ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Enterprise

apsara base

Apsara Uni-manager Management Console User Guide

Product Version: V3.18.1

Document Version: 20231222

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.What is the Apsara Uni-manager Management Console? 10
2.Getting started 14
2.1. Homepage 14
2.2. Initial configuration 15
2.2.1. Preset roles 15
2.2.2. Configuration description 23
2.2.3. Configuration process 24
2.3. Log on to the Apsara Uni-manager Management Console 25
3.Portal home pages 26
4.Enterprise 33
4.1. Overview 33
4.2. Resource management 33
4.2.1. Overview 33
4.2.2. Organization management 34
4.2.2.1. Create an organization 34
4.2.2.2. Manage organizations 35
4.2.2.3. Manage the AccessKey pairs of an organization 35
4.2.3. Resource set management 37
4.2.3.1. Create a resource set 37
4.2.3.2. Manage resource sets 37
4.2.4. Manage regions 38
4.2.5. Change management 39
4.3. User management 40
4.3.1. Overview 40
4.3.2. User management 40
4.3.2.1. Create a user 40

4.3.2.2. Manage system users 45
4.3.2.3. Manage historical users 48
4.3.3. User group management 48
4.3.3.1. Create a user group 48
4.3.3.2. Manage user groups 49
4.4. Permission management 51
4.4.1. Overview 51
4.4.2. Role permissions 52
4.4.2.1. Overview 52
4.4.2.2. RAM policies 53
4.4.2.3. Create a role 60
4.4.2.4. Manage roles 63
4.4.3. Data permissions 65
4.4.3.1. Overview 65
4.4.3.2. Set the data permissions on resource instances 65
4.4.3.3. Configure the data permissions of users 66
4.4.4. Access control 67
4.4.4.1. Create an access policy 67
4.4.4.2. Manage access policies 68
4.4.5. Permission boundary 69
4.4.5.1. Service boundary management 69
4.4.5.1.1. Create a service boundary 69
4.4.5.1.2. Manage service boundaries 71
4.4.5.2. Control policy management 72
4.4.5.2.1. Create an API control policy 75
4.4.5.2.2. Manage API control policies 77
4.4.6. AccessKey logs 77
5.Configurations 84

5.1. Overview 84
5.2. System configuration 85
5.2.1. Resource pool configuration 85
5.2.2. Region settings 86
5.2.3. Security policies 86
5.2.4. Authentication modes 89
5.2.4.1. Authentication modes 89
5.2.4.2. Configure an authentication mode 90
5.2.5. Specification management 90
5.2.5.1. Specification parameters 91
5.2.5.2. Create or synchronize specifications 96
5.2.5.3. Manage specifications 97
5.2.6. Manage menu items 99
5.2.7. Announcement management 101
5.2.8. Custom configurations 102
5.2.9. Third-party authentication 103
5.2.9.1. IdP settings 104
5.2.10. Message gateway 108
5.3. Other settings 111
5.3.1. Service-linked role management 111
5.3.2. Log storage settings 112
6.Operations 114
6.1. Service orders 114
6.1.1. Overview 114
6.1.2. Manage orders 115
6.2. Service management 118
6.2.1. Service catalog management 118
6.2.1.1. Overview 118

6.2.1.2. Create a service catalog 119
6.2.1.3. Manage service catalogs 121
6.2.1.4. FAQ 122
6.2.2. Process management 123
6.2.2.1. Overview 123
6.2.2.2. Create a process 124
6.2.2.3. Manage processes 125
6.3. Quota management 126
6.3.1. Quota parameters 128
6.3.2. Manage quotas 134
6.3.3. Quota alert management 136
6.3.3.1. Create a quota alert rule 136
6.3.3.2. Manage quota alerts 137
6.4. Statistical analysis 138
6.4.1. View reports 138
6.4.2. Report pushing 144
6.4.3. Download Center 145
6.4.4. Carbon emissions report 148
6.5. Usage statistics 148
6.5.1. View the usage statistics of cloud resources 148
6.6. Bills 152
6.6.1. View cloud service bills 152
6.6.2. View organization and resource set bills 153
6.7. Tag management 154
6.7.1. Query tags 154
6.7.2. Create a tag 155
6.8. Cloud resource optimization 155
6.8.1. Resource utilization rates 155

6.8.2. Resource health scores 161
6.8.3. Cloud optimization results 162
6.8.4. Resource optimization settings 162
6.9. Billing management 165
6.9.1. Billing overview 165
6.9.2. Billing configurations 165
6.9.2.1. Create a billing configuration 165
6.9.2.2. Manage billing configurations 168
7.Security 171
7.1. Overview 171
7.2. Operation logs 171
7.2.1. Supported cloud services 171
7.2.2. Manage operation logs 175
7.3. AccessKey 177
7.4. MFA 178
7.4.1. Overview 178
7.4.2. Manage MFA 178
8.User information management 181
9.Page integration 182
10.Cloud Monitor Service 185
10.1. Overview 185
10.2. Overview 185
10.3. Cloud service monitoring 186
10.3.1. Metrics 188
10.3.2. View the monitoring data 261
10.3.3. View alert rules 261
10.4. Host monitoring 262
10.4.1. Install a CloudMonitor agent 262

10.4.2. Process monitoring 263
10.5. Alert service 265
10.5.1. Alert history 265
10.5.2. Alert rules 266
10.5.2.1. Create an alert rule 266
10.5.2.2. Manage alert rules 269
10.5.3. Alert templates 271
10.5.3.1. Common alerting templates 271
10.5.3.2. Create an alert template 277
10.5.3.3. Manage alert templates 278
10.6. Statistical analysis 281
10.6.1. Sorting 281
10.7. Event monitoring 281
10.7.1. System events 282
10.7.1.1. View system events 282
10.7.1.2. Create a system event-triggered alert rule 285
10.7.1.3. Manage system event-triggered alert rules 286
10.8. Application group 288
10.8.1. Create an application group 288
10.8.2. Manage application groups 289
10.8.3. Availability monitoring 291
10.8.3.1. Create an availability monitoring task 291
10.8.3.2. Manage availability monitoring tasks 295
10.8.3.3. Status codes 296
10.9. Download a report 296

1.What is the Apsara Uni-manager Management Console?

