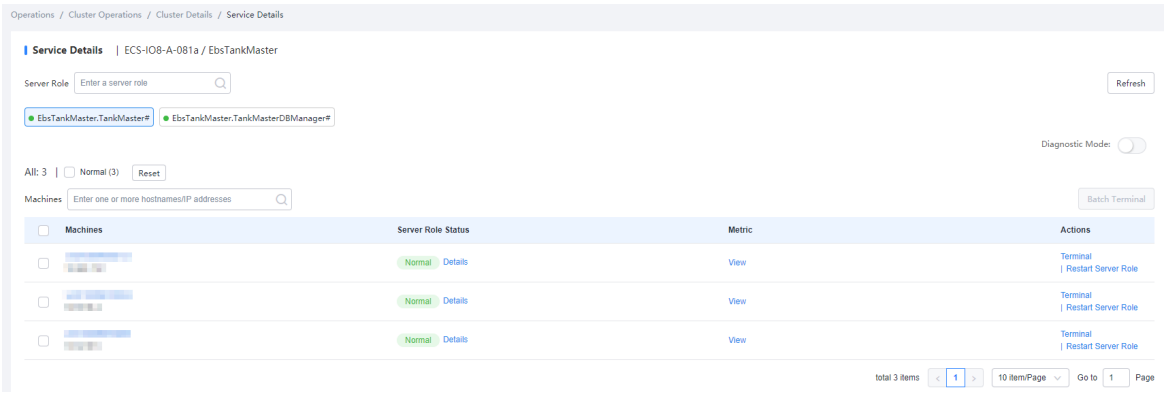

6. Enter a keyword in the search box to search for a server role. Then, the details of the server role are displayed in the list.

Parameter	Description
Machines	The machine to which the server role belongs. Click the machine name to view the machine details.
Server Role Status	The status of the server role. Click **Details** to view the basic information, application version information, application process information, and resources of the server role.
Metric	Click **View** to view the server role and machine metrics.
Actions	○ Click **Terminal** to log on to the machine and perform operations. ○ Click **Restart Server Role** to restart the server role.

5.3.2.7. Machine operations

This topic describes how to view the statistics of all machines.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. Use one of the following methods to go to the machine list:
   - On the **Homepage** page, move the pointer over the **Machines** section and click Details in the upper-right corner.
   - In the left-side navigation pane, choose **Operations > Machine Operations**.

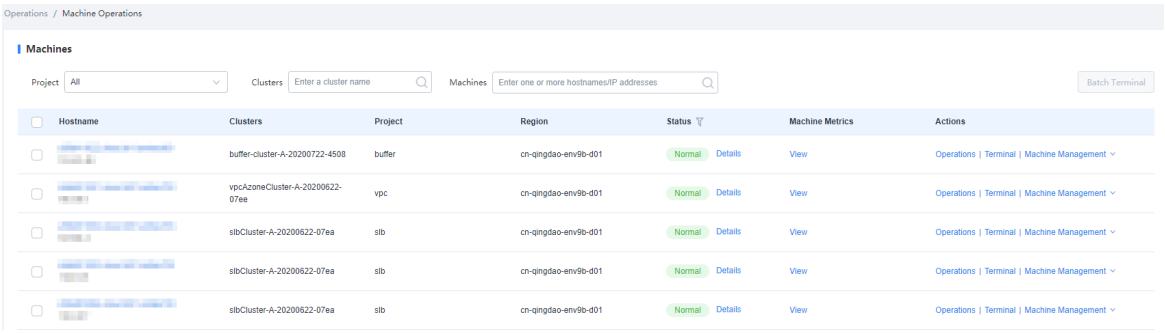

3. (Optional)Select a project from the drop-down list or enter a cluster or machine name to search for the machine.

Parameter	Description
Hostname	Click a hostname to go to the Machine Details page.
Status	The status of a machine. Click the icon to filter machines. Click **Details**. Then, the **Status Details of Machine** dialog box appears.
Machine Metrics	Click **View**. Then, the **Metrics** dialog box appears. 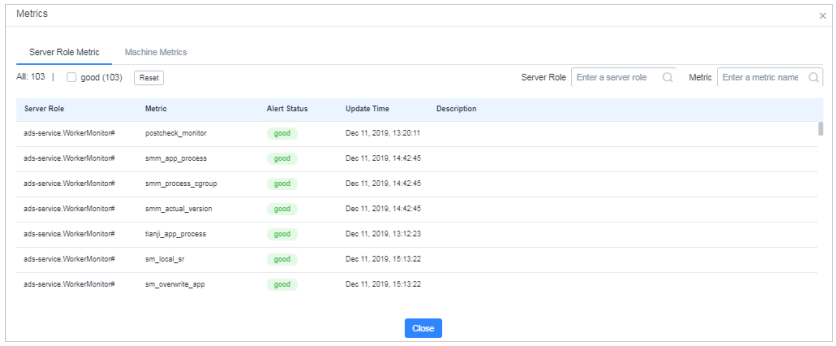 Metrics are displayed on the **Server Role Metric** and **Machine Metrics** tabs. You can view the status and update time of each metric. Enter a keyword in one of the search boxes in the upper-right corner to search for a server role or metric. You can also select the status to filter metrics.
Actions	○ Click **Operations** to go to the Machine Details page. ○ Click **Terminal** to log on to the machine and perform operations. You can select multiple machines and then click **Batch Terminal** in the upper-right corner to log on to multiple machines at a time. ○ Click **Machine Management** to perform an out-of-band restart operation on the machine.

5.3.2.8. Monitoring center

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

5.3.2.8.1. View the status of a metric

This topic describes how to view the status of a metric.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Operations > Service Operations**.
3. (Optional)Enter a service name in the search box to search for the service.
4. Click **Operations** in the **Actions** column.
5. On the **Clusters** tab, use filter conditions to find a cluster. Click **Operations** in the **Actions** column corresponding to the cluster.
6. On the **Services** tab, select a server role and click **View** in the **Metric** column corresponding to a machine to view the server role and machine metrics.

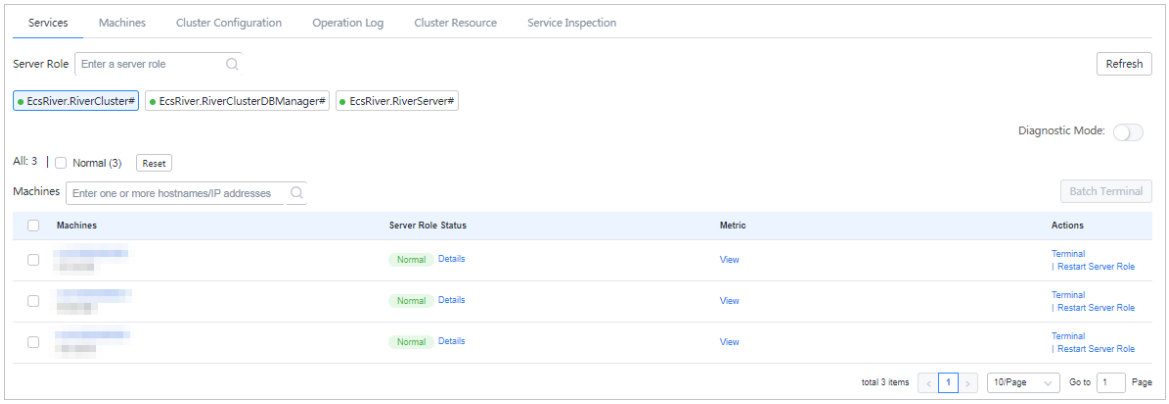

5.3.2.8.2. View the alert status

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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**.

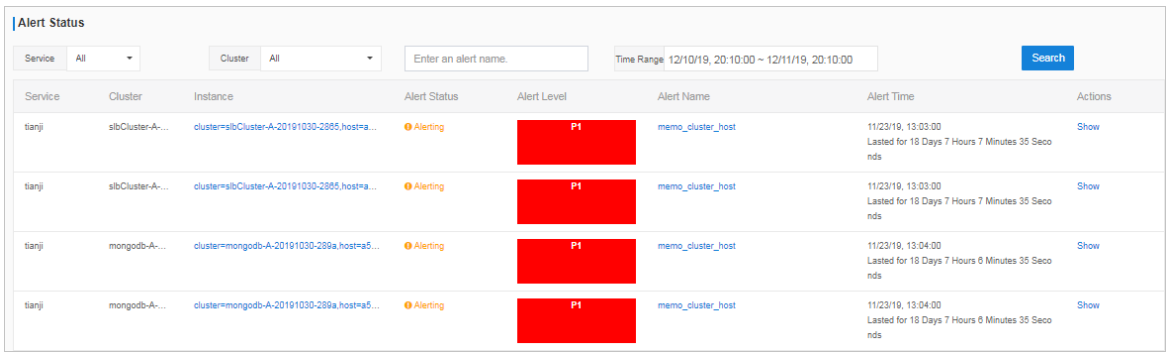

4. (Optional)Search for an alert by service name, cluster name, alert name, or alert time range.

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	Two alert states are available, which are **Normal** and **Alerting**.
Alert Level	Alerts are divided into four levels in descending order of severity: ○ P1 ○ P2 ○ P3 ○ P4
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 has lasted.
Actions	Click **Show** to view the data before and after the alert time.

5.3.2.8.3. View alert rules

This topic describes how to view alert rules.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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**.

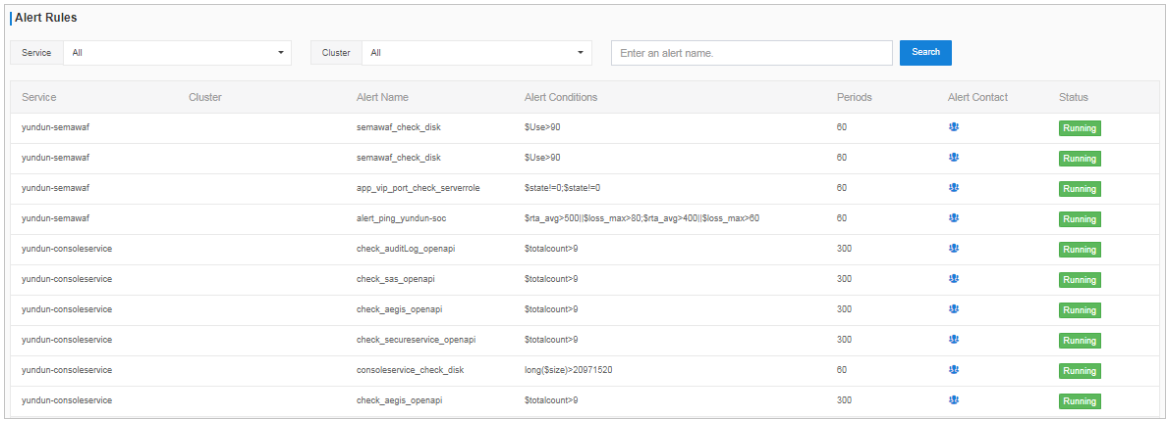

4. (Optional)Search for alert rules by service name, cluster name, or alert name.
5. View alert rules on the **Alert Rules** page. The following table describes the related parameters.

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.

5.3.2.8.4. 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 Apsara Infrastructure Management Framework.
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**.

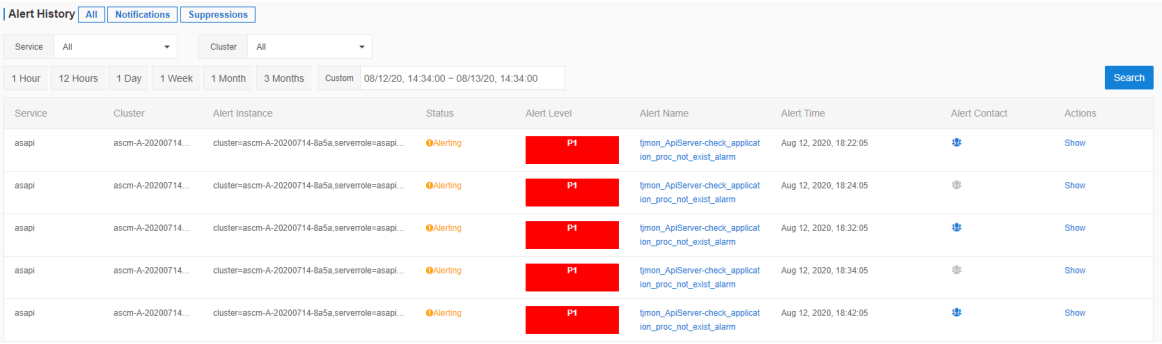

4. (Optional)Search for an alert by service name, cluster name, alert name, or alert time range.
5. View the alert history on the **Alert History** page. The following table describes the related parameters.

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	Two alert states are available, which are **Normal** and **Alerting**.
Alert Level	Alerts are divided into four levels in descending order of severity: ○ P1 ○ P2 ○ P3 ○ P4
Alert Name	The name of the alert. Click the name of an alert to view alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members to notify when the alert is triggered.
Actions	Click **Show** to view the data before and after the alert time.

5.3.2.9. View tasks

This topic describes how to view the submitted tasks and their status.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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. Click the icon in the **Status** column to filter tasks.
4. Find the target task 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.

5.3.2.10. Reports

5.3.2.10.1. View reports

The Reports module allows you to view the statistical data.

Context

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

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Reports**. On the Reports page, click **Go** to open the target page.

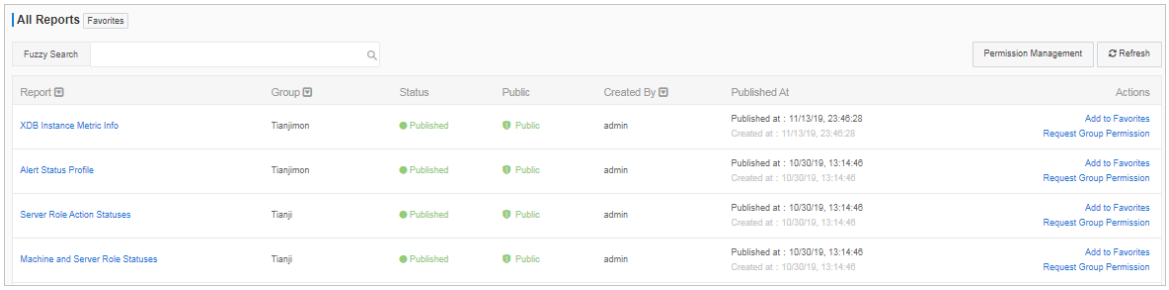

For more information about the report descriptions, see the following table.

Item	Description
Report	The report name. Move the pointer over the down-arrow button next to Report to search for reports by report name.
Group	The group to which the report belongs. Move the pointer over the down-arrow button next to Group to filter reports by group name.
Status	Indicates whether the report is published. ○ Published ○ Not published

Item	Description
Public	Indicates whether the report is public. ○ Public: All of the logon users can view the report. ○ Not public: Only the current logon user can view the report.
Created By	The person who creates the report.
Published At	The time when the report is published and created.
Actions	Click **Add to Favorites** to add this report to your favorites. Then, you can view the report by choosing **Reports > Favorites** in the top navigation bar.

3. (Optional)Enter the name of the report that you are about to view in the search box.
4. Click the report name to go to the corresponding report details page. For more information about the reports, see Appendix.

5.3.2.10.2. 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 Apsara Infrastructure Management Framework.
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**.

5.3.2.11. Tools

5.3.2.11.1. Machine tools

The Machine Tools module guides operations personnel to perform Operation & Maintenance (O&M) operations in common scenarios.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Tools > Operation Tools > Machine Tools**. On the Machine Tools page, click **Go** to open the target page.
3. Select the operation scene according to actual situations.

Operation scene	Description	Action
Scene 1: NC Scale-out (with existing machines)	Scales out an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 2: Host Scale-out (with existing machines)	Scales out the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 3: NC Scale-in	Scales in an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 4: Host Scale-in	Scales in the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 5: VM Migration	Migrates virtual machines (VMs) from a host to another host.	Select a source host and a destination host. Select the VMs to be migrated in the left-side section and then click Select> to add them to the right-side section. Click **Submit** and then click **Confirm** in the displayed dialog box.
Scene 6: Host Switching	Switches from a standby host to a primary host.	Select a source host and a destination host. Click **Submit** and then click **Confirm** in the displayed dialog box.

5.3.2.11.2. IDC shutdown

In some scenarios such as vehicle-mounted ones, you can shut down all machines of all clusters within an IDC with one click.

Prerequisites

The total number of machines of all clusters within an IDC is not more than 25.

Context

When you perform IDC shutdown, business clusters are shut down first, and then the base cluster is shut down.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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**. In the **Confirm Operation** message, enter *SHUTDOWN*.

   > **Warning** The IDC shutdown operation will shut down all services and machines and thus cause business interruption.

   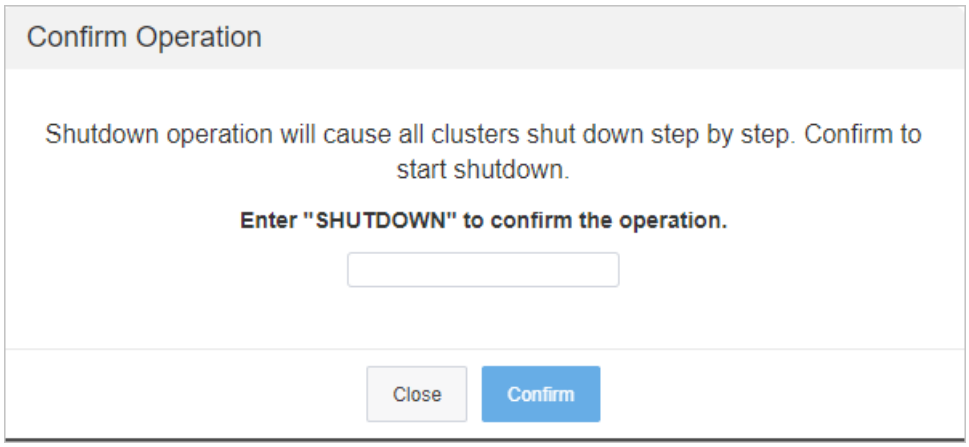

4. If you are sure that you want to perform IDC shutdown, click **Confirm**.

   > **Warning** Backend services need to communicate with the frontend shutdown page during the IDC shutdown process. Do not close the shutdown page until the shutdown is complete.

5. View the IDC shutdown progress and the status 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 states 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 during the shutdown of a cluster, machine, or server role.

You can perform the following operations:

- View the IDC shutdown progress: In the upper part of the **IDC Shutdown** page, view the IDC shutdown progress.
- View the cluster status: In the **Cluster List** section, view the status of each cluster, the total number of machines in each cluster, and the number of machines in each state.
- View the machine status: In the **Cluster List** section, click a status corresponding to a cluster. In the **Machine List** section, view all machines in the corresponding state in 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 status 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 a 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 status to filter all clusters or machines in the corresponding state.
- Refresh data: Click **Refresh** in the upper-right corner to refresh data.

If the status of all clusters in the **Cluster List** section is **shutdown**, the IDC 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 become inaccessible, go to the IDC and confirm that all machines are powered off.

What's next

If you want to use the machines in the future, power on all machines one by one in the IDC and wait until all services reach the desired state.

5.3.2.11.3. View the clone progress

This topic describes how to go to the OS Provision console (Corner Stone) by using Apsara Infrastructure Management Framework, which allows you to know the progress, status, and errors of the machine installation.

Prerequisites

You have obtained the username and password of the OS Provision console from the delivery personnel.

Context

Apsara Infrastructure Management Framework provides a quick entry of the OS Provision console, which allows you to view the machine installation details. The OS Provision console allows you to view the machine clone details and then you can know the progress and status of the machine installation and locate the installation faults.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
2. In the left-side navigation pane, choose **Tools > Clone Progress**.
3. On the logon page of the Corner Stone, enter the **Username** and **Password**, and then click **Submit**.

5.3.2.12. Metadata operations

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

5.3.2.12.1. 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.

Name	Type	Required	Description
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
  }
}
```

5.3.2.12.2. Connect to API operations

This topic describes how to connect to control-side and deployment-side API operations.

Procedure

1. Log on to Apsara Infrastructure Management Framework.
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 target machine.
4. Connect to API operations.
   - Connect to control-side API operations
     a. Find the target 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
```

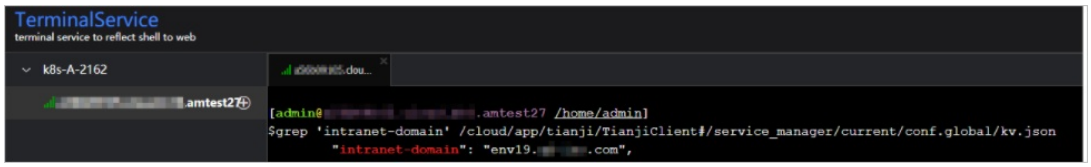

     c. Use one of the following methods to connect to control-side API operations. ListInstance is used in the example.
        - GET request

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

        - POST request

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

   - Connect to deployment-side API operations
     a. Find the target machine and record the IP address in the Hostname column.
     b. Use one of the following methods to connect to deployment-side API operations. CheckState is used in the example.