The Apsara Uni-manager Management Console is provided by Apsara Stack to help organizations perform cloud deployment. The platform improves IT management and operations, aiming to provide large-scale, cost-efficient end-to-end cloud computing and big data services for customers in sectors like public service, education, healthcare, and financial services.

Overview

The Apsara Uni-manager Management Console simplifies the management and deployment of physical and virtual resources, facilitating application development, improving resource utilization, and reducing O&M costs. The console is an integral part of the ecosystem built around cloud computing.

Procedure

Operations in the Apsara Uni-manager Management Console are classified into the following parts:

- Basic configurations for system initialization, such as creating organizations, resource sets, and users, creating basic resources such as virtual private clouds (VPCs), and creating contacts and contact groups in CloudMonitor.
- Creation of cloud resources (performed by administrators).
- Resource management, such as starting, using, and releasing resources, and changing resource configurations and resource quotas.

Emergency mode

The Apsara Uni-manager Management Console is switched to the emergency mode when faults occur to Apsara Stack.

In this mode, only emergency features are available.

Apsara Stack faults

When faults occur to Apsara Stack and Apsara Uni-manager Management Console enters the emergency mode, the following features become unavailable.

Module	Feature	Description

Enterprise	• Organization management: manage AccessKeys; create users and user groups. • Resource set management: grant permissions to users. • Resource set details: grant permissions to users. • User management: create, modify, disable, delete, and grant permissions to users; add users to and remove users from user groups; configure logon policies; configure multi-factor authentication (MFA). • User group management: create, modify, delete, and grant permissions to user groups; add users to and remove users from user groups. • Role permissions: create, modify, disable, delete, replicate, and customize roles. • Data permissions: modify data permissions. • Access control: create, modify, disable, and delete access control policies. • Permission boundary: create, modify, disable, and delete service boundaries; create, modify, and delete API control policies.	The features become unavailable.
Products	• Create cloud resources. • Manage cloud resources. • Redirect to consoles of cloud products.	The features become unavailable.

Central region faults

If your Apsara Stack is deployed across regions, when a fault occurs to the central region, the following features become limited or unavailable in the general regions.

> **Important**
>
> Data inconsistency between the central and general regions may exist when the services are recovered, but will be resolved within five minutes.

Module	Feature	Description
Console	Logon status	• When the console enters the emergency mode, users on the general regions are logged out. • This also happens when the console goes back to normal. Users need to log on again.
Homepage	• Overview: view resources loads. • Platform operations: all features. • Platform resources: all features.	The features become unavailable.

Enterprise	• Organization management: create, modify, and delete organizations; manage AccessKeys; create resource sets, users, and user groups. • Resource set management: create and delete resource sets; modify the names of resource sets; grant permissions to members. • Resource set details: grant permissions to users. • Region management: update associations. • Change management: all change management features. • User management: create, modify, disable, delete, and grant permissions to users; add users to and remove users from user groups; configure logon policies; configure MFA. • User group management: create, modify, delete, and grant permissions to user groups; add users to and remove users from user groups. • Role permissions: create, modify, disable, delete, replicate, and customize roles; grant permissions to users and user groups. • Data permissions: modify data permissions. • Access control: create, modify, disable, and delete access control policies. • Permission boundary: create, modify, disable, and delete service boundaries; create, modify, and delete API control policies.	The features become unavailable.
Operations	• Service orders: all features. • Services: all features. • Quotas: manage quotas and quota alerts. • Statistical analysis: all features. • Metering management: all features. • Bills: all features. • Tags: all features. • Cloud resource optimization: all features. • Billing management: all features. • Global search: search for organizations, resource sets, and resources.	The features become unavailable.
	Product overview	The feature becomes unavailable.

Products	Instance creation	The feature becomes limited. • Service catalogs cannot be used when you create instances. • If instances of a specific cloud product have never been created in a level-1 organization, instances of that cloud product cannot be created within that organization and its sub-organizations.

2.Getting started

2.1. Homepage

The homepage of the Apsara Uni-manager Management Console consists of the top navigation bar, the information section of the current logon user, and the operation section.

Homepage of the Apsara Uni-manager Management Console

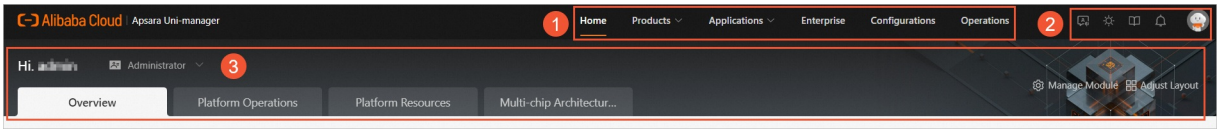

Table 1. Sections of the homepage

Section		Description
①	Top navigation bar	This section includes the following modules: • **Home**: You can view charts that show the usage and monitoring data of the resources in each region. • **Products**: You can manage the basic cloud products and resources. • **Applications**: Your can manage the cloud applications of your enterprise. • **Enterprise**: You can manage the organizations, resource sets, roles, users, regions, logon policies, user groups, and ownership . • **Configurations**: You can manage the resource pools, regions, password policies, authentication models, specifications, menus, announcements, personalization settings, third-party verification, usage of resources from other cloud service providers, adapters, service catalogs, and RAM roles. • **Operations**: You can perform management operations on your cloud resources by managing service orders, services, quotas, usage statistics, bills, tags, cloud resource optimization, and billing configurations. • **Security**: You can view operation logs, system logs, and create cloud security products.