        Assume that the IP address of the target 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}'
```

5.3.2.12.3. APIs on the control side

5.3.2.12.4. APIs on the deployment side

5.3.2.13. Appendix

5.3.2.13.1. 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.

5.3.2.13.2. 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.

5.3.2.13.3. 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.

5.3.2.13.4. Rolling info report

This report displays the running and completed rolling tasks and the task-related status.

Choose a rolling action

This list only displays the running rolling tasks. If no rolling task is running, no data is available in the list.

Item	Description
Cluster	The cluster name.
Git Version	The version of change that triggers the rolling task.
Description	The description about the change entered by a user when the user submits the change.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Submitted By	The ID of the user who submits the change.
Rolling Task Status	The current status of the rolling task.
Submitted At	The time when the change is submitted.

Server Role in Job

Select a rolling task in the **Choose a rolling action** section to display the rolling status of server roles related to the selected task. If no rolling task is selected, the server role statuses of all historical rolling tasks are displayed.

Item	Description
Server Role	The server role name.
Server Role Status	The rolling status of the server role.

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 that have the rolling task approved by the decider.
Failure Rate	The proportion of machines that have the rolling task failed.
Success Rate	The proportion of machines that have the rolling task succeeded.

Server Role Rolling Build Information

The source version and target version of each application under the server role in the rolling process.

Item	Description
App	The name of the application that requires rolling in the server role.
Server Role	The server role to which the application belongs.
From Build	The version before the upgrade.
To Build	The version after the upgrade.

Server Role Statuses on Machines

Select a server role in the **Server Role in Job** section to display the deployment status of this server role on the machine.

Item	Description
Machine Name	The name of the machine on which the server role is deployed.
Expected Version	The target version of the rolling.
Actual Version	The current version.
State	The status of the server role.
Action Name	The Apsara Infrastructure Management Framework action currently performed by the server role.

Item	Description
Action Status	The action status.

5.3.2.13.5. 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.

Item	Description
Action Status	The status of the action on the server role.
Actions	The approval button.

Machine Component

Displays the hard disk information of pending approval machines.

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

5.3.2.13.6. Registration vars of services

This report displays values of all service registration variables.

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

5.3.2.13.7. Virtual machine mappings

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

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

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

5.3.2.13.8. Service inspector report

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

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

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

5.3.2.13.9. 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.

Item	Description
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.

Item	Description
Reprocess Msg	The error message of the interaction with Business Foundation System during the VIP resource application.
Reprocess Result	The result of the interaction with Business Foundation System during the VIP resource application.
Refer Version List	The version that uses the resource.
Error Msg	The exception message.

5.3.2.13.10. Statuses of project components

This report displays the status of all server roles in an abnormal status on machines of the project, and the monitoring information (alert information reported by the server role to Apsara Infrastructure Management Framework monitor) of server roles and machines.

Error State Component Table

Only displays the information of server roles that are not in GOOD status and server roles to be upgraded.

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The machine name.
Need Upgrade	Whether the current version reaches the final status.
Server Role Status	The current status of the server role.
Machine Status	The current status of the machine.

Server Role Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

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

Item	Description
Server Role	The server role name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Machine Alert Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

Service Inspector Information

Select a row in the **Error State Component Table** section to display the corresponding information in the list.

Item	Description
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Monitored Item	The monitored item name of the server role.
Level	The alert level.
Description	The description about the alert contents.
Updated At	The updated time of the alert information.

5.3.2.13.11. Relationship of service dependency

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

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Dependent Service	The service on which the server role depends.
Dependent Server Role	The server role on which the server role depends.
Dependent Cluster	The cluster to which the dependent server role belongs.
Dependency in Final Status	Whether the dependent server role reaches the final status.

5.3.2.13.12. 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.

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

5.3.2.13.13. 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.

5.3.2.13.14. Auto healing/install approval pending report

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

5.3.2.13.15. 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.

Item	Description
Machine Name	The machine name.
Server Role Status	The running status of the server role.
Server Role Action	The action currently performed by the server role.
Server Role Action Status	The status of the action.
Error Message	The exception message.

Machine Statuses

Displays the running statuses of machines in the selected cluster.

Select a row in the **Statuses on Machines** section to display the information of the corresponding machine in the list.

Item	Description
Cluster	The cluster name.
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status of the machine.
Error Message	The exception message.

6.Products

The Products module allows you to click operations and maintenance services of other products on the cloud platform and ISV access configurations to go to the corresponding page.

6.1. Product list

On the Product List page, you can go to the corresponding operations and maintenance page of a product or ISV page by using Single Sign-On (SSO) and redirection.

Prerequisites

To access the ISV page, make sure that the ISV access information is configured on the **ISV Access Configurations** page. For more information about how to configure the ISV access information, see Configure the ISV access information.

Context

After you log on to the Apsara Stack Operations (ASO) console, you can view O&M icons of different products and different ISV icons on the **Product List** page based on your permissions. An operations system administrator can view all the O&M components of the cloud platform.

The read and write permissions for product O&M are separated. Therefore, the system can dynamically assign different permissions based on different roles.

Procedure

1. In the left-side navigation pane, choose **Products > Product List**.
2. On the **Product List** page, you can view the O&M icons of different products and ISV icons based on your permissions.

6.2. ISV access configurations

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

6.2.1. Configure the ISV access information

You can configure the ISV access information in the system based on business needs. Then, you can click an icon on the product list page to access the corresponding ISV page.

Procedure

1. In the left-side navigation pane, choose **Products > ISV Access Configuration**.
2. In the upper part of the page, click **Add**.
3. In the **Add** pane, configure the ISV access information.

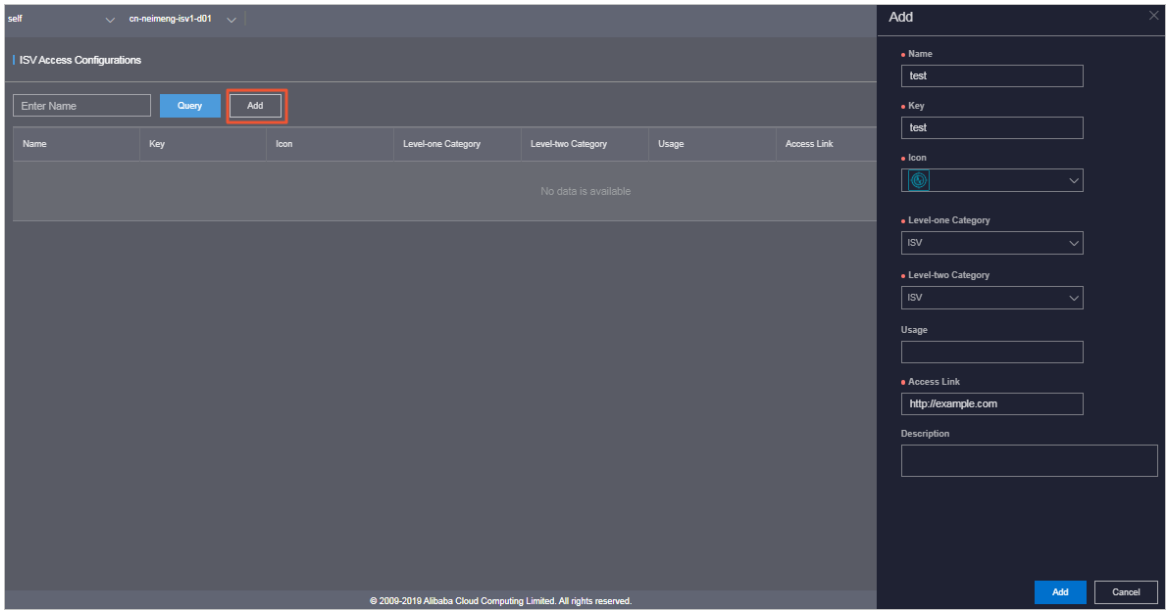

The following table describes the parameters.

Parameter	Description
Name	The name of the ISV to be accessed.
Key	Typically, enter an identifier related to the ISV business as the key.
Icon	Select the icon displayed on the **Product List** page for the ISV to be accessed.
Level-one Category and **Level-two Category**	The category to which the ISV to be accessed belongs on the **Product List** page.
Usage	The function of the ISV to be accessed.
Access Link	The address of the ISV to be accessed.
Description	The description related to the ISV to be accessed.

4. Click **Add**.

Result

You can view the added ISV icon in the Product List page by choosing **Products > Product List**. Click the icon and then you can go to the corresponding page.

6.2.2. Modify the ISV access information

If the ISV information is changed, you can modify the ISV access information.

Procedure

1. In the left-side navigation pane, choose **Products > ASV Access Configuration**.
2. (Optional)In the search box on the page, enter the ISV name, and then click **Query**. Fuzzy search is supported.

3. Find the ISV whose access information is to be modified. Click **Modify** in the **Actions** column.

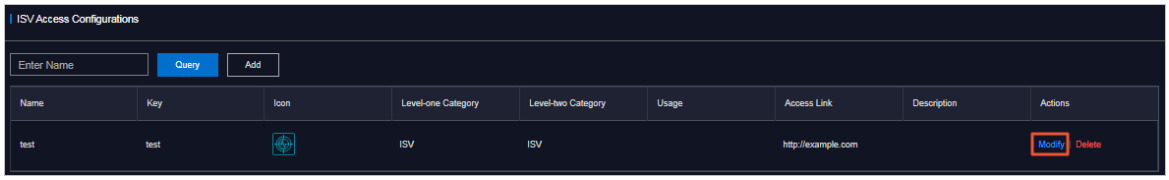

4. In the **Modify** pane, modify the name, key, icon, level-one category, level-two category, usage, access link, or description of the ISV.
5. Click **Modify**.

6.2.3. Delete the ISV access information

You can delete the ISV access information added in the system based on business needs.

Procedure

1. In the left-side navigation pane, choose **Products > ISV Access Configuration**.
2. (Optional)In the search box on the page, enter the ISV name, and then click **Query**. Fuzzy search is supported.
3. Find the ISV whose access information is to be deleted. Click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

Result

The deleted ISV will no longer be displayed in the **Product List**.

7.Log configurations

7.1. What is LogAgentconfig?

LogAgentconfig is a single-node service that is deployed on OPS1, and is one of the Apsara Stack Agility SE base modules. LogAgentconfig delivers and updates log collection configurations.

LogAgentconfig serves as a configuration delivery service on the Log Service client. After the log directories, name of the log file to be collected, Logtail configuration in full regex mode, and log parsing configuration are committed, LogAgentconfig delivers log collection configurations to base modules. After the configurations are delivered, the collected logs are written to the Kafka service. Then, the logs in Kafka are consumed by Logstash and stored in Elasticsearch. Finally, you can view logs in Kibana.

7.2. Log on to the LogAgentconfig console

This topic describes how to log on to the LogAgentconfig console.

Prerequisites

Before you log on to the LogAgentconfig console, confirm with the deployment personnel that the following requirements have been met:

- LogAgentconfig is deployed on the OPS1 server in the current deployment environment and reaches the desired state.
- logservicelite-Kafka and logservicelite-elk reach the desired state.

Procedure

1. Obtain the URL used to access the LogAgentconfig console by using Apsara Infrastructure Management Framework.

   > **Note** This topic describes how to Obtain the URL used to access the LogAgentconfig console by using Apsara Infrastructure Management Framework. If you have obtained the IP address of the OPS1 server, append port number 8888 to the IP address to obtain the URL of the LogAgentconfig console. Example: http://<OPS1 IP address>:8888.

   i. Log on to the Apsara Infrastructure Management Framework console. For more information about how to log on to the new version and the old version of the Apsara Infrastructure Management Framework console, see **Operations tools > Apsara Infrastructure Management Framework > New version**, or **Operations tools > Apsara Infrastructure Management Framework > Old version** in the "**Log on to Apsara Infrastructure Management Framework**" section of *Apsara Stack Agility SE Operations and Maintenance Guide*. This topic describes how to log on to the LogAgentconfig console by using the new version of the Apsara Infrastructure Management Framework console.

   ii. In the left-side navigation pane, click **Reports**.

   iii. Find and click **Resource Application Report** on the All Reports page.

iv. In the **Resource Status** section, click the More icon on the Service column, enter logservicelite-logagentconfig, and click Apply Filter.

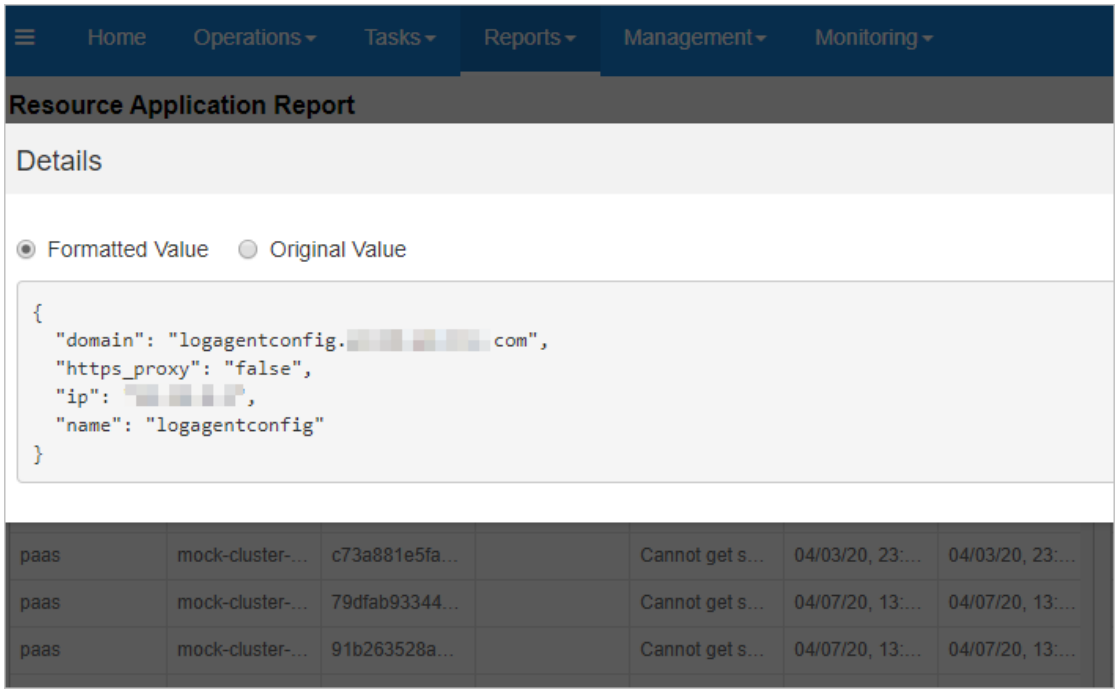

v. Append port number 8888 to the domain name that you obtained in the preceding step to obtain the URL of the LogAgentconfig console. Example: http://Logagentconfig.example.com:8888.

2. Enter the URL in the address bar of your browser, and press the Enter key to go to the logon page of the LogAgentconfig console.

7.3. Configure log collection

This topic describes how to configure log collection for base modules.

Procedure

1. Log on to the LogAgentconfig console.
2. In the **Configure** section, configure the parameters of the server role from which logs will be collected.

Parameter	Description
ConfigureName	The name of the configuration file. The name must be unique. Otherwise, the configuration file fails to be delivered.
ServiceName	The name of the service to which the target server role belongs.
ServiceRole	The name of the service role from which logs will be collected

Parameter	Description
ClusterName	The name of the cluster to which the target server role belongs.

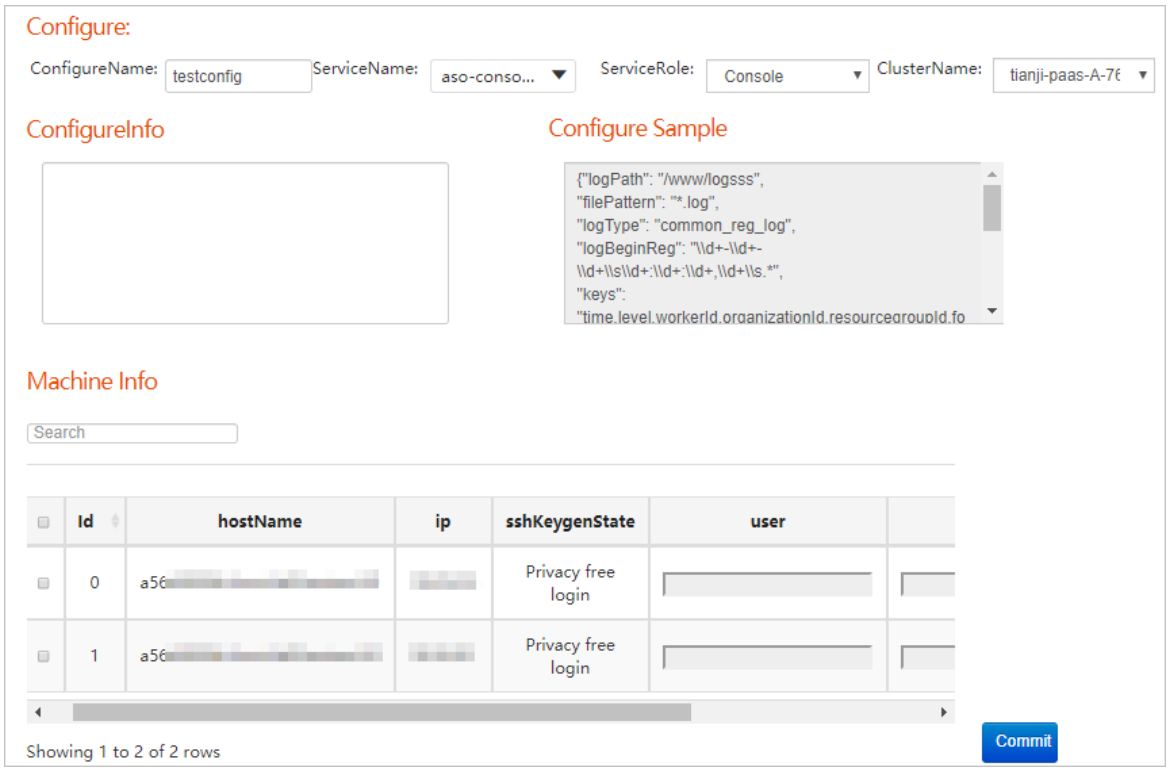

After you configure the preceding parameters, information about the machine where the server role resides is displayed in the **Machine Info** section.

3. Enter the collection configuration information in the **ConfigureInfo** section, as shown in the Configure Sample section. Based on the information, the system automatically generates the configuration file for log collection.
4. In the **Machine Info** section, select the machines to which you want to deliver the configuration file. Then, enter the username and password, and click **Check**.

> **Note** You can the enter a machine name in the search bar to search for the required machine.

The system checks whether the username and password are correct, and provides the check result.

5. In the message that appears, click **OK**.
6. Click **Commit**. The entered information is delivered to a backend service that dynamically generates the configuration file for log collection to be delivered to the specified machine. The entered information includes the configuration name and server role information specified in the **Configure** section, the configuration information specified in the **ConfigureInfo** section, and the machine information selected in the **Machine Info** section.
   After the configurations are delivered, the configuration information appears in the **Configure History** section.
7. After configurations are delivered, you can log on to Kibana Log O&M platform to view the log information collected. For more information about how to view log information on the Kibana Log

O&M platform, see the content in the "Kibana Log O&M" section of *Apsara Stack Agility SE Operations and Maintenance Guide* .

8.Log O&M

8.1. Kibana Log O&M

8.1.1. Overview of the Kibana log O&M platform

Kibana 7.1 is an open-source analytics and visualization platform. Logs for Apsara Stack Agility SE services such as ApsaraDB for RDS, Xnet2, Asapi, and POP are accessible to Elasticsearch, Logstash, and Kibana (ELK). You can use Kibana to view and retrieve related logs.

This topic describes only the common features of Kibana 7.1 for daily O&M. For more information, visit Kibana Guide.

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

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

Prerequisites

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

Procedure

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

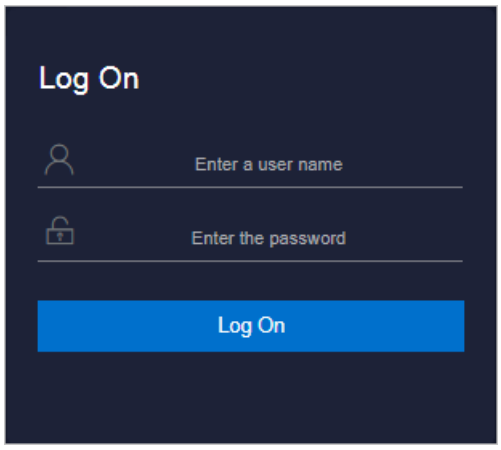

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

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

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, click **Products**.
6. In the **Apsara Stack O&M** section, click **Kibana Log O&M**.

8.1.3. Quick start

8.1.3.1. Create index patterns

You must create index patterns before you can view and retrieve logs.

Context

An index pattern can identify one or more indexes. Kibana allows you to search for indexes that match a specified index pattern.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon.
3. In the **Kibana** section, click **Index Patterns**.

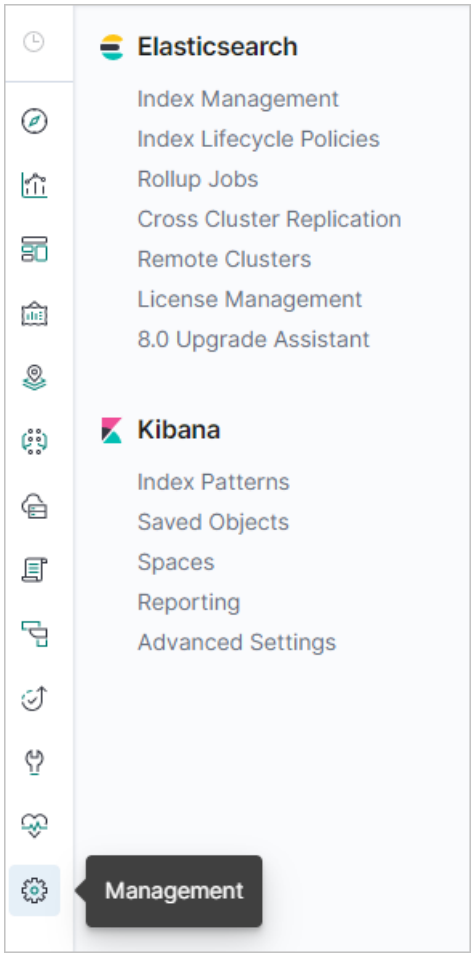

> **Note** When you create an index pattern for the first time, the **Create index pattern** page appears automatically. To create more index patterns, you must click **Create index pattern** in the upper-left corner of the page to open the **Create index pattern** page.

4. In the **Index pattern** field, enter the name of the index pattern as prompted. After you enter an index pattern name, the system matches indexes based on the index pattern. To match system indexes, turn on **Include system indices** in the upper-right corner of the page.

> **Note** To match a specified index, you can set the index pattern name to the index name without using wildcards.

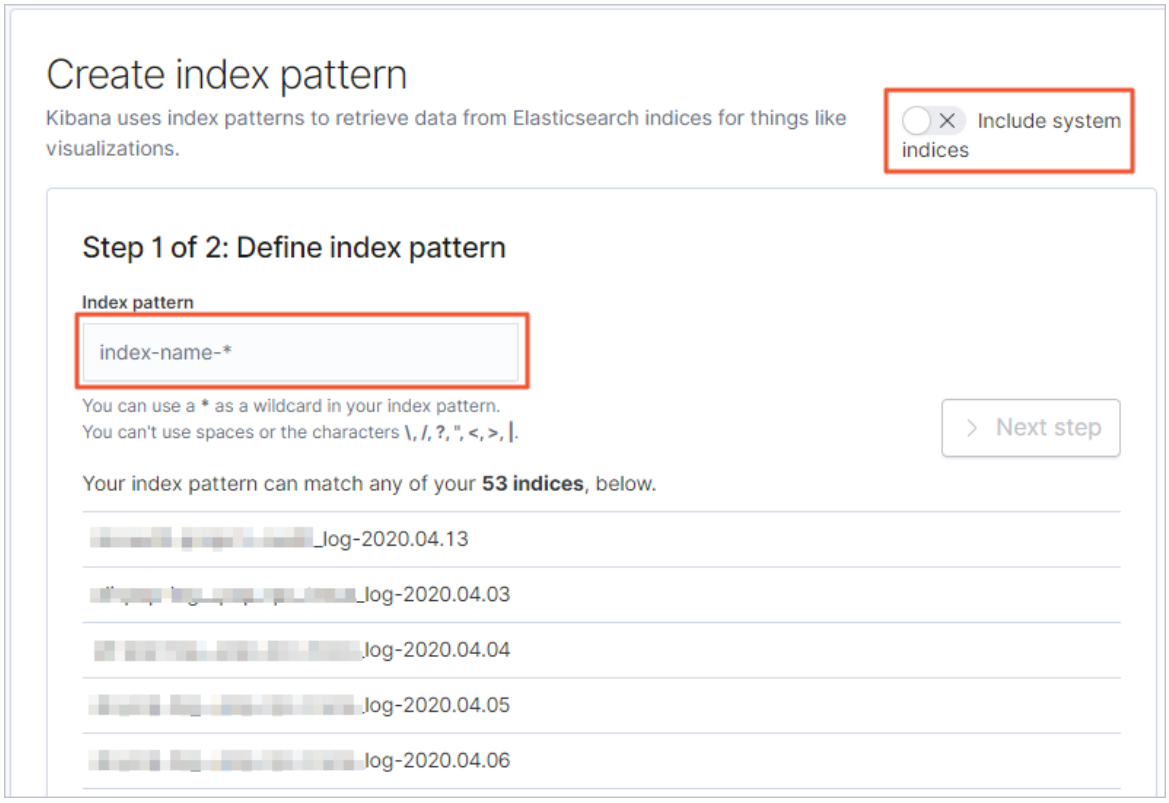

5. Click **Next step**.
6. On the right of **Time Filter field name**, click **Refresh**. Select a time field from the **Time Filter field name** drop-down list for visualization.

> Note
>
> - Before you select a time field from the **Time Filter field name** drop-down list, you can click **Refresh** in the upper-right corner.
> - If the time field is not contained in the matched indexes, this option is not required.

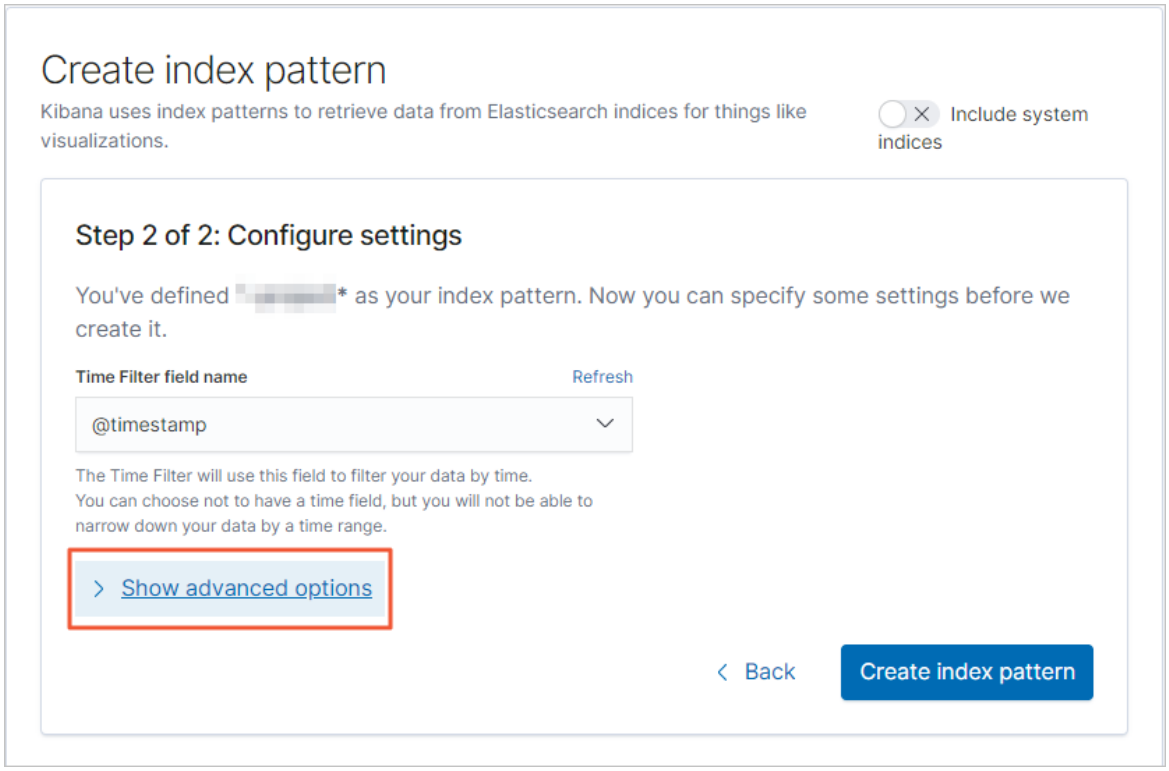

7. (Optional)Click **Show advanced options**, and enter a custom index pattern ID as prompted.
8. Click **Create index pattern**.

8.1.3.2. View data in documents

You can view the documents that match an index and all the data in the documents.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.

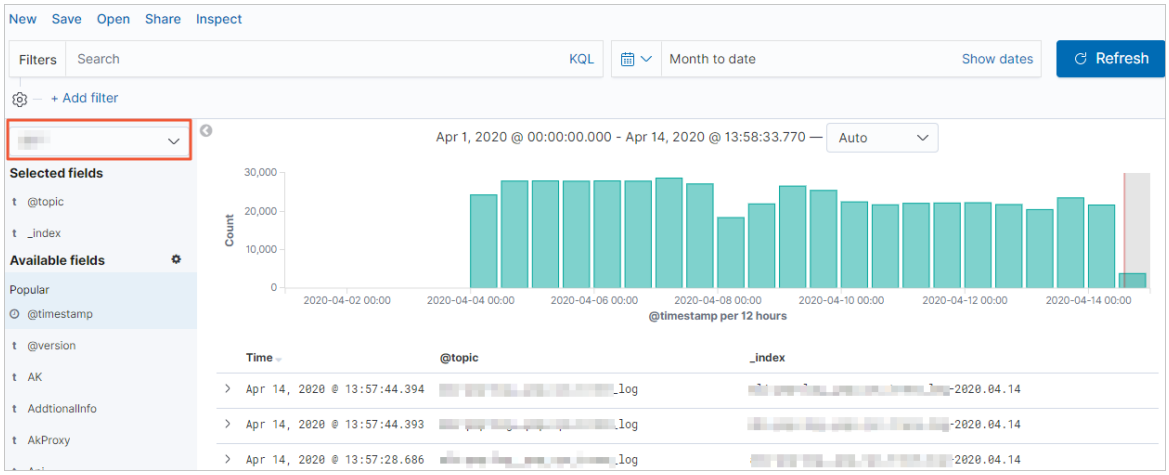

3. In the **Available fields** list, move the pointer over a field, and click **add** to add the field to the **Selected fields** list.
4. You can add multiple fields. All documents that match the fields listed in **Selected fields** are displayed on the right side of the page.

> **Note** If there is a field you do not want to view matches for, you can move the pointer over the field in the **Selected fields** list, and click **remove**. Alternatively, you can move the pointer over the field on the right side of the page, and click the Delete icon.

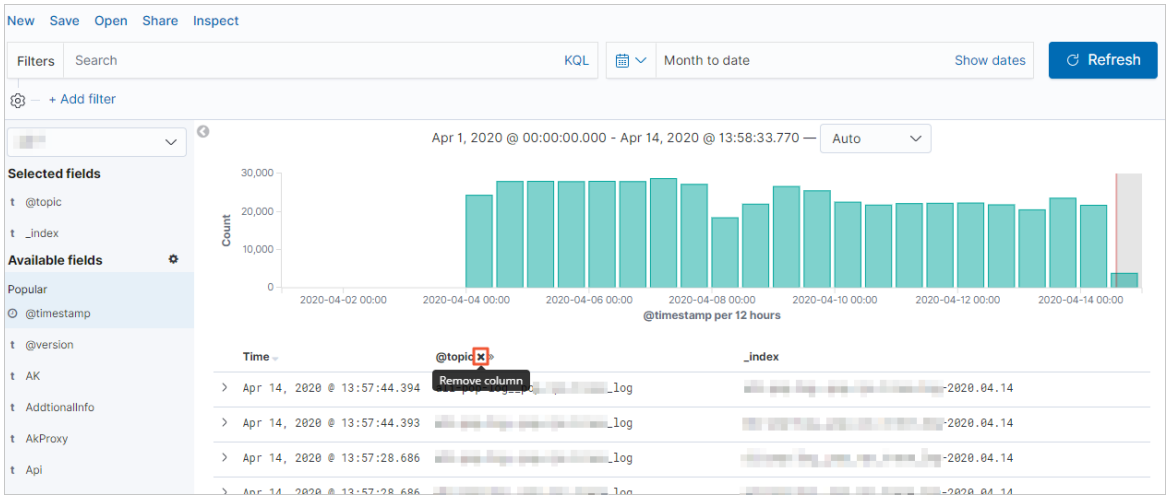

5. Click the Collapse icon in front of each document to view the details of each document. You can view data in the **Table** and **JSON** formats.

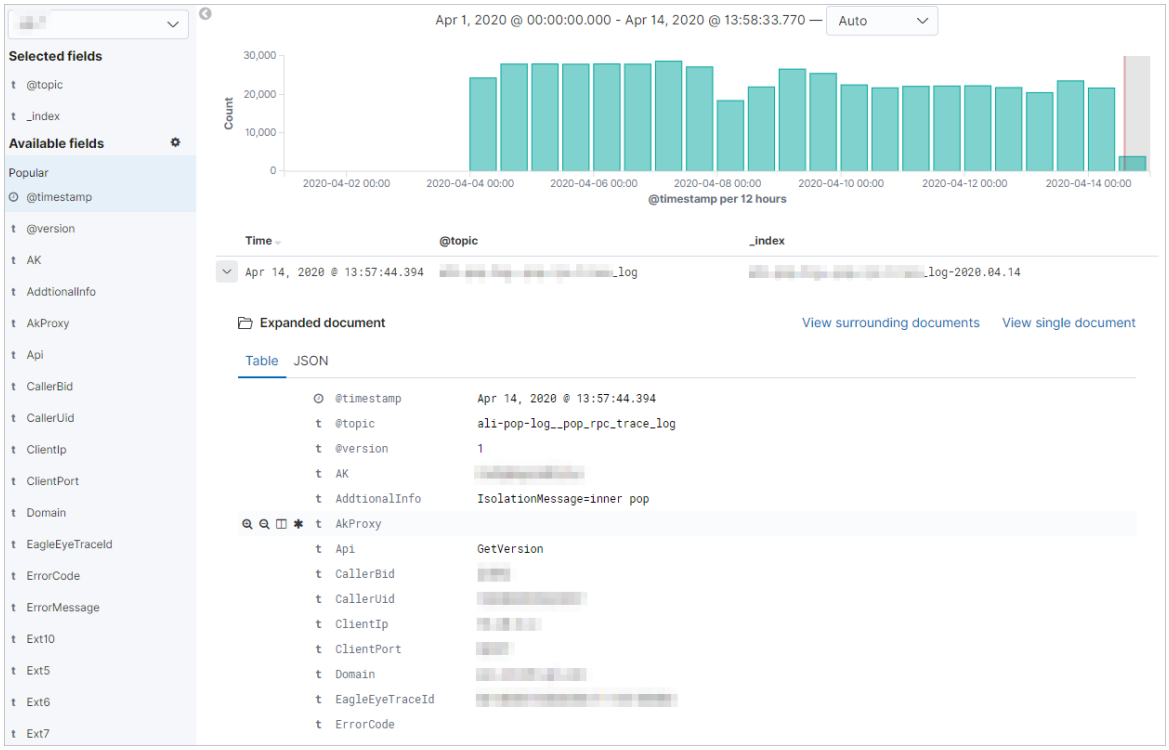

6. When you view documents, you can perform the following operations:
   - View a single document

You can click **View single document** to view a document on a new page. You can bookmark this page and share its link to enable direct access to the document.

- View surrounding documents

  For some applications, you may need to view documents before and after a document. You can only view surrounding documents for index patterns that contain time-based events.

  Click **View surrounding documents**. Documents listed before and after the specified document are displayed. The displayed documents are sorted by the time field specified in the index pattern configurations. If multiple documents have the same time field value, the documents are sorted in the internal order by default.

- Change the ordinal position of fields

  Move the pointer over the name of the target field and click the Move icon behind the field name to move the field to the left or to the right.

  > **Note** Only index fields can be sorted.

8.1.3.3. Filter data by using a time filter

A time filter allows you to search data generated within a specific time range. If a time field is configured for the selected index pattern, you can configure a time filter to filter data.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. On the **Discover** page, configure a time filter.

   By default, the time filter is set to the last 15 minutes. You can use the time selector to reset the time filter in the following ways:

   - Click the field on the right of the time selector to set the start time and end time in sequence. On the **Absolute** tab, set a start time and end time. On the **Relative** tab, set a start time and end time relative to the current time. The relative time can be in the past or future.

> **Note** To set the current time as the start time or end time, you can click **Set data and time to now** on the **Now** tab.

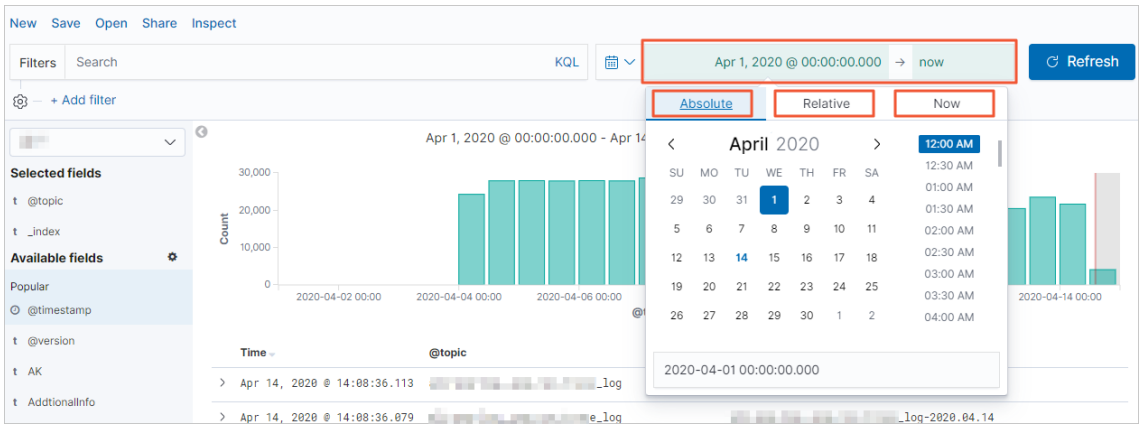

- Click the time selector icon. In the **Quick select** section, specify a time filter relative to the current time in years, months, weeks, days, hours, minutes, or seconds in the past or future. Click Apply.

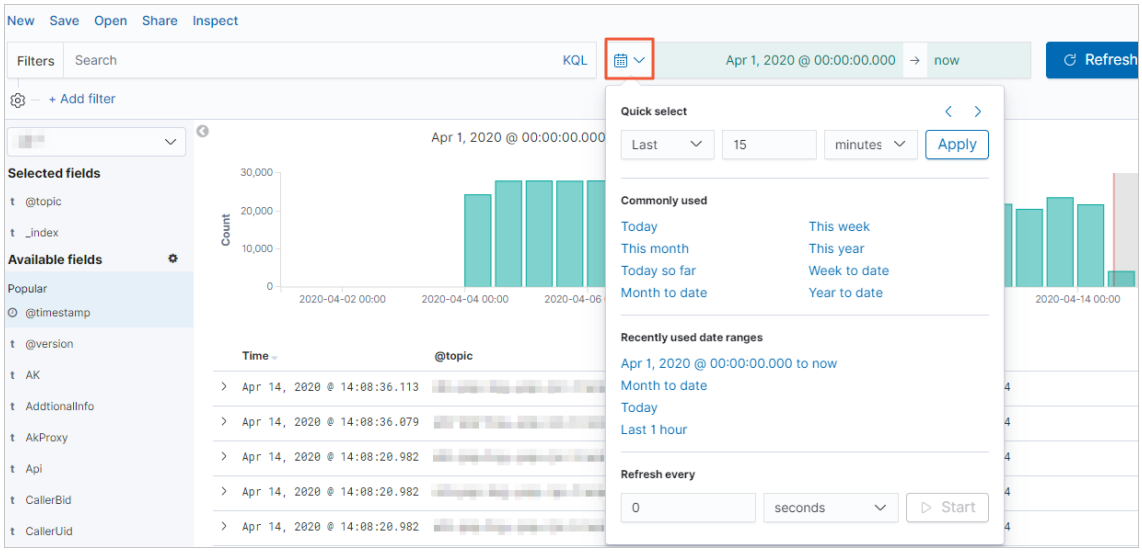

- Click the time selector icon. In the **Commonly used** section, select the desired time from the listed options, including Today, This week, This month, This year, Today so far, Week to date, Month to date, and Year to date. Click **Apply**.
- Click the time selector icon. In the **Recently used date ranges** section, select a recently used filter. Click **Apply**.

After a time filter is configured, data that meets the specified conditions are displayed.

8.1.3.4. Filter data by using column charts

You can use column charts to filter data if a time field is configured for the index pattern.

Procedure

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

2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. In column charts on the right side of the page, you can click a chart and move the pointer to view the data distribution in a specific time range.

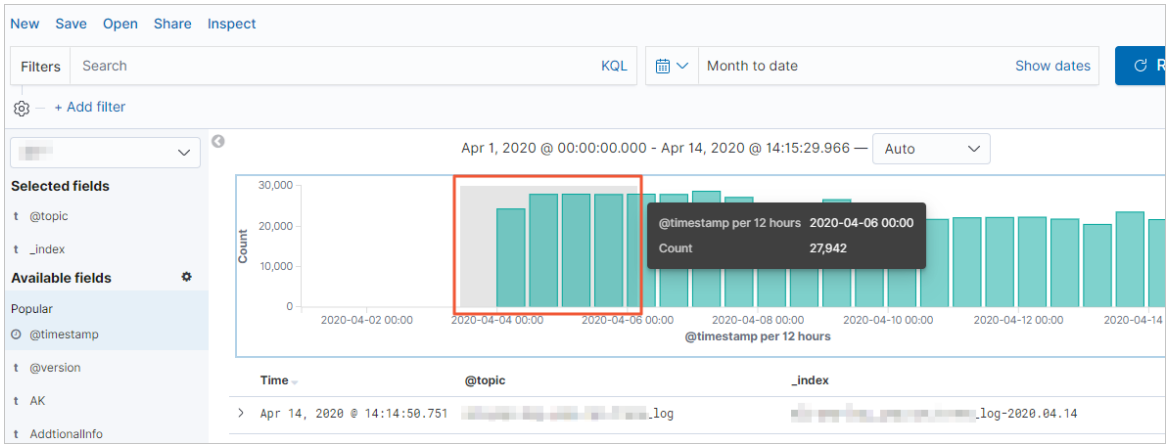

Note You can click a column chart multiple times to enlarge it.

8.1.3.5. Query data by using KQL

You can use the query syntax provided by Kibana Query Language (KQL) to query data.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. In the **Filters** field, enter the fields to be queried. When you enter the fields to be queried, you can select fields from the available fields, and then select filter conditions as prompted.

   The following table describes the symbols of filtering conditions.

Symbol	Description
:	The field value must be the specified value.
:*	The field value can be in any format.
and	The field value must meet the specified two conditions.
or	The field value must meet one or more conditions.

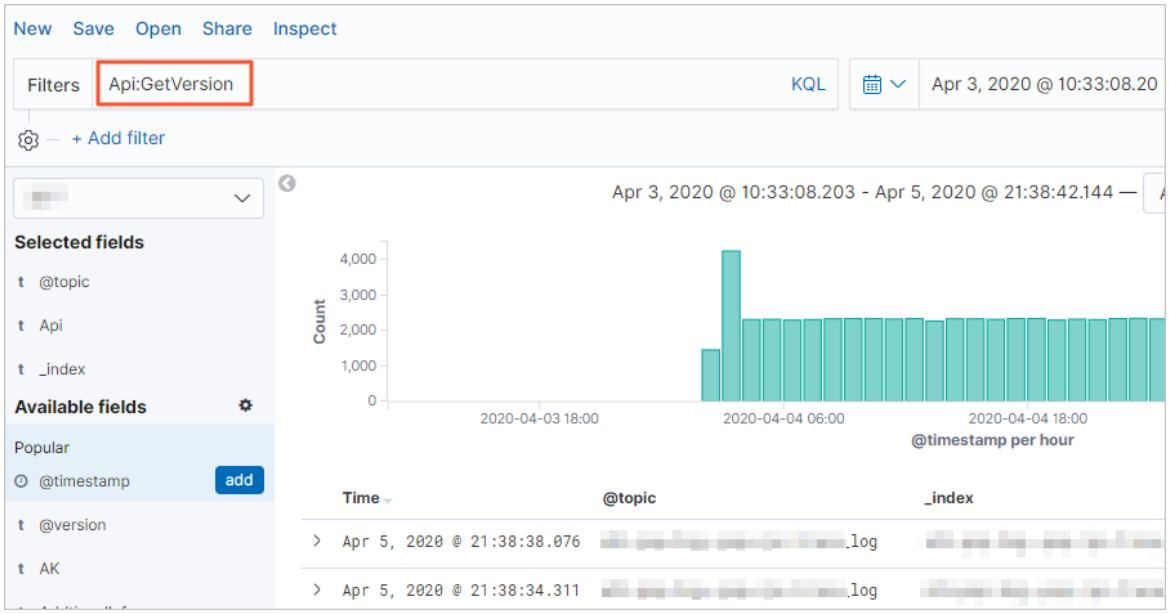

4. Press the Enter key to query data.

8.1.4. Explore data

8.1.4.1. Open a saved search

You can open a saved search on the **Discover** page.

Prerequisites

Make sure that a saved search exists. For more information about how to save a search, see Save a search.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon.
3. In the toolbar at the top of the page, click **Open**.
4. In the **Open Search** dialog box that appears, select a saved search.

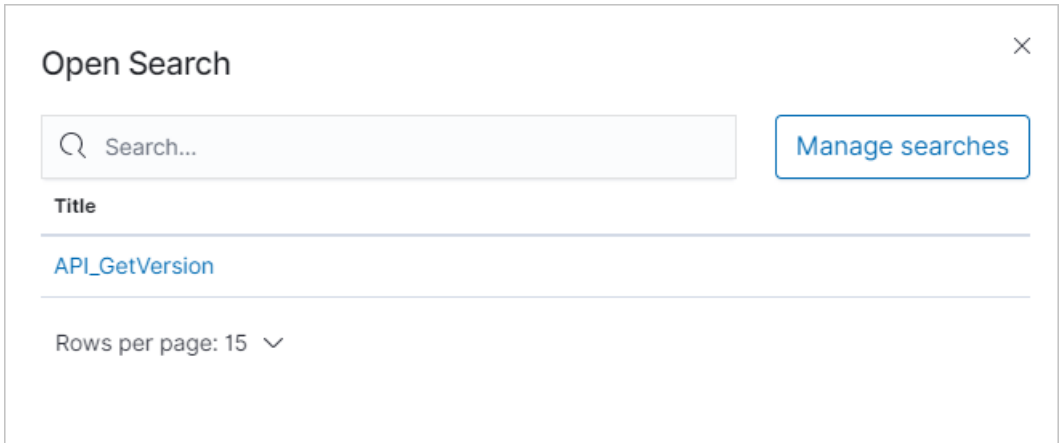

After a saved search is opened, if the index pattern specified in the saved search is different from the selected index pattern, the index pattern specified by the saved search is used. The query statements of the saved search are also used.

8.1.4.2. View statistics for field data

In the field list, you can view the statistics of each field, including the number of documents that contain the field, top five values in the field, and percentage of each value in these documents.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. In the **Available fields** list, click the name of the field that you want to view.

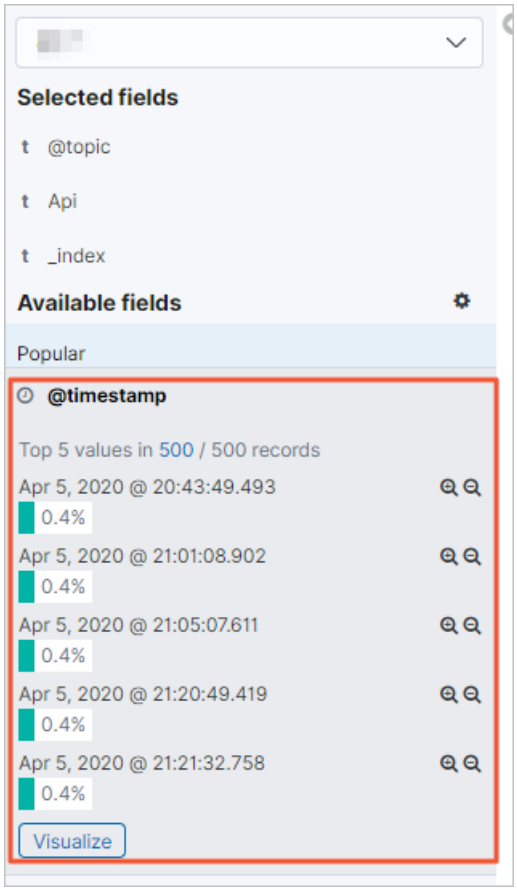

The following statistics are displayed below the field: the number of documents that contain the field, top five values in the field, and the percentage of each value in these documents.

4. (Optional)For aggregatable fields, you can click **Visualization** to view and analyze the visualized data.

8.1.4.3. Filter by fields

You can add a field filter from the field list or document list, or you can manually add a field filter to filter documents that contain specific field values. You can also create a negative filter to exclude documents that contain the specified field values.

Add a filter from the field list

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. In the **Available fields** list, click the name of the field for which you want to add a filter.

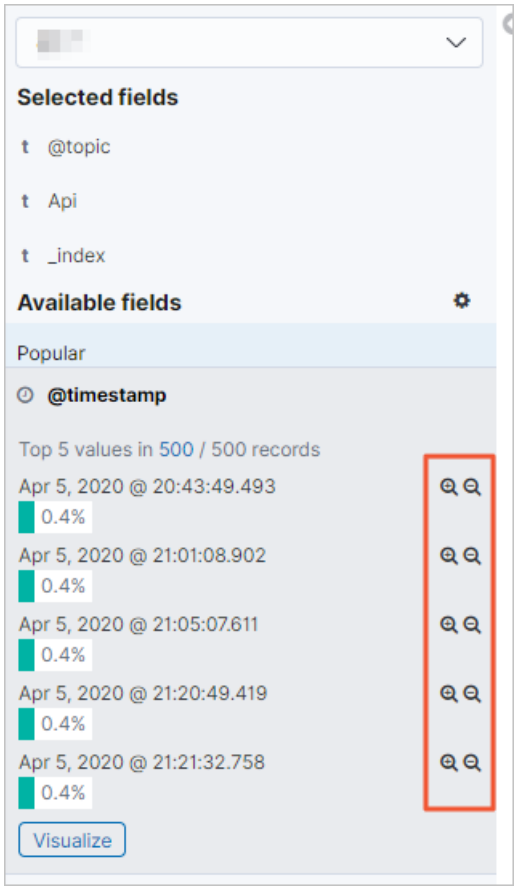

The statistics of the field are displayed.

4. You can perform the following operations:
   - Add a positive filter for a field

     Click the icon next to a field value. Documents that contain the specified field value are displayed on the right side of the page.
   - Add a negative filter for a field

     Click the icon next to a field value. Documents that do not contain the specified field value are displayed on the right side of the page.

Add a filter from the document list

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. In the **Available fields** list, move the pointer over a field, and click **add** to add the field to the **Selected fields** list.
   All documents that match the fields in the **Selected fields** list are displayed on the right side of the page.
4. Click the collapse icon in front of the target document to view all fields contained in the document.

5. You can perform the following operations:
   - Add a positive filter for a field

     Find the index field for which you want to add a filter, move the pointer over the field, and then click the icon. Documents that contain the specified field value are displayed on the right side of the page.
   - Add a negative filter for a field

     Find the index field for which you want to add a filter, move the pointer over the field, and then click the icon. Documents that do not contain the specified field value are displayed on the right side of the page.
   - Query documents that contain a specific field

     Find a field whose values are not empty and for which you want to add a filter, move the pointer over the field, and then click the icon. Documents that contain the field are displayed on the right side of the page.

Add filters manually

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

2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. In the toolbar at the top of the page, click **Add filter**.

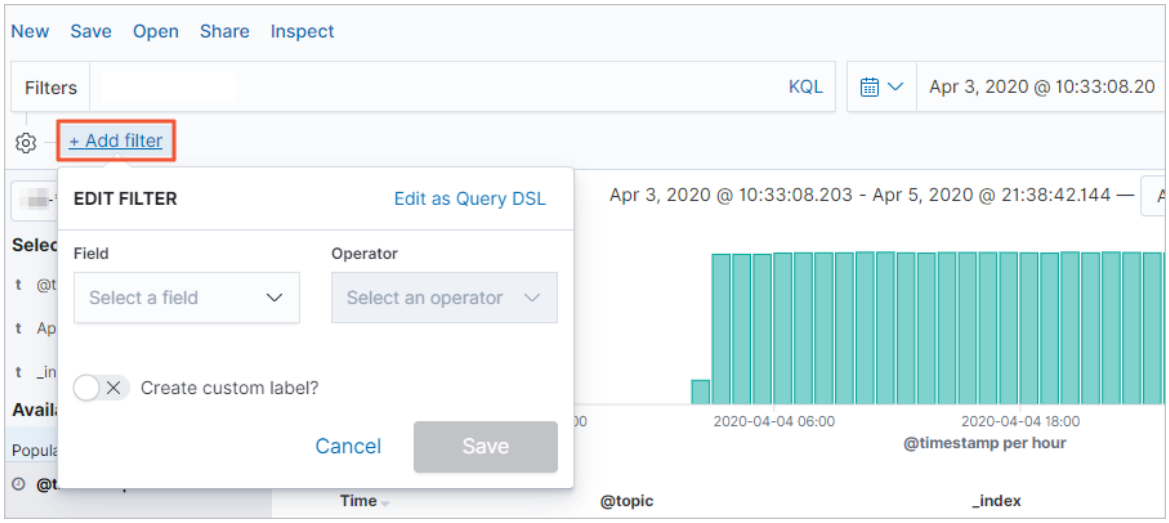

4. Select a field from the **Field** drop-down list, select an operator from the **Operator** drop-down list, and then enter a value in the **Value** field.

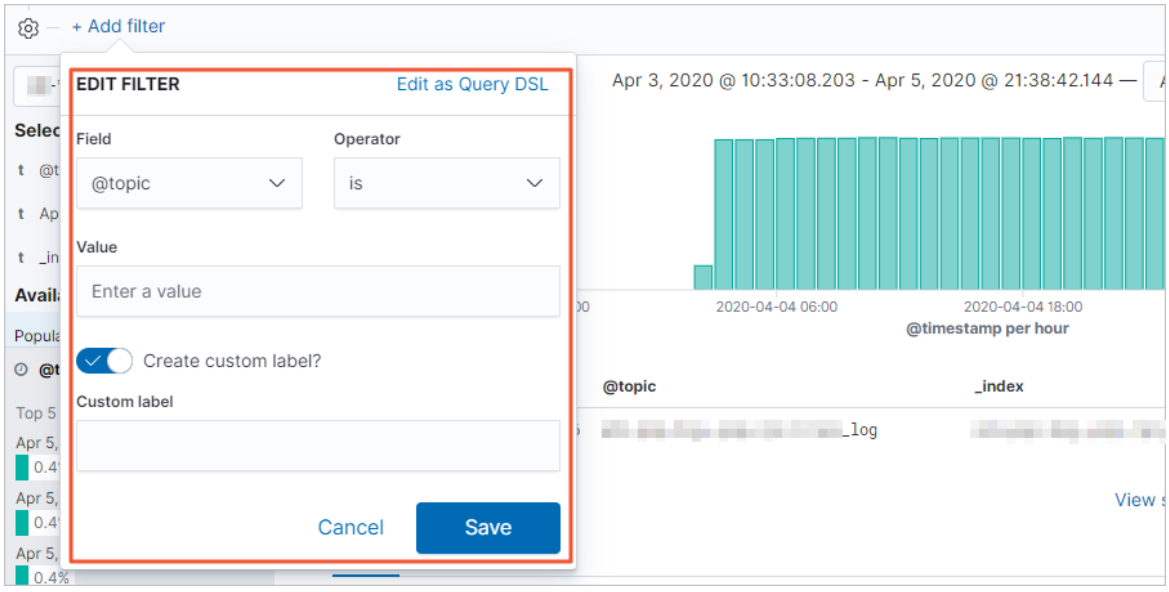

5. To create a custom label, turn on **Create custom label**, and enter the label information.
6. Click **Save**.
   Then, the system queries information of the specified field.

8.1.4.4. Configure a refresh interval

When index data changes, results displayed on the **Discover** page may become obsolete. To avoid such problems, you can configure a refresh interval to periodically resubmit your searches to retrieve the latest results.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon. The default index pattern is displayed below the filter bar. You can select an index pattern from the list as needed. Kibana matches the target indexes based on the specified index pattern.
3. On the **Discover** page, click the Time Selector icon in the Kibana toolbar.

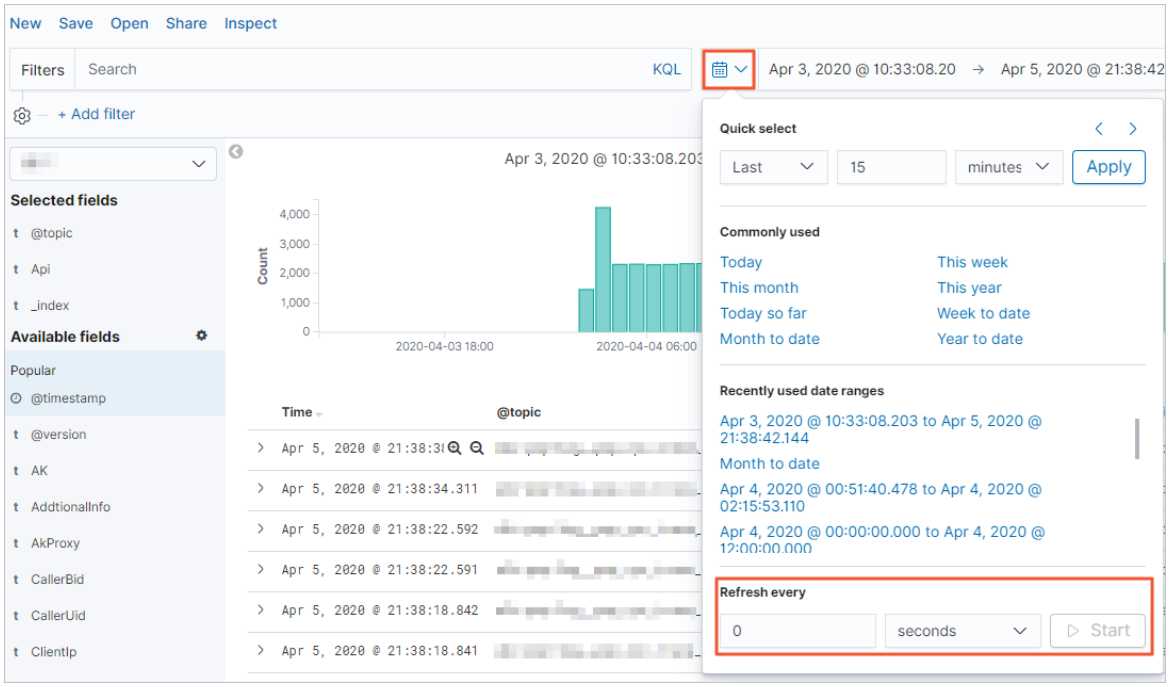

4. In the **Refresh every** field, enter a refresh interval. **0** indicates that the search results will not be refreshed.
5. Click **Start**. After you configure a refresh interval, your searches will be periodically resubmitted based on the specified interval to refresh the search results.

8.1.4.5. Save a search

After you query, filter, or search for specific data, you can save the search as needed.

Procedure

1. For more information about how to query, filter and search for data, see View data in documents, Filter data by using a time filter, Filter data by using column charts, or Query data by using KQL.
2. In the toolbar at the top of the page, click **Save**.

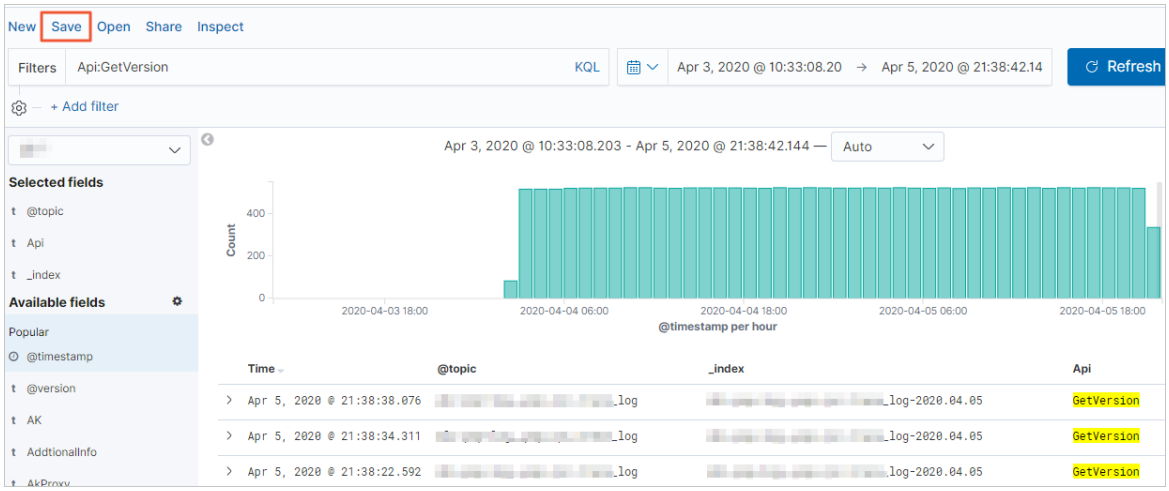

8.1.5. Use development tools to retrieve data

You can also use development tools to retrieve data.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the icon.
3. On the **Console** tab of the **Dev Tools** page, enter query statements to search for data that meets the specified conditions.

   Example:

   - Query all indexes: `GET /_cat/indices? pretty`
   - Query the specified index content: `Get indexname/_search`
   - Query indexes based on query statements:

```
Get /indexname/_doc/_search{
  "query":{
    "query_string":{
      "qurey":"condition"
    // The condition must be in the "field name: the data you want to search" format. You can use "an
d" to connect multiple queries.
        // Example: "query":"id:123 and name:Alice"
    }
  },"sort":{
    "@timestamp":{
      "order":"asc" // The sorting method.
    }
  },"from":?,"size": 100 //The number of logs to query each time.
}
```

8.1.6. Manage index patterns

8.1.6.1. Set the default index pattern

The first index pattern you create is automatically designated as the default pattern. You can click the star icon in the upper-right corner of the index pattern to reset the default pattern.

Prerequisites

An index pattern is created. For more information about how to create an index pattern, see Create index patterns.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the ⚙ icon.
3. In the **Kibana** section, click **Index Patterns**.
4. In the left-side index pattern list, select an index pattern that you want to set as the default pattern.

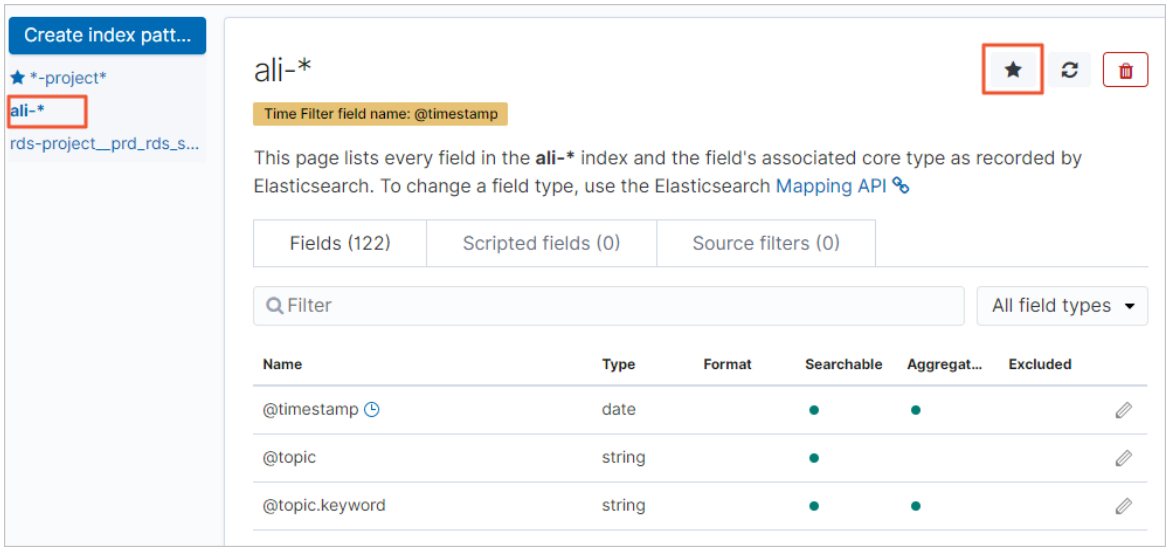

5. In the upper-right corner of the index pattern, click the ★ icon.

Result

After an index pattern is set as the default pattern, a star icon is displayed to the left of the default pattern name.

8.1.6.2. Delete index patterns

You can remove an index pattern from the list of Saved Objects in Kibana. After you delete an index pattern, you will not be able to restore field formatters, scripted fields, source filters, and field popularity data related to the index pattern.

Context

Deleting an index pattern will interrupt all visualizations, saved searches, and other saved objects that reference the pattern.

Deleting an index pattern will not remove any indexes or data documents from Elasticsearch.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the ⚙ icon.
3. In the **Kibana** section, click **Index Patterns**.
4. In the left-side navigation pane, select an index pattern to be deleted.
5. In the upper-right corner of the index pattern, click the 🗑 icon.
6. In the message that appears, click **Delete**.

8.1.7. Manage indexes

Index Management allows you to view index settings, mappings, and statistics and perform operations on indexes such as refreshing and freezing indexes, clearing the cache, and forcibly merging segments. Index Management also allows you to perform operations on multiple indexes at a time.

Procedure

1. Log on to the Kibana log O&M platform.
2. In the left-side navigation pane, click the ⚙ icon.
3. In the **Elasticsearch** section, click **Index Management**.
   Indexes are listed on the right side of the page. By default, the index list does not contain rollup indexes and system indexes. To list these indexes, you can turn on **Include rollup indices** or **Include system indices**.

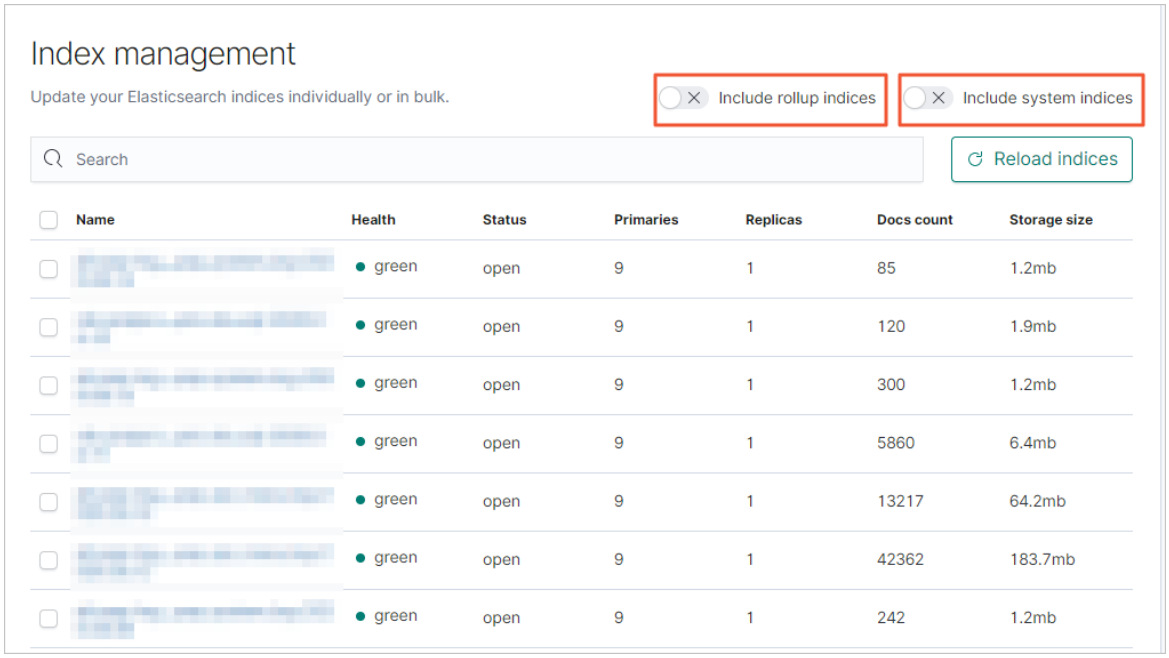

4. You can perform the following operations:

- View brief information about an index

  In the right-side index list, click the name of the index that you want to view. You can view brief information about the index on the **Summary** tab.

- View index settings

  In the right-side index list, select the index that you want to view, click **Manage index**, and then select **Show index settings** from the INDEX OPTIONS drop-down list. On the **Settings** tab that appears, view the configuration information of the index.

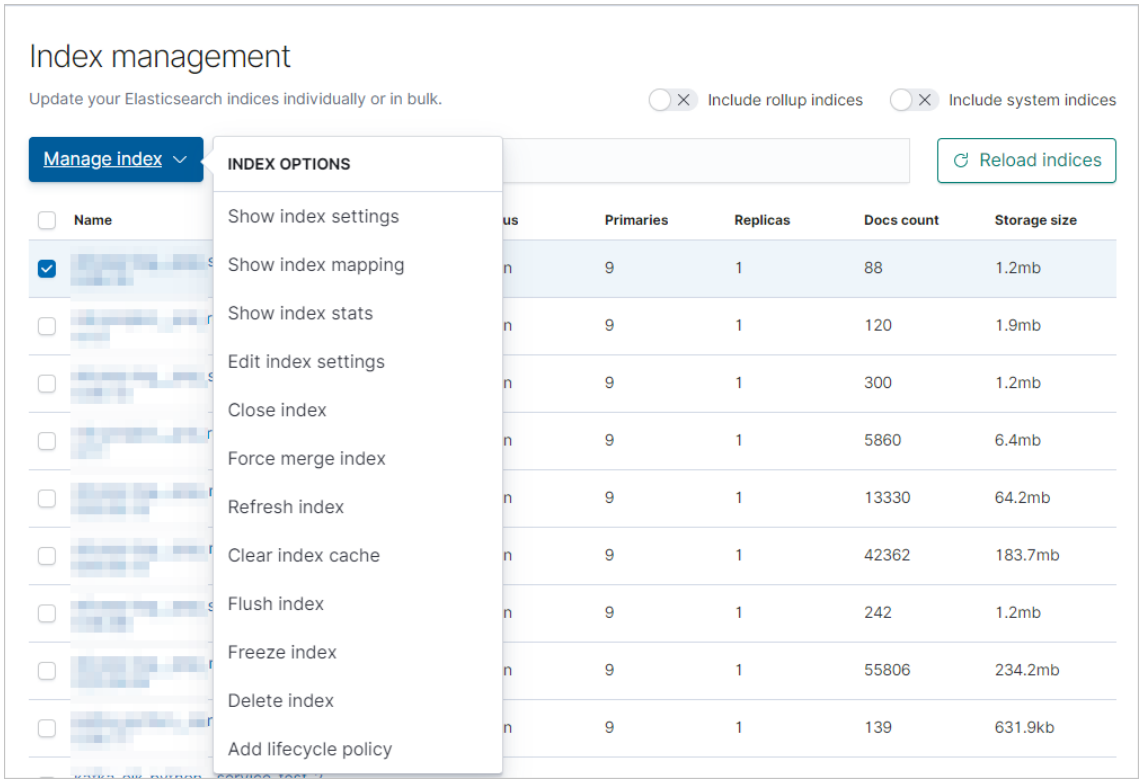

- View index mapping

  In the right-side Index list, select the index that you want to view, click **Manage index**, and then select **Show index mapping** from the INDEX OPTIONS drop-down list. On the **Mapping** tab that appears, view the mapping information of the index.

- View index statistics

  In the right-side index list, select the index that you want to view, click **Manage index**, and then select **Show index stats** from the INDEX OPTIONS drop-down list. On the **Stats** tab that appears, view the statistics of the index.

- Edit index settings

  In the right-side Index list, select the index that you want to view, click **Manage index**, and then select **Edit index settings** from the INDEX OPTIONS drop-down list. On the **Edit settings** tab that appears, modify the index settings and click **Save**.

- Close indexes

  When an index is closed, read and write operations are not allowed on the index. Closed indexes in clusters do not consume any resources other than disk space. If you reopen a closed index, it will go through the normal recovery process.

In the right-side index list, select one or more indexes to close, click **Manage index**, and then select **Close index** from the INDEX OPTIONS drop-down list.

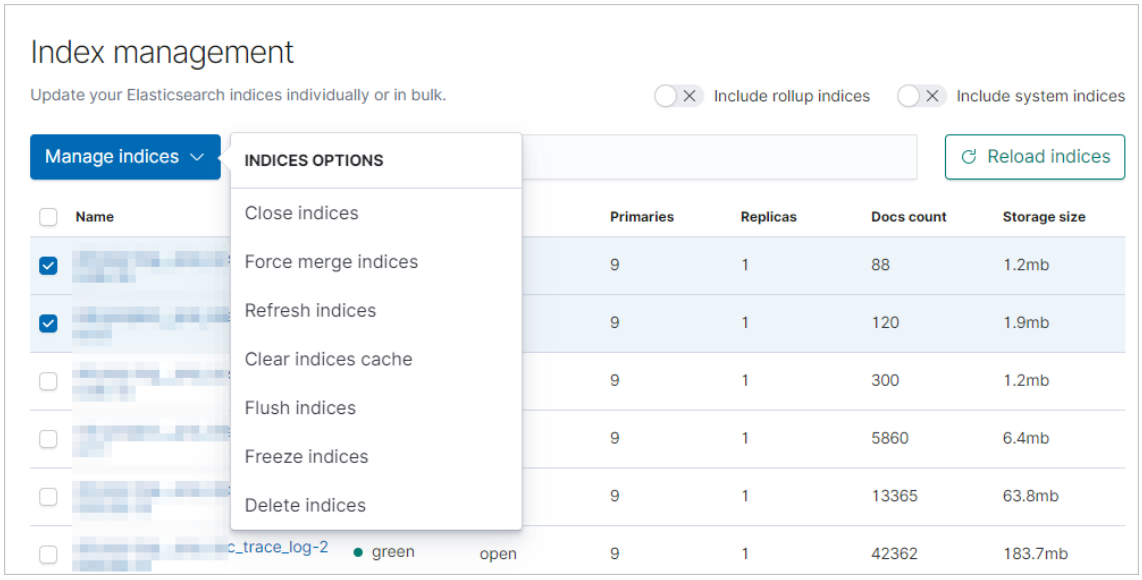

- Forcibly merge indexes

  You can forcibly merge an index by merging small files and clearing deleted files to reduce the number of segments in a shard. Only read-only indexes can be forcibly merged.

  In the right-side index list, select one or more indexes, click **Manage index**, and then select **Force merge index** from the INDEX OPTIONS drop-down list.

- Refresh indexes

  You can refresh indexes to write operations in the indexing buffer to the file system cache. This action is automatically performed once every second.

  In the right-side index list, select one or more indexes, click **Manage index**, and then select **Refresh index** from the INDEX OPTIONS drop-down list.

  > **Note** We recommend that you do not manually refresh indexes because this may affect performance.

- Clear index cache

  You can clear all caches related to indexes.

  In the right-side index list, select one or more indexes, click **Manage index**, and then select **Clear index cache** from the INDEX OPTIONS drop-down list.

- Flush indexes

  In the right-side index list, select one or more indexes to flush, click **Manage index**, and then select **Flush index** from the INDEX OPTIONS drop-down list.

- Freeze indexes

  Frozen indexes are read-only indexes whose shards have been moved to disks to reduce memory usage. Frozen indexes can still be queried, but will take longer.

  In the right-side index list, select one or more indexes to freeze, click **Manage index**, and then select **Freeze index** from the INDEX OPTIONS drop-down list.

- Delete indexes

  When you delete an index, the index and all of its documents are deleted permanently.

  In the right-side index list, select one or more indexes to delete, click **Manage index**, and then select **Delete index** from the INDEX OPTIONS drop-down list.

- Add a lifecycle policy

  > **Note** Before you add a lifecycle policy to an index, ensure that you have created an index lifecycle policy. For more information about how to create an index lifecycle policy, visit Creating an index lifecycle policy.

  a. In the right-side index list, select the index that you want to modify, click **Manage index**, and then select **Add lifecycle policy** from the INDEX OPTIONS drop-down list.

  b. In the dialog box that appears, select a lifecycle policy and click **Add policy**.

8.2. Kafka Manager

8.2.1. What is Kafka Manager?

This topic describes features of Kafka Manager.

Kafka Manager is a web-based management system for Kafka. You can use Kafka Manager to manage Kafka clusters that consist of the logs from base modules. You can use Kafka Manager to perform visualized O&M operations on Kafka clusters of Log Service. For example, you can manage topics, brokers, and consumers in Kafka clusters.

8.2.2. Log on to Kafka Manager

This topic describes how to log on to Kafka Manager.

Prerequisites

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

Procedure

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

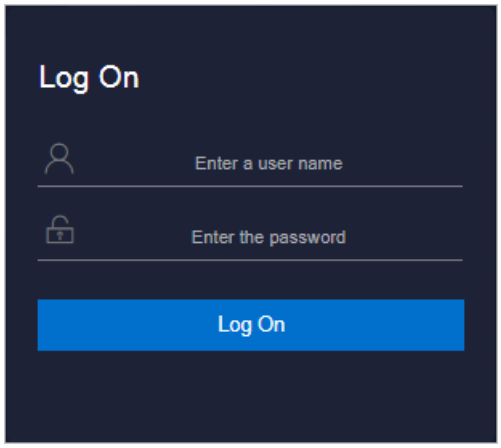

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**.
6. In the **Apsara Stack > Basic O&M** section, click **Platform Log Management**.

8.2.3. Quick start

8.2.3.1. Create a Kafka cluster

If you use Kafka Manager for the first time, you must create a Kafka cluster.

Context

Procedure

1. In the top navigation bar of the Kafka Manager homepage, select **Add Cluster** from the **Cluster** drop-down list.

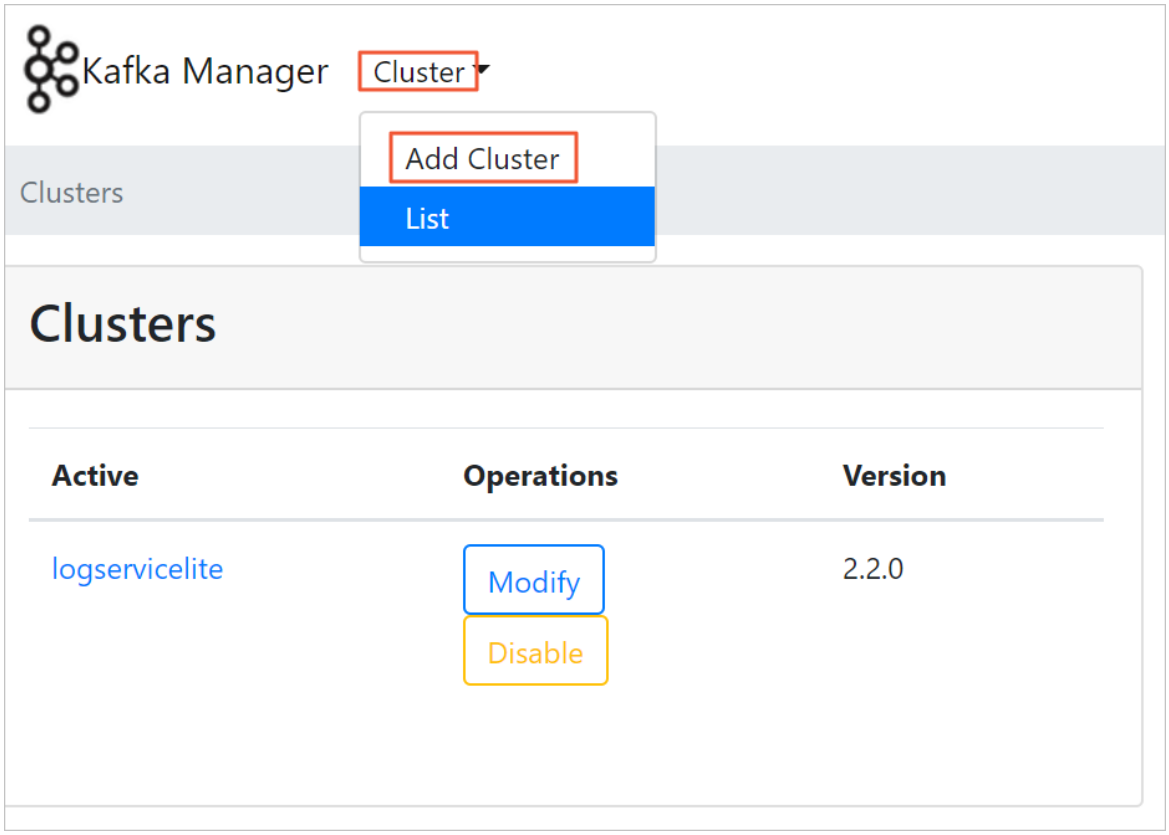

2. On the **Add Cluster** page, configure cluster parameters.

Field	Description
Cluster Name	Required. The name of the cluster to be created. After a cluster is created, its name cannot be changed.
ClusterZookeeper Hosts	Required. The connection string of a ZooKeeper node in the cluster. Format: ip:port.
Kafka Version	The version of Apache Kafka. We recommend that you select the latest version.
Enable JMX Polling	We recommend that you select this option. Java Management Extensions (JMX) is a Java technology that supplies tools for monitoring and managing system objects, such as clearing cache and reloading the configuration file.
JMX Auth Username	The username for JMX authorization.
JMX Auth Password	The password for JMX authorization.
JMX with SSL	You can select this option to enable JMX over SSL.
Enable Logkafka	You can select this option to enable Logkafka. Logkafka sends log file content to Kafka line by line.

Field	Description
Poll consumer information	You can select this option to view consumer information in a Kafka cluster on the **Consumers** page. We recommend that you do not select this option if the number of the consumers is large.
Filter out inactive consumers	You can select this option to filter out inactive consumers.
Enable Active OffsetCache	You can select this option to enable active offset cache.
Display Broker and Topic Size	You can select this option to view the numbers of brokers and topics in the cluster.
brokerViewUpdatePeriodSeconds	The update period for broker views. Unit: seconds.
clusterManagerThreadPoolsize	The size of the thread pool for the cluster manager.
clusterManagerThreadPoolQueueSize	The size of the thread pool queue for the cluster manager.
kafkaCommandThreadPoolSize	The size of the thread pool for Kafka commands.
kafkaCommandThreadPoolQueueSize	The size of the thread pool queue for Kafka commands.
logkafkaCommandThreadPoolSize	The size of the thread pool for Logkafka commands.
logkafkaCommandThreadPoolQueueSize	The size of the thread pool queue for Logkafka commands.
logkafkaUpdatePeriodSeconds	The update period for Logkafka. Unit: seconds.
partitionOffsetCacheTimeoutSecs	The timeout period of the partition offset cache. Unit: seconds.
brokerViewThreadPoolSize	The size of the thread pool for broker views.
brokerViewThreadPoolQueueSize	The size of the thread pool queue for broker views.
offsetCacheThreadPoolSize	The size of the thread pool for the offset cache.
offsetCacheThreadPoolQueueSize	The size of the thread pool queue for the offset cache.
kafkaAdminClientThreadPoolSize	The size of the thread pool for the Kafka administrator client.

Field	Description
kafkaAdminClientThreadPoolQueueSize	The size of the thread pool queue for the Kafka administrator client.
kafkaManagedOffsetMetadataCheckMillis	The check period for the offset metadata.
kafkaManagedOffsetGroupCacheSize	The size of the offset group cache.
kafkaManagedOffsetGroupExpireDays	The expiration period of the offset group.
Security Protocol	The security protocol.
SASL Mechanism (only applies to SASL based security)	The simple authentication and security layer (SASL) mechanism that only applies to SASL-based security.
SASL JAAS Config (only applies to SASL based security)	The SASL-based Java Authentication and Authorization Service (JAAS) configurations, including the username and password. Example: **org.apache.kafka.common.security.plain.PlanLoginModule required username=username password=password;** **Note** To mask sensitive information, developers configure this parameter within Apsara Stack Agility SE. You do not need to configure this parameter when you create a Kafka cluster.

3. Click **Save**.

Result

After a Kafka cluster is created, you can view the Kafka cluster in the cluster list of the Kafka Manager homepage.

What's next

To modify the configuration information of a Kafka cluster, perform the following steps: Go to the cluster list of the Kafka Manager homepage. Find the target cluster and click **Modify** in the **Operations** column.

8.2.3.2. View topics in a Kafka cluster

You can view the topics in a created Kafka cluster.

Prerequisites

A Kafka cluster is created. For more information about how to create a Kafka cluster, see Create a Kafka

cluster.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the target cluster and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.

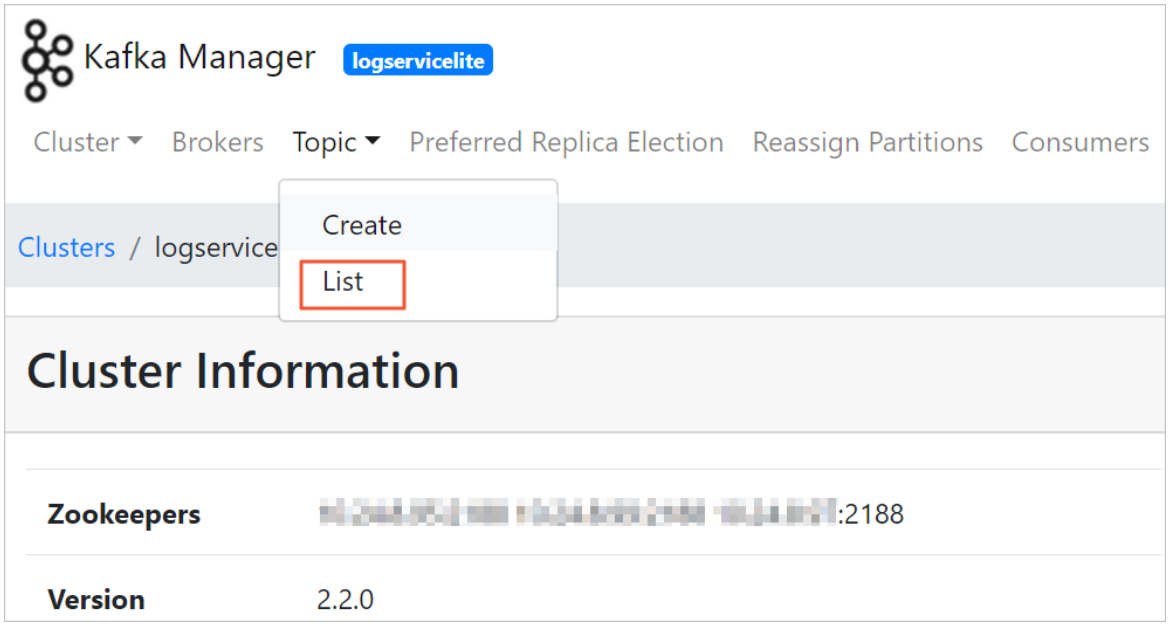

3. In the **Topics** section, view the number of topics in the cluster, the number of messages in each topic, the rate at which messages are generated by each topic, and the number of Kafka nodes in which messages are stored.
4. Click a topic name to view its details.

8.2.3.3. View consumers in a Kafka cluster

If you select **Poll consumer information** when creating a cluster, you can view the topics that each consumer consumes on the cluster details page.

Context

Kafka consumers read and consume messages from Kafka servers. Kafka consumers are also known as message subscribers or message consumers.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the target cluster.
2. Click the cluster name. The cluster details page appears.
3. In the top navigation bar, click **Consumers**.

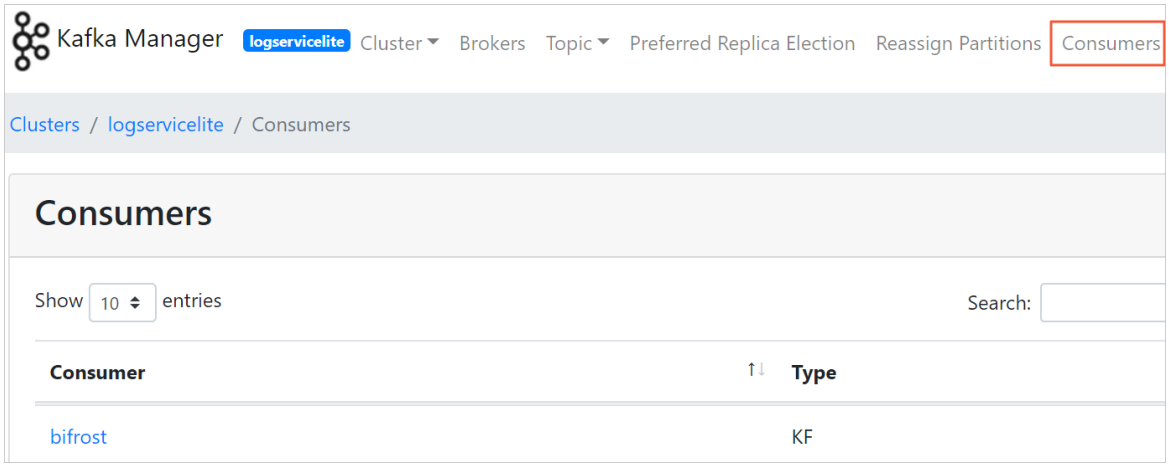

4. View all consumers in the cluster and their types.
5. Click a consumer and view the topics consumed by the consumer.

8.2.3.4. View brokers in a Kafka cluster

A broker is an independent Kafka server that runs in a Kafka Cluster. After a Kafka cluster is created, you can view the brokers in the Kafka cluster.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the target cluster.
2. Click the cluster name. The cluster details page appears.
3. In the top navigation bar, click **Brokers**.

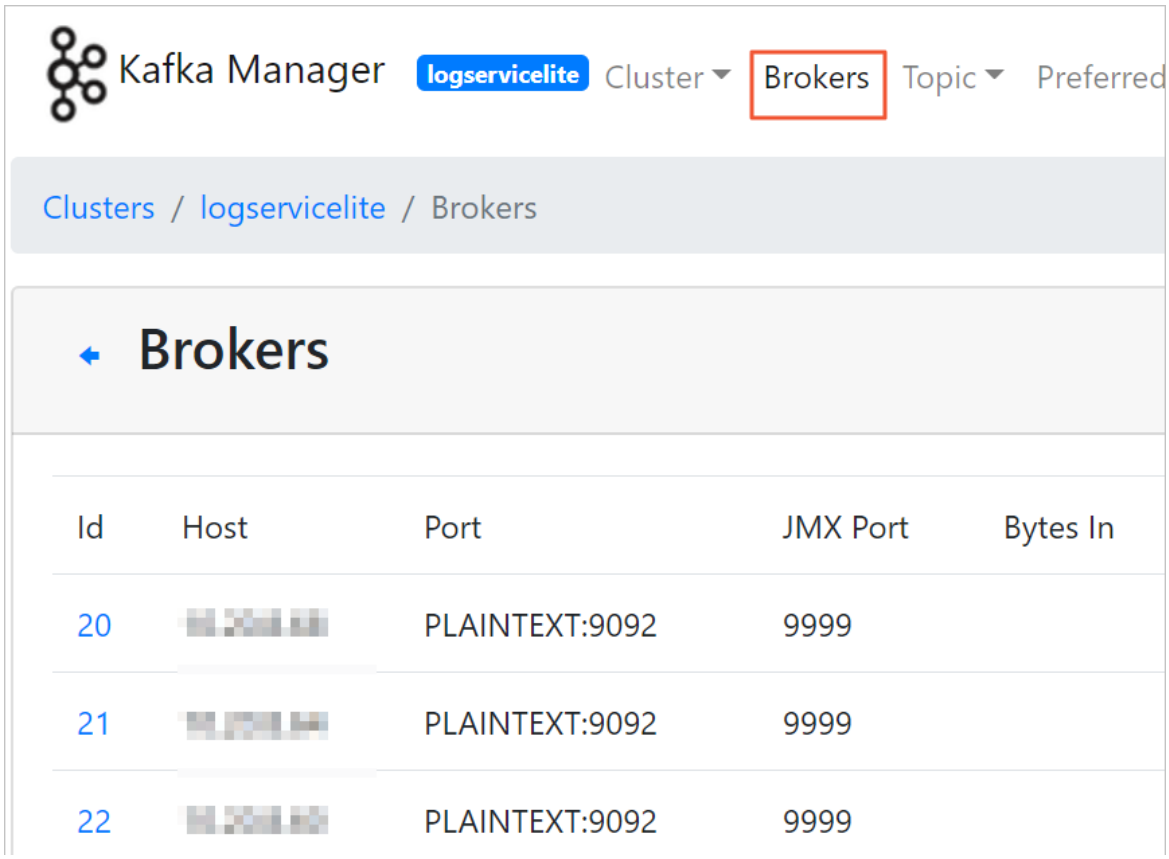

4. View information of brokers in a Kafka cluster, such as IDs, hosts, and ports.
5. Click a broker ID, and view details about the broker.

8.2.4. Kafka clusters

You can view, modify, disable, enable, and delete a Kafka cluster.

8.2.4.1. View a Kafka cluster

You can view topics and brokers in a Kafka cluster based on your requirements.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the target cluster.
2. Click the cluster name. The cluster details page appears.
3. In the **Cluster Summary** section, view the numbers of topics and brokers in the cluster.
4. Click the number following **Topics** or **Brokers** to view the information of all topics or brokers.

8.2.4.2. Disable or enable a Kafka cluster

This topic describes how to disable or enable a Kafka cluster based on your requirements.

Disable a Kafka cluster

After you create a Kafka cluster, the cluster is enabled by default. You can disable the Kafka cluster if you do not need to use it.

1. In the cluster list of the Kafka Manager homepage, find the target cluster.
2. Click **Disable** in the **Operations** column.

Enable a Kafka cluster

If you want to use a disabled cluster, enable it first.

1. In the cluster list of the Kafka Manager homepage, find the target cluster.
2. Click **Enable** in the **Operations** column.

8.2.4.3. Delete a Kafka cluster

This topic describes how to delete a Kafka cluster that you no longer use.

Prerequisites

The Kafka cluster is disabled. For more information about how to disable a Kafka cluster, see Disable a Kafka cluster.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster you want to delete.
2. Click **Delete** in the **Operations** column.

8.2.5. Topics

A Kafka topic is a unique string associated with a message type. Messages in Kafka are classified by topic. A topic is divided into one or more partitions distributed across one or more brokers.

8.2.5.1. Create a Kafka topic

This topic describes how to create a Kafka topic.

Context

> **Note** In Apsara Stack Agility SE, topic information is generated when Log Service is used. This function is not common in O&M scenarios and is for reference only.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster for which you want to create a topic, and click the cluster name.
2. In the top navigation bar, select **Create** from the **Topic** dropdown list.

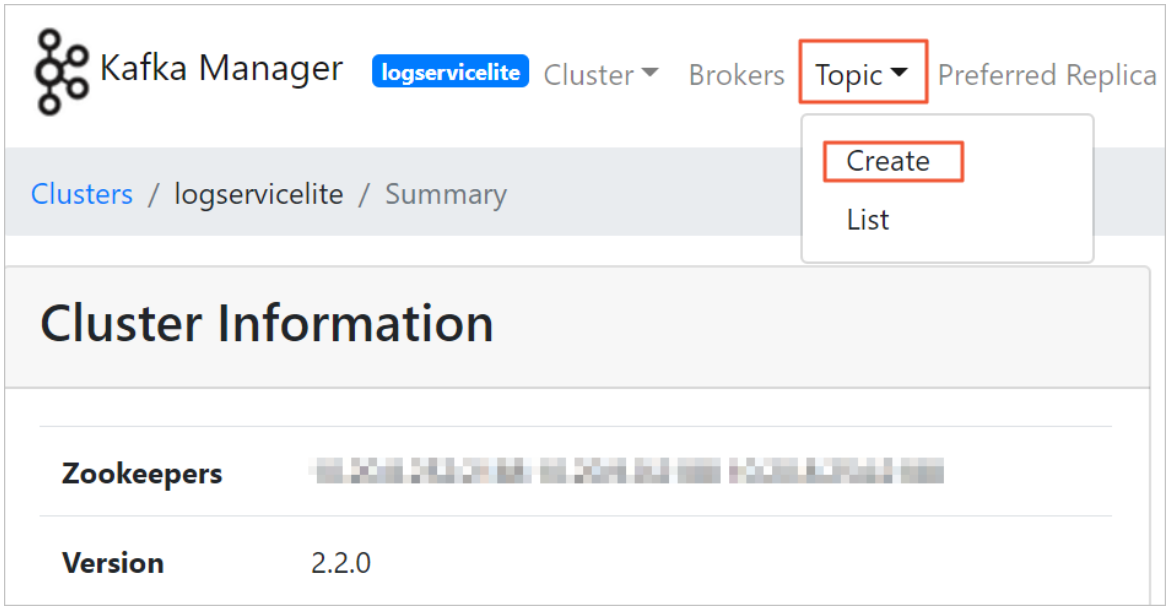

3. On the **Create Topic** page, configure the topic parameters.

Field	Description
Topic	The name of the topic.
Patitions	The number of partitions within the topic. The value must be an integer greater than 1. An appropriate partitioning strategy leads to higher throughput.
Replication Factor	The number of replicas. Replicas ensures the high availability of Kafka clusters.

4. Click **Create**.

8.2.5.2. Generate partition assignments

This topic describes how to generate partition assignments for one or more topics. You can move multiple partitions at a time to reassign partitions. For example, you can reassign partitions to specified brokers when the number of brokers increases in a cluster or the number of partitions increases in a topic.

Context

Each topic has one or more partitions.

> **Note** In Apsara Stack Agility SE, topic information is generated when Log Service is used. This function is not common in O&M scenarios and is for reference only.

Reassign partitions for a single topic

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the topic for which you want to reassign partitions and click its name.
4. In the **Operations** section, click **Generate Partition Assignments**.
5. On the **Confirm Assignment** page, select a broker.

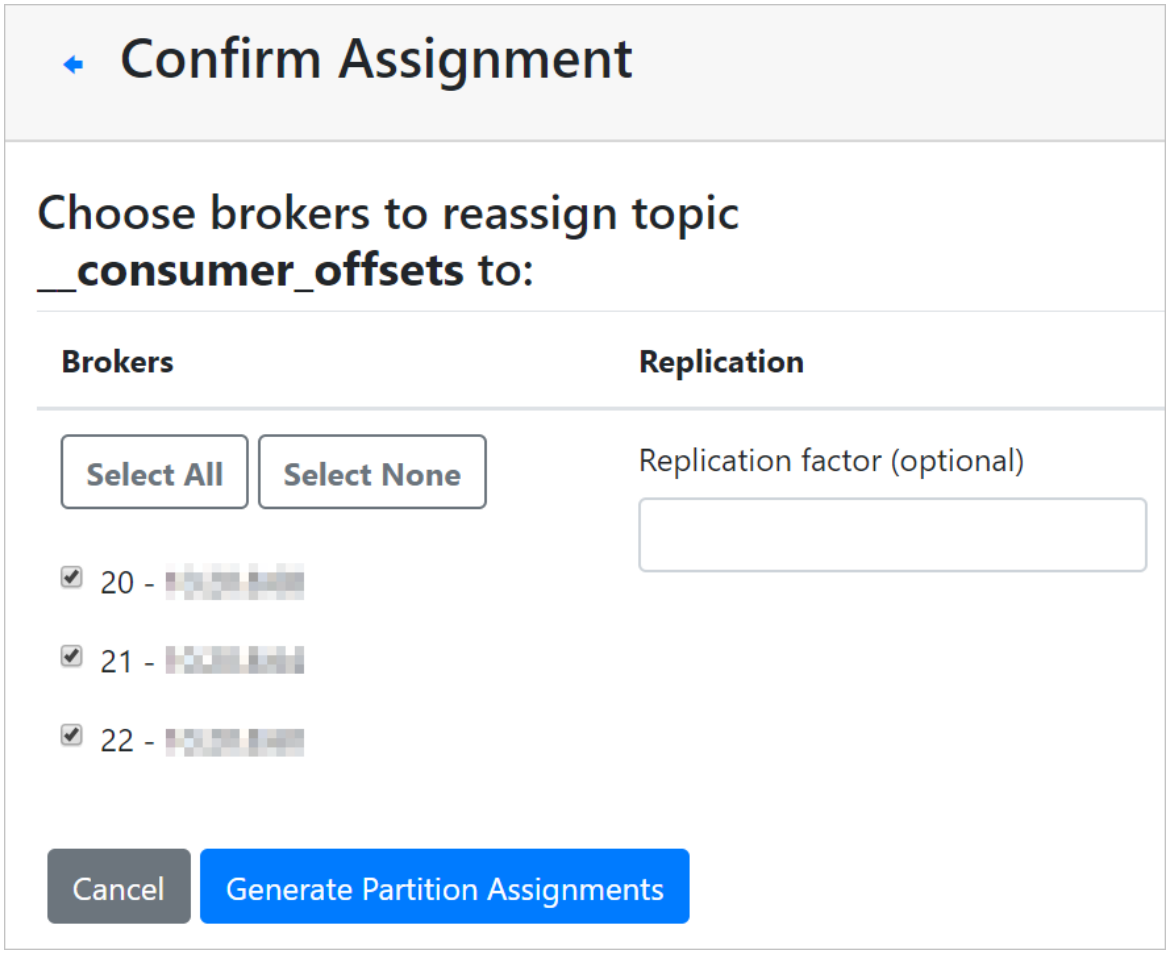

6. Click **Generate Partition Assignments**.

Reassign partitions for multiple topics at a time

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Operations** section, click **Generate Partition Assignments**.
4. On the **Confirm Assignments** page, select the topics for which you want to reassign partitions from the Topics list and corresponding brokers from the Brokers list.
5. Click **Generate Partition Assignments**.

8.2.5.3. Add partitions

You can add partitions for one or more topics.

Context

> **Note** In Apsara Stack Agility SE, topic information is generated when Log Service is used. This function is not common in O&M scenarios and is for reference only.

Add partitions for a single topic

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the topic to which you want to add partitions and click its name.
4. In the **Operations** section, click **Add Partitions**.

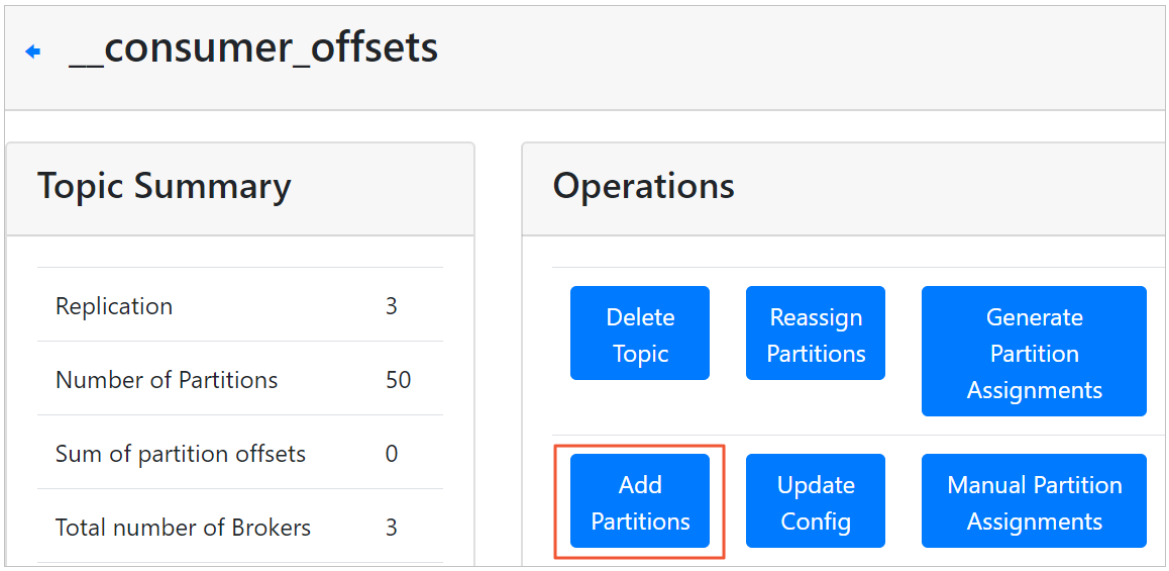

5. On the **Add Partitions** page, select the broker and modify the number of partitions.

> **Note** The **Partitions** field specifies the total number of partitions. The new number of partitions must be greater than the original number.

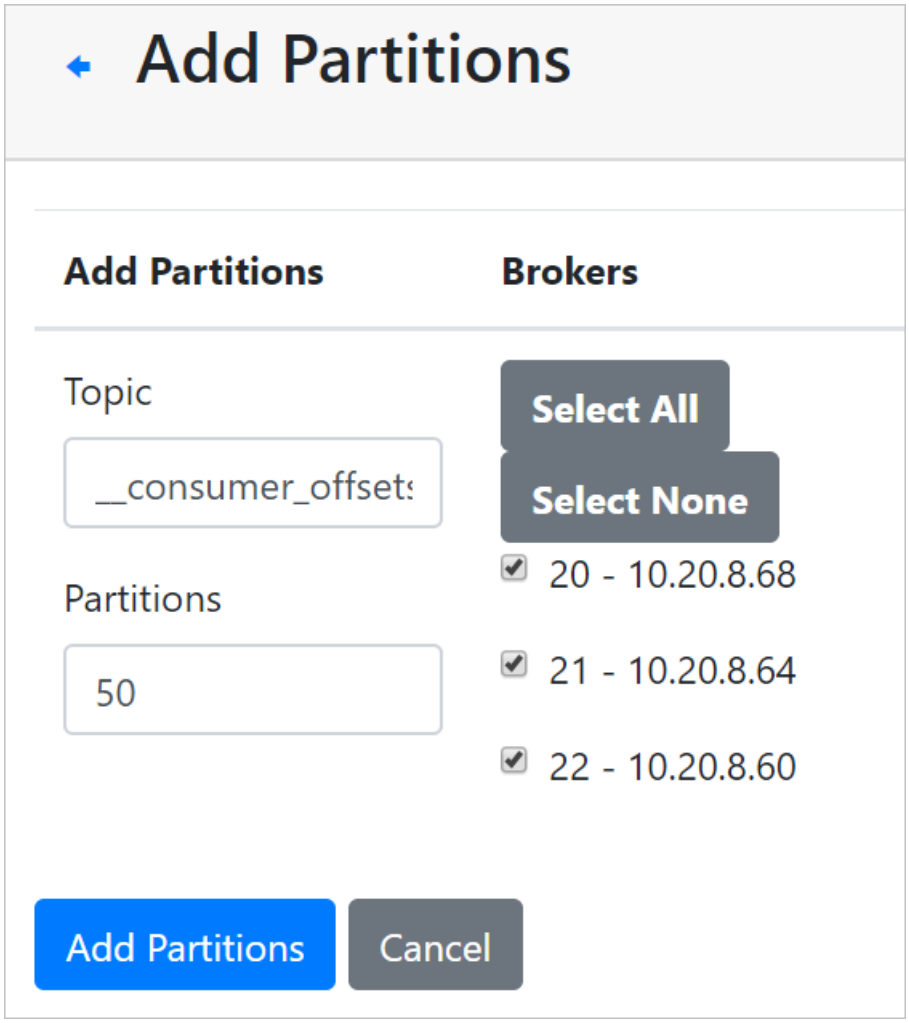

6. Click **Add Partitions**.

Add partitions for multiple topics

You can add partitions for multiple topics at a time. The new number of partitions must be greater than the maximum number of partitions of a topic among all topics.

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Operations** section, click **Add Partitions**.

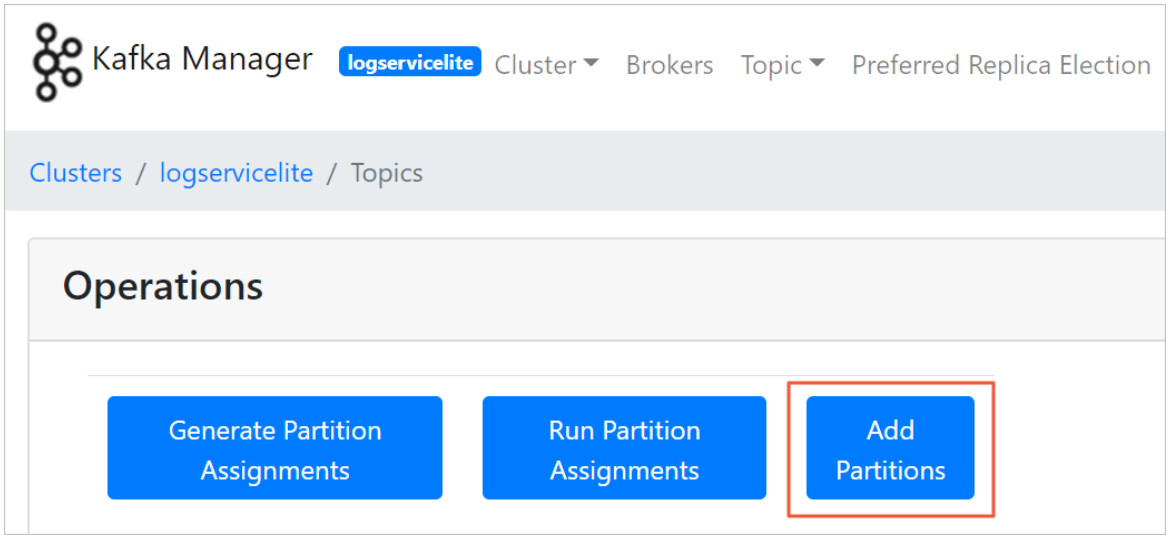

4. In the **Partitions** field on the **Add Partitions** page, enter the number of partitions and select the topics and brokers for which you want to add partitions.

   > **Note** The new number of partitions must be greater than the maximum number of partitions of a topic among all topics.

5. Click **Add Partitions**.

8.2.5.4. Run partition assignments

This topic describes how to run partition assignments.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Operations** section, click **Run Partition Assignments**.
4. On the **Run Assignments** page, select the topic for which you want to reassign partitions. Click **Run Partition Assignments**.

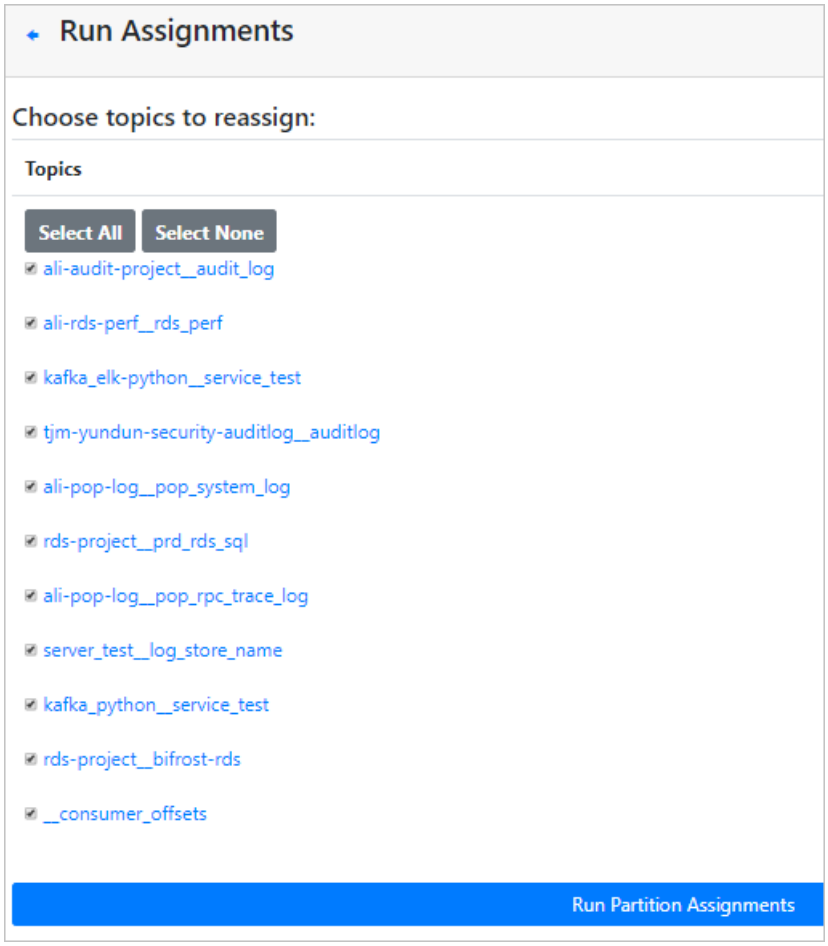

8.2.5.5. Reassign partitions

This topic describes how to reassign partitions. To balance cluster loads, you can reassign partitions to elect a new leader for assigned replicas.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the target topic and click its name.
4. In the **Operations** section, click **Reassign Partitions**.

8.2.5.6. Update configurations for a topic

This topic describes how to update configurations for a Kafka topic.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the target topic and click its name.

4. In the **Operations** section, click **Update Config**.
5. On the **Update Config** page, update the configurations for the topic.

Field	Description	Default value
Topic	The name of the topic.	N/A
cleanup.policy	The cleanup policy on old log segments. Valid values: ○ delete: deletes log segments. ○ compact: compresses log segments.	delete
compression.type	The compression type specified for the topic. Valid values: ○ **gzip** ○ **snappy** ○ **lz4** ○ **uncompressed** ○ **producer**: specifies the original compression encoder set by the producer.	producer
delete.retention.ms	The maximum amount of time to retain compressed log data. Unit: milliseconds.	86400000, which is one day
file.delete.delay.ms	The time to wait before deleting a file from the file system. Unit: milliseconds.	60000, which is one minute
flush.messages	The number of messages written to a log partition before you force an `fsync` on the log. We recommend that you do not modify the default value.	9223372036854775807
flush.ms	The time interval at which you force an fsync of data written to the log. We recommend that you do not modify the default value.	9223372036854775807
follower.replication.throttled.replicas	The list of replicas for which log replication must be throttled on the follower side.	N/A

Field	Description	Default value
index.interval.bytes	The frequency at which Kafka adds an index entry to its offset index. We recommend that you do not modify the default value.	4096
leader.replication.throttled.replicas	The list of replicas for which log replication must be throttled on the leader side.	N/A
max.message.bytes	The largest record batch size allowed by Kafka.	1000012, which is about 1 MB