②	Information of the current logon user	• English: You can switch between English, simplified Chinese, and traditional Chinese. • /: allows you to switch between light mode and dark mode for the console. • : redirects you to **Help and Open Platform** where you can search for documentation such as user guides, API references, and SDK references. • : displays message notifications. You can click **More** to go to **Message Center** for details of the messages. • User Information: displays the **User Information**, **View Version**, and **Log Out** menu items when you click the profile picture. If you click **User Information**, you can perform the following operations on the User Information page: ◦ View basic information. ◦ Modify user information. ◦ Change the logon password. ◦ View and create AccessKey pairs. ◦ Switch to other roles. ◦ View the current role policy. ◦ View the resource sets of the user. ◦ View or bind virtual multi-factor authentication (MFA) devices. ◦ Enable or disable messages sent by email or DingTalk. If you click **View Version**, you can view the authorization status and build number of the Apsara Uni-manager Management Console.
③	Operation section	You can adjust the **operation section** by performing the following operations: • Click **Mange Module** in the upper right corner to select the modules that are displayed on the page. • Click **Adjust Layout** in the upper right corner and drag the modules to adjust their positions.

2.2. Initial configuration

2.2.1. Preset roles

To implement separation of privilege, the Apsara Uni-manager Management Console provides several preset roles based on common use cases in cloud platforms. Each role comes with a collection of permissions that fit their definition.

Before you use the Apsara Uni-manager Management Console, make sure that you are familiar with the preset roles, including their management permissions and authorized permissions to properly manage the cloud platform.

- Management permissions refers to the scope of a role's permissions:
  - All Organizations: The role has permissions on all organizations.
  - Specified Organization and Subordinate Organizations: The role has permissions on an organization and its subordinate organizations.
  - Resource set: The role has permissions on the resource sets of the associated user.
- Authorized permissions refers to the operations that a role is allowed to perform on the Apsara Uni-manager Management Console.

> **Note**
> - The preset roles and their permissions cannot be modified or deleted.
> - Basic Permissions in the following tables refers to a set of permissions to perform basic operations like viewing the homepage and managing user information.

Platform administrator

Initializes the configurations of the Apsara Uni-manager Management Console, such as the technical component configurations, language configurations, and customization settings.

Management permissions	Authorized permissions
All Organizations	User management: view users, create users, delete users, activate users, disable users, operate users, and view the AccessKey pairs of users. Access policy: view access policies, create and delete access policies, and operate access policies. Menu management: manage menus. Announcement: manage announcements. System configuration: configure systems. Basic Permissions

Operations administrator

Creates and manages organizations and resource sets, adds and manages users, assigns and manages roles, and activates and manages resources in the Apsara Stack console.

Management permissions	Authorized permissions

All Organizations	API audit: view API audit, view API directory, perform online API debugging. Billing management: view the bills of cloud products, view the bills of organizations and resource sets, export bills, and modify bills. Change management: view changes, change organizations, change resource sets, change users, and change resources. Billing management: viewing the pricing, configure the pricing, create and delete the pricing, obtain the billing overview, view billing configurations, create and modify billing configurations, and delete billing configurations. Cloud resource optimization: view optimization tasks, modify optimization tasks, view optimization configurations, modify optimization configurations, view algorithm configurations, modify algorithm configurations, view alert configurations, and modify alert configurations. Download center: view tasks, create tasks, delete tasks, and download reports. Access policy: view access policies, create and delete access policies, and operate access policies. Usage statistics: view usage statistics and export usage statistics. Announcement: manage announcements. Operation logs: view the log storage mode and configure the log storage mode. Organization management: view organizations, create and delete organizations, and operate organizations. AccessKey management: view organization AccessKey pairs, create organization AccessKey pairs, delete organization AccessKey pairs, enable organization AccessKey pairs, and disable organization AccessKey pairs. Quota management: view quotas and operate quotas. Resource set management: view resource sets, create and delete resource sets, and operate resource sets. Role management: view roles, create and delete roles, operate roles, lock and unlock roles, and configure authorization of cloud services. Service catalog management: view service catalogs, create and modify service catalogs, delete service catalogs, publish and unpublish service catalogs, and grant permissions on service catalogs. Service-linked role: view service-linked roles and create service-linked roles. Tags: operate tags. User management: view users, create users, delete users, enable users, disable users, lock and unlock users, operate users, and view user AccessKey pairs. Process management: view processes, create and modify processes, publish and withdraw processes, delete processes, and view applications. Basic Permissions

Organization administrator

Manages the organization, creates and manages subordinate organizations, creates and manages resource sets, adds and manages users, allocates and manages roles, activates and manages resources, assigns administrators and quotas to subordinate organizations and resource sets, applies for resources or quotas, and examines and approves resource or quota applications.

Management permissions	Authorized permissions

Specified Organization and Subordinate Organizations	API tool: view API audit, view API directories, and perform online API debugging. Billing management: view organization and resource set bills, and export bills. Change management: view changes, change organizations, change resource sets, change users, and change resources. Cloud resource optimization: view optimization tasks, view algorithm configurations, and view alert configurations. Download center: view tasks, create tasks, delete tasks, and download reports. Access policy: view access policies. Usage statistics: view and export usage statistics. Organization management: view organizations, create and delete organizations, and operate organizations. AccessKey management: view organization AccessKey pairs, create organization AccessKey pairs, delete organization AccessKey pairs, enable organization AccessKey pairs, and disable organization AccessKey pairs. Quota management: view quotas and operate quotas. Resource set management: view resource sets, create and delete resource sets, and operate resource sets. Role management: view roles, create and delete roles, and operate roles. Service catalog management: view service catalogs, create and modify service catalogs, delete service catalogs, publish and unpublish service catalogs, and grant permissions on service catalogs. Service-linked role: view service-linked roles and create service-linked roles. Tags: operate tags. User management: view users, create users, delete users, enable users, disable users, operate users, and view the AccessKey pairs of users. User group management: view user groups, create and delete user groups, and operate user groups. Process management: view processes and view applications. Basic Permissions

Resource set administrator

Manages the resource set, adds and manages users, activates and manages resources, examines and approves resource applications, and applies for resources or quotas.

Management permissions	Authorized permissions

Resource Sets	Cloud resource optimization: view optimization tasks. Download center: view tasks, create tasks, delete tasks, and download reports. Usage statistics: view and export usage statistics. Organization management: view organizations. Quota management: view quotas. Resource set management: view resource sets and operate resource sets. Tags: operate tags. User management: view users. User group management: view user groups. Basic Permissions

Resource user

Uses the cloud resources created and allocated by the operations administrator.

Management permissions	Authorized permissions
Resource Sets	Cloud resource optimization: view optimization tasks. Organization management: view organizations. Resource set management: view resource sets. Service-linked role: view service-linked roles. Tags: operate tags. Basic Permissions

Resource auditor

Has read-only permission on all resources in Apsara Stack.

Management permissions	Authorized permissions

All Organizations	API tool: view API audit. Billing management: view organization and resource set bills, and export bills. Change management: view changes. Billing management: view the pricing. Cloud resource optimization: view optimization tasks, view algorithm configurations, and view alert configurations. Download center: view tasks and download reports. Access policy: view access policies. Usage statistics: view and export usage statistics. Operation logs: view operation logs, download operation logs, and view the storage mode of operation logs. Organization management: view organizations. Quota management: view quotas. Resource set management: view resource sets. Role management: view roles. Service catalogs: view the list of service catalogs. User management: view users. User group management: view user groups. Process management:view applications. Basic Permissions

Organization resource auditor

Has read-only permissions on all resources in the organization to which it belongs.

Management permissions	Authorized permissions

Specified Organization and Subordinate Organizations	API tool: view API audit. Billing management: view organization and resource set bills, and export bills. Change management: view changes. Cloud resource optimization: view optimization tasks, view algorithm configurations, and view alert configurations. Download center: view tasks and download reports. Access policy: view access policies. Usage statistics: view and export usage statistics. Operation log: view operation logs and download operation logs. Organization management: view organizations. Quota management: view quotas. Resource set management: view resource sets. Role management: view roles. Service catalogs: view the list of service catalogs. User management: view users. User group management: view user groups. Process management:view applications. Basic Permissions

Security auditor

Audits the security of Apsara Stack.

Management permissions	Authorized permissions
All Organizations	Operation log: view operation logs and download operation logs. Basic Permissions

Platform security administrator

Uses Cloud Security Operation Center (SOC) to manage the security of Apsara Stack.

Management permissions	Authorized permissions
All Organizations	Basic Permissions

Global organization security administrator

Uses SOC to provide security management for Apsara Stack and all tenants. This role also manages security products related to ECS.

Management permissions	Authorized permissions
All Organizations	Organization management: view organizations. Resource set management: view resource sets. Role management: view roles, create and delete roles, and operate roles. Service-linked role: view service-linked roles and create service-linked roles. User management: view users, enable users, disable users, and operate users. Basic Permissions

Organization security administrator

Uses SOC to manage the security of hosts, applications, and networks for an organization. This role also manages security products related to Elastic Compute Service (ECS).

Management permissions	Authorized permissions
Specified Organization and Subordinate Organizations	Operation log: view operation logs and download operation logs. Organization management: view organizations. Resource set management: view resource sets. Service-linked role: view service-linked roles and create service-linked roles. Basic Permissions

Security system configuration administrator

Configures system security features such as the upgrade center and global configurations.

Management permissions	Authorized permissions
All Organizations	Basic Permissions

Platform security auditor

Uses SOC to check the security conditions of Apsara Stack.

Management permissions	Authorized permissions

All Organizations	Operation log: view operation logs and download operation logs. Basic Permissions

Global organization security auditor

Uses SOC to check the security conditions of all organizations.

Management permissions	Authorized permissions
All Organizations	Operation log: view operation logs and download operation logs. Basic Permissions

Platform security configuration administrator

Configures security services in Apsara Stack, such as Server Guard and Web Application Firewall (WAF).

Management permissions	Authorized permissions
All Organizations	Basic Permissions

2.2.2. Configuration description

Before you use the Apsara Uni-manager Management Console, you must complete a series of basic configuration operations as an administrator, such as creating organizations, resource sets, users, and roles and initializing resources. This is the initial system configuration.

The Apsara Uni-manager Management Console manages the organizations, resource sets, users, and roles of cloud data centers in a centralized and service-oriented manner to grant different resource access permissions to different users.

- Organization

  After the Apsara Uni-manager Management Console is deployed, a root organization is automatically generated. You can create other organizations under the root organization.

  Organizations are displayed in a hierarchical structure. You can create subordinate organizations under each organization level.

- Resource Set

  A resource set is a container used to store resources. Each resource must belong to a resource set.

- User

  A user is a resource manager and user.

- Role

A role is a set of access permissions. You can assign different roles to different users to implement system access control to meet a variety of different requirements.

The following table describes the relationships among organizations, resource sets, users, roles, and cloud resources.

Relationship between two items	Relationship type	Description
Organization and resource set	One-to-many	An organization can have multiple resource sets, but each resource set can belong to only a single organization.
Organization and user	One-to-many	An organization can have multiple users, but each user can belong to only a single organization.
Resource set and user	Many-to-many	A user can have multiple resource sets, and a resource set can be assigned to multiple users under the same level-1 organization.
User and role	Many-to-many	A user can have multiple roles, and a role can be assigned to multiple users.
Resource set and resource	One-to-many	A resource set can have multiple resources, but each cloud resource can belong to only a single resource set.

2.2.3. Configuration process

Before users can the Apsara Uni-manager Management Console, administrators must complete the initial system configuration based on the process shown in the following figure.

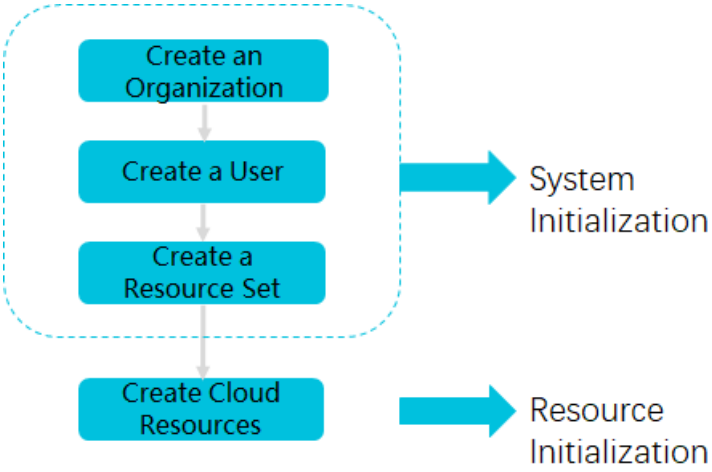

1. Create an organization

   Create an organization to store resource sets and their resources.