message.downconversion.enable	Specifies whether down-conversion of message formats is enabled to satisfy the requests from consumers. When this parameter is set to false, the broker will not perform down-conversion for consumers that are configured to receive older message format. The broker returns the UNSUPPORTED_VERSION error to these the requests from consumers. This configuration does not apply if the message conversion format is required to enable replication to followers.	true
message.format.version	Specifies the message format version that the broker uses to append messages to the log. The value must be a valid API version.	N/A
message.timestamp.difference.max.ms	The maximum difference allowed between the timestamp when the broker receives a message and the timestamp specified in the message. When message.timestamp.type is set to CreateTime, a message will be rejected if the difference in timestamp exceeds this threshold. When message.timestamp.type is set to LogAppendTime, this parameter is ignored.	9223372036854775807

Field	Description	Default value
message.timestamp.type	The timestamp type in the message. The value can be CreateTime or LogAppendTime.	CreateTime
min.cleanable.dirty.ratio	The frequency at which the log compactor will attempt to clean the log.	0.5
min.compaction.lag.ms	The minimum amount of time a message will remain uncompressed in the log. This parameter is only applicable to logs that are being compressed.	0
min.insync.replicas	The minimum number of in-sync replicas that acknowledge a write operation to be considered successful.	1
preallocate	Specifies whether to preallocate a file on the disk when you create a new log segment.	false
retention.bytes	The maximum size of a partition. If the maximum size is reached, you can discard old log segments to free up space by using the "delete" retention policy.	-1
retention.ms	The maximum amount of time to retain a log. If the maximum time is reached, you can discard old log segments by using the "delete" retention policy. Unit: milliseconds.	604800000, which is seven days
segment.bytes	The segment file size for the log.	1073741824, which is 1 GB
segment.index.bytes	The size of the index that maps offsets to file positions.	10485760, which is 10 MB
segment.jitter.ms	The maximum random jitter subtracted from the scheduled segment roll time to avoid thundering herds of segment rolling.	0

Field	Description	Default value
segment.ms	The period of time after which Kafka will force the log to roll even if the segment file is not full. This ensures that retention can delete or compress old data. Unit: milliseconds.	604800000, which is seven days
unclean.leader.election.enable	Specifies whether an out-of-sync replica is elected as the leader when there is no live in-sync replica (ISR).	false

6. Click **Update Config**.

8.2.5.7. Manually assign partitions

The system allows you to manually assign brokers for replicas in each partition. You can assign partitions if a broker or a broker leader is skewed.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the target topic and click its name.
4. In the **Operations** section, click **Manual Partition Assignments**.
5. On the **Manual Partition Assignments** page, reassign brokers for replicas in each partition.

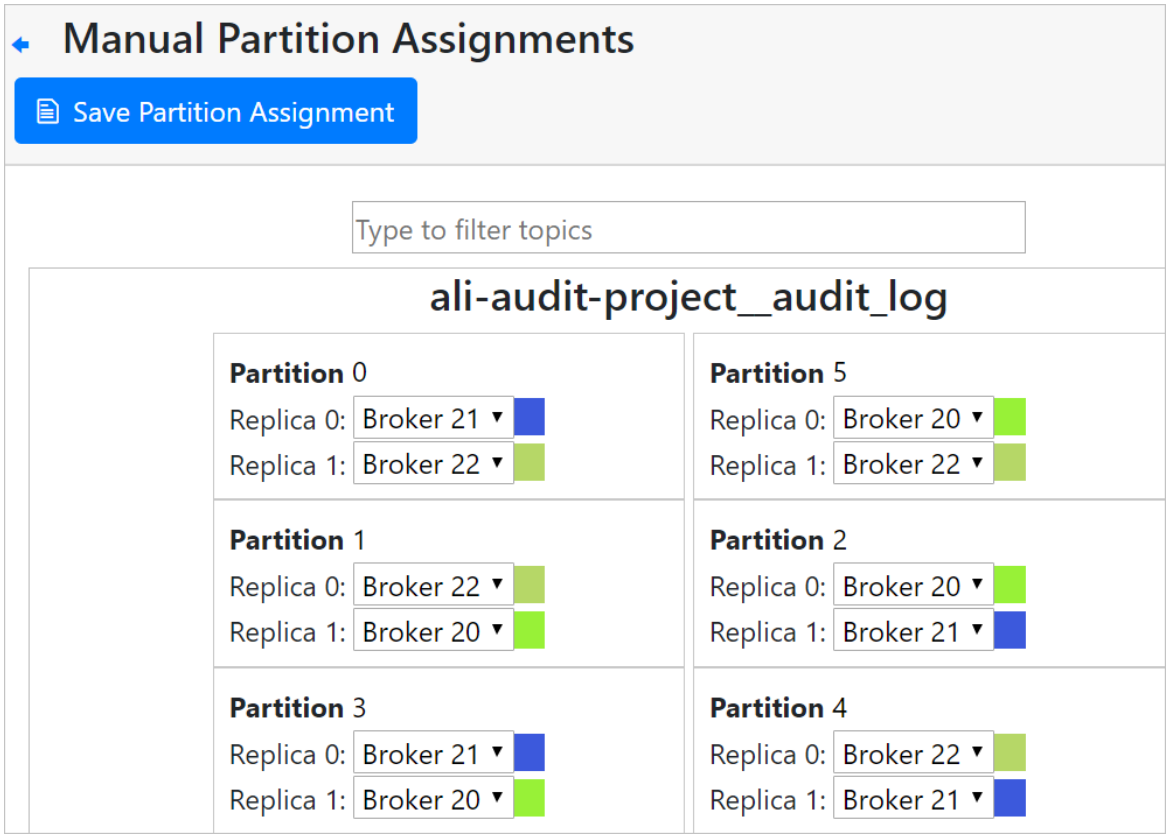

6. Click **Save Partition Assignments**.

8.2.5.8. Configure automatic partition assignment

The system automatically assigns brokers for replicas in each partition.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the topic and click its name.
4. In the **Operations** section, click **Create Partition Assignments**.

8.2.5.9. Delete a topic

This topic describes how to delete a topic that you no longer need.

Procedure

1. In the cluster list of the Kafka Manager homepage, find the cluster to which the target topic belongs, and click the cluster name.
2. In the top navigation bar, select **List** from the **Topic** dropdown list.
3. In the **Topics** section, find the topic you want to delete and click its name.
4. In the **Operations** section, click **Delete Topic**.

5. On the confirmation page, click **Delete Topic**.

9.PaaS operations and maintenance

9.1. PaaS console

9.1.1. 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.

9.1.2. Log on to the PaaS console

This topic describes how to log on to the PaaS console.

Prerequisites

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

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

Procedure

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

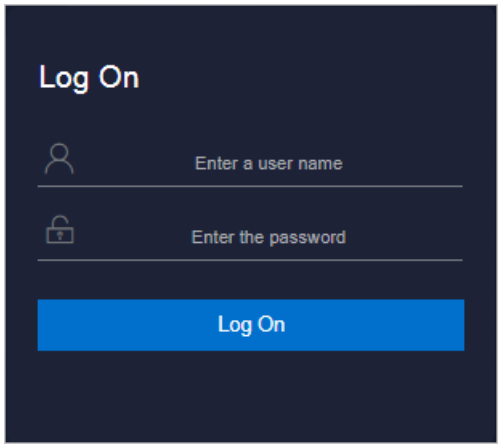

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

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

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**.
6. In the **Apsara Stack O&M > Basic O&M** section, click **PaaS Console**.

9.1.3. Platform overview

On the **Overview** page, view status statistics of platform products and clusters.

Procedure

1. Log on to the PaaS console. The **Overview** page appears by default.
2. On the **Overview** page, view status statistics of platform products and clusters.

   The page contains two sections:

   - Platform Product Overview: displays status statistics of products deployed in the PaaS console. Click the **Total** value to go to the **Product Center > Products** page.
   - Cluster Overview: displays status statistics of nodes and pods in clusters. Click a status value of **Node Status** to go to the **Clusters > Nodes** page. Click the **Not Ready** value of **Pod Status** to view error messages of abnormal pods.

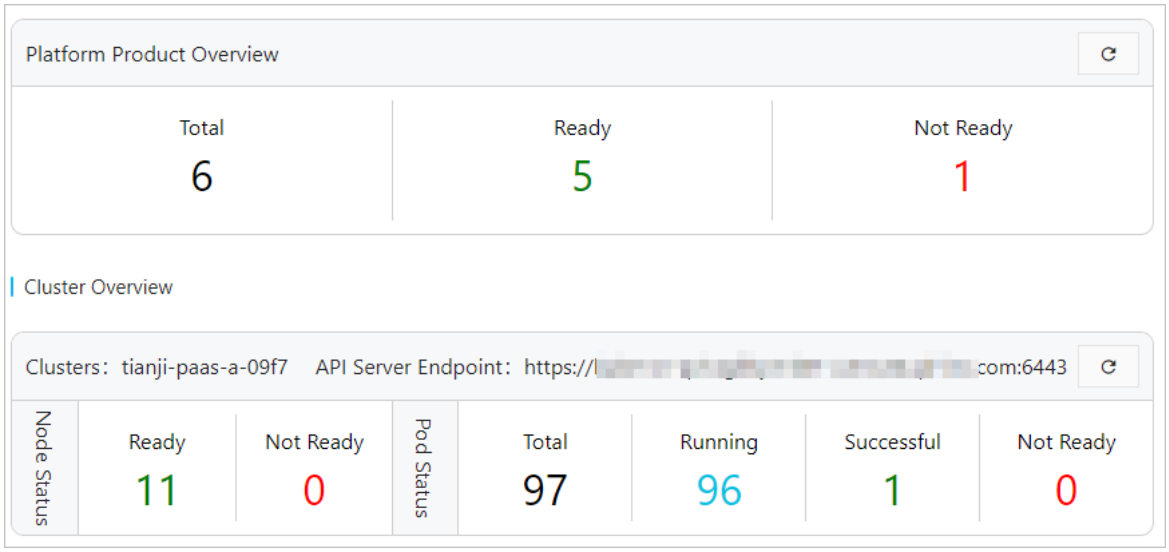

9.1.4. Clusters

9.1.4.1. View the cluster list

On the Clusters page, you can view the status and kubeconfig connection information of the clusters managed by PaaS.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Clusters** from the **Clusters** drop-down list.
2. On the **Clusters** page, view all clusters managed by PaaS.

3. Find a cluster, and then click **View** in the **Actions** column to view the kubeconfig connection information of the cluster.

9.1.4.2. Node management

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

9.1.4.2.1. Add tags

You can add tags to nodes for subsequent cluster scheduling, configuration, and behavior customization.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Select one or more nodes to which you want to add a tag. Click **Add Label** in the lower-left corner.
4. Perform the following operations:
   - Add a built-in tag

     In the **Add Label to Node** dialog box, click a tag in the Built-in Labels field. The tag name is automatically filled into the Key field. Set **Value** and then click **OK**.

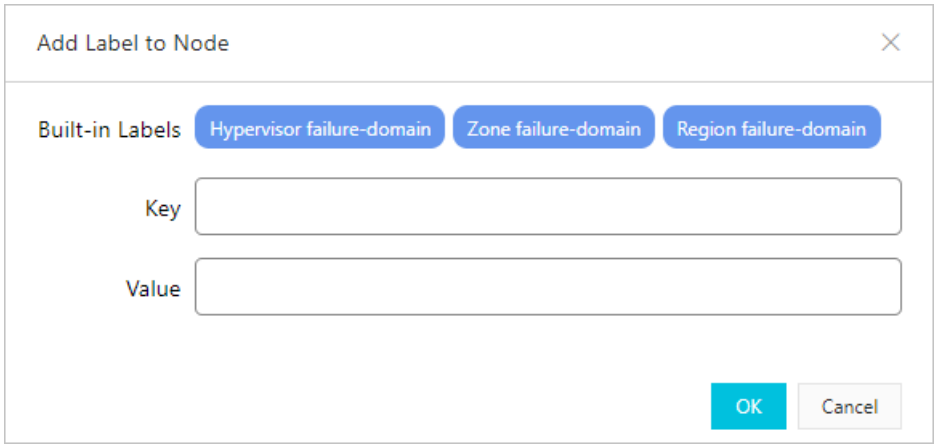

The following table describes the parameters.

Parameter	Description
Built-in Labels	Existing tags in the system. Valid values: ■ **Hypervisor failure-domain**: During output virtualization, virtual machines are distributed across different physical machines. This tag can be used to distribute pods to different physical machines. ■ **Zone failure-domain**: distributes Kubernetes nodes to different zones. ■ **Region failure-domain**: distributes Kubernetes nodes to different regions.
Key	After you click a tag in the Built-in Labels field, the tag name is automatically filled into the **Key** field. You can also set **Key** to specify a custom tag.
Value	The custom tag value.

- Add a custom tag

  In the **Add Label to Node** dialog box, set Key and Value, and then click **OK**.

9.1.4.2.2. Add taints

You can add taints to nodes for subsequent pod scheduling.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Select one or more nodes to which you want to add a taint. Click **Add Taint** in the lower-left corner.
4. Perform the following operations:
   - Add a built-in taint

In the **Add Taint** dialog box, click a taint in the Built-in Taints field. The taint name is automatically filled into the Key field. Specify **Value** and **Effect**, and then click **OK**.

The following table describes the parameters.

Parameter	Description
Built-in Taints	Existing taints in the system.
Key	After you click a taint in the Built-in Taints field, the taint name is automatically filled into the **Key** field. You can also set **Key** to specify a custom taint.
Value	The custom taint value.
Effect	The effects of the taint. Valid values: ■ **PreferNoSchedule**: indicates that if possible, pods will not schedule the node. ■ **NoSchedule**: indicates that pods will not be allowed to schedule the node. ■ **NoExecute**: indicates that pods will not be allowed to schedule the node and that pods that are running on the node will be evicted.

- Create a custom taint

  In the **Add Taint** dialog box, specify **Key**, **Value**, and **Effect**, and then click **OK**.

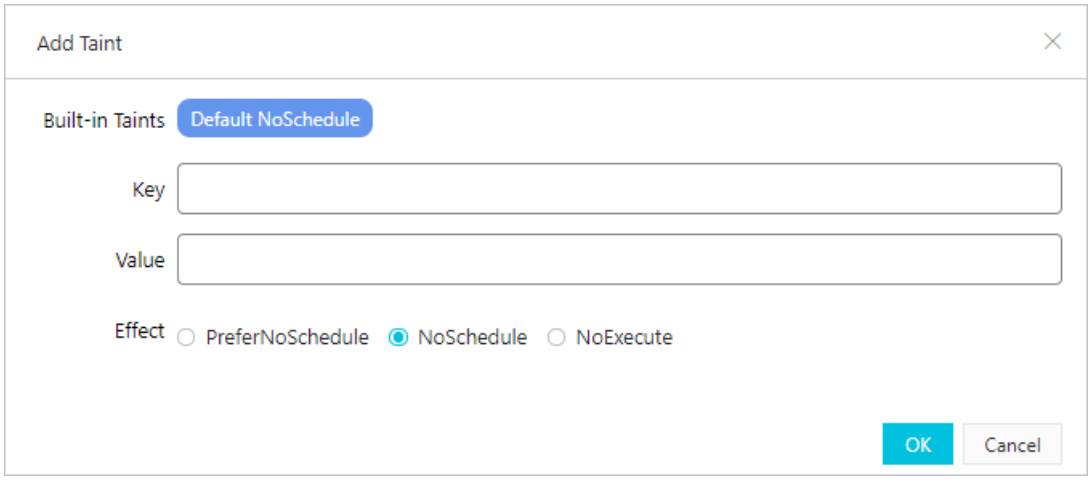

9.1.4.2.3. Query nodes by tag

You can filter nodes by tag to find nodes that have a specified tag.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. In the upper-right corner of the page, enter a tag name or specify a tag in the **key=value** format in the search box and then click the Search icon.

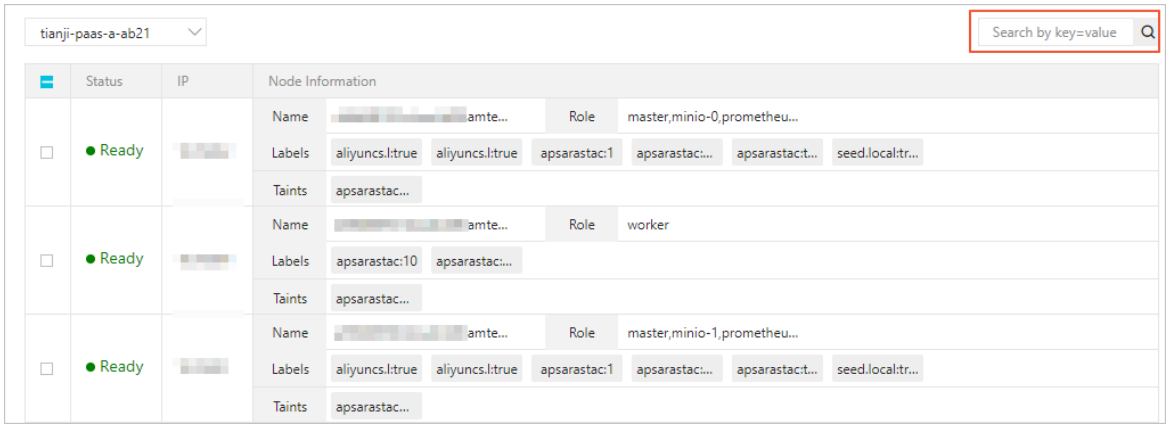

9.1.4.2.4. Delete a tag

You can delete a built-in or custom tag from a node. Kebernetes-defined tags of nodes cannot be deleted.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional)In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Find the target node and the target tag in the **Labels** row, and then click the icon.

> **Note** You can select the **Hide Kebernetes-defined Labels** check box to hide Kebernetes-defined tags, so that you can quickly find the target tag.

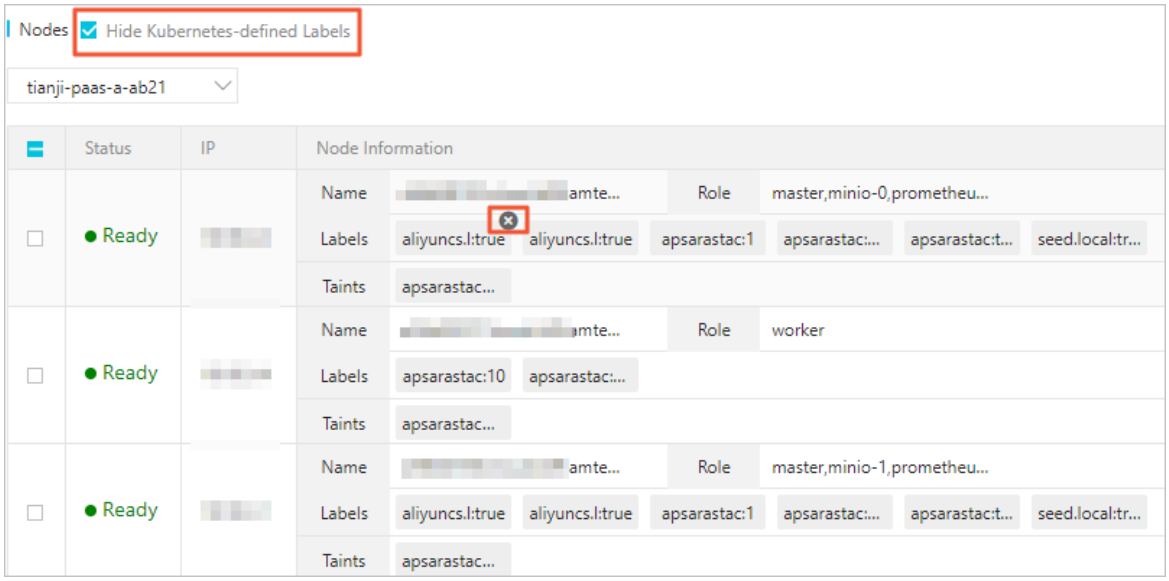

4. In the message that appears, click **OK**.

9.1.4.2.5. Delete a taint

You can delete a taint from a node.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Nodes** from the **Clusters** drop-down list.
2. (Optional) In the upper-left corner of the Nodes page, select the target cluster from the drop-down list.
3. Find the target node and the target taint in the **Taints** row, and then click the icon.

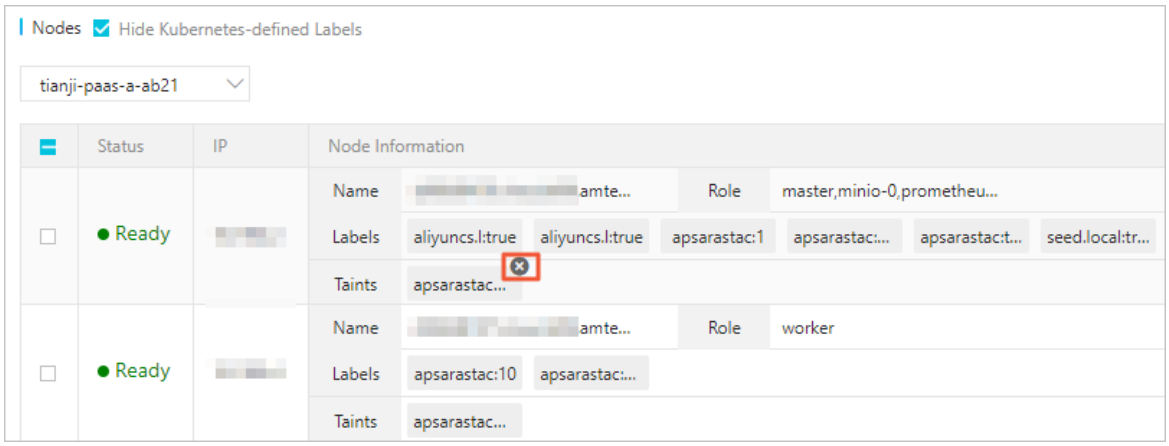

4. In the message that appears, click **OK**.

9.1.5. Product center

9.1.5.1. Product list

The product list displays the information about all products deployed in the PaaS console, including their names and versions. In the product list, you can perform O&M operations and view product resources or register variables. You can also remove products that are no longer needed.

9.1.5.1.1. View product details

You can view the details of products deployed in the PaaS console, including their names, versions, and components.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.

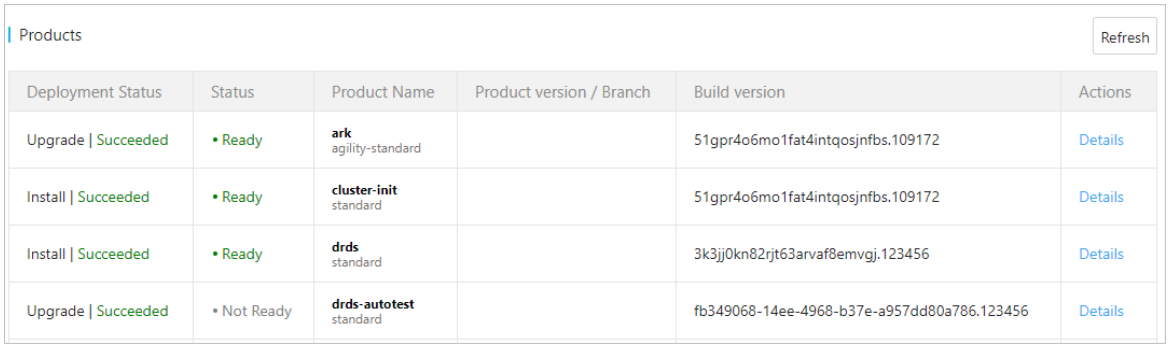

3. On the **Overview** page, view the name, version information, and components of the product.

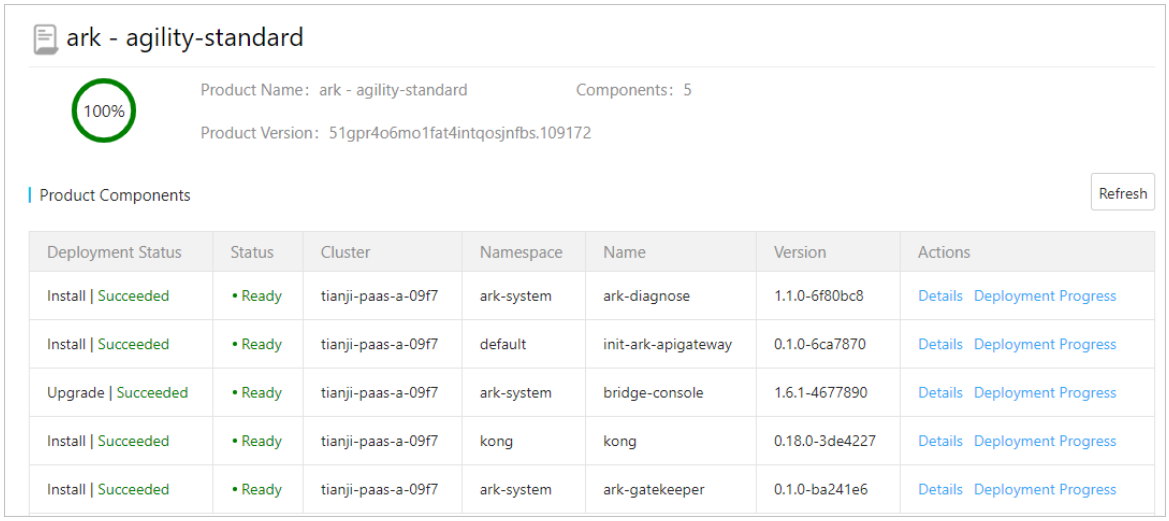

9.1.5.1.2. View component information

You can view the component details in the Product Components section of the Overview page of a product.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the Product Components section of the **Overview** page, view the deployment information of components, such as the deployment status, component status, cluster, namespace, component name, and component version.
4. Find a component and click **Details** in the **Actions** column to view details of the component.
5. The **Component Details** page contains the following tabs: StatefulSets, Deployments, DaemonSets, Jobs, Services, and Persistence Volume Claims.

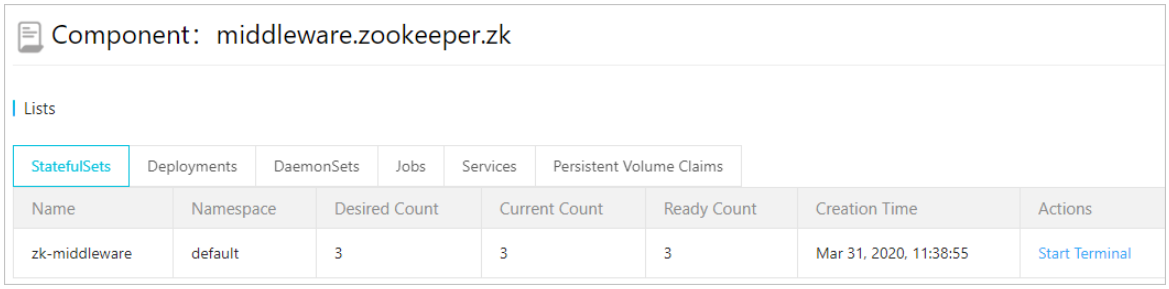

9.1.5.1.3. View the deployment progress of product components

You can view the deployment progress of product components.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the Product Components section of the **Overview** page, find the target component and click **Deployment Progress** in the **Actions** column.

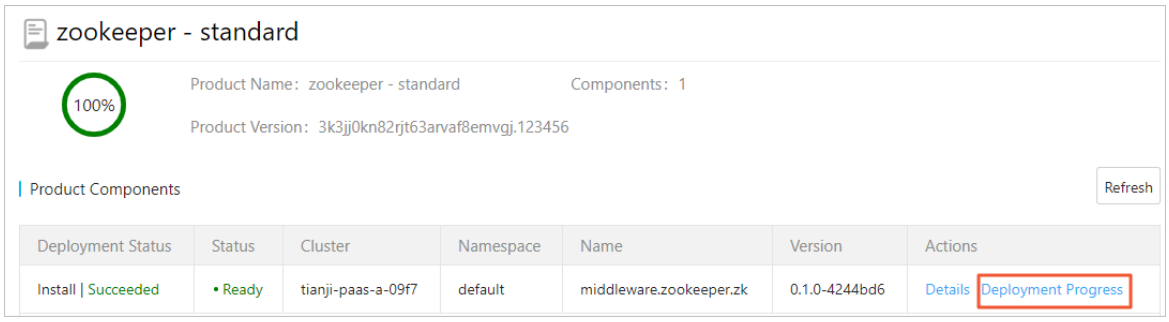

4. On the **Product Component Deployment Progress** page, click the deployment nodes in sequence to view the deployment progress and logs of the current component.

> **Note** You can click **LOGS** in the lower-left corner of the Summary tab to view the deployment logs.

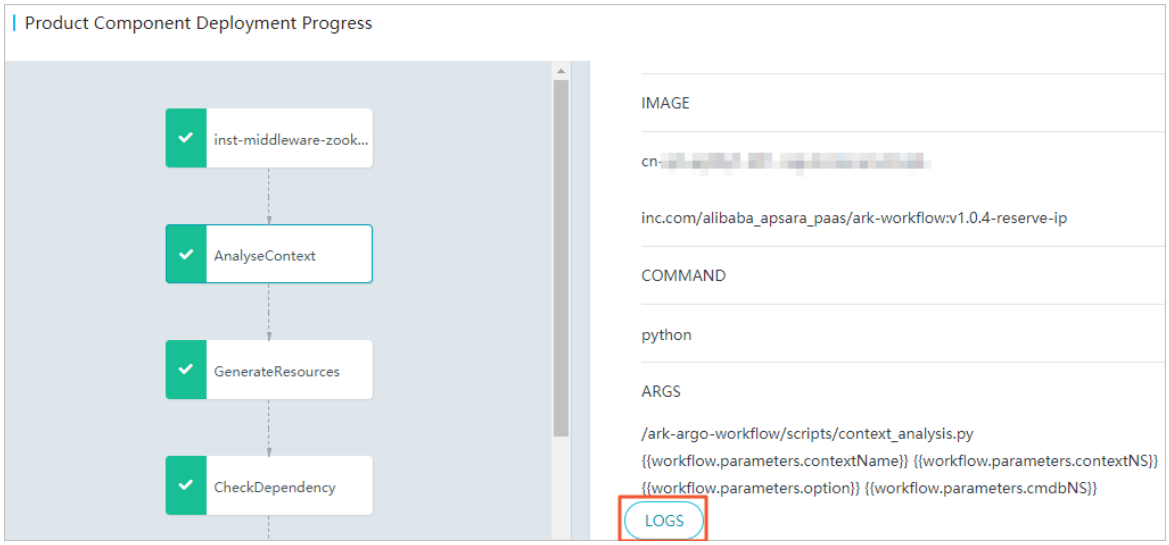

9.1.5.1.4. Log on to a web terminal

The StatefulSets and Deployments tabs of the Component Details page list available terminals. Browser-based terminals are used for O&M management and troubleshooting.

Procedure

1. In the left-side navigation pane of the PaaS console, select Products from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the Product Components section of the **Overview** page, find the target component and click **Details** in the **Actions** column.
4. On the **Component Details** page, click the **StatefulSets** or **Deployments** tab.
5. Find the target component, and then click **Start Terminal** in the **Actions** column. Available containers that are based on the number of replicas are displayed in the pane.

6. Select the target container and then click **OK** to start the terminal process.

9.1.5.1.5. Perform O&M operations

The O&M Actions page displays the O&M operations that are available to a product. You can also perform O&M operations on this page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.

3. In the left-side navigation pane, click **O&M Actions**.
4. Perform O&M operations that are available to the product.

9.1.5.1.6. View a resource report

The Resource Report page displays the information of all resources that a product has requested from the PaaS console. The resource type can be cni (ip), db, vip, dns, and accesskey.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.
2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the left-side navigation pane, click **Resource Report**.

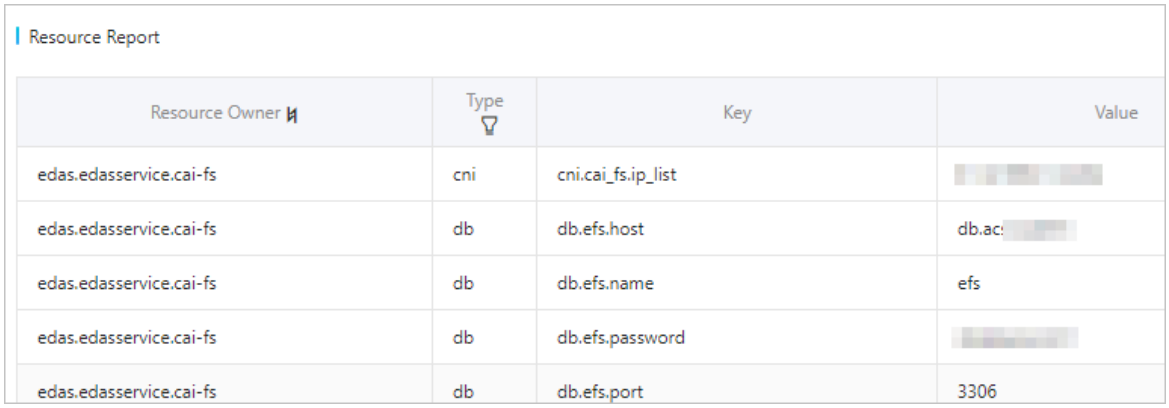

4. View the information of resources.

   By default, all resources are displayed. You can click the up and down arrows next to **Resource Owner** to sort resources. You can also click the icon next to **Type** to filter resources.

Field	Description
Resource Owner	The name of the component to which the resource belongs.
Type	The type of the resource.
Key	The attribute name of the resource.
Value	The attribute value of the resource.

9.1.5.1.7. View service registration variables

The Service Registration Variables page displays the values of all service registration variables. You can view the service registration variables of a product. The service registration variables report for a product lists the variables that the product can deliver to other products or components.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Products** from the **Product Center** drop-down list.

2. In the product list, find the target product and click **Details** in the **Actions** column.
3. In the left-side navigation pane, click **Service Registration Variables**.

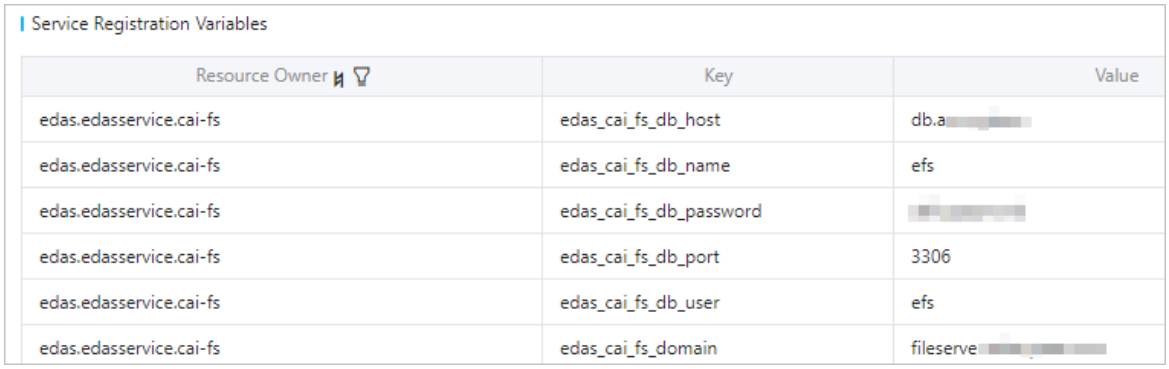
Service Registration Variables

Resource Owner	Key	Value
edas.edasservice.cai-fs	edas_cai_fs_db_host	db.a
edas.edasservice.cai-fs	edas_cai_fs_db_name	efs
edas.edasservice.cai-fs	edas_cai_fs_db_password	
edas.edasservice.cai-fs	edas_cai_fs_db_port	3306
edas.edasservice.cai-fs	edas_cai_fs_db_user	efs
edas.edasservice.cai-fs	edas_cai_fs_domain	fileserve

4. View the information of service registration variables.

   By default, all service registration variables are displayed. You can click the up and down arrows next to **Resource Owner** to sort service registration variables. You can also click the icon next to **Resource Owner** to filter service registration variables.

   The following table describes the fields for service registration variables.

Field	Description
Resource Owner	The name of the component to which the resource belongs.
Key	The variable name that is registered on CMDB and can be used by this product or other product components.
Value	The variable value that is registered on CMDB.

9.1.5.2. Deployment and upgrade

This topic describes how to perform batch upgrade and incremental deployment. You can deploy a product by product feature. If the product supports custom configuration, the system automatically goes to the custom configuration page.

Prerequisites

The deployment upgrade package is imported to the PaaS console.

You can follow the import the deployment upgrade package in the following way:

1. Upload the installation disk used for deployment and upgrade to the bootstrap node in the on-site environment.
2. Log on to the bootstrap node over SSH.
3. Run the following command to import deployment packages and generate a deployment package list:

   sh upgrade.sh *{packages -path}.iso*

   Replace *{packages -path}.iso* with the actual storage path of the iso file on the installation disk.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Deploy&Upgrade** from the **Products** drop-down list. The **System Packages** page displays deployment packages that have been imported to the PaaS console.
2. Find the target deployment package on the **System Packages** page.

> **Note** If multiple deployment upgrade packages exist, you can enter the system ID in the search box to query a deployment package that meets the condition.

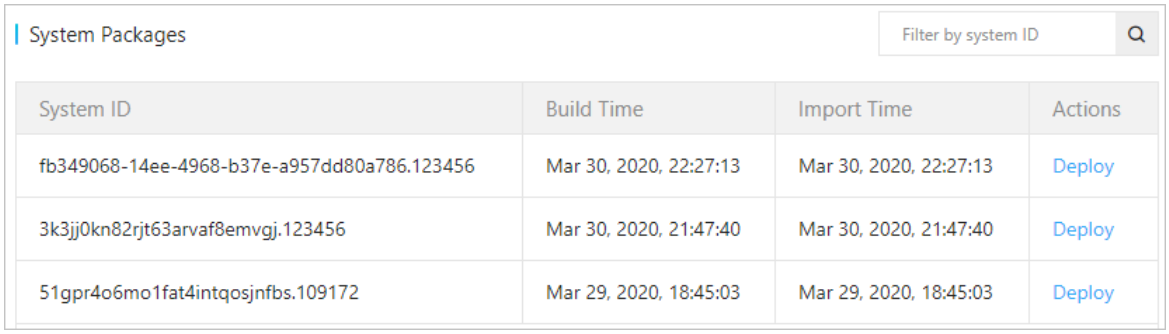

3. Click **Deploy** in the **Actions** column to start the deployment or upgrade process.
4. (Optional)In the **Select Products** step, click the number in the **Components** column to view the components and versions of the current product.

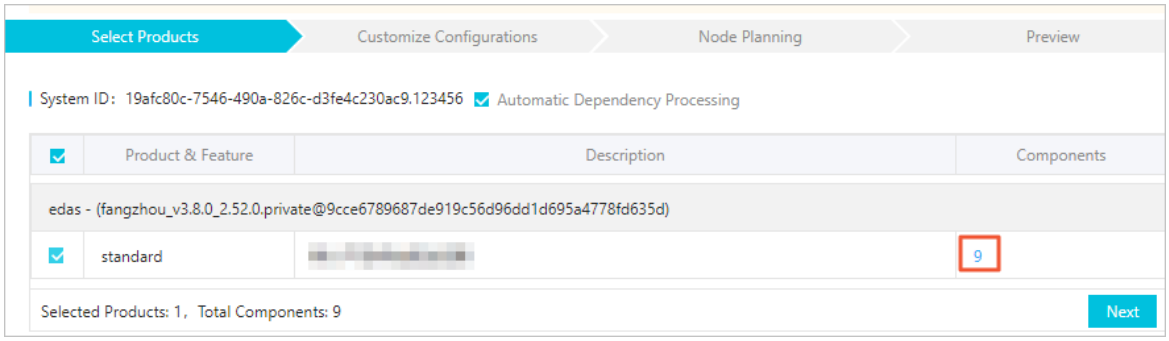

5. Select the required features and click **Next**.

> **Note** The system can automatically parse dependencies among products. When the Automatic Dependency Processing check box is selected, the system automatically checks whether dependencies exist between the deployed products and the products to be deployed and then select dependent products. If you want to manually select the products to be deployed, you can clear the **Automatic Dependency Processing** check box. If you select products that have been deployed, the system upgrades these products. If you select products that have not been deployed, the system performs incremental deployment on these products.

If the custom configuration feature is enabled for a selected product, the Customize Configurations page is displayed. Otherwise, the **Preview** page is displayed.

6. In the **Customize Configurations** step, configure the parameters as prompted and then click **Save**. Click **Next**.
7. In the **Node Planning** step, verify that the node planning is correct and click **Next**.

> **Note** If you want to modify the node planning, click **Reselect**.

8. In the **Preview** step, check the information of the products to be deployed.

   Icons before **Product & Feature** indicate different states of products:

   - : indicates that the product is newly deployed.
   - : indicates that the product has been deployed and does not need to be updated.
   - : indicates that the product has been updated. You can click the icon to check the differences.

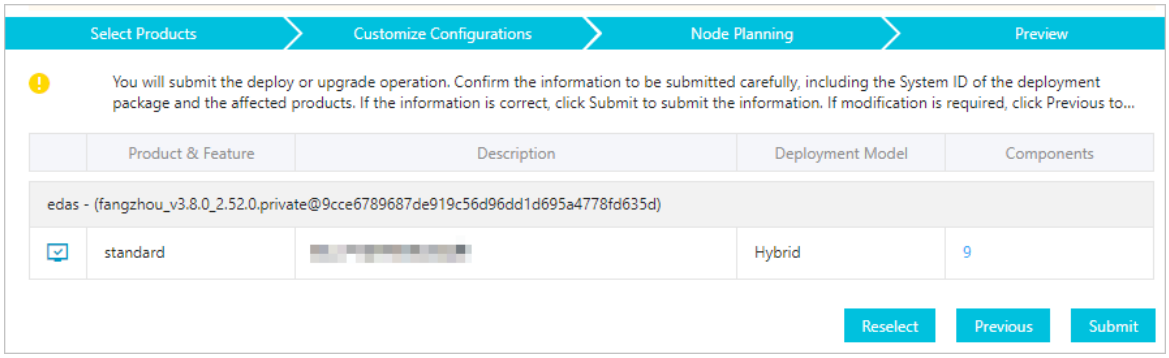

9. Click **Submit** to start the deployment or upgrade process.

   After the deployment process starts, you can view the progress on the Task Instances page by choosing **Task Center > Task Instances**.

9.1.6. Task center

The Task Center module provides general task management capabilities. You can view and run task templates, and view, suspend, resume, terminate, and delete tasks.

9.1.6.1. Task templates

The Task Templates page lists all task templates, both imported and preset.

9.1.6.1.1. View a task template

You can view information of all task templates on the Task Templates page, such as the name, description, parameters, and workflow definition.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Templates** from the **Task Center** drop-down list.
2. Find the target task template. Click **View** in the **Actions** column.

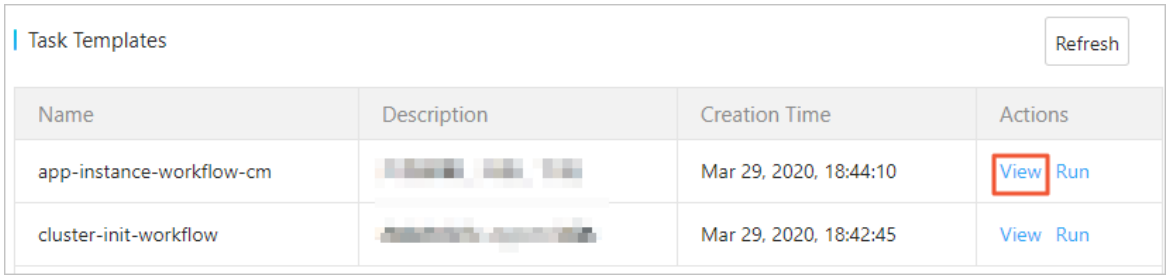

3. In the pane that appears, view the name, description, parameters, and workflow definition of the task template.

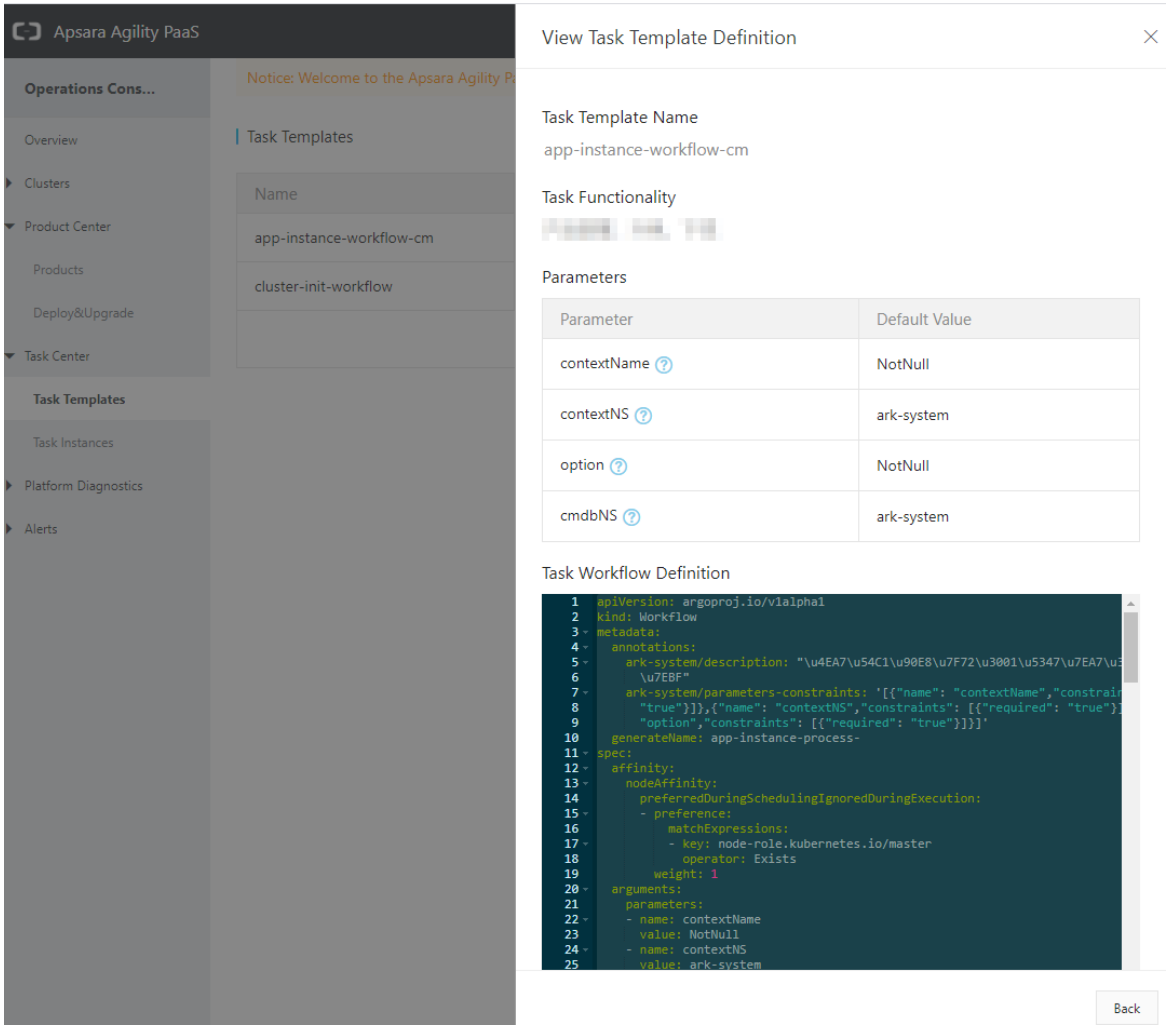

9.1.6.1.2. Run a task

You can run a task on the Task Templates page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Templates** from the **Task Center** drop-down list.
2. Find the target task template. Click **Run** in the **Actions** column.
3. In the pane that appears, set Task Instance Name and Action Parameters.

> **Note** If the task instance name is not specified, the system automatically generates a task instance name. We recommend that you enter a recognizable name for easy query.

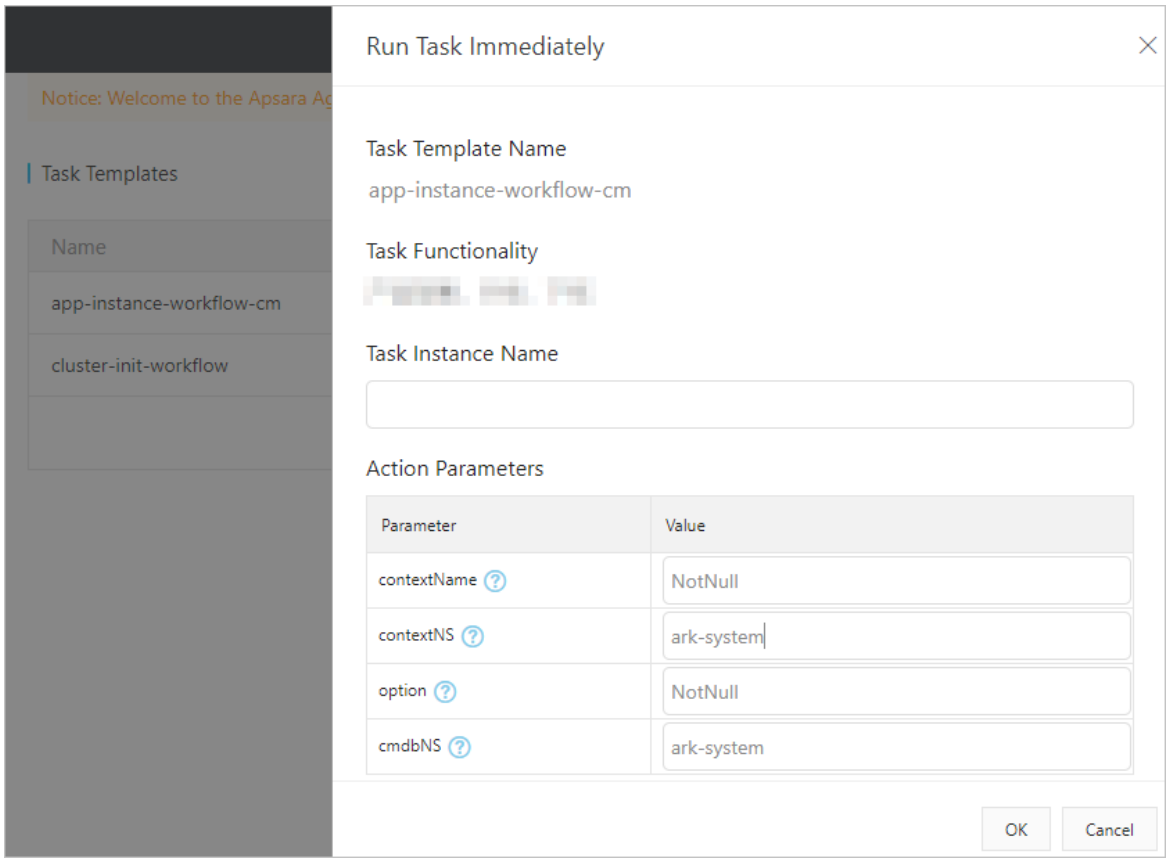

4. Click **OK**.

9.1.6.2. Task instances

The Task Instances page displays information of all tasks. On this page, you can view, suspend, resume, terminate, retry, and delete tasks.

9.1.6.2.1. View task details

After you run a task, you can view the progress, logs, and parameters of the task on the Task Instances page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, view the status of all tasks.

   Valid values of the task status:

   - **Succeeded**: indicates that the task has been executed.
   - **Running**: indicates that the task is being executed.
   - **Running (Suspended)**: indicates that the task has been suspended.

- **Failed**: indicates that the task has failed.
- **Failed (Terminated)**: indicates that the task has been terminated.

3. Find the target task. Click **View** in the **Actions** column. Then, you are redirected to the **Task Instance Details** page.

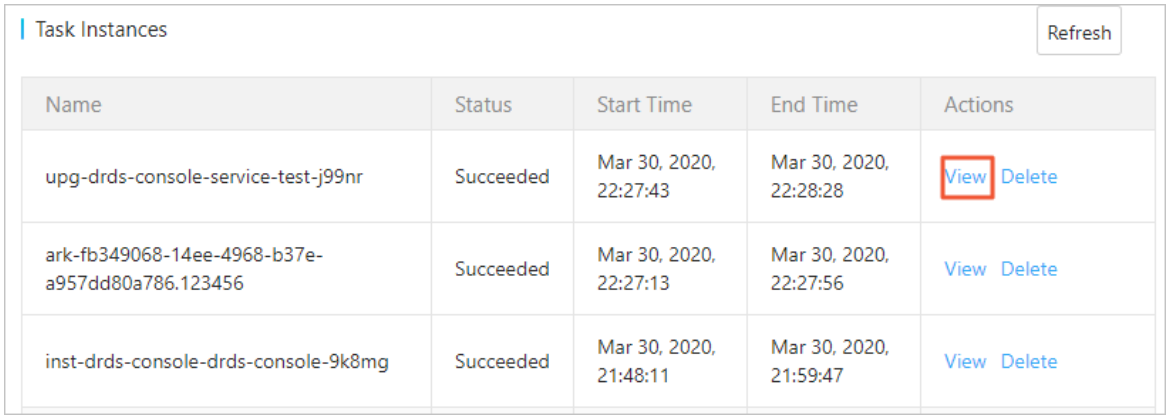

4. On the Task Instance Details page, click the task nodes in sequence to view the information and logs of the current task.

> **Note** You can click **LOGS** in the lower-left corner of the Summary tab to view task logs.

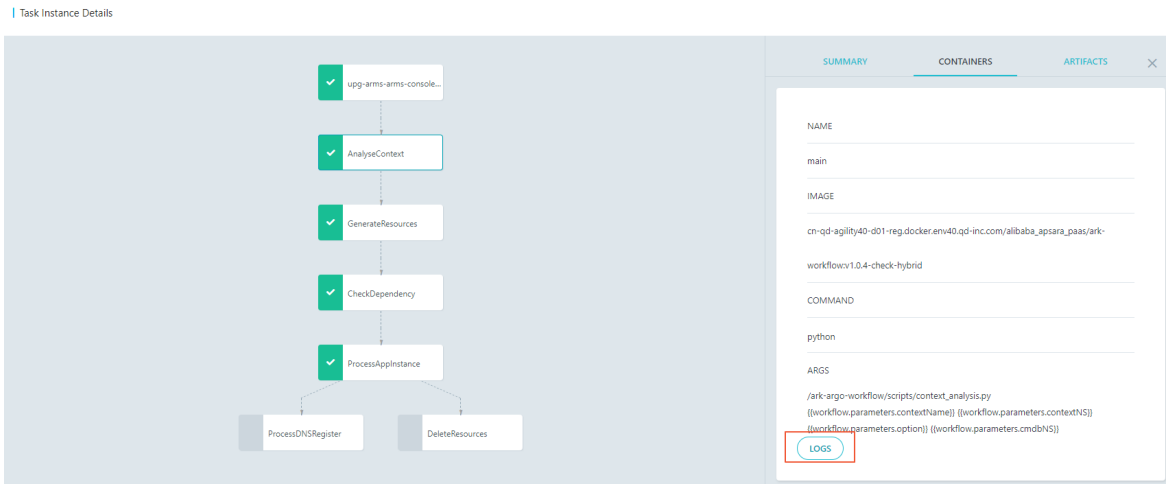

9.1.6.2.2. Suspend a task

You can suspend a task in the Running state. Then, the task status becomes Running (Suspended).

Prerequisites

The task is in the **Running** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find the task in the **Running** state that you want to suspend. Click **Suspend** in the **Actions** column. After a successful operation, the task status changes from