2. Create a user

   Create users and assign the users different roles to meet different requirements for system access control.

3. Create a resource set

   Create a resource set before you apply for resources.

Users that have the permissions on a resource set can use the resources in the resource set. For information about how to authorize a user to use a resource set, see the Grant users the permissions on a resource set section of the "Manage resource sets" topic.

4. Create cloud resources

   Create instances in the consoles of cloud services based on project requirements. For more information about how to create cloud service instances, see the user guide of each cloud service.

2.3. Log on to the Apsara Uni-manager Management Console

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

Prerequisites

- The URL of the console that is obtained from the deployment personnel.
- A web browser. We recommend that you use Google Chrome.

Procedure

1. Enter the URL of the Apsara Uni-manager Management Console in the address bar and press the Enter key.
2. Enter your username and password.

   You need to obtain the username and password from an operations administrator.

   > **Note**
   > - **First logon**
   >
   >   The first time that you log on to the Apsara Uni-manager Management Console, you need to change the password of your account. The password must be 10 to 32 characters in length. It must contain all of the following character types: uppercase letters, lowercase letters, digits, and special characters. Supported special characters include: ! @ # $ %
   > - **Forget password**
   >
   >   If you have forgotten your password, click Forgot Password. On the page that appears, enter the username of your account, the email address that was used to create the account, and the CAPTCHA code. Then, the system sends a link for resetting the password to the specified email address.

3. Click **Log On**.
4. If multi-factor authentication (MFA) is enabled for your account, perform the corresponding operations in the following scenarios:
   - You log on to the Apsara Uni-manager Management Console for the first time after MFA is enabled by the administrator.
     a. On the Bind Virtual MFA Device page, bind an MFA device.
     b. Enter the username and password again as in Step 2 and click **Log On**.
     c. Enter a six-digit MFA verification code and click **Authenticate**.
   - You have enabled MFA and bound an MFA device:

     Enter a six-digit MFA verification code and click **Authenticate**.

3.Portal home pages

Portal home pages show the usage and monitoring status of system resources in different regions from different dimensions in charts. This helps you understand the overall resource status and provides supporting data for refined operations and statistical analysis.

The modules and layouts of portal home pages vary based on the user role. The following sections describe the **home page for platform administrators**, **home page for operations administrators**, **home page for organization administrators**, and **home page for resource set administrators** that are supported by default. Portal home pages, modules, and layouts for other roles or custom roles are automatically matched based on the defined scope of permissions and authorized permissions.

> **Note**
>
> You can perform resource lifecycle management operations, such as creating and deleting resources, in the service consoles. To manage the lifecycle of resources, click **Products** in the upper part of the Apsara Uni-manager Management Console, and create or manage services based on your business requirements.

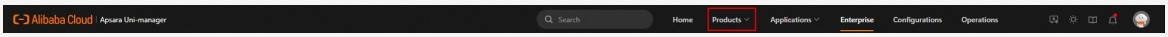

Home page for platform administrators

The home page shows the overview of quotas, organizations, and resource loads from the cloud platform dimension.

Tab	Parameter	Description
Overview	Quota Overview	Shows the quota, assigned quota, and used quota of cloud services in the current region from the cloud platform dimension.
	Organization Overview	Shows the number of organizations, resource sets, and users from the cloud platform dimension.
	Announcement	Shows related O&M, operations, and upgrade notifications. Click More to go to the announcement list and view more announcements.
	Resource Load	Shows the top 5 memory usage, top 5 disk usage, and top 5 CPU usage of Elastic Compute Service (ECS) and ApsaraDB RDS instances from the cloud platform dimension.
	Number of instances	Shows the number of instances of cloud services from the cloud platform dimension and helps you understand the usage of the cloud services.
	Historical trends	Shows recent usage trends of cloud services from the cloud platform dimension and helps you understand the overall usage of the cloud services.

	Alarm Statistics	Shows the alarms generated on the cloud platform, including the total alarms, alarm severities, and alarm information. The statistics help you quickly locate alarms and provide effective data support for quickly removing alarms.

Home page for operations administrators

The home page shows the overview of quotas, organizations, and resource loads. The home page also shows information about platform operations and platform resources from the cloud platform dimension.

Tab	Parameter	Description
Overview	Quota Overview	Shows the quota, assigned quota, and used quota of cloud services in the current region from the cloud platform dimension.
	Organization Overview	Shows the number of organizations, resource sets, and users from the cloud platform dimension.
	Announcement	Shows related O&M, operations, and upgrade notifications. Click More to go to the announcement list and view more announcements.
	Resource Load	Shows the top 5 memory usage, top 5 disk usage, and top 5 CPU usage of ECS and ApsaraDB RDS instances from the cloud platform dimension.
	Number of instances	Shows the number of instances of cloud services from the cloud platform dimension and helps you understand the usage of the cloud services.
	Historical trends	Shows recent usage trends of cloud services from the cloud platform dimension and helps you understand the overall usage of the cloud services.
	Alarm Statistics	Shows the alarms generated on the cloud platform, including the total alarms, alarm severities, and alarm information. The statistics help you quickly locate alarms and provide effective data support for quickly removing alarms. The alarm severities include Critical, Warn, and Info. You can view more alarm information in CloudMonitor.
	Organization Overview	Shows the number of organizations, resource sets, and users from the cloud platform dimension.
	Service Orders	Shows the number of pending orders and all orders on the cloud platform.

Platform Operations	Regions	Shows the regions from the cloud platform dimension and helps you understand the regional distribution.
	Resources	Shows the number of resources used by each organization on the cloud platform, and shows the total number of resources in descending order.
	Activated Services	Shows the number of cloud services activated by each organization on the cloud platform, and shows the total number of activated services in descending order.
	Resource Sets	Shows the number of resource sets of each organization on the cloud platform.
Platform Resources	Virtual Resource Overview	Shows the number of virtual resources (cloud services and instances) from the cloud platform dimension.
	Physical Resource Overview	Shows the number of physical resources from the cloud platform dimension, such as servers, racks, network devices, and data centers.
	Platform Resource Health	Shows the overall resource health score in the last 30 days from the cloud platform dimension. The health score is calculated based on the utilization, idle status, and compliance analysis of each cloud resource on the cloud platform. • 0 - 50 points: poor • 51 - 70 points: general • 71 - 90 points: moderate • 91 - 100 points: healthy
	Resource Health Analytics	Supports resource health analysis and display by organization or resource set. Click Sort by Health Score in Ascending Order or Sort by Health Score in Descending Order to view health scores in ascending or descending order.
	Capacity Analysis	Predicts resource capacity based on the resource availability in the last 6 months. Helps users understand the resource capacity usage. The prediction results are for reference only.
	Server distribution	Shows the total number of physical servers that are based on different chips from the data center dimension. You can click a chip name in the legend to filter servers by chip.

Multi-chip Architecture in One Cloud	Instance distribution	Shows the number of instances that have been created for each cloud service from the chip dimension. The data source is the total number of resources within the management scope of the current user. A delay of more than 10 minutes occurs when the data is updated.
	Resource Load	Shows the top 5 memory usage, top 5 disk usage, and top 5 CPU usage of ECS and ApsaraDB RDS instances from the cloud platform dimension and the chip types of instances.
	Alarm Statistics	Similar to the "Alarm statistics" section on the Overview tab, this section shows the alarm information from the cloud platform dimension. The difference is that this section shows chip type information of the instance. The alarm severities include Critical, Warn, and Info. You can view more alarm information in CloudMonitor.

> **Note**
> The dashboard of multi-chip architecture in one cloud is displayed only when multiple chip models are deployed on the Apsara Stack platform.

Home page for organization administrators

The home page shows the overview, resource operations, and organization operations of the organization and its subordinate organizations.

Tab	Parameter	Description
Overview	Quota Overview	Shows the quota, assigned quota, and used quota of cloud services in different regions from the organization dimension.
	Organization Overview	Shows the number of organizations, resource sets, and users from the organization dimension.
	Announcement	Shows related O&M, operations, and upgrade notifications. Click More to go to the announcement list and view more announcements.
	Resource Load	Shows the top 5 memory usage, top 5 disk usage, and top 5 CPU usage of ECS and ApsaraDB RDS instances from the organization dimension.
	Number of instances	Shows the number of instances of cloud services from the organization dimension and helps you understand the usage of the cloud services.
	Historical trends	Shows recent usage trends of cloud services from the organization dimension and helps you understand the overall usage of the cloud services.

	Alarm Statistics	Shows the alarms generated in the organization, including the total alarms, alarm severities, and alarm information. The statistics help you quickly locate alarms and provide effective data support for quickly removing alarms. The alarm severities include Critical, Warn, and Info. You can view more alarm information in CloudMonitor.
Resource Operations	Organization Resources	Shows the number of resources in the last 3 months from the organization dimension. **Note** You can select current organization or All Organizations. If the current organization has sub-organizations and you select All Organizations, resource statistics of the current organization and all its sub-organizations are displayed.
	All Cloud Service Categories	Shows the number of resources by all cloud service categories, and categories such as elastic computing, networking, big data, storage, and security from the organization dimension.
	Service Orders	Shows the number of service orders from the organization dimension.
	Organization Resource Health	Shows the resource health score in the last 30 days from the organization dimension. The health score is calculated based on the utilization, idle status, and compliance analysis of each cloud resource in the organization. 0 - 50 points: poor 51 - 70 points: general 71 - 90 points: moderate 91 - 100 points: healthy
	Organization Resource Health Analytics	Shows the resource health analytics from the organization dimension. Click Sort by Health Score in Ascending Order or Sort by Health Score in Descending Order to view health scores in ascending or descending order.
	Regions	Shows the regions from the organization dimension and helps you understand the regional distribution.
	Organization Overview	Shows the number of organizations, resource sets, users, and cloud services from the organization dimension.
	Shortcut	Allows you to quickly go to corresponding pages. You can click the Edit icon on the right side of the Shortcut section to customize a shortcut. You can specify up to six shortcuts.

Organization Operations	Statistics for Current Organization	Shows the fees in the last 12 months and the estimated fees for the next 3 months. Estimated fees are calculated based on the average growth rate in the last 3 months and are for reference only.
	Fee Ranking and Composition	Shows the organization fees of cloud service categories, cloud services, and resource sets by Current Month, Last Month, and Accumulated Since Its Establishment.
	Cost and resource health	Shows the costs and resource health status by organization and resource set. You can perform statistical analytics based on the data.
	Cost and resource retention	Shows the costs and inventory by organization and resource set. You can perform statistical analytics based on the data.
	Cloud product classification fees	Shows categories of cloud services and corresponding fees by This Month, Last Month, and Accumulated Since Its Establishment.
	Cloud Service Fees	Shows specific cloud services and corresponding fees by This Month, Last Month, and Accumulated Since Its Establishment.

Home page for resource set administrators

Shows the resource overview and resource operations from the resource set dimension.

Tab	Parameter	Description
Overview	Quota Overview	Shows the quota, assigned quota, and used quota of cloud services in different regions.
	Organization Overview	Shows the number of organizations, resource sets, and users.
	Announcement	Shows related O&M, operations, and upgrade notifications. Click More to go to the announcement list and view more announcements.
	Resource Load	Shows the top 5 memory usage, top 5 disk usage, and top 5 CPU usage of ECS and ApsaraDB RDS instances from the resource set dimension.
	Number of instances	Shows the number of instances of cloud services and helps you understand the usage of the cloud services.
	Historical trends	Shows recent usage trends of cloud services and helps you understand the usage trends of the cloud services.