Running to **Running (Suspend)** in the **Status** column.

9.1.6.2.3. Resume a task

After a task is suspended, the task is in the Running (Suspended) state. Then, you can click Resume in the Actions column to resume the task.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find the task in the **Running (Suspended)** state that you want to resume. Click **Resume** in the **Actions** column. After a successful operation, the task status changes from **Running (Suspend)** to **Running** in the **Status** column.

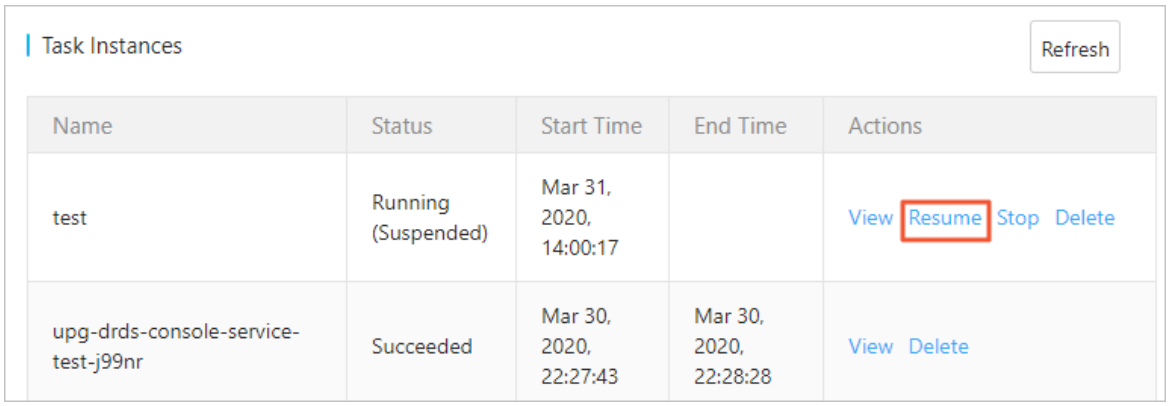

9.1.6.2.4. Terminate a task

You can terminate a task in the Running (Suspended) or Running state.

Prerequisites

The task is in the **Running (Suspended)** or **Running** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find a task in the **Running (Suspended)** or **Running** state. Click **Stop** in the **Actions** column. For a task in the Running the system immediately terminates the task and the task status becomes **Failed (Terminated)**. For a task whose **Status** is **Running (Suspended)**, the system immediately terminates the task when the task status becomes Running again. Then the task status becomes **Failed (Terminated)**.

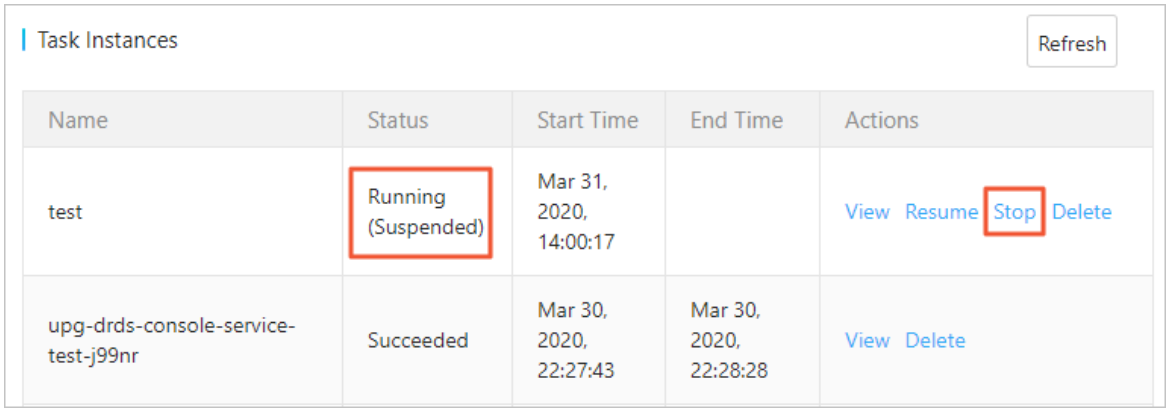

9.1.6.2.5. Retry a task

You can retry a task in the Failed or Failed (Terminated) state. When a task is retried, the task restarts from the failed or terminated task node.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find a task in the **Failed** or **Failed (Terminated)** state. Click **Retry** in the **Actions** column.

9.1.6.2.6. Delete a task

You can delete a task in any state. If a task is in the Running state, this operation enables the system to immediately terminate the task and delete the task record. If a task is in a state other than Running, this operation enables the system to immediately delete the task record.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Task Instances** from the **Task Center** drop-down list.
2. In the task instance list, find the target task. Click **Delete** in the **Actions** column.

9.1.7. Platform diagnostics

The PaaS console provides platform-level diagnostics. This module collects information about the console and products deployed in the console, presents summary diagnostic results, and allows you to download detailed diagnostic results. The module aims to improve user experience of diagnostics.

9.1.7.1. Diagnostic items

The Diagnostic Items page displays all diagnostic items in the PaaS console. On this page, you can view, execute, and delete diagnostic items.

9.1.7.1.1. View a diagnostic item

You can view details about the current diagnostic item, such as the name, type, description, start time, deletion protection, and definition.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Items** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic item. Click **View** in the **Actions** column.
3. In the pane that appears, view details of the diagnostic item.

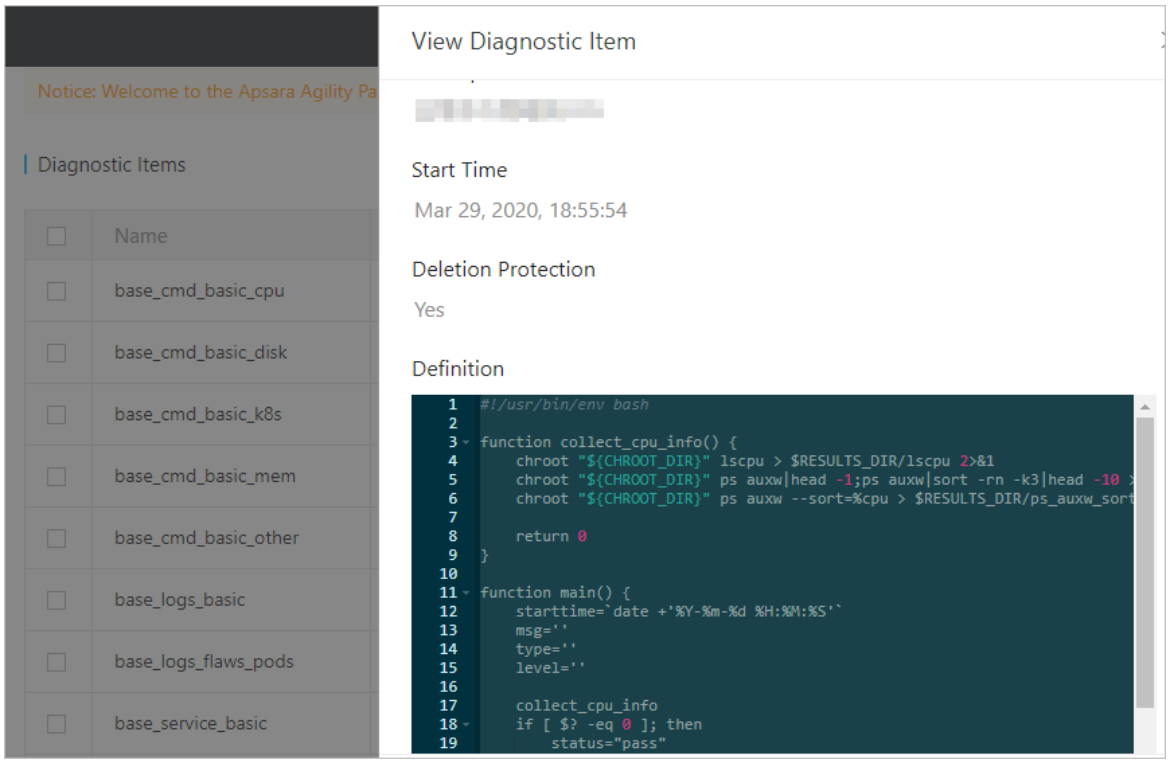

9.1.7.1.2. Execute diagnostic items

You can execute diagnostic items on the Diagnostic Items page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Items** from the **Platform Diagnostics** drop-down list.
2. Select one or more diagnostic items and click **Submit Diagnosis**.

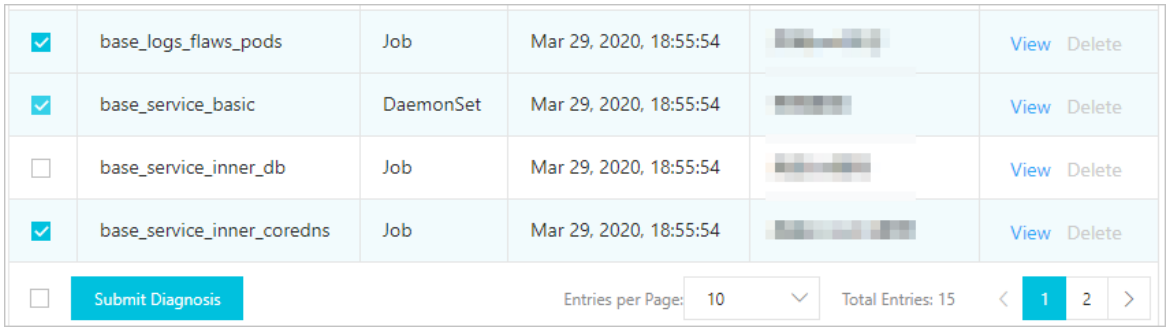

3. In the message that appears, click **OK**.

9.1.7.1.3. Delete a diagnostic item

You can delete a diagnostic item. You can only delete imported diagnostic items, but not the diagnostic items preset by the system.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Items** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic item. Click **Delete** in the **Actions** column.
3. In the message that appears, click **OK**.

9.1.7.2. Diagnostic tasks

The Diagnostic Tasks page displays all diagnostic tasks. On this page, you can view diagnostic progress, view diagnostic reports, download diagnostic reports, terminate diagnostic tasks, and delete diagnostic tasks.

9.1.7.2.1. View diagnostic progress

After you start a diagnostic task, you can view its diagnostic progress on the Diagnostic Tasks page.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Click **Diagnostic Progress** in the **Actions** column.
3. On the **Diagnostic Progress** page, click the task nodes in sequence to view the diagnostic progress and logs of the current diagnostic task.

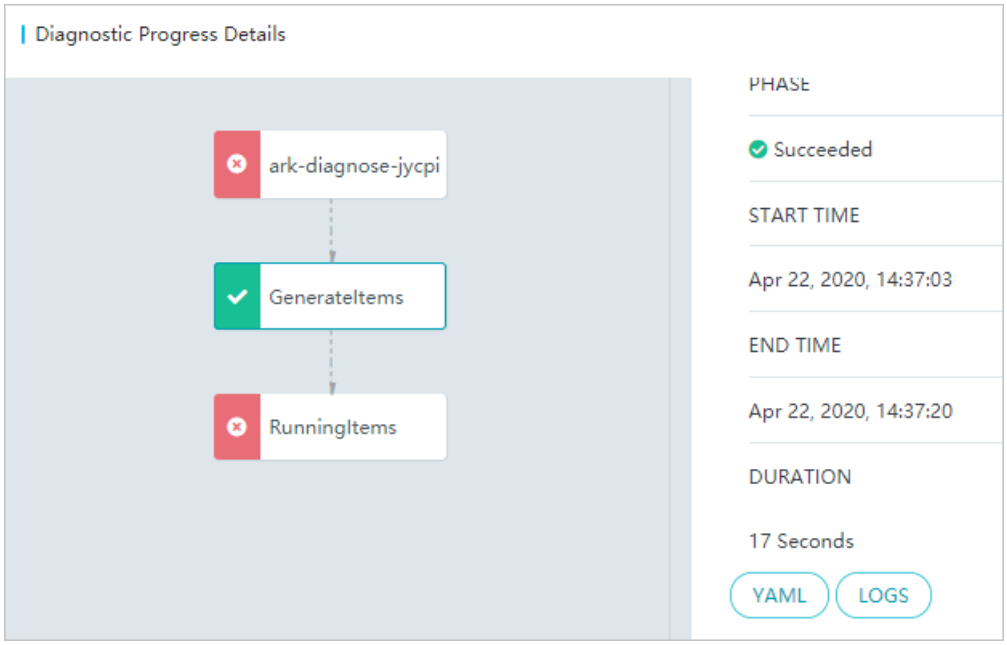

9.1.7.2.2. View a diagnostic report

After a diagnostic task is complete, you can view its diagnostic report.

Prerequisites

You can view the diagnostic report only for a diagnostic task in the **Succeeded** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Click **View Report** in the **Actions** column.
3. In the pane that appears, view the diagnostic results, such as the name, status, and details.

9.1.7.2.3. Download a diagnostic report

After a diagnostic task is complete, you can download its diagnostic report to your on-premises machine for offline query and analysis.

Prerequisites

You can download the diagnostic report only for a diagnostic task in the **Succeeded** state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Click **Download** in the **Actions** column.

9.1.7.2.4. Terminate a diagnostic task

You can terminate a diagnostic task in the Running state.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Choose **More > Terminate** in the **Actions** column.
3. In the message that appears, click **OK**.

9.1.7.2.5. Delete a diagnostic task

You can delete a diagnostic task that is no longer needed.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Diagnostic Tasks** from the **Platform Diagnostics** drop-down list.
2. Find the target diagnostic task. Choose **More > Delete** in the **Actions** column.
3. In the message that appears, click **OK**.

9.1.8. Alerts

The **Alerts** module implements unified management of alerts in the PaaS console. You can view alert rules, notification channels, and alert events. You can also configure alert rules and notification channels in the **Alerts** module.

9.1.8.1. Alert rule groups

An alert rule must belong to an alert rule group. You can create alert rule groups and add alert rules to alert rule groups.

9.1.8.1.1. Create an alert rule group

You can create an alert rule group. When you create an alert rule group, you must add an alert rule to the group.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Create Rule Group**.
4. In the **Create Rule Group** dialog box, configure the parameters.

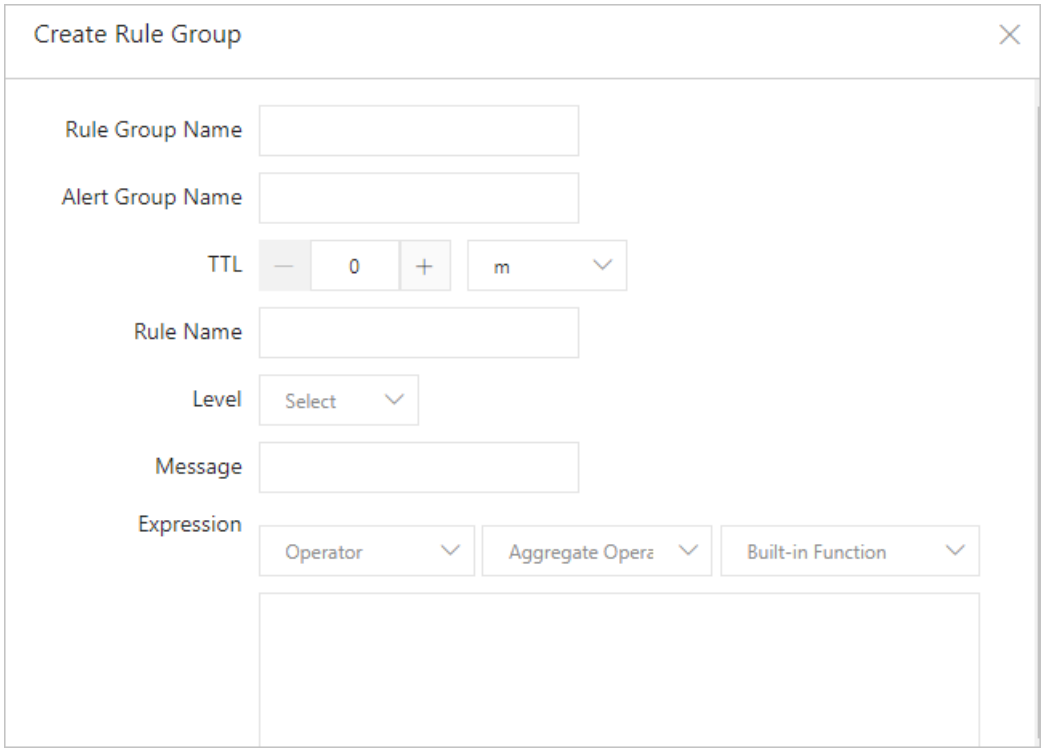

Parameter	Description
Rule Group Name	The globally unique name of the alert rule group.
Alert Group Name	The globally unique name of the alert group. An alert rule group must have an alert group.
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: indicates hours. ○ **m**: indicates minutes. ○ **s**: indicates seconds.
Rule Name	The globally unique name of the alert rule.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

5. Click **Submit**.

9.1.8.1.2. Create an alert rule

After you create an alert rule group, you can add an alert rule to the group.

Prerequisites

An alert rule group is created. For more information about how to create an alert rule group, see Create an alert rule group.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. Find the target rule group. Click **Modify Rule** in the **Actions** column.
   The **Rules** page is displayed. You can view all alert rules in the alert rule group.

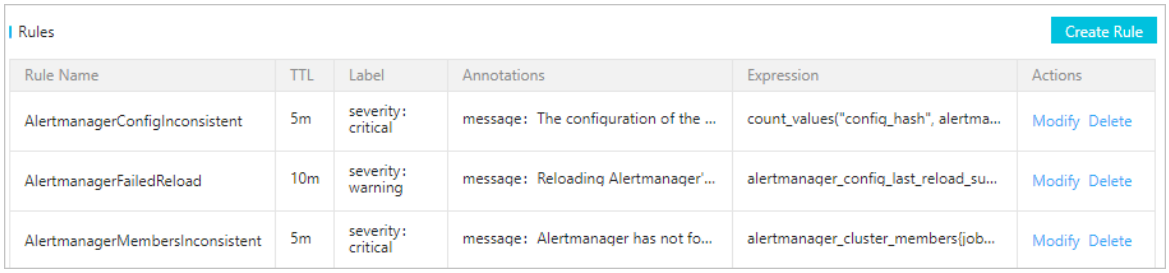

4. In the upper-right corner of the page, click **Create Rule**.
5. In the **Create Rule** dialog box, configure the parameters.

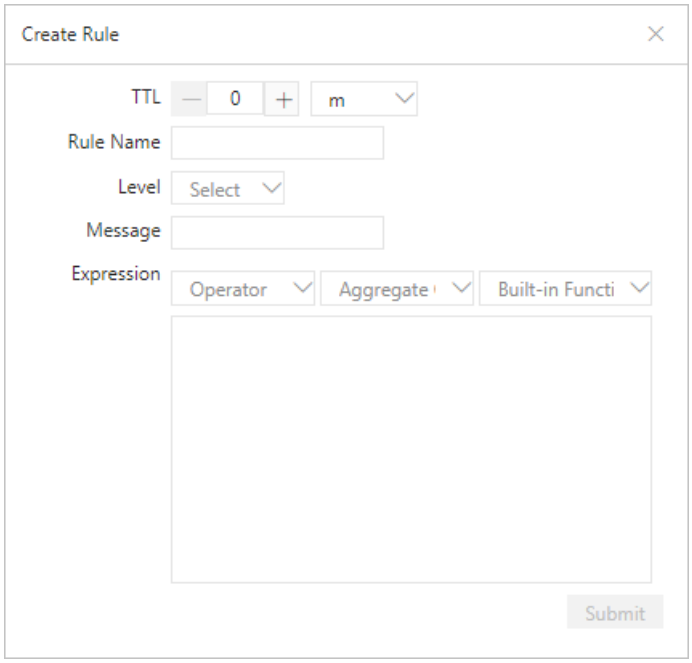

Parameter	Description

Parameter	Description
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: indicates hours. ○ **m**: indicates minutes. ○ **s**: indicates seconds.
Rule Name	The globally unique name of the alert rule.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

6. Click **Submit**.

9.1.8.1.3. Modify an alert rule

You can modify an alert rule.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. On the **Rule Groups** page, view all alert rule groups defined in the system.
4. Find the rule group for the target rule. Click **Modify Rule** in the **Actions** column.
5. On the **Rules** page, view all alert rules in the rule group.
6. Find the target rule. Click **Modify** in the **Actions** column.
7. Modify the TTL, Level, Message, and Expression parameter settings of the alert rule.

Parameter	Description

Parameter	Description
TTL	Specifies the time period that an error lasts for before an alert is sent. ○ **h**: indicates hours. ○ **m**: indicates minutes. ○ **s**: indicates seconds.
Level	The severity of the alert. Valid values: ○ **Warning**: indicates a warning alert. ○ **Critical**: indicates a critical alert.
Message	The description of the alert.
Expression	The criteria to trigger the alert. **Note** We recommend that you select operators, aggregate operations, or built-in functions from the drop-down lists if you need to use them in the expression.

8. Click **Submit**.

9.1.8.1.4. Delete an alert rule

You can delete an alert rule that is no longer needed from an alert rule group.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. On the **Rule Groups** page, view all alert rule groups defined in the system.
4. Find the rule group for the target rule. Click **Modify Rule** in the **Actions** column.
5. On the **Rules** page, view all alert rules in the rule group.
6. Find the target rule. Click **Delete** in the **Actions** column.
7. In the message that appears, click **OK**.

9.1.8.1.5. Delete an alert rule group

You can delete an alert rule group that is no longer needed.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Groups** from the **Alerts** drop-down list. The **Rule Groups** page is displayed.
2. In the upper part of the page, select the target cluster from the drop-down list.

3. Find the target rule group. Click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

9.1.8.2. Notification channels

You can view and modify notification channel settings on the Notification Channels page.

9.1.8.2.1. View notification channel settings

You can view the current notification channel settings.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the **Global Settings**, **Routing**, and **Receiver** sections, view the relevant information.

9.1.8.2.2. Modify notification channel settings

You can modify notification channel settings such as global settings, routing, and receivers.

9.1.8.2.2.1. Modify global settings

You can modify global settings, such as the resolve_timeout, smtp_info, and notifications settings.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Edit**.
4. In the **Global Settings** section, modify the resolve_timeout, smtp_info, and notifications settings.

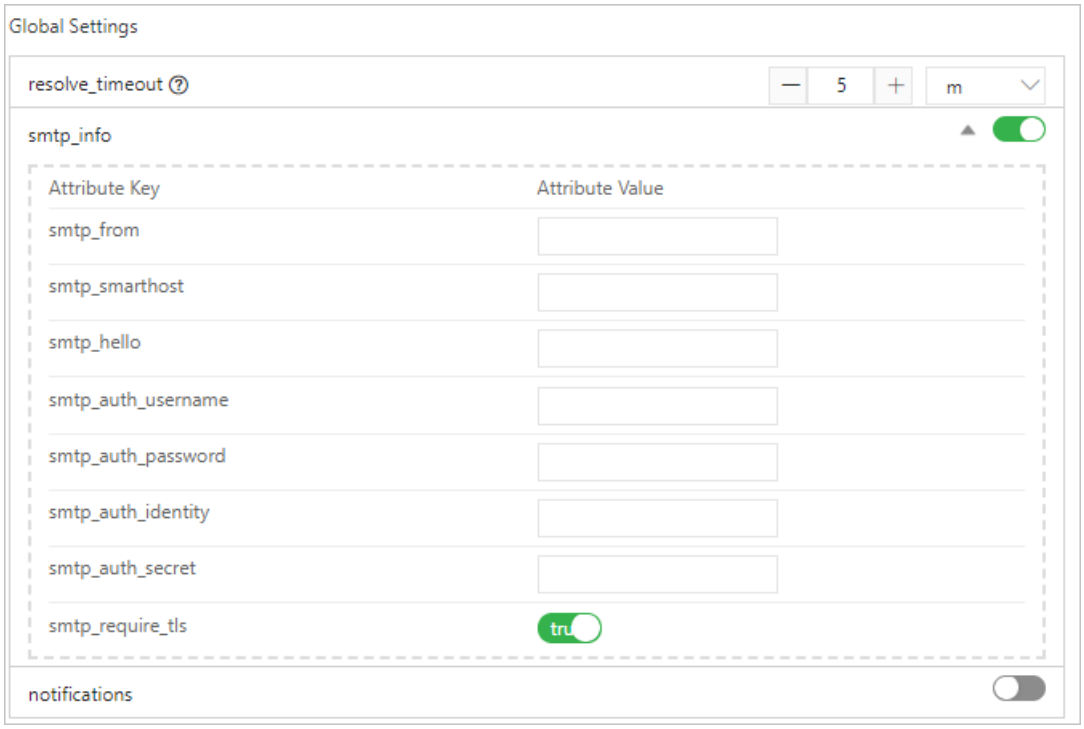

Parameter	Description
resolve_timeout	Specifies the time period before an alert is marked as resolved if the Alertmanager does not receive further notifications of the alert.
smtp_info	Specifies global SMTP information. To modify this item, turn on the switch on the right and then click the Show icon. You can configure the following parameters: ○ **smtp_from**: the source email address used to send alerts. ○ **smtp_smarthost**: the SMTP server address and port number for the source email address used to send alerts. Example: smtp_smarthos t:smtp.example.com:465 ○ **smtp_hello**: the default hostname that identifies the SMTP server. ○ **smtp_auth_username**, **smtp_auth_password**: the username and password for the source email address used to send alerts. ○ **smtp_auth_identity**: specifies the PLAIN SMTP authentication method. ○ **smtp_auth_secret**: specifies the CRAM-MD5 SMTP authentication method. ○ **smtp_require_tls**: the default SMTP TLS configuration. Although the default value is **true**, the parameter is typically set to **false** to avoid starttls errors that occur if the parameter is set to **true**.

Parameter	Description
notifications	The Slack configuration. To modify this item, turn on the switch on the right and then click the Show icon. You can configure the following parameters: ○ **slack_api_url**: the API URL for Slack notifications. ○ **victorops_api_key**: the VictorOps API key. ○ **victorops_api_url**: the VictorOps API URL. ○ **pagerduty_url**: the destination URL for API requests. ○ **opsgenie_api_key**: the Opsgenie API key. ○ **opsgenie_api_url**: the destination URL for Opsgenie API requests. ○ **hipchat_api_url**: the source URL for API requests. ○ **hipchat_auth_token**: the authentication token. ○ **wechat_api_url**: the WeChat API URL. ○ **wechat_api_secret**: the WeChat API key. ○ **wechat_api_corp_id**: the WeChat API corporate ID.

5. In the upper-right corner of the page, click **Save**.

9.1.8.2.2.2. Modify routing settings

You can modify global routing settings, and create or delete sub-routes.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Edit**.
4. In the **Routing** section, perform the following operations:
   - Modify routing settings

     You can modify the default route or subroutes.

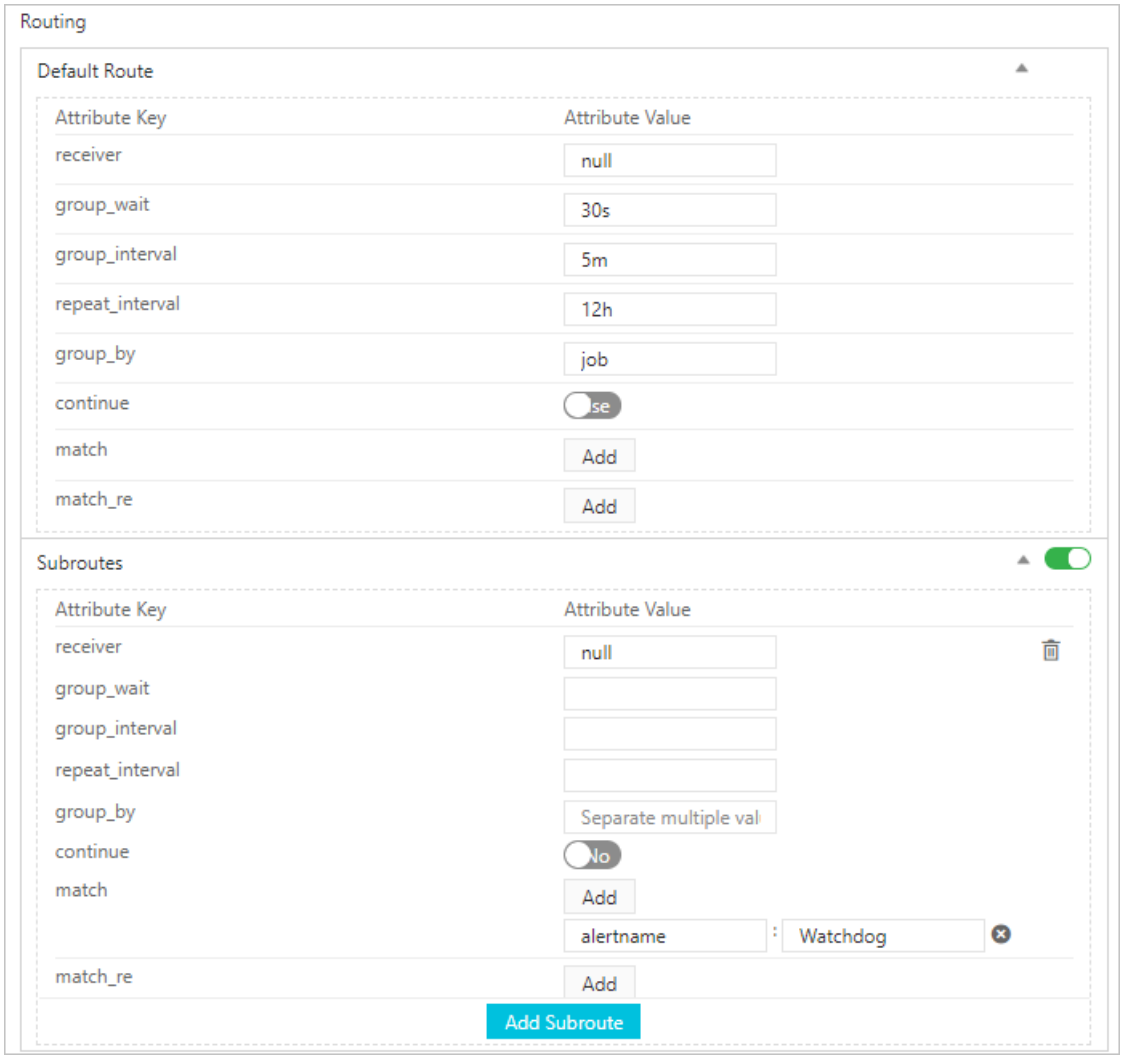

Parameter	Description

Parameter	Description
Default Route	The global route. You can configure the route information based on the actual environment. ■ **receiver**: the name of the alert receiver. ■ **group_wait**: specifies the waiting time to initialize a message when a new alert group is created. This method ensures that the system can have enough time to obtain multiple alerts for the same alert group, and then trigger an alert message. ■ **group_interval**: specifies the waiting time to send a new alert message. ■ **repeat_interval**: specifies the waiting time to resend an alert message. ■ **group_by**: the tag list. It is the regrouping tag list after alert messages are received. For example, all received alert messages that contain the `cluster=A` and `alertname=Latncy High` tags are aggregated into a group. ■ **continue**: specifies whether an alert matches subsequent nodes. ■ **match**: Click **Add** and specify a receiver for matched alerts. ■ **match_re**: Click **Add**. Enter a regular expression and specify a receiver for alerts that match the regular expression.
Subroutes	Configure subroutes in a similar way to the global route, so that you can export an alert type to another location. ■ **receiver**: the name of the alert receiver. ■ **group_wait**: specifies the waiting time to initialize a message when a new alert group is created. This method ensures that the system can have enough time to obtain multiple alerts for the same alert group, and then trigger an alert message. ■ **group_interval**: specifies the waiting time to send a new alert message. ■ **repeat_interval**: specifies the waiting time to resend an alert message. ■ **group_by**: the tag list. It is the regrouping tag list after alert messages are received. For example, all received alert messages that contain the `cluster=A` and `alertname=Latncy High` tags are aggregated into a group. ■ **continue**: specifies whether an alert matches subsequent nodes. ■ **match**: click **Add**. Enter the key and value of a tag and specify a receiver for alerts that match the tag. ■ **match_re**: click **Add**. Enter a regular expression based on the key and value of a tag and specify a receiver for alerts that match the regular expression.

- Create a subroute

  To export an alert type to another location, you can click **Add Subroute** in the lower part of the **Routing** section to configure a new subroute.

- Delete a subroute

  In the **Routing** section, find a subroute that is no longer needed and click the Delete icon to delete the subroute.

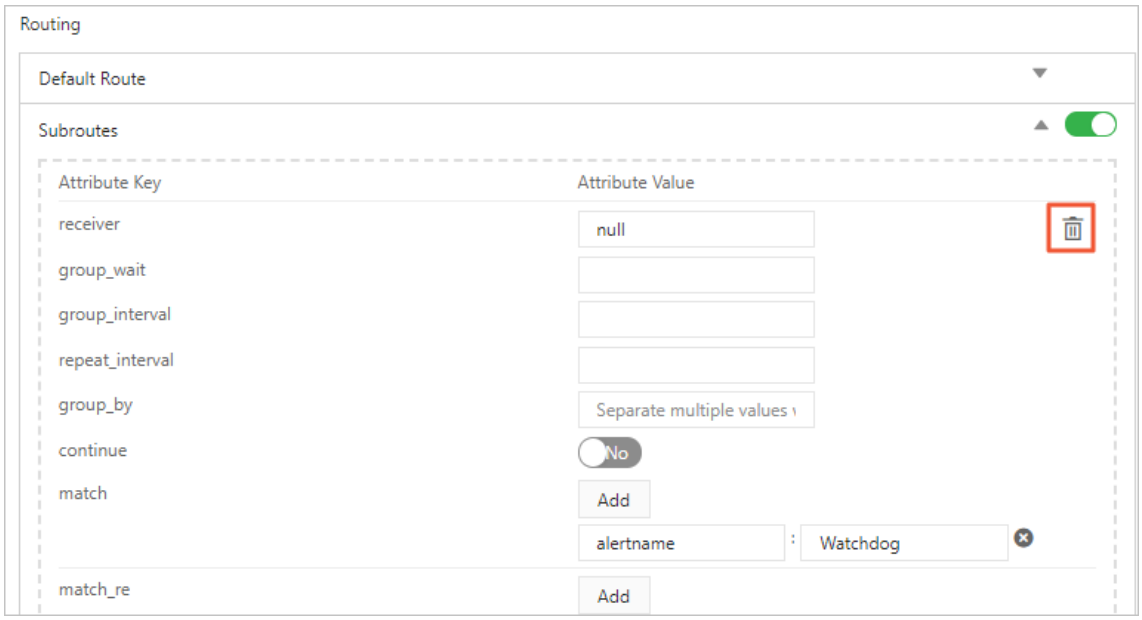

9.1.8.2.2.3. Modify receiver settings

You can create, modify, or delete alert receiver settings.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Notification Channels** from the **Alerts** drop-down list.
2. In the upper part of the page, select the target cluster from the drop-down list.
3. In the upper-right corner of the page, click **Edit**.
4. In the **Receivers** section, perform the following operations:

   - Modify receiver settings

     Modify the name and type of a receiver.

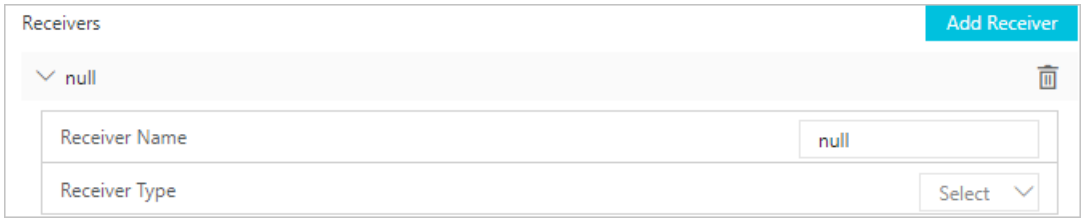

Parameter	Description
Receiver Name	The name of the alert receiver.

Parameter	Description
Receiver Type	Valid values for **Receiver Type**: **webhook** and **email**. If Receiver Type is set to **webhook**, you must configure the following parameters: ■ **url**: the URL of the alert receiver. ■ **send_resolved**: specifies whether to send messages for resolved alerts. Default value: **No**. If Receiver Type is set to **email**, you must configure the following parameters: ■ **send_resolved**: specifies whether to send messages for resolved alerts. Default value: **No**. ■ **to**: the destination email address for alerts. ■ **from**: the source email address used to send alerts. ■ **smarthost**: the server address and port number for the source email address used to send alerts. ■ **hello**: the default hostname that identifies the email server. ■ **auth_username**: the username for the source email address used to send alerts. ■ **auth_password**: the password for the source email address used to send alerts. ■ **auth_secret**: specifies the CRAM-MD5 authentication method. ■ **auth_identity**: specifies the PLAIN authentication method. ■ **require_tls**: the default TLS configuration. Although the default value is **Yes**, the parameter is typically set to **No** to avoid starttls errors that occur if the parameter is set to **Yes**.

- Add a receiver

  In the upper-right corner of the **Receivers** section, click **Add Receiver**. Configure the parameters.

- Delete a receiver

  In the **Receivers** section, find the target receiver and click the Delete icon to delete a receiver that is no longer needed.

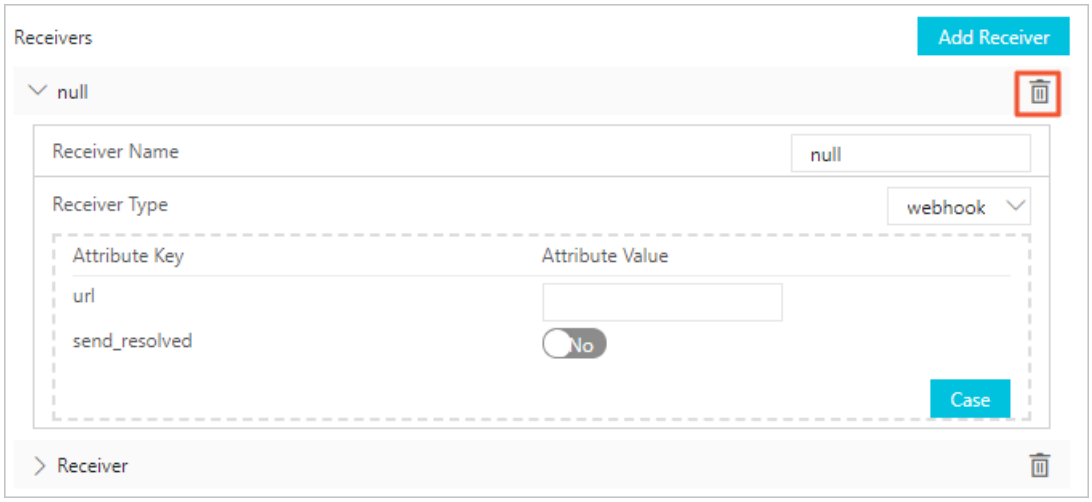

9.1.8.3. Alert events

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

9.1.8.3.1. View aggregated alert events by alert name

You can view aggregated alert events by alert name on the Alert Aggregation tab.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Events** from the **Alerts** drop-down list.
   The **Alert Aggregation** tab is displayed by default.
2. In the upper part of the page, select the target cluster from the drop-down list. By default, all alert events that are aggregated by alert name are displayed.
3. In the alert name view, the Alert Aggregation tab displays all aggregated alert events by alert name. The aggregated alert event list includes the following columns: Alert Name, Details, Total Alerts, Severity, and Actions.

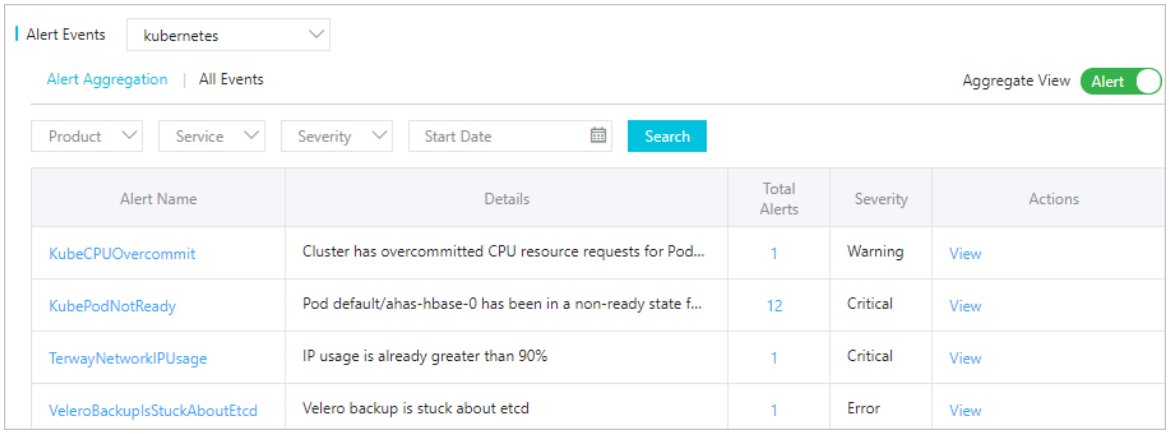

4. (Optional)In the search box at the top of the tab, set Product, Service, Severity, and Start Date, and then click **Search** to query aggregated alert events that meet the conditions.
5. Find the target aggregated alert events. Click the name in the **Alert Name** column and the number

in the **Total Alerts** column, or click **View** in the **Actions** column to view details of individual alert event within the aggregated alert events.

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

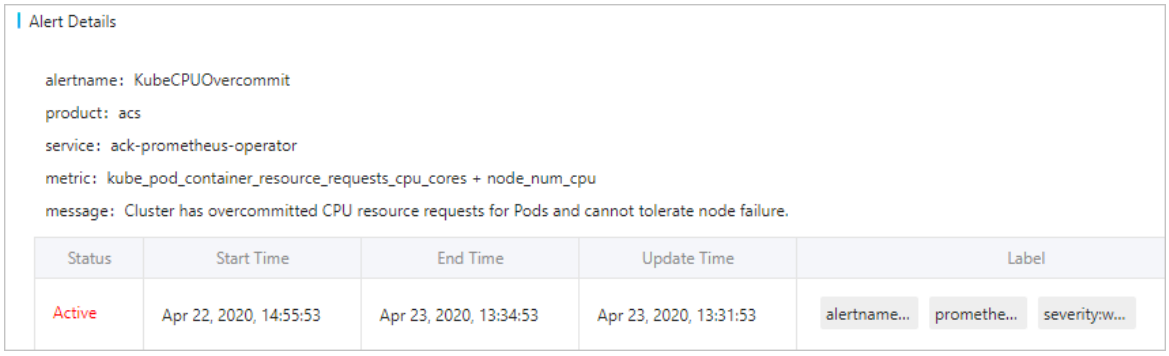

9.1.8.3.2. View aggregated alert events by product name

You can view aggregated alert events by product name on the Alert Aggregation tab.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Events** from the **Alerts** drop-down list.
   The **Alert Aggregation** tab is displayed by default.
2. In the upper part of the page, select the target cluster from the drop-down list. By default, all alert events that are aggregated by alert name are displayed.
3. Turn off the **Aggregate View** to switch to the product name view.

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

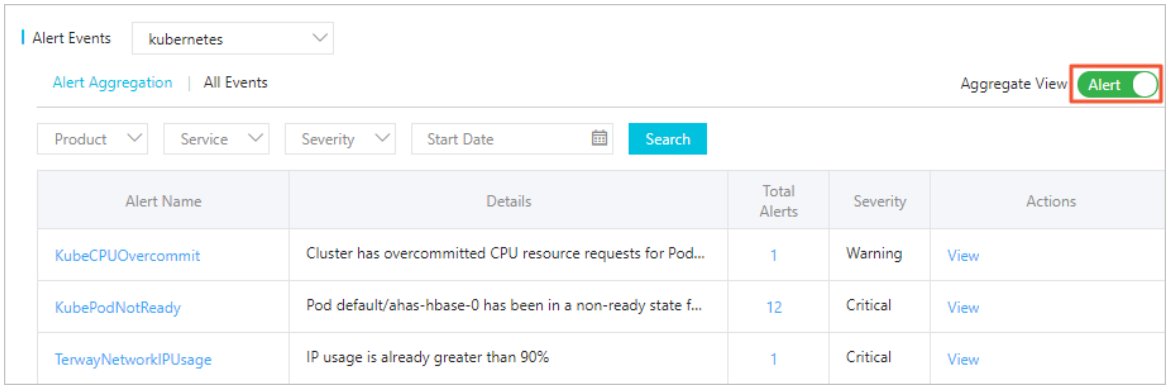

4. (Optional)In the search box at the top of the tab, set Product, Service, Severity, and Start Date, and then click **Search** to query aggregated alert events that meet the conditions.
5. Find the target aggregated alert events. Click the name in the **Alert Name** column and the number in the **Total Alerts** column, or click **View** in the **Actions** column to view details of individual alert events within the aggregated alert events. The alert details include the following columns: Status, Start Time, End Time, Update Time, and Label.

9.1.8.3.3. View all alert events

On the All Events tab, you can view all alert events generated in the PaaS console.

Procedure

1. In the left-side navigation pane of the PaaS console, select **Alert Events** from the **Alerts** drop-down list.
2. Click the **All Events** tab.
3. All alert events are displayed on the tab. The alert event list includes the following columns: Alert Name, Start Time, End Time, Update Time, Status, Details, Severity, and Label.

9.2. Harbor-based image repository console

9.2.1. Overview

This topic describes the features and purposes of Harbor.

The Harbor service is integrated into Apsara Stack Agility PaaS Kubernetes to provide a high-availability image repository and support image permission control, security scanning, and synchronization.

You can use Docker to push and pull images, as well as grant different image repository permissions to different roles.

Harbor 1.9.3 is an open-source tool. This documentation only provides basic instructions on image pushing and permission control. For more information about the features of Harbor, visit https://github.com/goharbor/harbor/blob/release-1.9.0/docs/user_guide.md.

9.2.2. Preparations

Before you use the Harbor-based image repository, you must obtain the domain name of the Harbor-based image repository and configure the host information.

9.2.2.1. Query the domain name of the Harbor-based image repository

After the environment is deployed, you can perform the following steps to query the domain name of the Harbor-based image repository:

Procedure

1. Log on to the master1 node.
2. Run the following command: **kubectl get ingress -n acs-harbor**

```
[root@node2 ~]# kubectl get ingress -n acs-harbor
NAME                    HOSTS                    ADDRESS   PORTS   AGE
harbor-harbor-ingress   harbor.myk8s.paas.com              80      3d2h
[root@node2 ~]#
```

You can obtain the domain name of the Harbor-based image repository based on the returned result.

For example, if the returned result is `harbor.myk8s.paas.com`, the domain name of the Harbor-based image repository is harbor.myk8s.paas.com:80.

9.2.2.2. Configure host information

Before you use the Harbor-based image repository, you must perform a series of configurations to ensure that each machine in the cluster can access the Harbor-based image repository.

Add master1 node information

If the DNS of the machine accessing the Harbor-based image repository cannot resolve to the domain name of the repository, you must add the following content to the */etc/hosts* file of the machine:

```
xx.xx.xx.xx harbor.myk8s.${domain}
```

In this example, the domain name of the Harbor-based image repository is `harbor.myk8s.paas.com`. You must add the following content to the /etc/hosts file:

```
xx.xx.xx.xx harbor.myk8s.paas.com
```

In the content, `xx.xx.xx.xx` is the IP address of the master1 node.

Configure insecure-registry for Docker

You must access the Harbor-based image repository over HTTP because Harbor does not have TLS enabled. If the current machine encounters HTTPS problems when attempting to use Docker to access the Harbor-based image repository, you can solve these problems by configuring insecure-registry.

This example shows how to configure insecure-registry for Linux machines. Follow these steps:

1. Log on to each node of the PaaS cluster separately.
2. Find the *daemon.json* file in the */etc/docker/daemon.json* directory.
3. Add the domain name of the Harbor-based image repository to the *daemon.json* file.

   An example of the result is as follows:

   ```
   {
    "insecure-registries";[
    "harbor.myk8s.paas.com:80"
    ]
   }
   ```

4. After the modification, run the following command to restart Docker: **systemctl restart docker**

9.2.3. Log on to the Harbor-based image repository console

This topic describes how to log on to the Harbor-based image repository console.

Prerequisites

Before you log on to the Harbor-based image repository console, make sure that the following requirements are met:

- You have obtained the URL of the Harbor-based image repository console. For more information, see Query the domain name of the Harbor-based image repository.
- You have obtained the username and password that are used to log on to the Harbor-based image repository console from the deployment personnel or administrator.
- We recommend that you use the Google Chrome browser.

Procedure

1. Enter the access URL of the Harbor-based image repository console in the address bar: harbor.myk8s.${domain}:80. Press the Enter key.

   > **Note** If you cannot access the Harbor-based image repository console, see Configure host information.

2. Enter your username and password.

   > **Note** We recommend that you change your password immediately after you log on to the Harbor-based image repository console.
   >
   > Method: In the upper-right corner of the page, click **Change Password**. Enter the current password and new password, and confirm the password. Click **OK**.

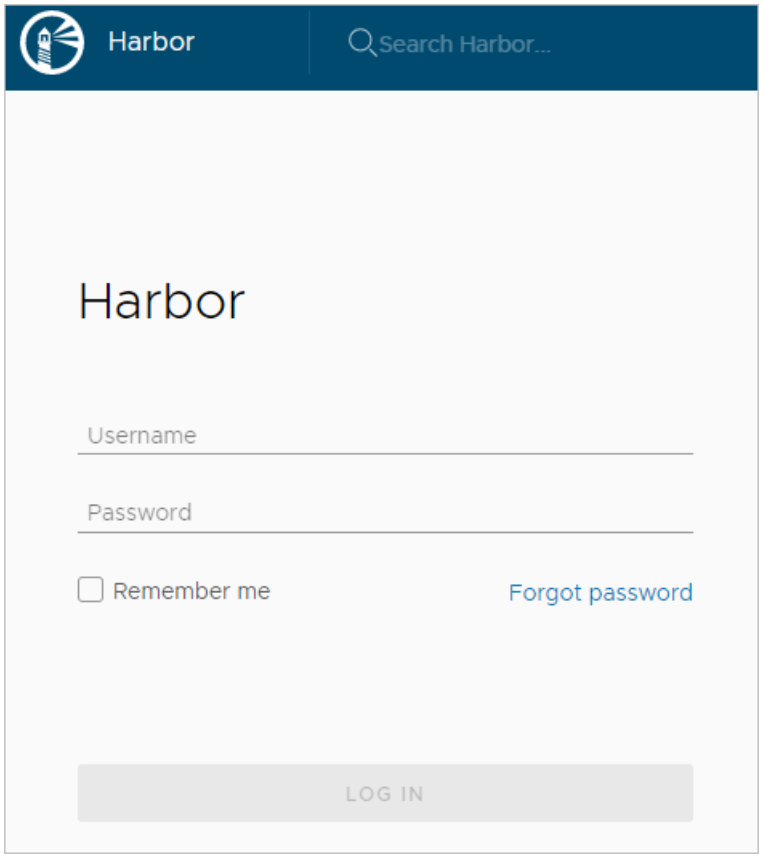

3. Click **LOG IN**.

9.2.4. Create users

After logging on to the Harbor-based image repository console, an administrator must create users to meet different requirements for access control.

Procedure

1. Log on to the Harbor-based image repository console as an administrator.
2. In the left-side navigation pane, choose **Administration > Users**.
3. On the **Users** page, click **NEW USER**.
4. In the **New User** dialog box that appears, set Username, Email, First and last name, Password, Confirm Password, and Comments. Click **OK**.

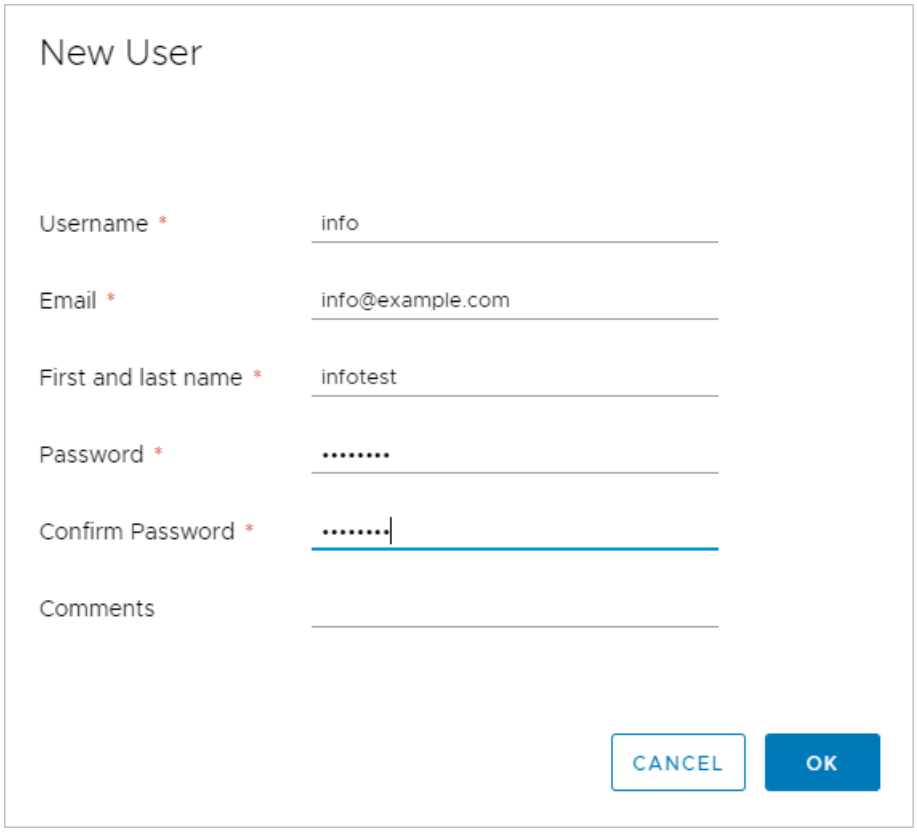

9.2.5. Create projects

Projects are containers that store image repositories. You must create a project before you can push or pull images.

Procedure

1. Log on to the Harbor-based image repository console as an administrator.
2. In the left-side navigation pane, click **Projects**.
3. On the **Projects** page, click **NEW PROJECT**.
4. In the **New Project** dialog box that appears, configure the following parameters.

Parameter	Description
Project Name	The name of the project to be created.
Access Level	The access level of the project. Whether the project is public determines whether permissions are required to pull images from the image repository. If **Public** is selected, all users including unlogged users are allowed to use Docker to pull images.
Count quota	The maximum number of images that can be stored in the image repository. The default value is **-1**, indicating that no upper limit is set on the number of images that can be stored in the image repository.

Parameter	Description
Storage quota	The storage capacity of the image repository. The default value is -1, indicating that no upper limit is set on the storage capacity of the image repository.

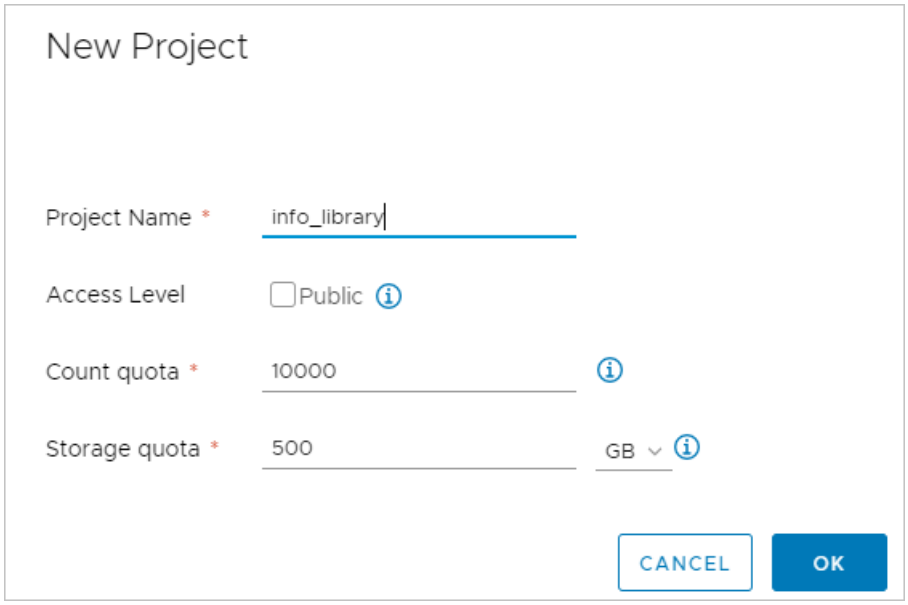

5. Click **OK**. By default, the creator of the project is the project administrator.

9.2.6. Grant project permissions

After creating a user and a project, you must grant the user project permissions so that the user can access the project.

Prerequisites

- A user is created. For more information about how to create a user, see Create users.
- A project is created. For more information about how to create a project, see Create projects.

Procedure

1. Log on to the Harbor-based image repository console as an administrator.
2. In the left-side navigation pane, click **Projects**.
3. In the project list, find the project you just created and click the project name.
4. Click the **Members** tab.
5. Click **+ USER**.
6. In the **New Member** dialog box that appears, configure the following parameters.

Parameter	Description
Name	Enter the name of the user to whom you want to grant project permissions. For example, you can enter the name of the user you just created.
Role	Valid values: ○ **Project Admin**: This role has all project permissions such as those to push images, pull images, and configure the project. ○ **Master**: This role has the permissions to push images and pull images. ○ **Developer**: This role has the permissions to push images and pull images. ○ **Guest**: This role has the permission to pull images.

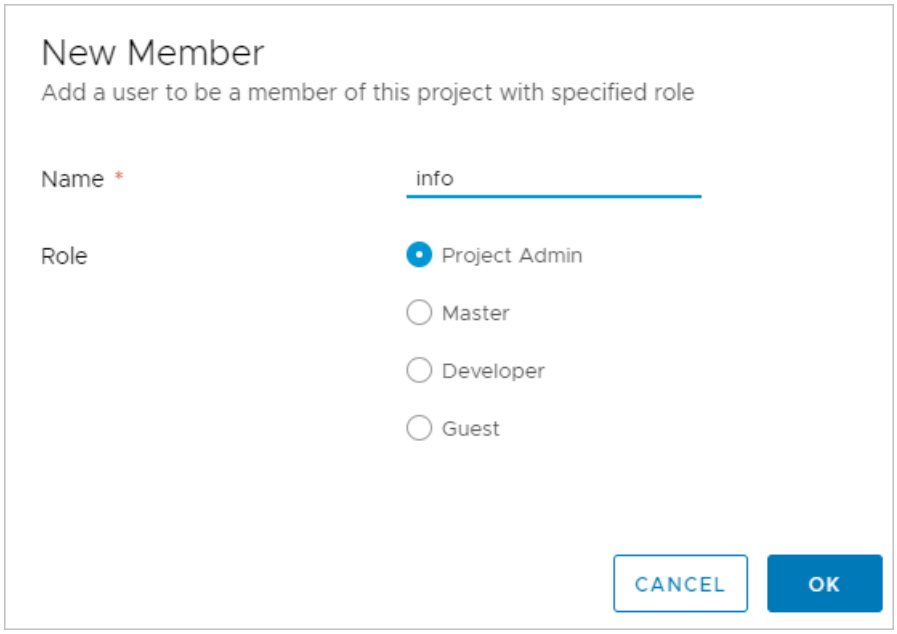

7. Click **OK**.

9.2.7. Push images

To deploy your applications on the cloud, you must push images to the Harbor-based image repository.

Prerequisites

Before you push images, make sure that the following requirements are met:

- You have the permission to push images as a project administrator, maintenance personnel, or developer.
- The image to be pushed is available. It can be a local image or an image downloaded from another image repository.

> **Note**
>
> - The Harbor-based image repository does not support images with manifest v2 schema 1.
> - The Harbor-based image repository limits the size of an image layer to 5,000 MB. If this limit is exceeded, the image will fail to be uploaded. If an image layer is larger than 5,000 MB in size, reduce the size of the image.

Procedure

1. In the left-side navigation pane of the Harbor-based image repository console, click **Projects**.
2. In the project list, find the project to which you want to push the image and click the project name.
3. Click the **Repositories** tab.
4. In the upper-right corner of the page, click **PUSH IMAGE**. The command used to push images is displayed.

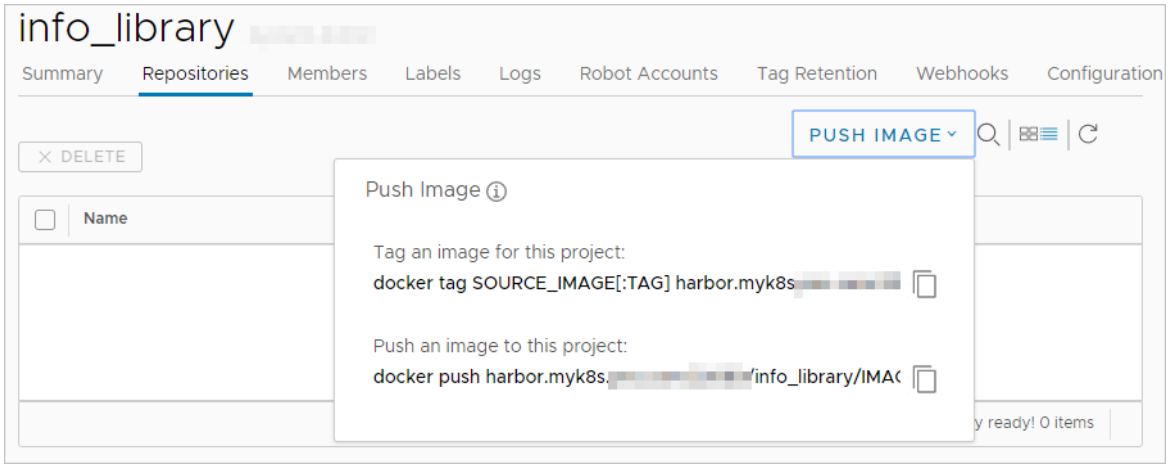

> **Note** The specific domain name of the image repository to which you push images varies with the environment. This topic uses the domain name harbor.myk8s.paas.com:80 of the Harbor-based image repository as an example.

5. Tag and push the image by using the displayed commands.
   i. Log on to the master1 node.
   ii. Prepare the image to be pushed in the local environment.
   iii. Run the following command to tag the image.

   In this example, the *1.13.3-k8s* tag is added to the *info_library/nginx* image repository.

   docker tag *info_library/nginx:1.13.3-k8s* **harbor.myk8s.paas.com:80** */info_library/nginx:1.13.3-k8s*

iv. Use the Docker command to log on to the image repository.

a. Run the following command to log on to the image repository:

docker login harbor.myk8s.paas.com:80

b. Enter the administrator account and password of the image repository.

After you log on to the image repository, `Login Succeeded` is displayed.

> **Notice** If the system prompts that the certificate verification failed during logon, you must refer to Configure host information to complete the configuration of insecure-registry.

v. After you log on to the image repository, run the following command to push the tagged image to the current project:

docker push harbor.myk8s.paas.com:80/info_library/nginx:1.13.3-k8s

6. Wait a few minutes and then log on to the Harbor-based image repository console again. On the **Projects** page, the value of **Repositories Count** corresponding to the **info_library** project is displayed as 1.

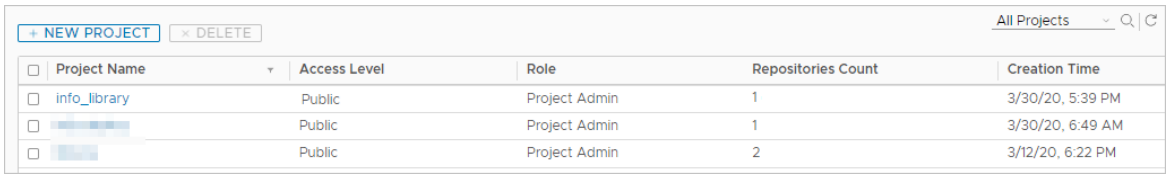

7. Click the project name. On the page that appears, click the **Repositories** tab to view the pushed image.

8. Click the image name to view the image tag.

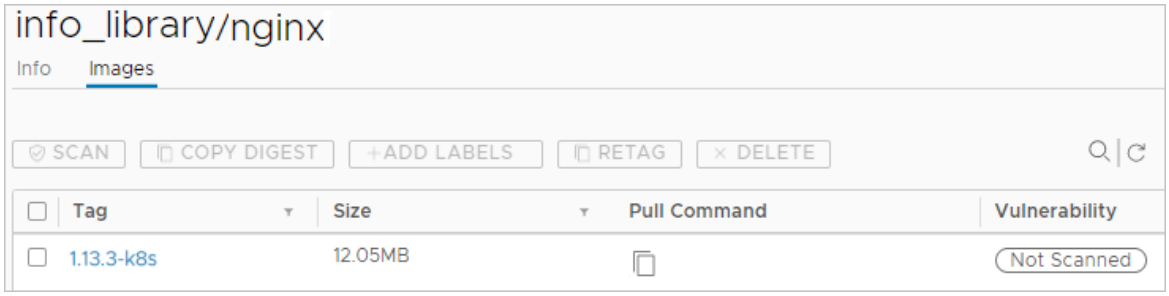

10.Operations of basic cloud products

10.1. ApsaraDB for RDS

10.1.1. Architecture

10.1.1.1. System architecture

10.1.1.1.1. Backup system

ApsaraDB for RDS can back up databases at any time and restore them to any point in time based on the backup policy, which makes the data more traceable.

Automatic backup

ApsaraDB RDS for MySQL supports both physical and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can create temporary backup files when necessary. Temporary backup files are retained for seven days.

Log management

ApsaraDB RDS for MySQL automatically generates binlogs and allows you to download them for local incremental backup.

Instance cloning

A cloned instance is a new instance with the same content as the primary instance, including data and settings. This feature allows you to restore data of the primary instance or create multiple instances that are the same as the primary instance.

10.1.1.1.2. Monitoring system

RDS provides multi-dimensional monitoring services across the physical, network, and application layers to ensure business availability.

Performance monitoring

RDS provides nearly 20 metrics for system performance monitoring, such as disk capacity, IOPS, connections, CPU utilization, network traffic, TPS, QPS, and cache hit rate. You can obtain the running status information for any instances within the past year.

SQL auditing

The system records the SQL statements and related information sent to RDS instances, such as the connection IP address, database name, access account, execution time, and number of records returned. You can use SQL auditing to check instance security and locate problems.

Threshold alerts

RDS provides alert SMS notifications if status or performance exceptions occur in the instance.

These exceptions can be involved in instance locking, disk capacity, IOPS, connections, and CPU. You can configure alert thresholds and up to 50 alert recipients (of which five are effective at a time). When an instance exceeds the threshold, an SMS notification is sent to the alert recipients.

Web operation logs

The system logs all modification operations in the RDS console for administrators to check. These logs are retained for a maximum of 30 days.

10.1.1.1.3. Control system

If a host or instance does not respond, the RDS high-availability (HA) component checks for exceptions and fails over services within 30 seconds to guarantee that applications run normally.

10.1.1.1.4. Task scheduling system

You can use the RDS console or API operations to create and delete instances, or switch instances between the internal network and Internet. All instance operations are scheduled, traced, and displayed as tasks.

10.1.2. Log on to the Apsara Stack Operations console

Prerequisites

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

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

Procedure

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

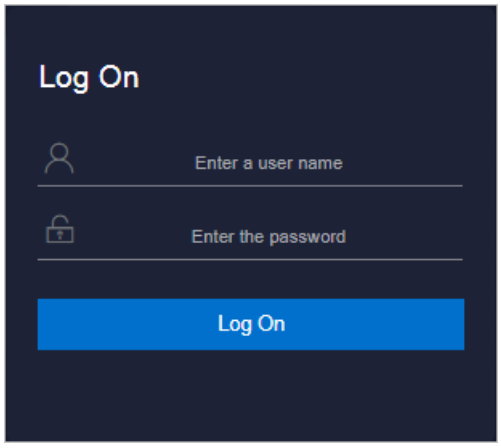

> **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

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

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

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

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.

10.1.3. Manage instances

You can view instance details, logs, and user information.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > RDS**.
3. On the **Instance Management** tab of the **RDS** page, you can perform the following operations:
   - View instances

     View instances that belong to the account on the **Instance Management** tab, as shown in Instances.

     Instances

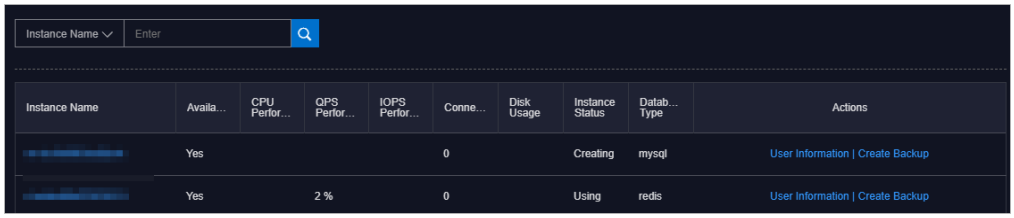

- View instance details

  Click the ID of an instance to view details, as shown in Instance details. You can switch your service between primary and secondary instances and query historical operations on this page.

  > **Note** We recommend that you do not perform forced switchover, because it may result in data loss if data is not synchronized between the primary and secondary instances.

  Instance details

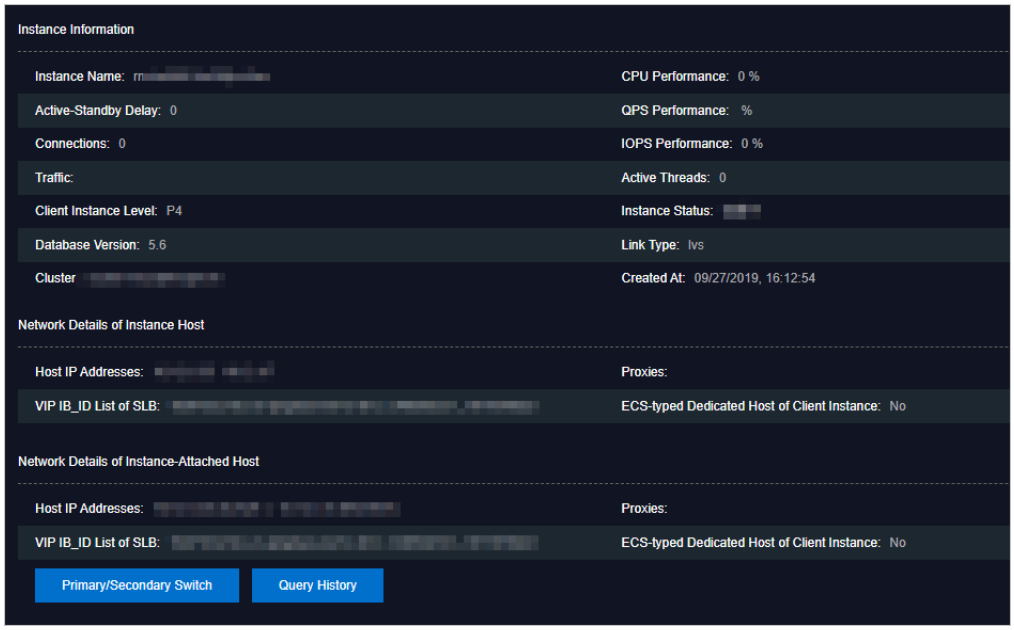

- **View user information**

  Click **User Information** in the **Actions** column corresponding to an instance, as shown in User information.

  User information

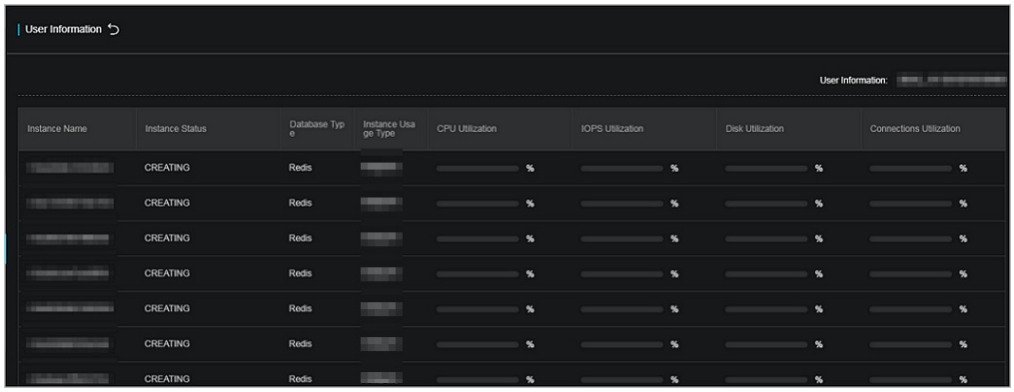

- **Create backups**

  For ApsaraDB RDS for MySQL instances, click **Create Backup** in the **Actions** column to view the backup information, as shown in Backup information. You can also click **Create Single Database Backup** on the Backup Information page to back up a single database.

  Backup information

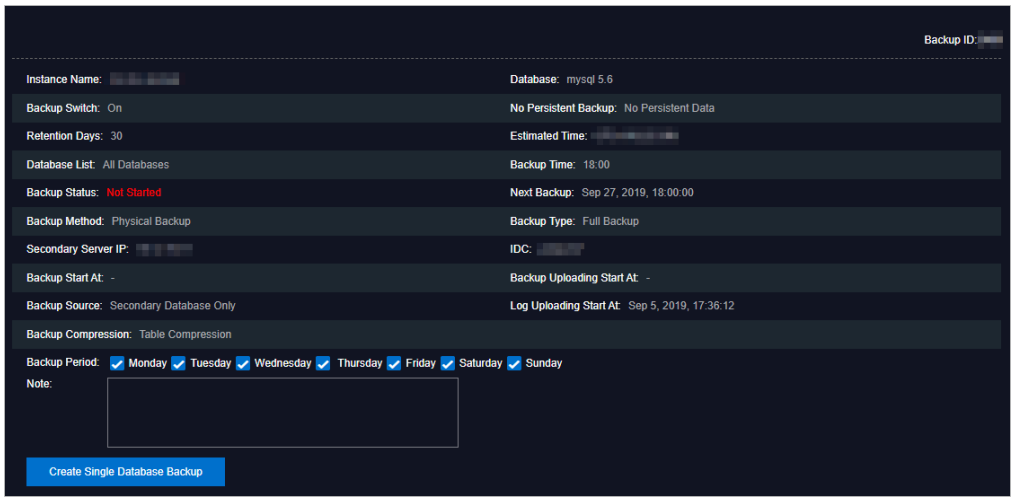

10.1.4. Manage hosts

You can view and manage hosts.

Procedure

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > RDS**.
3. On the **Host Management** tab of the **RDS** page, you can view the information of all hosts.

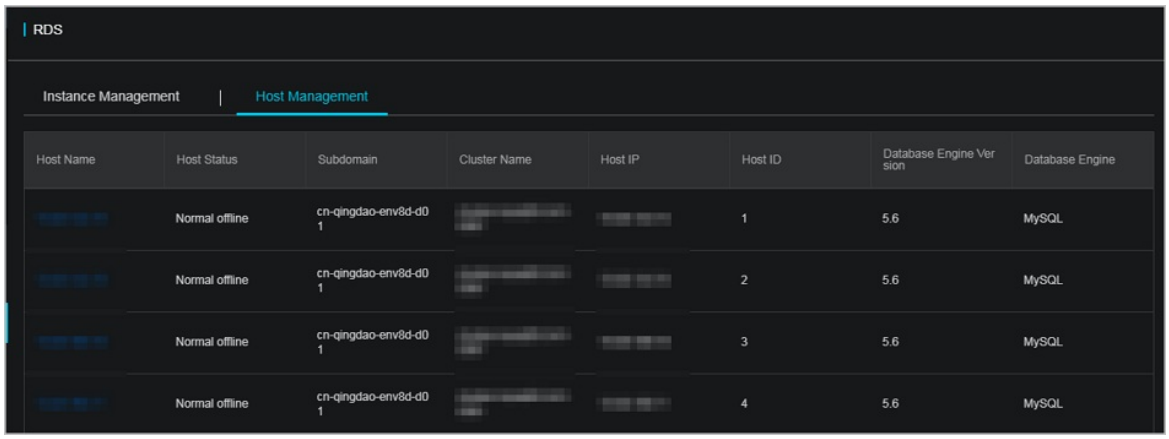

4. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.

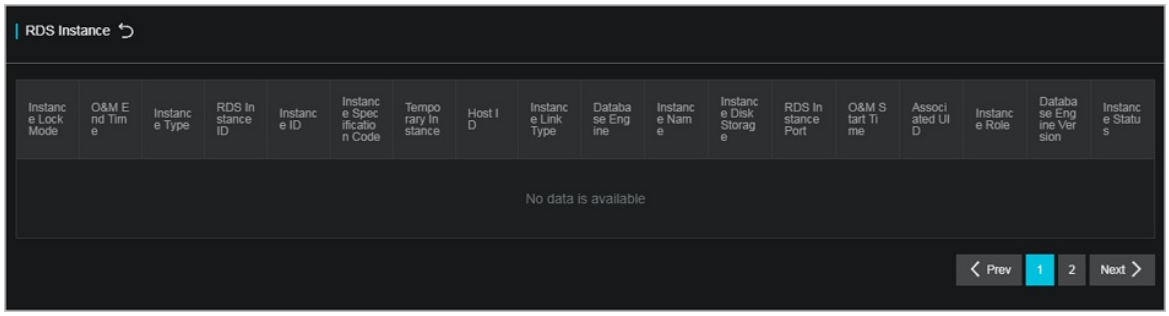

10.1.5. Security maintenance

10.1.5.1. 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.

10.1.5.2. 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.

11.Apsara Opsapi Management system

11.1. Apsara Opsapi Management system overview

This topic describes the features and infrastructure of the Apsara Opsapi Management system (Opsapi).

The Apsara Opsapi Management system is a platform that manages O&M APIs and SDKs in the Apsara Stack environment in a centralized manner. This system also manages API and SDK versions.

Most Apsara Stack products provide APIs for tenants. Only a few Apsara Stack products provide APIs for O&M. To address the business needs at the O&M level and meet custom development requirements of users such as developing their own O&M consoles or obtaining O&M data, Alibaba Cloud provides the Apsara Opsapi Management system.

The Apsara Opsapi Management system has the following features:

- Provides APIs at the system level and typical APIs for resource usage, monitoring, and alerting.
- Manages APIs, including querying, editing, testing, and deleting APIs.
- Provides an API designer to customize an API flow based on the existing API, which facilitates custom business.
- Manages versions and relationships between these versions. These versions include Apsara Stack versions, product versions, SDK versions, and API versions.
- Supports SDKs. The Apsara Opsapi Management system provides SDKs for Java and Python to call O&M APIs.

Infrastructure of the Apsara Opsapi Management system shows the infrastructure of the Apsara Opsapi Management system.

Infrastructure of the Apsara Opsapi Management system

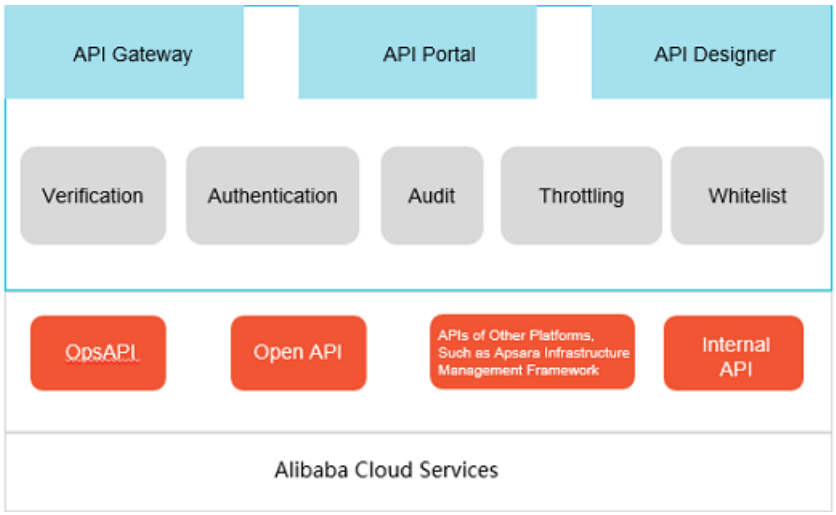

The Apsara Opsapi Management system contains the following components:

- api-server: contains O&M APIs that are available in SDKs.

- API Portal: the O&M console used to manage Opsapis.
- api-node: the API designer.

11.2. Log on to the Apsara Opsapi Management system

The Apsara Opsapi Management system provides basic platform management functions for O&M engineers. These functions include API management, version management, test management, and system management. This topic describes how to log on to the Apsara Opsapi Management system.

Prerequisites

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

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

Procedure

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

   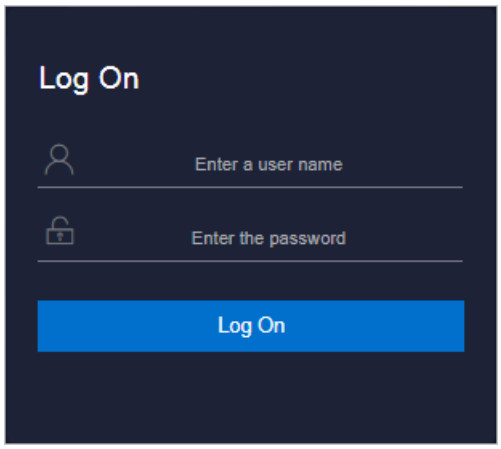

   Note You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

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

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

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

   - It must contain uppercase and lowercase letters.
   - It must contain digits.

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

4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, click **Products**.
6. In the **Apsara Stack O&M** section, click **Opsapi**.

11.3. API management

The Apsara Opsapi Management system provides APIs of various products in the Apsara Stack environment. You can manage these APIs, such as uploading, querying, editing, testing, and deleting APIs. You can also use the API designer to customize APIs.

11.3.1. Register APIs

An API can be defined in an XML file. Each API corresponds to one XML file. You can upload an XML file to register an API in the Apsara Opsapi Management system.

Context

The following fields must be defined in an XML file:

- API name
- namespace (or product name)
- API type
- Parameter

When you upload APIs, you can upload one or more XML files simultaneously. You must compress the XML files into a ZIP file before you upload these files.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **API Platform > APIs**. The **APIs** page appears.
3. In the upper-right corner, click **Upload API**. Select the XML or ZIP file to be uploaded. After you upload the XML or ZIP file, you can view the uploaded APIs in the API list on the **APIs** page.

11.3.2. Modify information about APIs

You can modify basic information and specific parameters of an API.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **API Platform > APIs**. The **APIs** page appears.
3. (Optional)Select a product from the **Select Product** drop-down list. Enter an API name in the search box.

   You can enter a full or partial API name to search for APIs.

4. Click the icon in the **Actions** column corresponding to the API to be modified.

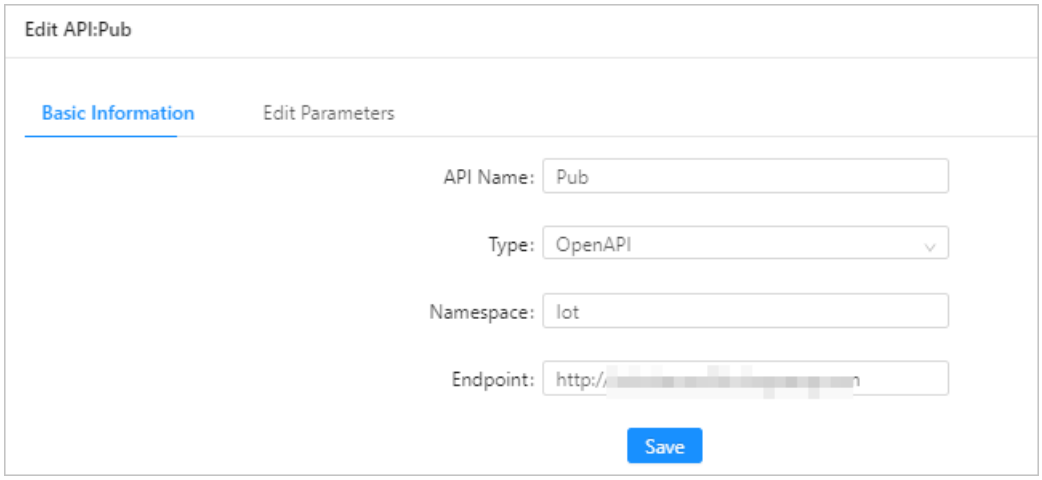

5. In the Edit API dialog box that appears, click the **Basic Information** tab, modify the basic information, and then click **Save**.

   The following table describes parameters of an API.

Parameter	Description
API Name	The name of the API.
Type	The type of the API. Different products have different API types. The types are as follows: ○ opsAPI: the API for O&M ○ OpenAPI: the API for product operations ○ customAPI: the API customized through the API designer
Namespace	The namespace of the API that corresponds to the product name.
Endpoint	The domain name of the product that corresponds to the API.

6. Click the **Edit Parameters** tab to modify configuration information such as the request parameters, response parameters, and error handling mechanism. Follow these steps: Set the request parameters in the `Parameters` section, response parameters in the `ResultMapping` section, and error handling mechanism in the `ErrorMapping` section.
7. After the modification is completed, click **Save**.

11.3.3. Test APIs

The Apsara Opsapi Management system allows you to test APIs online to check whether an API is available. During the test, you can save input parameters as a test case for subsequent execution.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **API Platform > APIs**. The **APIs** page appears.
3. On the APIs page, click the icon in the **Actions** column corresponding to the API to be tested.
4. In the Test API dialog box, set **Request Parameters**.

   Request parameters may vary with APIs. The following table describes the typical request parameters.

Parameter	Required	Description
regionId	Yes	The region ID of the test environment.
accessKeyId and **accessKeySecret**	Yes	The identification of the visitor. You can obtain them from the Apsara Stack console.

5. After request parameters are configured, click **Send**.

   The Apsara Opsapi Management system sends a corresponding test request to the configured domain name. The response appears in the **Responses** section.
6. (Optional)After the test is complete, click **Save As Test Case** for subsequent execution of this test case. You can do this by choosing **Testing Platform > Test Cases** on the Test Cases page.

11.3.4. Remove information about APIs

You can remove information about an API that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **API Platform > APIs**. The **APIs** page appears.
3. (Optional)Select a product from the **Select Product** drop-down list. Enter an API name in the search box.

   You can enter a full or partial API name to search for APIs.
4. Click the icon in the **Actions** column corresponding to the API to be removed.
5. In the message that appears, click **OK**.

11.3.5. API design

The Apsara Opsapi Management system provides an API designer to help you customize APIs.

11.3.5.1. API designer

This topic describes the API designer.

If Apsara Stack Opsapis do not match the APIs you are using, or if you need to customize APIs to meet the requirements of specific projects, you can use an API designer to assemble and create desired APIs in the flow design process.

The API designer is built based on the open-source project Node-RED. Node-RED is a powerful tool launched by IBM to build Internet of Things (IoT) applications. It uses the visual programming method that allows developers to connect predefined code blocks (nodes) to perform tasks. Connected nodes are a combination of input nodes, processing nodes, and output nodes. When they are connected to form a flow, they are able to process requests such as HTTP requests.

Node-RED is highly capable of customizing flows and processing HTTP messages. These capabilities can be easily expanded.

To design an API, follow these steps:

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **API Platform > API Design**. The **API Designer** page appears.
3. Drag and drop the components on the left to the flow chart section. You can combine these components to complete a specific flow chart and design an API.

11.3.5.2. Designer nodes

This topic describes typical nodes that are used in the designer.

To customize Opsapis, the Apsara Opsapi Management system adds some nodes through the mechanism provided by Node-RED.

Typical nodes are described as follows:

- api-request: used to create a request to access an Opsapi. There is a small icon before each node. After you click the icon, a request is sent to the flow that contains the node.
- api response: used to provide responses and format the returned data.
- api selector: used to select and execute an existing API.
- db exec: used to execute a specified SQL operation.
- new api: used to create an API that contains a specified endpoint and specific input parameters.
- sync msg: used to synchronize multiple requests. Multiple responses are merged into one response.
- py function: Python is used in some modules during the API design process.
- Input components: detailed operations involved in a request process. For example, set the protocol type of a request to HTTP, TCP, or UDP, and specify the status code and created link for the request.
- Output components: the returned data, status code, and protocol in the response. Output components are used to describe fixed output modes such as request and response methods and formats of returned data.

11.3.5.3. Design an API flow

This topic describes how to design an API flow.

Each customized API has its own API flow. Each API flow consists of multiple connected nodes, including one input node, several processing nodes, and one return node (or output node).

Among the nodes:

- Typical input nodes are **api request** and **http in**.
- Typical return nodes are **api response** and **http response**.
- Typical processing nodes are **function**, **api selector**, and **db exec**. The **function** node is used to convert parameters and process simple logic.

An API flow is designed as follows:

1. Select an input node and an output node, and add processing nodes to the flow.
2. Specify the name and configurations of each node, such as the endpoint of the input node.
3. Connect the nodes as needed to form a flow.
4. In the upper-right corner of API Designer, click **Deploy** to publish the flow.
5. Access this flow in the browser. You can obtain the response.

11.4. Version management

11.4.1. Apsara Stack version management

Apsara Stack has multiple versions that vary with projects.

11.4.1.1. Add information about versions

You can add information about Apsara Stack versions as needed to manage the relationships among Apsara Stack versions, products, and product versions.

Context

Each Apsara Stack version can have either one release version or one snapshot of the on-premises environment or deployment environment. It can be distinguished by its version name and description.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack Versions** page appears.
3. In the upper-right corner of the page, click **Add Version**.
4. In the **Add Version** dialog box that appears, set **Apsara Stack Version**, **Version**, and **Release Notes**.

   We recommend that you enter information that is related to the current version for Release Notes.
5. Click **Submit**.

11.4.1.2. Select products for an Apsara Stack version

After you add information about an Apsara Stack version, you can select products that are supported in an Apsara Stack version based on version output conditions.

Procedure

1. Log on to the Apsara Opsapi Management system.

2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack Versions** page appears.
3. In the version list, click **Configure Products** in the Product column corresponding to the specified Apsara Stack version.
4. In the Configure Products dialog box that appears, select a version from the **Version** drop-down list and select the check box in the **Output** column corresponding to the product.

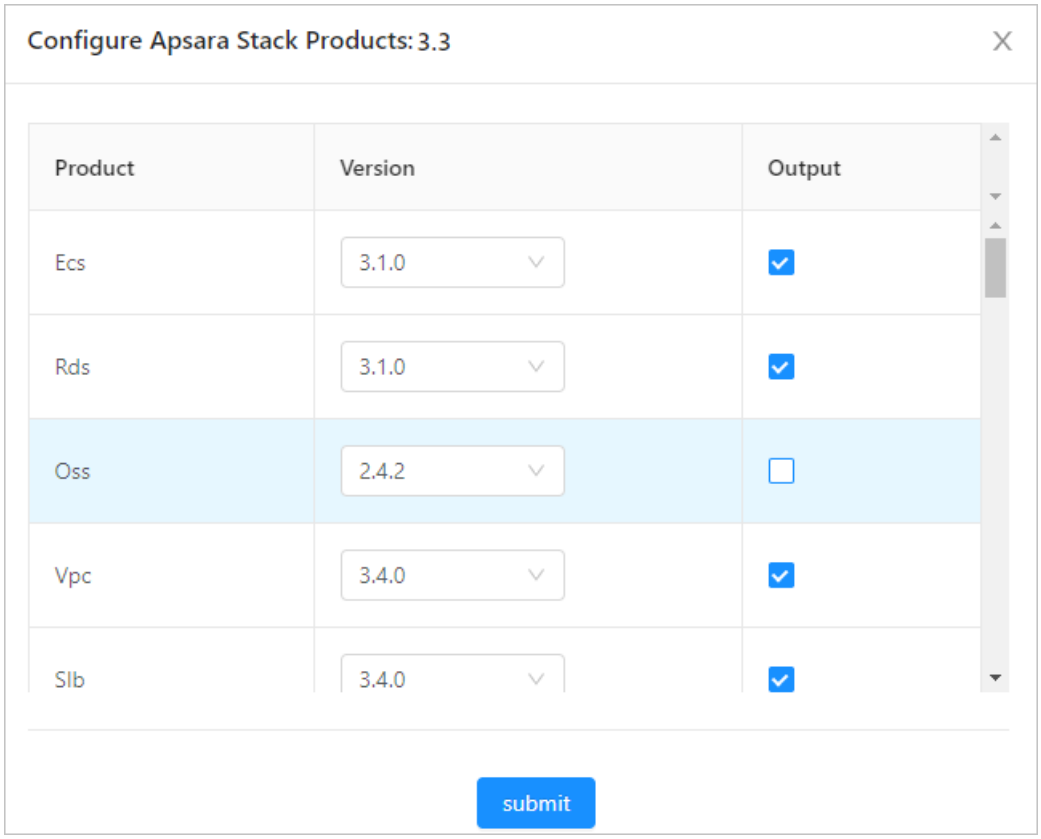

5. Click **Submit** to generate information of the products of the specified Apsara Stack version.

11.4.1.3. Compare versions

You can use the version comparison function to compare the product differences between two Apsara Stack versions. Based on these product differences, you can further learn about the differences of their APIs as well as of the definitions and parameters of these APIs.