	Alarm Statistics	Shows current alarms, including the total alarms, alarm severities, and alarm information. The statistics help you quickly locate alarms and provide effective data support for quickly removing alarms.
Resource Operations	Shortcut	Allows you to quickly go to corresponding pages. You can click the Edit icon on the right side of the Shortcut section to customize a shortcut. You can specify up to six shortcuts.
	Service Orders	Shows the number of service orders from the resource set dimension.
	Resource Set Resource Statistics	Shows the number of resources in the last 3 months from the resource set dimension.
	Resource Set Resource Health	Shows the resource health score in the last 30 days from the resource set dimension. The health score is calculated based on the utilization, idle status, and compliance analysis of each cloud resource in the resource set. 0 - 50 points: poor 51 - 70 points: general 71 - 90 points: moderate 91 - 100 points: healthy
	Resource Set Resource Health Analytics	Shows the resource health analytics from the resource set dimension. Click Sort by Health Score in Ascending Order or Sort by Health Score in Descending Order to view health scores in ascending or descending order.
	Recently Accessed Resources	Shows information about recently accessed resource instances from the resource set dimension.

4.Enterprise

4.1. Overview

Most modern enterprises are in hierarchical structures that are constantly evolving. To allow enterprises to manage cloud resources in a tiered manner, Apsara Stack provides the Enterprise component in Apsara Uni-manager Management Console.

You can organize your resources on the platform by organizations and projects to implement cost-effective and refined resource management.

The Enterprise component consists of the following modules.

Module	Description
Resources	You can manage organizations, resource sets, regions, and changes. This way, you can build a resource structure that suits your enterprise. You can easily manage the ownership of resources to improve resource utilization.
Users	The user management and user group management features allow you to centrally manage the structure and permissions of users and user groups. This can improve management efficiency and meet the various needs for system and resource access.
Permissions	The role permissions, data permissions, and access control features allow you to implemented refined permission management and improve system security.

4.2. Resource management

4.2.1. Overview

In the Resources module, you can manage organizations, resource sets, regions, and changes. This way, you can build a resource structure that suits your enterprise. You can easily manage the ownership of resources to improve resource utilization.

Organizations

You can create organizations of different levels in the console to map the actual structure of your organization.

In Apsara Uni-manager Management Console, organizations are where resources and resource sets belong. A root organization is created by default. It does not accommodate any resource. Organizations that are directly subordinate to the root organization are level-1 organizations, and those that are directly subordinate to level-1 organizations are level-2 organizations, and so on. The lowest-level organizations are the level-5 organizations.

After a level-1 organization is created, a cloud account and an AccessKey pair that belong to the organization are automatically generated. They are used to create resources in the system.

Operations for managing organizations:

- Basic operations: create organizations, modify organizations, delete organizations, change the sorting of organizations, and change the relationship between organizations.

- Manage AccessKey pairs of organizations: create, disable, activate, and delete the AccessKey pairs of organizations.

Resource sets

Each time an organization is created, the system automatically creates a resource set for that organization.

A resource set is a collection of resources. Resource sets are isolated from each other. You can grant permissions to users on resource sets to allow them to use the resources. An organization can have multiple resource sets, but each resource set can belong to only a single organization.

When a user creates a resource, they need to configure the resource set to which the resource belongs. This way, resources can be managed in groups. Users that are granted permissions on a resource set can use resources in that resource set.

Operations for managing resource sets:

- Create resource sets
- Change the names of resource sets
- Grant users the permissions on resource sets
- Delete resource sets

Regions

The Apsara Uni-manager Management Console allows you to manage the regions in which resources are created by associating regions with organizations.

Regions can only be associated with level-1 organizations. An organization can be associated with multiple regions, and a region can be associated with multiple organizations.

Changes

The Apsara Uni-manager Management Console allows you to modify the organizations, resources, and users in accordance with the changes of your organization and project. Changes cannot be made across level-1 organizations.

- Change organizations: You can change the hierarchical relationships of organizations at level-2 and lower. Level-1 organizations do not support changes of hierarchical relationships.
- Change resources: You can change the resource sets to which resources belong.
- Change users: You can change the organizations and roles of users, including the users that belong to level-1 organizations.

4.2.2. Organization management

4.2.2.1. Create an organization

You can create organizations to store resource sets and their resources. This topic describes how to create an organization.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources > Organizations**.
4. In the organization navigation tree, click the name of the parent organization under which you want to create an organization. Click **Create Organization** in the upper-right corner of

the page.

5. In the Add Organization dialog box, enter an organization name and click **OK**.

4.2.2.2. Manage organizations

You can modify or delete organizations in the console as your organization changes.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources** > **Organizations**.
4. (Optional) Search for the organization in the search box above the organization navigation tree. Fuzzy search is supported.
5. Click the name of the organization and perform operations as described in the following table.

Operation	Procedure
View the information of the organization	Click the **Resource Sets**, **Users**, or **User Groups** tab to view the resource sets, users, and user groups of the organization.
Change the name of the organization	i. Click **Editing**. ii. In the dialog box that appears, specify the new name of the organization. iii. Click **OK**.
Delete the organization	**Important** ◦ Before you delete an organization, make sure that the organization does not contain users, resource sets, or subordinate organizations. Otherwise, the organization cannot be deleted. ◦ After the organization is deleted, the data of the organization cannot be recovered. Proceed with caution. i. Click **Delete**. ii. In the message that appears, click **OK**.

4.2.2.3. Manage the AccessKey pairs of an organization

When a level-1 organization is created, an AccessKey pair is automatically generated for that organization. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. AccessKey pairs are used to authenticate request senders. The AccessKey ID is used to identify a user, and the AccessKey secret is used to encrypt the signature string.

A maximum of two AccessKey pairs can be created for each level-1 organization. Only operations administrators and level-1 organization administrators can obtain AccessKey pairs of an organization.

- Operations administrators can manage the AccessKey pairs for all level-1 organizations.
- Level-1 organization administrators can manage only the AccessKey pairs of their respective organization.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources** > **Organizations**.
4. In the organization list, click the name of the level-1 organization whose AccessKey pairs you want to manage.
5. On the right side of the page, click **Management Accesskey**.