Context

- Apsara Stack version: Each Apsara Stack version can have either one release version or one snapshot of the on-premises environment or deployment environment. Versions are distinguished by its version name and description.
- Product version: the specific version of a product when each Apsara Stack version is released, such as RDS 3.7.0. An Apsara Stack release version can have only one version of a specific product.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack**

Versions page appears.

3. In the upper-left corner of the page, click **Compare Versions**.
4. On the **Select Version** tab, select two versions to be compared, and click **Next**.

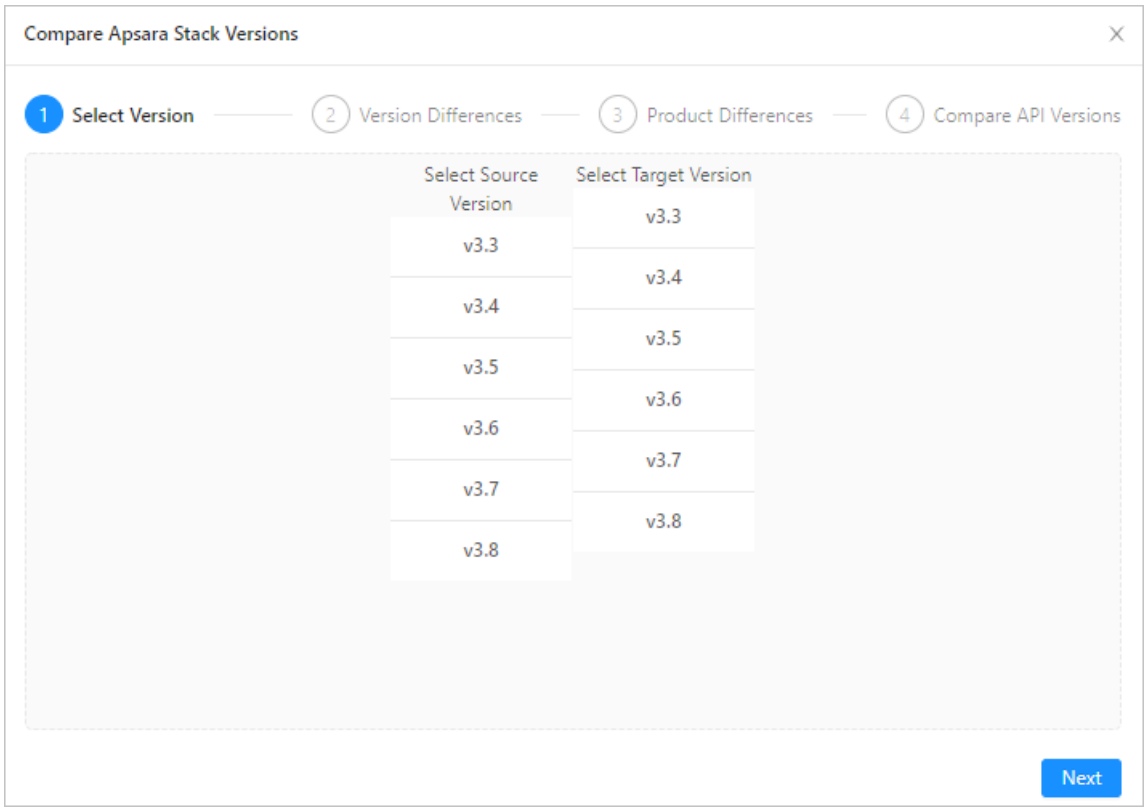

5. On the **Version Differences** tab, you can compare the product differences between the two versions. For example, you can view versions for which product information has been added or removed.
6. Click a product. Click **Next** to go to the **Product Differences** tab. You can compare the differences in product APIs between these two versions. You can view functions for which APIs have been added or removed, and APIs remain the same in these two versions.

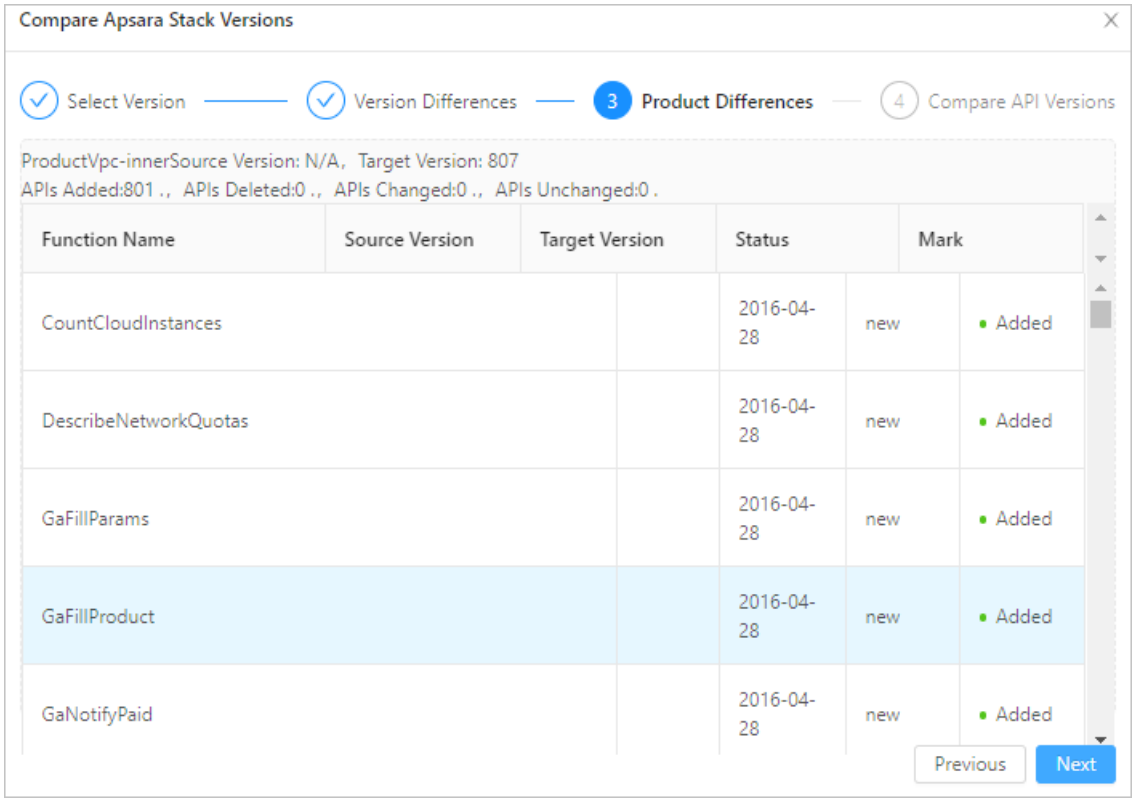

7. Click an API. Click **Next** to view the changes that are made to this API.

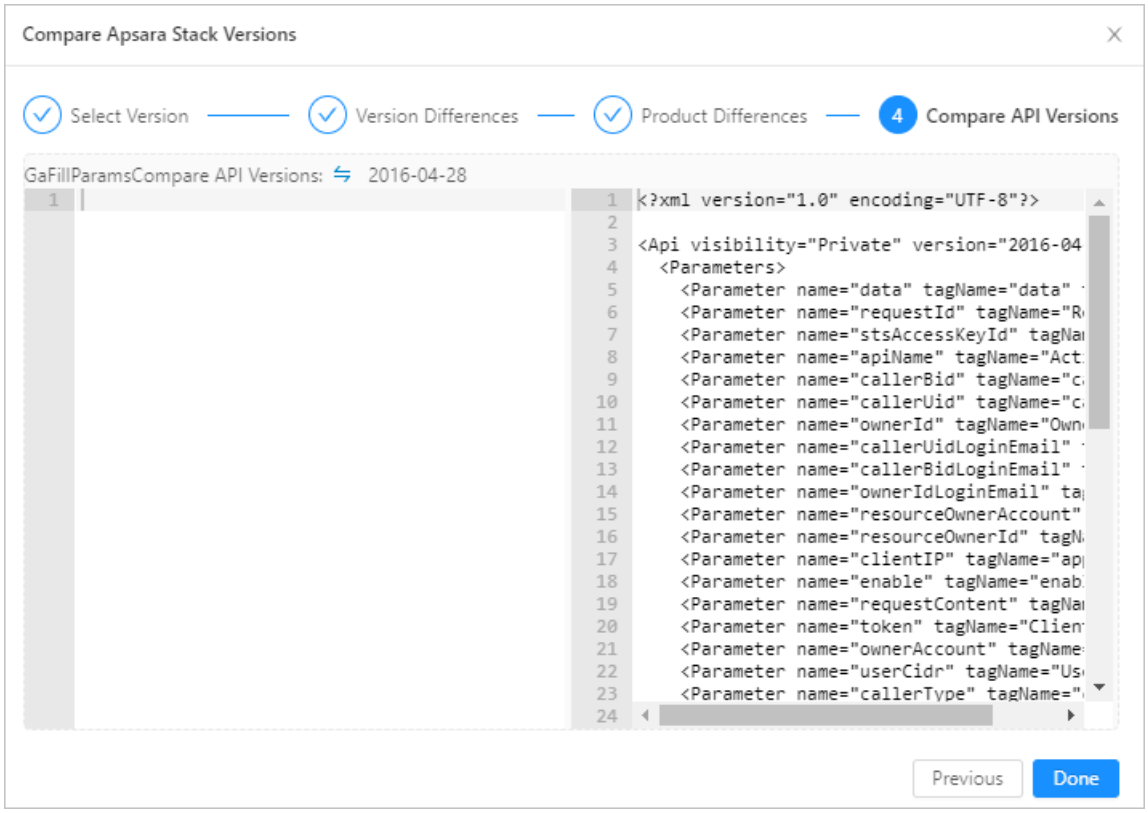

11.4.1.4. Remove information about Apsara Stack versions

If you no longer need the Apsara Stack version, you can remove its version and output information.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack Versions** page appears.
3. Click the icon in the **Actions** column corresponding to the Apsara Stack version information about which is to be removed.
4. In the message that appears, click **OK**.

11.4.2. Product baseline management

Product baselines are a set of configurations used by Apsara Stack products to define products, services, service roles, and applications. The Apsara Opsapi Management system provides basic information about products, services, and service roles. During initialization, the Apsara Opsapi Management system automatically scans all product baseline information in the Apsara Stack environment. You can use the system to scan the metadatabases and servers of services and service roles.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Product Baselines**. The **Product Baselines** page appears.
3. Follow these steps:
   - In the upper-right corner of the page, click **Scan Apsara Stack Environment** to scan metadatabases and servers that correspond to all products and update their information in the system.
   - Select a product from the drop-down list to query the service and service roles of the product.
   - Click the icon in the **Actions** column corresponding to the service role to scan the metadatabases of the service role.
   - Select a service role. click the icon in the **Actions** column corresponding to the service role to scan the server of the service role.

11.4.3. Product management

Operations and maintenance engineers can manage information of current Apsara Stack versions and product versions in real time.

Context

- Apsara Stack version and product version: Each Apsara Stack version can have only one specific version of products.
- Product version and SDK version: Each product version can have one SDK version.

11.4.3.1. Add information about products

You can add information about a product you need to manage.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. In the upper-right corner of the page, click **Add Product**.
4. In the **Add Product** dialog box that appears, set **Product Name** and **Product Description**.
5. Click **Submit** to add information about a product.

11.4.3.2. Add information about product versions

After you add information about a product, you need to add its product version and API version information for subsequent version management.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the Products page, click the icon in the **Actions** column corresponding to the product about which the version information is to be added.
4. In the **Add Version** dialog box that appears, set **Version** and **API Version**.

Parameter	Description
Version	The version of the current product.
API Version	The API version of the current product.

5. Click **Submit**.

11.4.3.3. Import information about APIs

You can import information about a preset API to the Apsara Opsapi Management system.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. In the product list, click the icon in the **Actions** column corresponding to the product to be managed.
4. In the dialog box that appears, select the product version from the left drop-down list.

If information about the API has been imported for the product version, this API is displayed in the APIs section.

5. In the upper-right corner, click **Import API**.
6. In the Import API dialog box that appears, set **Apsara Stack Version** and **API File**.

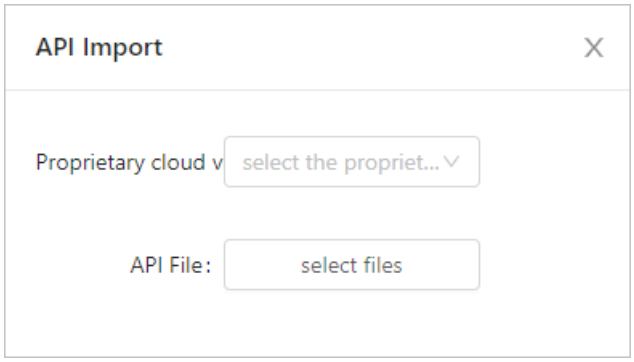

7. Click **OK** to import the API to the system.

11.4.3.4. Set SDK versions

A product has multiple SDK versions. You can set the SDK version of a product to obtain the SDK version of the product in the Apsara Stack release version.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. In the product list, click the icon in the **Actions** column corresponding to the product of which the SDK version is to be modified.
4. In the **SDK Settings** dialog box that appears, click **Modify** in the **Actions** column corresponding to the product version.
5. Select the specified SDK version from the drop-down list. Click **Submit** in the **Actions** column corresponding to the SDK version. The SDK version is modified.

11.4.3.5. Modify product names and descriptions

You can modify the name and description of a product.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the Products page, click the icon in the **Actions** column corresponding to the product information about which is to be modified.
4. In the dialog box that appears, modify the product name or description, and click **Submit**.

11.4.3.6. View information about product versions

When you need to learn about how to use a product, you can view information about the product version and API version.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the Products page, click the icon in the **Actions** column corresponding to the product about which the version information is to be viewed. You can view information about the product version and API version.

11.4.3.7. Modify information about product versions

You can modify information about a product version or API version as needed.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the **Products** page, click the icon in the **Actions** column corresponding to the product about which the version information is to be modified.
4. In the **View Version** dialog box that appears, click the icon in the **Actions** column corresponding to the version information about which is to be modified.
5. In the dialog box that appears, modify information about the product version and API version.

11.4.3.8. Remove information about product versions

You can remove information about a product version that is not applicable.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the **Products** page, click the icon in the **Actions** column corresponding to the product about which the version information is to be removed.
4. In the **View Version** dialog box that appears, click the icon in the **Actions** column corresponding to the version information about which is to be removed.
5. In the message that appears, click **OK**.

11.4.3.9. Remove information about products

You can remove information about a product that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.

2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. Click the icon in the **Actions** column corresponding to the product information about which is to be removed.

11.4.3.10. Remove information about product APIs

You can remove information about APIs that are not applicable to a product.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. In the product list, click the icon in the **Actions** column corresponding to the product to be managed.
4. In the dialog box that appears, select the product version from the left drop-down list.
   Imported APIs are displayed in the APIs section.
5. Click the icon corresponding to the API information about which is to be removed.
6. In the message that appears, click **Yes**.

11.4.4. SDK management

The Apsara Opsapi Management system enables you to customize SDKs. You can customize an SDK as needed to export APIs of Apsara Stack products of a specific version. You can also modify and delete the customized SDK.

11.4.4.1. Customize SDKs

The Apsara Opsapi Management system provides a tool to customize SDKs. The tool enables you to customize multiple combinations of SDKs for APIs within and across Apsara Stack products of specified versions.

Context

Each product has corresponding SDKs for different programming languages. The Apsara Opsapi Management system supports only SDKs for Java and Python.

Each SDK consists of an SDK core and an SDK model. The SDK core is the framework of the SDK. It is used to generate HTTP requests or requests of other protocols. The SDK core is fixed. You do not need to generate it each time. The SDK model defines the request parameters and responses of each API.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > SDK Tools**. The **SDK Tools** page appears.
3. In the upper-right corner of the page, click **Customize SDK**.
4. Set Apsara Stack Version, Product Name, Product Version, SDK Version, API Version, and Language. The corresponding APIs are displayed in the following APIs section.

5. Select APIs and click **Create SDK**.

   After an SDK is created, you can view the created SDK in the SDK list on the **SDK Tools** page.

6. (Optional)Click the link in the **Download** column corresponding to the product to download the SDK.

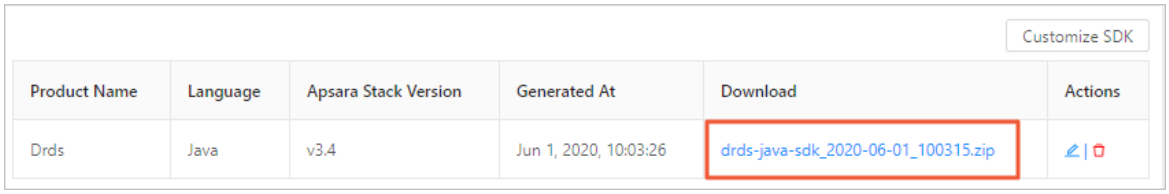

> **Note** The SDK generated in the Apsara Opsapi Management system is the SDK model. To use this SDK, you need to download the SDK core. You can download the SDK core from the Alibaba Cloud official website or obtain the SDK core from the Apsara Stack after-sales service.

11.4.4.2. Modify SDKs

When you need to update an SDK, you can upload an SDK to replace the original SDK.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Versions > SDK Tools**. The **SDK Tools** page appears.
3. In the product list, click the icon in the **Actions** column corresponding to the product of which the SDK is to be modified.
4. In the dialog box that appears, upload an SDK as prompted and click **Submit**.

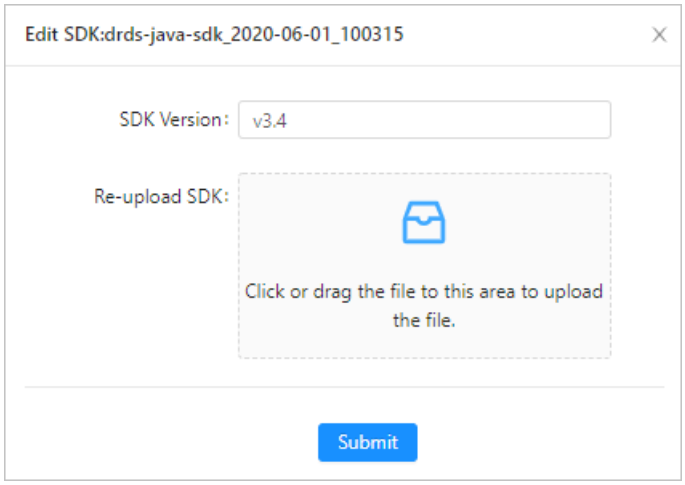

11.4.4.3. Delete SDKs

You can delete an SDK that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.

2. In the left-side navigation pane, choose **Versions > SDK Tools**. The **SDK Tools** page appears.
3. In the product list, click the icon in the **Actions** column corresponding to the product of which the SDK is to be deleted.
4. In the message that appears, click **Yes**.

11.5. Test management

To facilitate API tests in the Apsara Opsapi Management system, the system provides the test management function. Each API can be saved as a test case during the test. A test case contains the request parameters of these APIs. You can associate multiple test cases to create a test set. You can choose to run one test case and one test set at a time. You can view execution results on the Execution History page.

11.5.1. Test cases

A test case is used to test a specified API.

11.5.1.1. Modify test cases

You can modify the request parameters of a test case as needed.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.
3. On the Test Cases page, click the icon in the **Actions** column corresponding to a test case.
4. In the Edit Test Case dialog box that appears, modify values of the regionId, accessKeyId, accessKeySecret, Product, apiId, and apiVersion parameters.

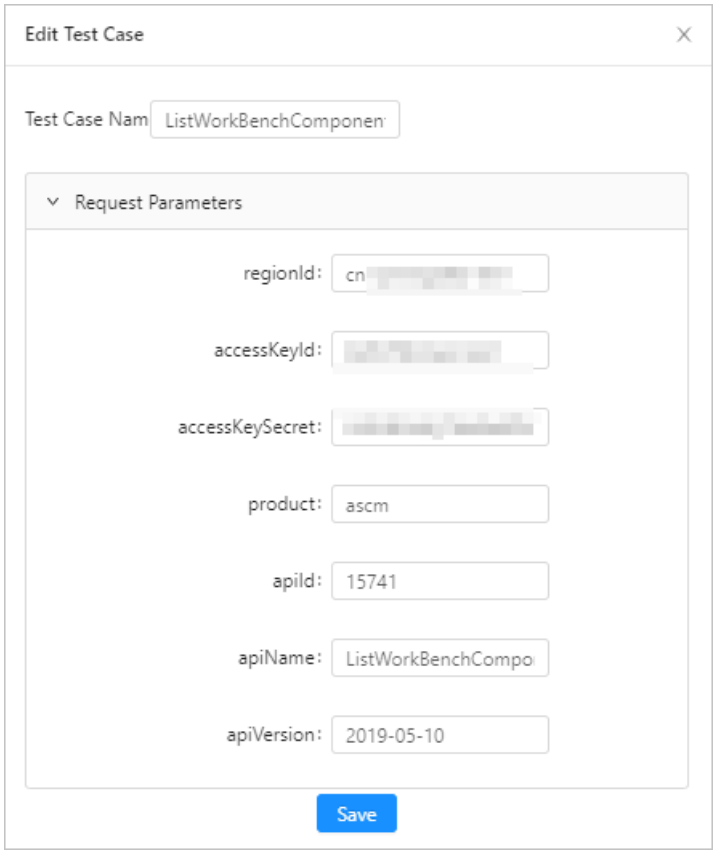

5. Click **Save**.

11.5.1.2. Run test cases

You can run a test case as needed.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.
3. On the Test Cases page, click the

   icon in the **Actions** column corresponding to a test case. After the test case is run, you can choose **Testing Platform > Execution History** to view the execution results on the Execution History page.

11.5.1.3. Delete test cases

You can delete a test case that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.

3. On the Test Cases page, click the icon in the **Actions** column corresponding to the test case to be deleted.
4. In the message that appears, click **Yes**.

11.5.2. Test sets

A test set consists of multiple associated test cases.

11.5.2.1. Create test sets

You can create a test set based on test requirements.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
3. On the Test Sets page, click **Create Test Set**.
4. In the dialog box that appears, enter the test set name and description. Click **Save**. We recommend that you enter a test set name that is easily identified.

11.5.2.2. Associate test cases

You can associate test cases with a test set to manage test cases in a unified manner.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
3. On the Test Sets page, click the icon in the **Actions** column corresponding to the test set.
4. (Optional)You can update the name and description of the test set and click **Save**.
5. Click **Relate to Test Case**.
6. In the dialog box that appears, search for and select the test case to be associated. You can select multiple test cases and add them to the test set.
7. Click **Save**.

11.5.2.3. Run test sets

You can run a test set to check whether the APIs in the test set are available.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page

appears.

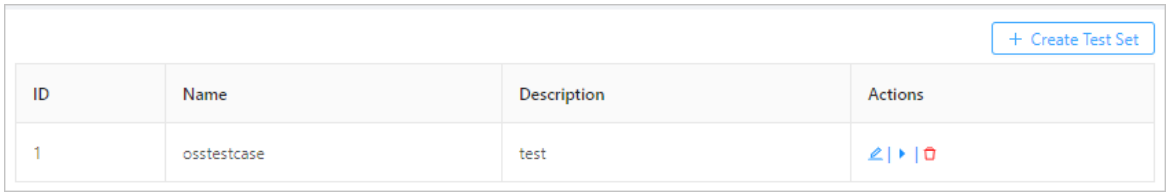

3. On the Test Sets page, click the ▶ icon in the **Actions** column corresponding to a test set. The test cases in the test set start to run. After the test set is run, you can choose **Testing Platform > Execution History** to view the execution results on the Execution History page.

11.5.2.4. Delete test sets

You can delete test sets that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
3. On the Test Sets page, click the 🗑 icon in the **Actions** column corresponding to the test case to be deleted.
4. In the message that appears, click **Yes**.

11.5.3. View execution history of test cases

You can view information about the API for which a test case was executed, including the corresponding product, version information, execution time, and execution status.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **Testing Platform > Execution History**. The **Execution History** page appears.
3. On the Execution History page, click the details icon in the **Details** column corresponding to the API to be viewed.
4. In the **Execution Details** dialog box that appears, view the execution details of the test case, including request parameters and responses.

11.6. System management

11.6.1. Metadatabase management

You can add or remove information about metadatabase in the Apsara Opsapi Management system.

11.6.1.1. View information about added metadatabases

The Apsara Opsapi Management system automatically scans all metadatabases in the Apsara Stack environment during initialization. The Apsara Opsapi Management system allows you to scan all metadatabases to view information about the added metadatabases. You can also manually add information about the metadatabases.

Context

The metadatabase information contains the domain name, database name, port, and server that are used in Apsara Stack products.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **System Management > Metabase**. The **Metabase** page appears.
3. Use one of the following methods to view the connection information about metadatabases:
   - Scan metadatabases

     Click **Scan Metabase** to scan all added metadatabases in the Aspara Stack environment.
   - Add information about metadatabases

     Click **Add Metabase**. In the **Add Metabase** dialog box that appears, set Product Name, Metabase Name, Metabase Server, Metabase Port, Username, and Password. Click **Submit**.

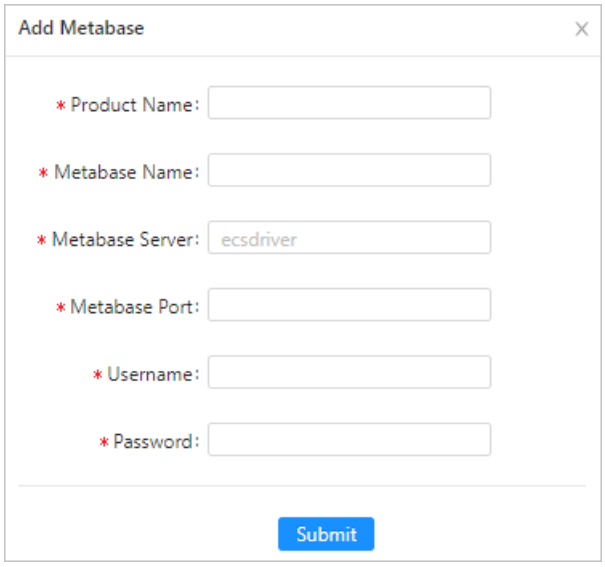

Parameter	Description
Product Name	The product to which the metadatabase belongs.
MetaBase Name	The name of the metadatabase.
MetaBase Server	The name of the server where the metadatabase is located.
MetaBase Port	The access port of the metadatabase.
Username	The username used to log on to the metadatabase.

Parameter	Description
Password	The password used to log on to the metadatabase.

You can view information about the added metadatabases in the metadatabase list.

11.6.1.2. View connection information about metadatabases

You can view connection information about a metadatabase on the **Metabase** page.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **System Management** > **Metabase**. The **Metabase** page appears.
3. In the upper-left corner of the page, select the product to view from the drop-down list.
4. On the Metabase page, click the

   icon in the **Actions** column corresponding to a metadatabase. In the message that appears, you can view the metadatabase connection information.

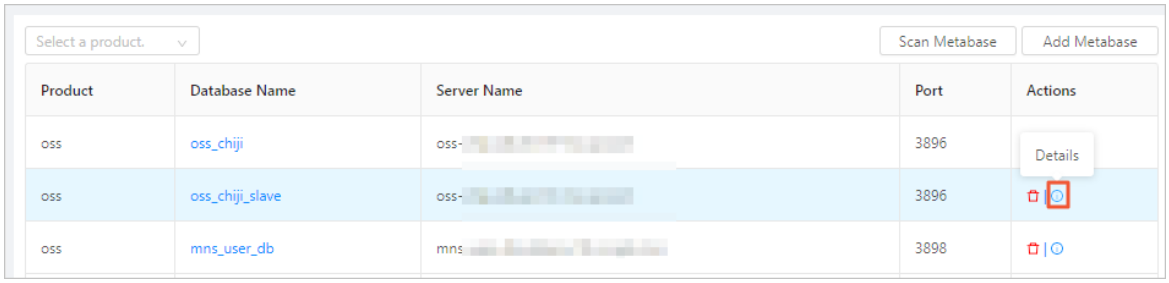

11.6.1.3. Remove information about metadatabases

To facilitate management, you can remove information about metadatabases that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **System Management** > **Metabase**. The **Metabase** page appears.
3. On the Metabase page, click the icon in the **Actions** column corresponding to the metadatabase information about which is to be removed.
4. In the message that appears, click **Yes**.

11.6.2. Server management

You can add or remove information about servers in the Apsara Opsapi Management system.

11.6.2.1. View information about added servers

The Apsara Opsapi Management system automatically scans all servers (including physical servers and VMs) in the Apsara Stack environment during initialization. When new servers are added to the Apsara Stack environment, you can scan servers to view information about the added servers. You can also add information about the added servers.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **System Management > Server Management**. The **Server Management** page appears.
3. Use either of the following methods to view information about added servers:
   - Scan servers

     Click **Scan Server** to scan the information about all servers in the Apsara Stack environment.
   - Add information about new servers

     Click **Add Server**. In the **Add Server** dialog box that appears, set Name, Server Name, Server IP, SSH Port, and SSH User. Upload the SSH private key. Click **Submit**.

     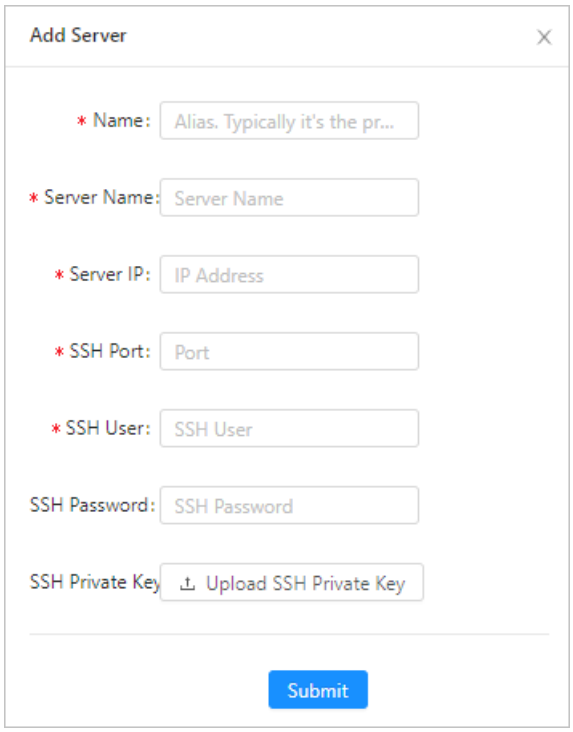

     You can view information about the added servers in the server list.

11.6.2.2. Remove server information

To facilitate management, you can remove information about servers that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.

2. In the left-side navigation pane, choose **System Management > Server Management**. The **Server Management** page appears.
3. On the Server Management page, click the icon in the **Actions** column corresponding to the server information about which is to be removed.

11.6.3. Audit APIs

You can view call records of all Opsapis. The records contain the specific API, statuses, time, and result of each call.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **System Management > API Audit**. The **API Audit** page appears.
3. On the API Audit page, click the icon in the **Actions** column corresponding to the API. You can view the call result of this API.

11.6.4. View logs

You can view API logs, error logs, and Egg logs to better maintain the back end.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose **System Management > Log Management**. The **Log Management** page appears.
3. View the details of all logs on the **API Log**, **Error Log**, and **Egg Log** tabs.

> **Note** You can modify log levels in the Apsara Opsapi Management system in real time. The level you set is valid only during the active service operating period. If you restart the service, the default level remains.

12.Appendix

12.1. Operation Access Manager (OAM)

12.1.1. OAM introduction

Overview

Operation Access Manager (OAM) is a centralized permission management platform of Apsara Stack Operations (ASO). OAM uses a simplified role-based access control (RBAC) model. Administrators can use OAM to assign roles to operations personnel, granting them corresponding operation permissions to operations systems.

OAM permission model

In RBAC, administrators do not directly grant system operation permissions to users. Instead, they create a collection of roles between a collection of users and a collection of permissions. Each role corresponds to a group of permissions. If a role is assigned to a user, the user is granted all the operation permissions of that role. Therefore, when creating a user, administrators are only required to assign a role to the user, saving the trouble to grant specific permissions to the user. In addition, the frequency of role permission changes is less than that of user permission changes, simplifying the user permission management and reducing the system overhead.

See the OAM permission model as follows.

Permission model

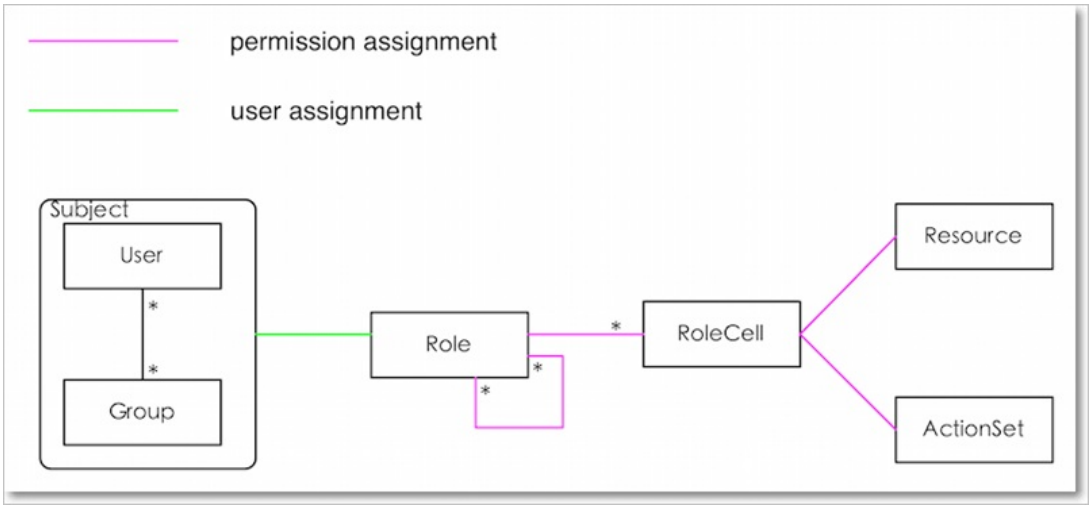

12.1.2. Instructions

Before using Operation Access Manager (OAM), you must know the following basic concepts about permission management.

subject

Operators of the access control system. OAM has two types of subjects: users and groups.

user

Administrators and operators of operations systems.

group

A collection of users.

role

The core of the role-based access control (RBAC) system.

Generally, a role can be regarded as a collection of permissions. A role can contain multiple RoleCells or roles.

RoleHierarchy

In the OAM system, a role can contain other roles to form RoleHierarchy.

RoleCell

The specific description of a permission. A RoleCell consists of resources, ActionSets, and available authorizations.

resource

The description of an authorized object. For more information about resources of operations platforms, see Permission lists of operations platforms.

ActionSet

The description of authorized actions. An ActionSet can contain multiple actions. For more information about actions of operations platforms, see Permission lists of operations platforms.

available authorizations

The maximum number of authorizations in the cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.

For example, if administrator A sets **Available Authorizations** to 5 when granting a permission to administrator B, the permission can be granted for another five times at most. When administrator B grants the permission to administrator C, the value of **Available Authorizations** cannot be greater than 4. If **Available Authorizations** is set to 0 when administrator B grants the permission to operator D, operator D can only use the permission but cannot grant it to others.

> **Note** Currently, OAM does not support the cascaded revocation for cascaded authorization. Therefore, administrator C and operator D still have the permission even if the permission is revoked for administrator B.

12.1.3. Quick Start

By completing the steps in this guide, you will learn how to create and assign roles for O&M.

12.1.3.1. Log on to OAM

This topic describes how to log on to Operation Administrator Manager (OAM).

Prerequisites

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

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

Procedure

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

   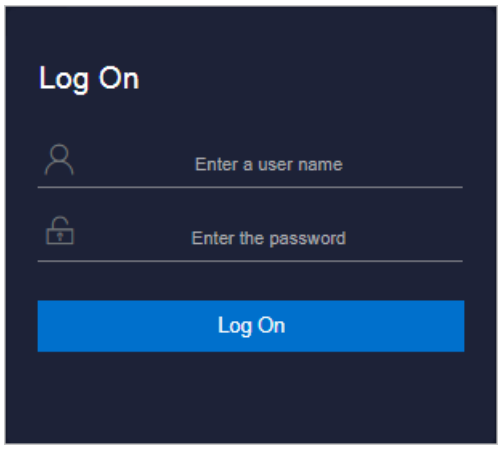

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.
3. Enter your username and password.

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

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

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

   - It must contain uppercase and lowercase letters.
   - It must contain digits.
   - It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%).
   - It must be 10 to 20 characters in length.
4. Click **Log On** to go to the **ASO** console.
5. In the left-side navigation pane, choose **Products > Product List**.
6. In the **Apsara Stack O&M > Basic O&M** section, click **OAM**.

12.1.3.2. Create groups

You can create user groups for centralized management.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. On the **Owned Groups** page, click **Create Group** in the upper-right corner. In the **Create Group** dialog box that appears, set **Group Name** and **Description**.

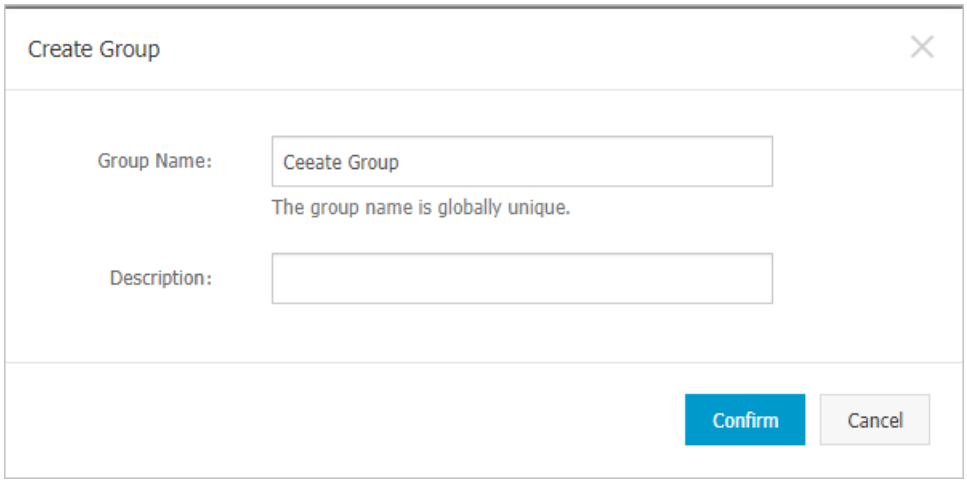

4. Click **Confirm**.After the group is created, it is displayed on the **Owned Groups** page.

12.1.3.3. Add group members

You can add members to an existing group to grant permissions to the group members in a centralized manner.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.
4. In the upper-right corner of the **Group Member** section, click **Add Member**.

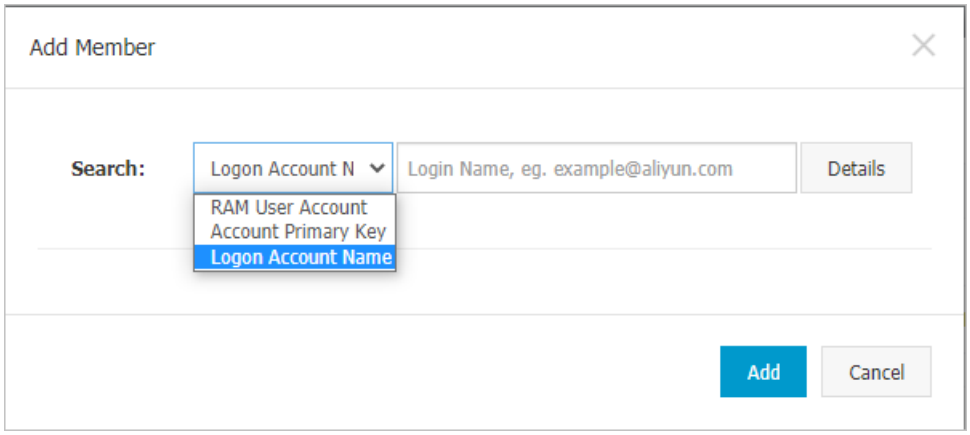

5. Select a search mode, enter the corresponding information, and click **Details**. Details of the specified user are displayed.

   Three search modes are available:

   - **RAM User Account**: Search in the format of RAM user@Apsara Stack tenant account ID.
   - **Account Primary Key**: Search by using the unique ID of the Apsara Stack tenant account.
   - **Logon Account Name**: Search by using the logon name of the Apsara Stack tenant account.
6. Click **Add**.
7. You can repeat the preceding steps to add multiple group members.To remove a member from a group, click **Remove** in the Actions column corresponding to the member.

12.1.3.4. Add group roles

You can add roles to an existing group.

Prerequisites

- The role to be added is created. For more information, see Create roles.
- You are the owner of the group and the role.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.
4. In the upper-right corner of the **Role List** section, click **Add Role**.

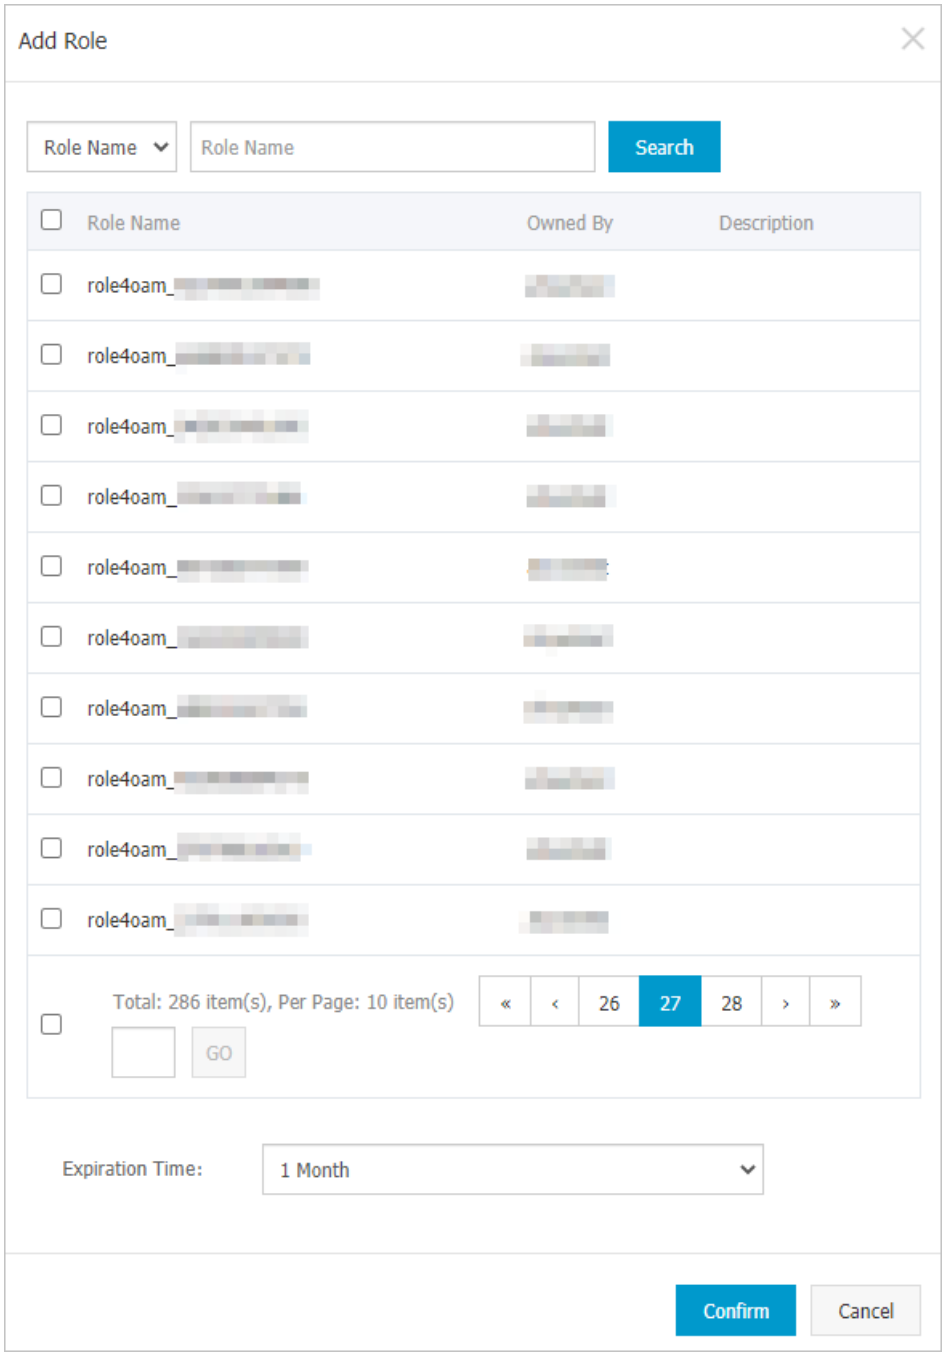

5. Search for roles by **Role Name**. Select one or more roles and set Expiration Time.
6. Click **Confirm**.

   To remove a role from a group, find the role in **Role List**, and click **Remove** in the **Actions** column.

12.1.3.5. Create roles

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. On the **Owned Roles** page, click **Create Role** in the upper-right corner.

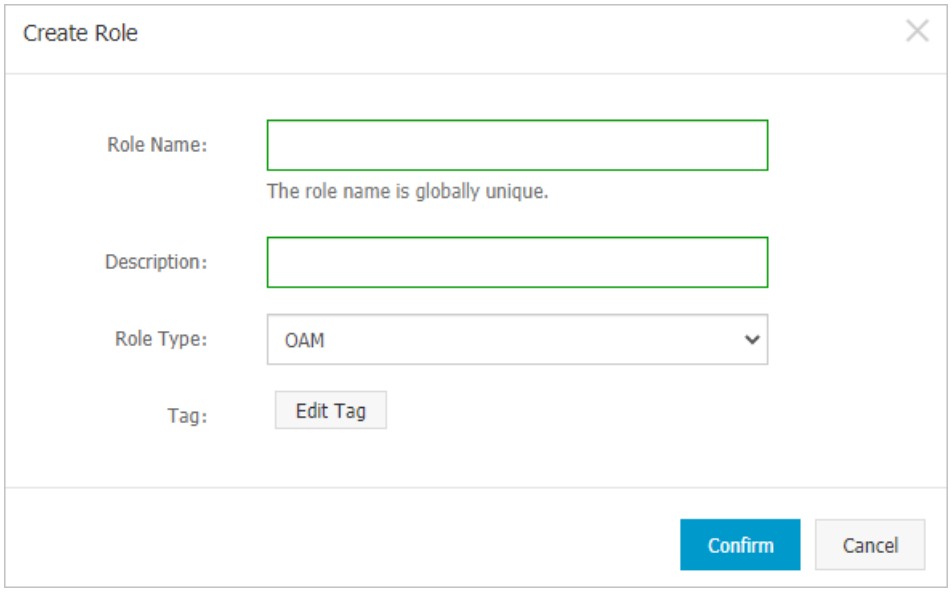

4. In the Create Role dialog box that appears, set **Role Name**, **Description**, and **Role Type**.
5. (Optional)Configure the role tags, which can be used to filter roles.
   i. Click **Edit Tags**.

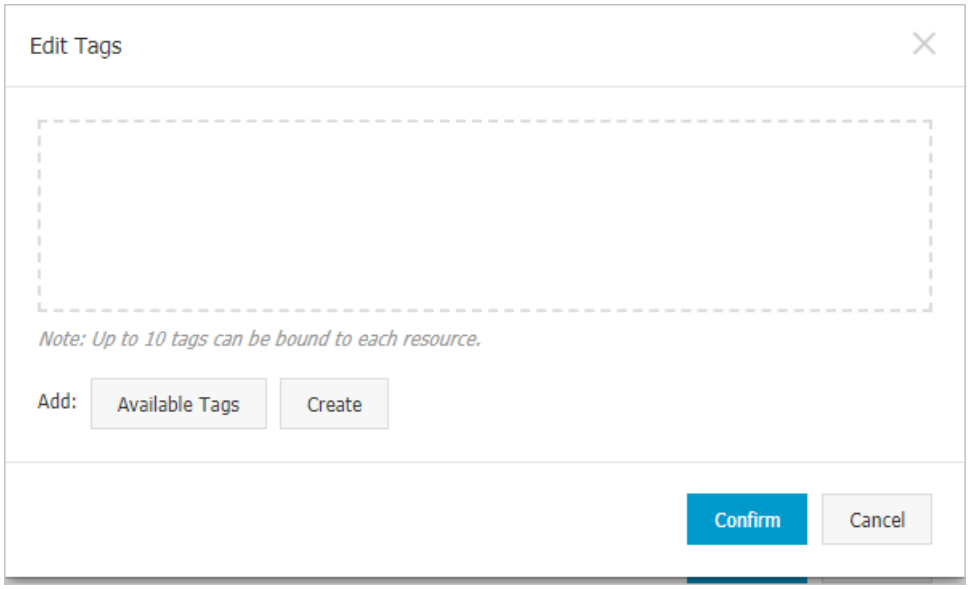

   ii. In the **Edit Tags** dialog box that appears, click **Create**.

iii. Set **Key** and **Value** for the tag and click **Confirm**.

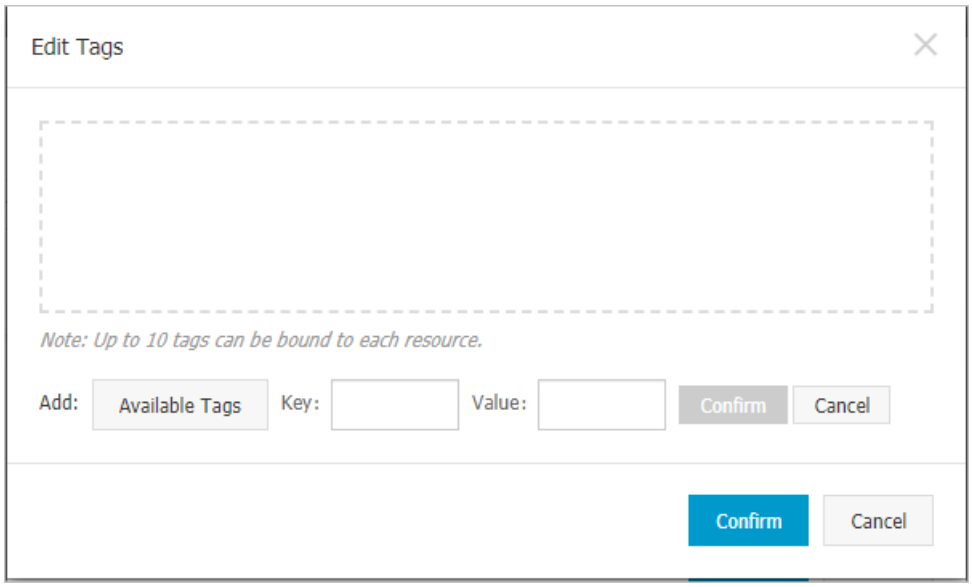

iv. Repeat the preceding step to create more tags.

The created tags are displayed inside the dotted box.

v. Click **Confirm** to create the tags and exit the **Edit Tags** dialog box.

6. Click **Confirm** to create the role.

12.1.3.6. Add inherited roles to a role

You can add inherited roles to a role to grant the permissions of the inherited roles to the role.

Prerequisites

You are the owner of the current role and the inherited role to be added.

For more information about how to query your owned roles, see Query roles.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add an inherited role and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Inherited Role** tab.
5. Click **Add Role**. In the **Add Role** dialog box that appears, search for roles by **Role Name**. Select one or more roles.

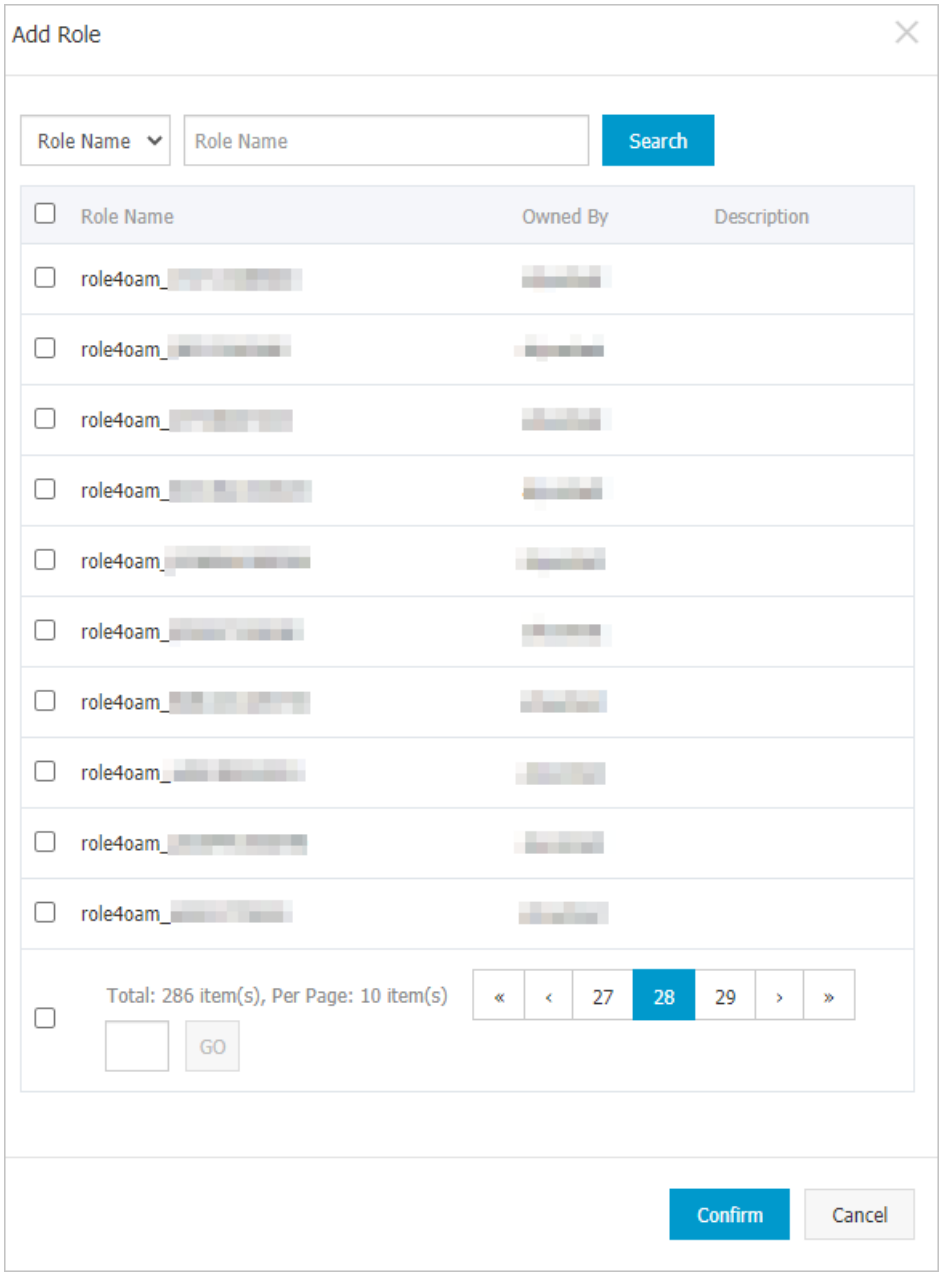

6. Click **Confirm**.

12.1.3.7. Add resources to a role

You must add resources to a created role.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add a resource and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Resource List** tab.
5. Click **Add Resource**.

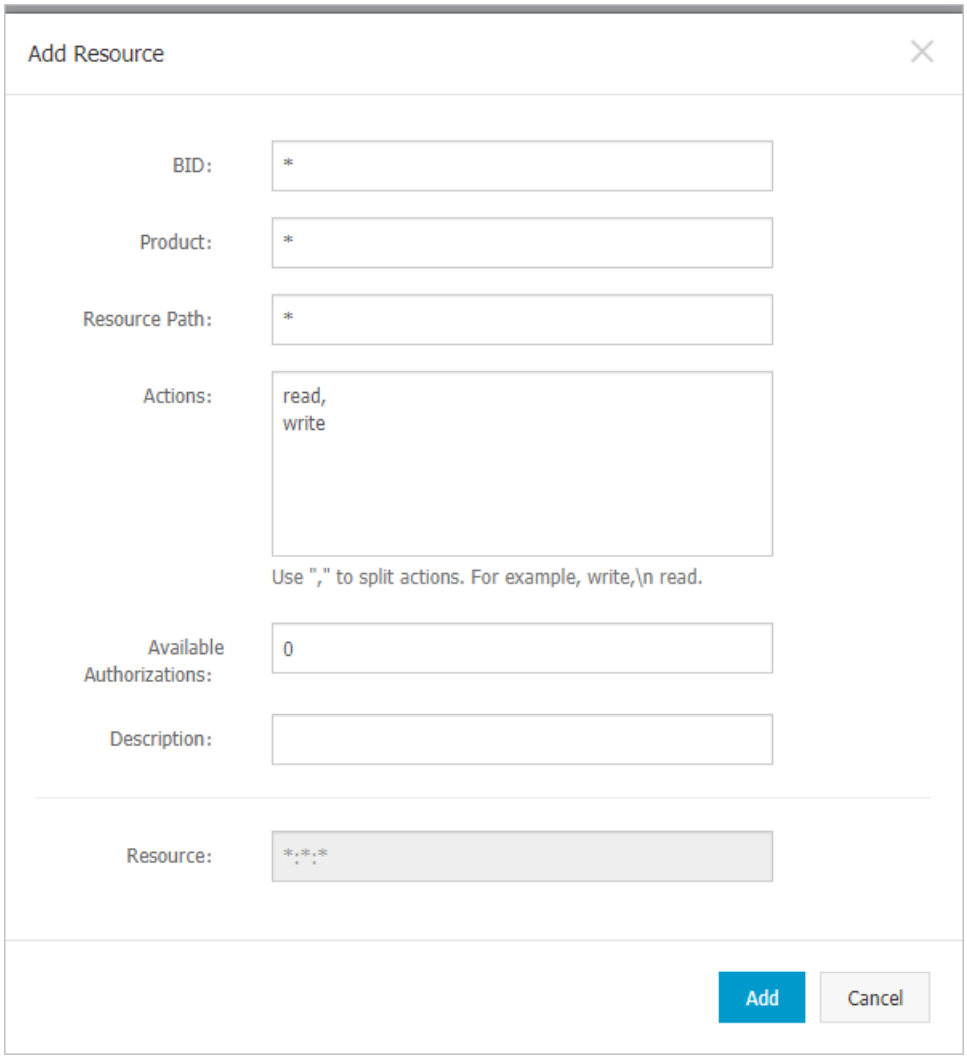

6. In the **Add Resource** dialog box, complete the configurations. For more information, see Parameters.

Parameters

Parameter	Description
BID	The deployment region ID.
Product	The cloud product to be added, such as rds. **Note** The cloud product name must be lowercase. For example, enter **rds** instead of **RDS**.
Resource Path	The resources of the cloud product. For more information about resources of the O&M platforms, see Permission lists of operations platforms.

Parameter	Description
Actions	An action set, which can contain multiple actions. For more information about actions on the O&M platforms, see Permission lists of operations platforms.
Available Authorizations	The maximum number of authorizations in cascaded authorization, which must be an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.
Description	The description of the resource.

7. Click **Add**.

12.1.3.8. Add authorized users to a role

You can assign an existing role to users or user groups.

Prerequisites

The corresponding users or user groups are created. Users are created in the Apsara Stack Cloud Management (ASCM) console. For more information about how to create a user group, see Create groups.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to which you want to add an authorized user and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Authorized Users** tab.
5. Click **Add User** in the upper-right corner.

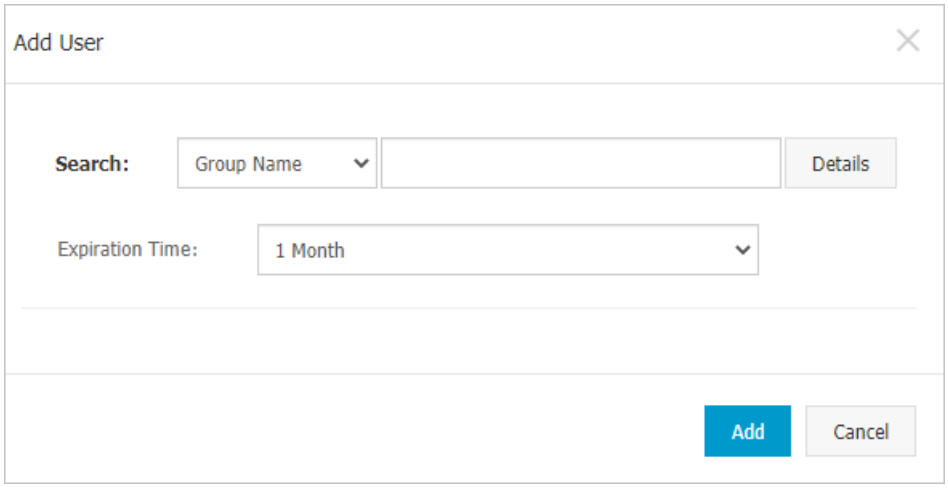

6. Select a search mode and enter the corresponding information.

   Four search modes are available:

- **RAM User Account**: Search in the format of *RAM user@Apsara Stack tenant account ID*.
- **Account Primary Key**: Search by using the unique ID of the Apsara Stack tenant account.
- **Logon Account Name**: Search by using the logon name of the Apsara Stack tenant account.
- **Group Name**: Search by group name.

> **Note** You can search for a single user or user group. For more information about how to create a user group, see Create groups.

7. Set Expiration Time.When the specified expiration time is due, the user no longer has the permissions of the role. To authorize the user again, the role creator must click **Renew** in the Actions column corresponding to the authorized user on the **Authorized Users** tab to modify the expiration time.
8. Click **Add** to assign the role to the user.To cancel the authorization, click **Remove** in the Actions column corresponding to the authorized user on the **Authorized Users** tab.

12.1.4. Manage groups

Group Management allows you to view, modify, or delete groups.

12.1.4.1. Modify group information

After you create a group, you can modify the group name and description on the Group Information page.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.
4. On the **Group Information** page, click **Modify** in the upper-right corner.
5. In the **Modify Group** dialog box that appears, modify the group name and description.
6. Click **Confirm**.

12.1.4.2. View group role details

You can view information about the inherited roles, resource list, and inheritance tree of a group role.

Prerequisites

A role is added to the group.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group whose name and description you want to modify and click **Manage** in the **Actions** column.

4. In **Role List** section, click **Details** in the Actions column corresponding to a role.
5. On the **Role Information** page, perform the following operations:
   - Click the **Inherited Role** tab to view the information about the inherited roles of the role.

     To view the detailed information of an inherited role, click **Details** in the **Actions** column corresponding to the inherited role.
   - Click the **Resource List** tab to view the resource information of the role.

     For information about how to add other resources to this role, see Add resources to a role.
   - Click the **Inheritance Tree** tab to view the basic information and resource information of the role and its inherited roles by using the inheritance tree on the left.

12.1.4.3. Delete groups

You can delete groups that are no longer needed.

Prerequisites

The group to be deleted does not contain members.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group to be deleted and click **Delete** in the **Actions** column.

12.1.4.4. View authorized groups

You can view the groups to which you are added on the Authorized Groups page.

Context

You can view only the groups to which you belong, but cannot view groups of other users.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Group Management > Authorized Groups**.
3. On the **Authorized Groups** page, view the name, owner, description, and modification time of the group to which you belong.

12.1.5. Manage roles

Role Management allows you to view, modify, transfer, or delete roles.

12.1.5.1. Query roles

You can view your owned roles on the Owned Roles page.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Enter a role name in the Role Name field and click **Search** to search for roles that meet the search criteria.

   > **Note** If the role you want to search for has a tag, you can click **Tag** and select the tag key to search for the role based on the tag.

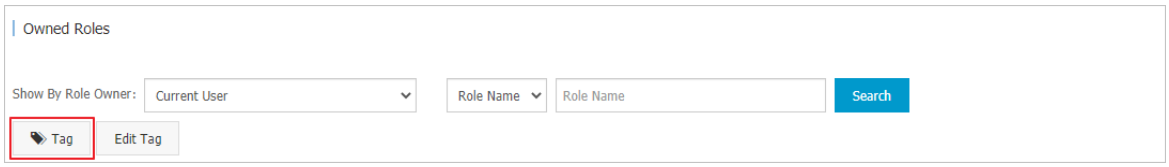

12.1.5.2. Modify role information

After you create a role, you can modify the role information.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role whose information you want to modify and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click **Modify** in the upper-right corner.
5. In the **Modify Role** dialog box that appears, set **Role Name**, **Description**, **Role Type** and **Tag**.
6. Click **Confirm**.

12.1.5.3. View the role inheritance tree

You can view the role inheritance tree to learn about the basic information and resource information of a role and its inherited roles.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role whose information you want to modify and click **Manage** in the **Actions** column.
4. On the **Role Information** page, click the **Inheritance Tree** tab.

   View the basic information and resource information of the role and its inherited roles by using the inheritance tree on the left.

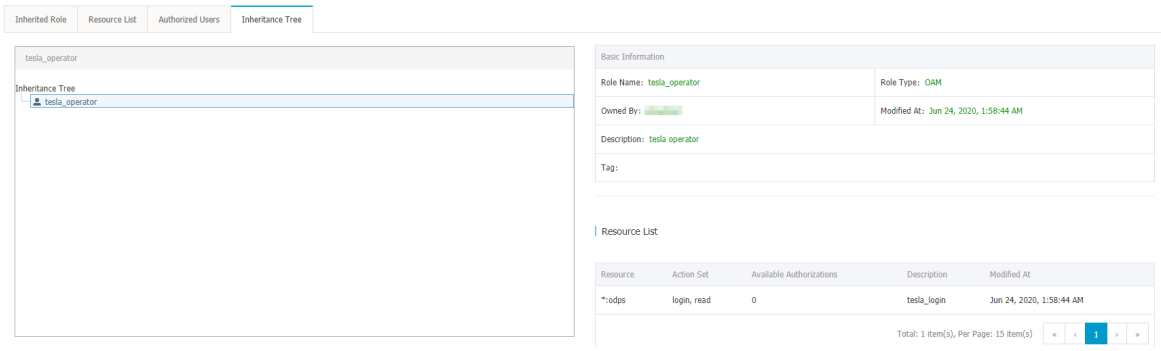

12.1.5.4. Transfer roles

You can transfer roles to other groups or users based on business requirements.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. On the **Owned Roles** page, configure the search condition and search for the roles to be transferred.
4. Select one or more roles in the search results and click **Transfer** in the lower-left corner.
5. In the **Transfer** dialog box that appears, select a search mode, enter the corresponding information, and then click **Details**. Details of the user or group are displayed.

   Four search modes are available:

   - **RAM User Account**: Search in the format of RAM user@Apsara Stack tenant account ID.
   - **Account Primary Key**: Search by using the unique ID of the Apsara Stack tenant account.
   - **Logon Account Name**: Search by using the logon name of the Apsara Stack tenant account.
   - **Group Name**: Search by group name.

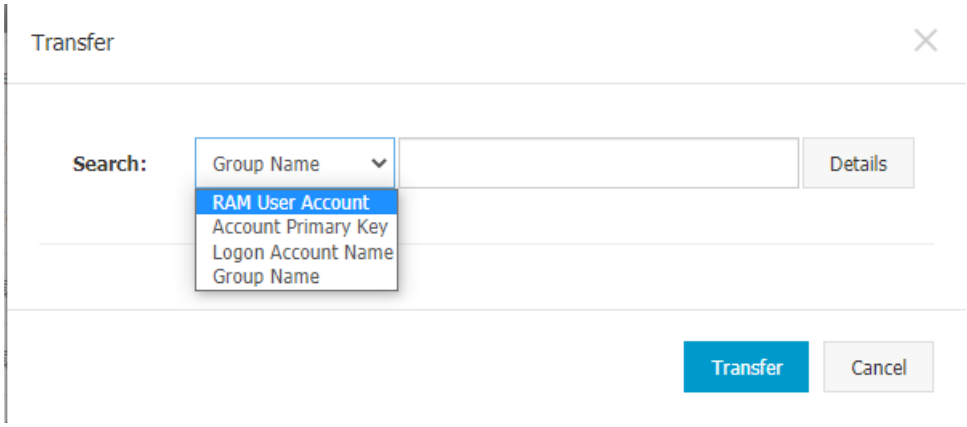

6. Click **Transfer**.

12.1.5.5. Delete a role

You can delete a role that is no longer in use according to business requirements.

Prerequisites

The role to be deleted does not contain inherited roles, resources, or authorized users.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. At the right of the role to be deleted and then click **Delete**.

12.1.5.6. View assigned roles

You can view the roles assigned to you and permissions granted to the roles.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > Authorized Roles**.
3. On the **Authorized Roles** page, you can view the name, owner, description, modification time, and expiration time of each role assigned to you.You can also click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

12.1.5.7. View all roles

You can view all roles in Operation Administrator Manager (OAM) on the All Roles page.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, choose **Role Management > All Roles**.
3. On the **All Roles** page, view all the roles in the system.You can search for roles by **Role Name**.
4. Click **Details** in the Actions column corresponding to a role to view the inherited roles, resources, and inheritance tree of the role.

12.1.6. Search for resources

You can search for resources to view the roles to which the resources are assigned.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, click **Search Resource**.
3. Set **Resource** and **Action**, and click **Search** to search for the roles that meet the specified conditions.

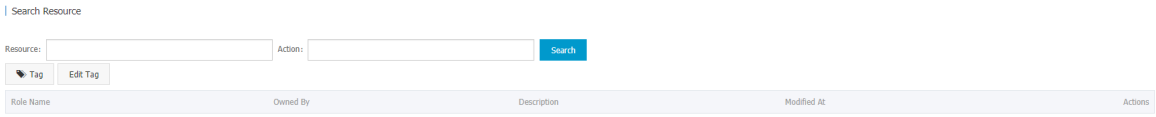

4. Click **Details** in the **Actions** column corresponding to a role to view the inherited roles, resources,

and inheritance tree of the role.

12.1.7. View personal information

You can view the personal information of the current user on the Personal Information page and test the user permissions.

Procedure

1. Log on to OAM.
2. In the left-side navigation pane, click **Personal Information**.
3. In the **Basic Information** section, you can view the username, type, creation time, AccessKey ID, and AccessKey secret of the current user.

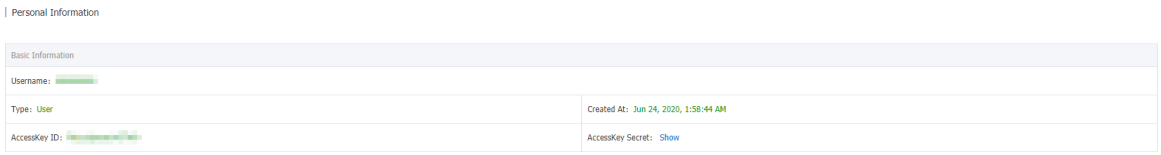

> **Note** You can click **Show** or **Hide** to show or hide the AccessKey secret.

4. In the **Test Permission** section, you can check whether the current user has a specific permission.
   i. Enter the resource information in the **Resource** field.

   > **Note** Use the English input method when you enter values in the **Resource** and **Action** fields.

   ii. Enter the permissions such as create, read, and write in the **Action** field. Separate multiple permissions with commas (,).

12.1.8. Default roles and permissions

12.1.8.1. Default roles and their functions

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

12.1.8.1.1. Default role of OAM

This topic describes the default role of Operation Access Manager (OAM) and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Super administrator	An administrator with root permissions	*:*	*	10

12.1.8.1.2. Default roles of Apsara Infrastructure Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Tianji_Project read-only	Has the read-only permission to Apsara Infrastructure Management Framework projects, which allows you to view the configurations and statuses of all projects and clusters	*:tianji:projects	["read"]	0
Tianji_Project administrator	Has all the permissions to Apsara Infrastructure Management Framework projects, which allows you to view and modify the configurations and statuses of all projects and clusters	*:tianji:projects	["*"]	0
Tianji_Service read-only	Has the read-only permission to Apsara Infrastructure Management Framework services, which allows you to view the configurations and templates of all services	*:tianji:services	["read"]	0

Role name	Role description	Resource	Actions	Available authorizations
Tianji_Service administrator	Has all the permissions to Apsara Infrastructure Management Framework services, which allows you to view and modify the configurations and templates of all services	*:tianji:services	["*"]	0
Tianji_IDC administrator	Has all the permissions to Apsara Infrastructure Management Framework data centers, which allows you to view and modify the data center information	*:tianji:idcs	["*"]	0
Tianji administrator	Has all the permissions to Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastructure Management Framework configurations	*:tianji	["*"]	0

12.1.8.1.3. Default role of Tianjimon

This topic describes the default role of Tianjimon and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations

Role name	Role description	Resource	Actions	Available authorizations
Tianjimon operations	Has all Tianjimon permissions, which allows you to perform basic monitoring and operations	26842:tianjimon:*	["*"]	0

12.1.8.1.4. Default roles of Opsapi

This topic describes the default roles of the Apsara Opsapi Management (Opsapi) system and the corresponding grant options.

Opsapi is a platform that manages O&M APIs and SDKs in the Apsara Stack environment in a centralized manner. This system also manages API and SDK versions.

The following table describes the default roles of Opsapi and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
Opsapi platform administrator	Has all the permissions on Opsapi.	*:opsapi:*	["read","write"]	0
Opsapi platform developer	Has the read permissions on Opsapi and the permissions to call API operations.	*:opsapi:*	["read","invoke"]	0
Opsapi common user	Has the read permissions on Opsapi.	*:opsapi:*	["read"]	0

12.1.8.1.5. Default roles of ASO

This topic describes the default roles of the Apsara Stack Operations (ASO) system and the corresponding grant options.

ASO is a centralized O&M management system that is developed for the Apsara Stack O&M personnel to perform centralized O&M operations.

The following table describes the default roles of ASO and the corresponding grant options.

Role	Role description	Resource	Action	Grant option

Role	Role description	Resource	Action	Grant option
ASO system administrator	Has the permissions to manage platform nodes, physical devices, and virtual resources, back up, restore, and migrate product data, and query and back up system logs.	*:aso:api-adapter:*	["read","write"]	0
		:aso:auth:	["read"]	0
		:aso:backup:	["read","write"]	0
		:aso:cmdb:	["read","write"]	0
		:aso:doc:	["read","write"]	0
		:aso:fullview:	["read","write"]	0
		:aso:init:	["read","write"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read","write"]	0
		:aso:lock:	["read","write"]	0
		:aso:physical:	["read","write"]	0
		:aso:psm:	["read","write"]	0
		:aso:scm:	["read","write"]	0
		:aso:serviceWhitelist:	["read","write"]	0
		:aso:slalink:	["read","write"]	0
		:aso:task:	["read","write"]	0
ASO security officer	Has the permissions to manage permissions, security polices, and network security, and review and analyze security logs and activities of security audit officers.	*:aso:auth:*	["read","write"]	0
		:aso:plat-access:	["read","write"]	0
		:aso:twoFactorAuth:	["read","write"]	0

Role	Role description	Resource	Action	Grant option
ASO security auditor	Has the permissions to audit, track, and analyze the activities of the system administrator and security officer.	*:aso:audit:*	["read","write"]	0
		:aso:auth:	["read"]	0
		:aso:serviceWhitelist:	["read"]	0
ASO product O&M officer	Has the permissions to perform O&M operations such as data import and export, modification, configuration, upgrade, and troubleshooting coordination.	*:aso:api-adapter:*	["read"]	0
		:aso:backup:	["read"]	0
		:aso:cmdb:	["read"]	0
		:aso:doc:	["read"]	0
		:aso:fullview:	["read","write"]	0
		:aso:init:	["read"]	0
		:aso:inventory:	["read","write"]	0
		:aso:itil:	["read"]	0
		:aso:lock:	["read"]	0
		:aso:physical:	["read","write"]	0
		:aso:psm:	["read"]	0
		:aso:scm:	["read"]	0
		:aso:slalink:	["read"]	0
		:aso:task:	["read"]	0
ASO common O&M	Has the permissions to perform daily health checks, query service	*: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

Role	Role description	Resource	Action	Grant option
officer	status query, inventory information, and query product usage.	*:aso:itil:*	["read"]	0
		:aso:lock:	["read"]	0
		:aso:physical:	["read","write"]	0
		:aso:psm:	["read"]	0
		:aso:scm:	["read"]	0
		:aso:slalink:	["read"]	0
		:aso:task:	["read"]	0
ASO duty observer	Has the permissions to view and monitor the dashboard, and monitor system alerts.	*:aso:doc:*	["read"]	0
		:aso:fullview:	["read"]	0

12.1.8.1.6. Default roles of PaaS

This topic describes the default roles of the Platform as a Service (PaaS) console and the corresponding grant options.

The PaaS console is an O&M platform designed for the PaaS platform and products, and is used to view, manage, and upgrade the products deployed on the PaaS platform.

The following table describes the default roles of the PaaS console and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
PaaS_Operation_Manager	Has all the permissions in the PaaS console.	*:paas-ops:*	["*"]	0

12.1.8.1.7. Default roles of ZStack

This topic describes the default roles of the ZStack Operations and Maintenance System console and the corresponding grant options.

The ZStack Operations and Maintenance System console allows ZStack users to manage and schedule the compute, storage, and network resources.

The following table describes the default roles of the ZStack Operations and Maintenance System console and the corresponding grant options.

Role	Role description	Resource	Action	Grant option
ZStack administrator	Has all the permissions in the ZStack Operations and Maintenance System console and can manage all ZStack resources.	*:zstack:*	["*"]	0

12.1.8.2. Operation permissions on O&M platforms

This topic describes the operation permissions on O&M platforms.

12.1.8.2.1. Permissions on Apsara Infrastructure Management Framework

This topic describes the operation permissions on Apsara Infrastructure Management Framework.

Resource	Action	Description
*:tianji:services:[sname]:tjmontemplates:[tmplname]	delete	DeleteServiceTjmonTmpl
*:tianji:services:[sname]:tjmontemplates:[tmplname]	write	PutServiceTjmonTmpl
*:tianji:services:[sname]:templates:[tmplname]	write	PutServiceConfTmpl
*:tianji:services:[sname]:templates:[tmplname]	delete	DeleteServiceConfTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTjmonTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:tssessions	terminal	CreateTsSessionByService
*:tianji:services:[sname]:serviceinstances:[siname]:template	write	SetServiceInstanceTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:template	delete	DeleteServiceInstanceTmpl

Resource	Action	Description
*:tianji:services:[sname]:serviceinstances:[siname]:template	read	GetServiceInstanceTmpl
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanceProductTagInService
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	write	AddServiceInstanceProductTagInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	GetServerroleResourceInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	write	OperateSRMachineInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	read	GetMachineSRInfoInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	delete	DeleteSRMachineActionInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	read	GetMachinesSRInfoInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	delete	DeleteSRMachinesActionInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	write	OperateSRMachinesInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:resources	read	GetAppResourceInService

Resource	Action	Description
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:apps: [app]:machines: [machine]:tianjilogs	read	TianjiLogsInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles	read	GetServiceInstanceServerrolesInService
*:tianji:services: [sname]:serviceinstances: [siname]:schema	write	SetServiceInstanceSchema
*:tianji:services: [sname]:serviceinstances: [siname]:schema	delete	DeleteServiceInstanceSchema
*:tianji:services: [sname]:serviceinstances: [siname]:rollings:[version]	write	OperateRollingJobInService
*:tianji:services: [sname]:serviceinstances: [siname]:rollings	read	ListRollingJobInService
*:tianji:services: [sname]:serviceinstances: [siname]:resources	read	GetInstanceResourceInService
*:tianji:services: [sname]:serviceinstances: [siname]:machines:[machine]	read	GetMachineAllSRInfoInService
*:tianji:services: [sname]:serviceinstances: [siname]	write	DeployServiceInstanceInService
*:tianji:services: [sname]:serviceinstances: [siname]	read	GetServiceInstanceConf
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:files:name	read	GetMachineAppFileListInService
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:download	read	GetMachineAppFileDownloadInService

Resource	Action	Description
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:content	read	GetMachineAppFileContentInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:filelist	read	GetMachineFileListInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:dockerlogs	read	DockerLogsInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:debuglog	read	GetMachineDebugLogInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps	read	GetMachineAppListInService
*:tianji:services:[sname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	DockerInspect
*:tianji:services:[sname]:schemas:[schemaname]	write	PutServiceSchema
*:tianji:services:[sname]:schemas:[schemaname]	delete	DeleteServiceSchema
*:tianji:services:[sname]:resources	read	GetResourceInService
*:tianji:services:[sname]	delete	DeleteService
*:tianji:services:[sname]	write	CreateService
*:tianji:projects:[pname]:machinebuckets:[bname]:machines:[machine]	read	GetMachineBucketMachineInfo
*:tianji:projects:[pname]:machinebuckets:[bname]:machines	read	GetMachineBucketMachines

Resource	Action	Description
*:tianji:projects:[pname]:machinebuckets:[bname]	write	CreateMachineBucket
*:tianji:projects:[pname]:machinebuckets:[bname]	write	OperateMachineBucketMachines
*:tianji:projects:[pname]:machinebuckets:[bname]	delete	DeleteMachineBucket
*:tianji:projects:[pname]:machinebuckets:[bname]	read	GetMachineBucketMachinesLegacy
*:tianji:projects:[pname]:machinebuckets	read	GetMachineBucketList
*:tianji:projects:[pname]:projects:[pname]:clusters:[cname]:tssessions:[tssessionname]:tsses	terminal	UpdateTsSessionTssByCluster
*:tianji:projects:[pname]:projects:[pname]:clusters:[cname]:tssessions	terminal	CreateTsSessionByCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTjmonTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	delete	DeleteServiceInstanceTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	write	SetServiceInstanceTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:template	read	GetServiceInstanceTmplInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	write	AddServiceInstanceProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanceProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	GetServerroleResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:name	read	GetMachineAppFileList
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:download	read	GetMachineAppFileDownload
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:content	read	GetMachineAppFileContent
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:filelist	read	GetMachineFileList
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:dockerlogs	read	DockerLogsInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:debuglog	read	GetMachineDebugLog
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps	read	GetMachineAppList
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	read	GetMachineSRInfoInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	write	OperateSRMachineInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	delete	DeleteSRMachineActionInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	write	OperateSRMachinesInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	delete	DeleteSRMachinesActionInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	read	GetAllMachineSRInfoInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:resources	read	GetAppResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	TianjiLogsInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	DockerInspectInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles	read	GetServiceInstanceServerrolesInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	delete	DeleteServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	write	SetServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:resources	read	GetInstanceResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	delete	DeleteServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	write	CreateServiceInstance

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	read	GetServiceInstanceConfInCluster
*:tianji:projects:[pname]:clusters:[cname]:rollings:[version]	write	OperateRollingJob
*:tianji:projects:[pname]:clusters:[cname]:rollings	read	ListRollingJob
*:tianji:projects:[pname]:clusters:[cname]:resources	read	GetResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:quota	write	SetClusterQuotas
*:tianji:projects:[pname]:clusters:[cname]:machinesinfo	read	GetClusterMachineInfo
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	read	GetMachineAllSRInfo
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	write	SetMachineAction
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	delete	DeleteMachineAction
*:tianji:projects:[pname]:clusters:[cname]:machines	write	OperateClusterMachines
*:tianji:projects:[pname]:clusters:[cname]:difflist	read	GetVersionDiffList
*:tianji:projects:[pname]:clusters:[cname]:diff	read	GetVersionDiff
*:tianji:projects:[pname]:clusters:[cname]:deploylogs:[version]	read	GetDeployLogInCluster
*:tianji:projects:[pname]:clusters:[cname]:deploylogs	read	GetDeployLogListInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:builds:[version]	read	GetBuildJob
*:tianji:projects:[pname]:clusters:[cname]:builds	read	ListBuildJob
*:tianji:projects:[pname]:clusters:[cname]	write	OperateCluster
*:tianji:projects:[pname]:clusters:[cname]	delete	DeleteCluster
*:tianji:projects:[pname]:clusters:[cname]	read	GetClusterConf
*:tianji:projects:[pname]:clusters:[cname]	write	DeployCluster
*:tianji:projects:[pname]	write	CreateProject
*:tianji:projects:[pname]	delete	DeleteProject
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	CreateRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	SetRackunitAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	delete	DeleteRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	SetRackAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	CreateRack
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	delete	DeleteRack
*:tianji:idcs:[idc]:rooms:[room]	write	CreateRoom
*:tianji:idcs:[idc]:rooms:[room]	delete	DeleteRoom
*:tianji:idcs:[idc]:rooms:[room]	write	SetRoomAttr
*:tianji:idcs:[idc]	delete	DeleteIdc

Resource	Action	Description
*:tianji:idcs:[idc]	write	SetIdcAttr
*:tianji:idcs:[idc]	write	CreateIdc

12.1.8.2.2. 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