Operation	Procedure
View an AccessKey pair	View the AcessKey ID and AccessKey secret. ◦ Click the icon to show the hidden AccessKey secret. ◦ Click the icon to hide the AccessKey secret.
Create an AccessKey pair	In the Management AccessKey dialog box, click**Create AccessKey**. **Note** A maximum of two AccessKey pairs can be created for each level-1 organization.
Disable an AccessKey pair	**Important** Disabling an AccessKey pair may cause errors in the applications that depend on it. Proceed with caution. i. In the Management AccessKey dialog box, find the target AccessKey pair and click **Disable** in the Actions column. ii. In the message that appears, click**OK**.
Enable an AccessKey pair	In the Management AccessKey dialog box, find the target AccessKey pair and click **Enable** in the Actions column.

Delete an AccessKey pair	**Important** Deleting an AccessKey pair may cause errors in the applications that depend on it. Proceed with caution. i. In the Management AccessKey dialog box, find the target AccessKey pair and click **Enable** in the Actions column. ii. In the message that appears, click **OK**.

4.2.3. Resource set management

4.2.3.1. Create a resource set

You must create a resource set before you apply for resources.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources > Resource Sets**.
4. In the upper-left corner, click **Create Resource Set**.
5. In the **Create Resource Set** dialog box, set **Name** and **Organization**.
6. Click **OK**.

4.2.3.2. Manage resource sets

You can change the name of a resource set, authorize users to use a resource set, or delete a resource set.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources** > **Resource Sets**.
4. (Optional) In the upper part of the page, select an organization or search for the target resource set by using the search box.
5. Find the target resource set and perform operations as described in the following table.

Operation	Procedure
View the information of a resource set	i. Click the name of a resource set to view its details. ii. On the Resource Set Details page, click the **Resources** and **Users** tabs to view all resources and users of the resource set.

Change the name of a resource set	i. Find the target resource set and click **Edit Name** in the Actions column. ii. In the Modify Resource Set dialog box, specify a new name for the resource set. iii. Click **OK**.
Grant users the permissions on a resource set	i. Find the target resource set and click **Member authorization** in the Actions column. ii. (Optional) In the Member authorization dialog box, select the role type of the user or the user group whose permissions you want to manage. iii. Grant or revoke permissions. ▪ Grant the permissions to a user or a user group: In the lower-left section, select the user or the user group that you want to authorize. ▪ Revoke the permissions of a user or a user group: In the lower-right section, click **Remove** in the Actions column corresponding to the target user or user group. iv. Click **OK**.
Delete a resource set	**Important** ◦ The default resource set cannot be deleted. ◦ Before you delete a resource set, make sure that the resource set does not contain resources and that no users or user groups are authorized to use it. Otherwise, the resource set cannot be deleted. i. Find the resource set that you want to delete and click **Delete** in the Actions column. ii. In the message that appears, click **OK**.

4.2.4. Manage regions

You can manage the regions associated with organizations. Regions determine the geographical locations in which resources are created.

Regions can be associated only with level-1 organizations, and resources of organizations at lower levels can only be created in the regions associated with their respective level-1 organization.

An organization can be associated with multiple regions, and a region can be associated with multiple organizations.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources > Regions**.
4. In the left-side organization navigation tree, click the name of the target organization.
5. In the right-side **Regions** section, select the names of the regions to be associated with the organization.
6. Click **Update Association**.

4.2.5. Change management

You can modify the organizations, resources, and users in the console according to the changes of your organizations or projects.

Changes cannot be made across level-1 organizations, and names of organizations and users within a level-1 organization must be unique.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Resources > Changes**.
4. (Optional) Search for the target organization in the search box above the organization navigation tree. Fuzzy search is supported.
5. Find the target organization and perform the operations described in the following table.

Operation	Procedure
Change the parent organization of an organization	i. Click the name of the subordinate organization. ii. On the right side of the page, click **Change Organization**. iii. Select a new parent organization from the **Change Organization To** drop-down list. iv. Click **OK**. **Note** After the subordinate organization is moved to the new parent organization, the resource sets and users of the subordinate organization are also moved.
Change the ownership of resources	i. Click the + icon to the left of the target organization and click a resource set under the organization. ii. On the right side of the page, select a **product type** and a **resource type**. iii. Change the ownership of resources. ▪ Change the ownership of a single resource: In the instance list, find the target resource and click **Change Ownership** in the Actions column. ▪ Change the ownership of multiple resources: In the instance list, select the target resources and click **Batch Change Ownership** at the bottom of the page. iv. In the dialog box that appears, select a resource set from the **Change Resource Set To** drop-down list. v. Click **OK**. **Note** Changing the ownership of a resource does not affect the organization and users.

Change the organization to which a user belongs	i. Click the name of the target organization. ii. On the right side of the page, click **Users**. iii. Find the target user and click **Change** in the Actions column. iv. In the dialog box that appears, select an organization from the **Change Organization To** drop-down list and select roles from the **Assigned Roles** drop-down list. v. Click **OK**. **Note** Changing the organization of a resource does not affect the organization and users.

4.3. User management

4.3.1. Overview

The user management and user group management features allow you to centrally manage the structure and permissions of users and user groups. This can improve management efficiency and meet the various needs for system and resource access.

User management

You can perform management operations on users such as managing lifecycle for users, attaching roles to users, and controlling logons.

- On the System Users tab of the Users page, you can create, modify, disable, enable, and delete users. You can also grant permissions, configure logon policies, view user information, add users to user groups, remove users from user groups, configure multi-factor authentication (MFA), view the initial passwords of users, and view the resources sets of users.
- On the Historical Users tab, you can restore deleted users or permanently delete users.

User group management

You can add users to user groups to implement batch permission management. You can:

- Create user groups
- Modify user groups
- Add users to user groups
- Grant permissions to user groups
- Remove users from user groups
- Delete user groups

4.3.2. User management

4.3.2.1. Create a user

Administrators can create users and assign the users different roles to meet different requirements for system access control.

Create a single user

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users** > **Users**.
4. Click the **System Users** tab. Click **Create a user**.
5. In the **Create a user** dialog box, configure the parameters.

Parameter	Required	Description
User name	Yes	The name of the user. The name must be 2 to 64 characters in length and can contain letters, digits, hyphens (-), underscores (_), and periods (.). **Note** User names must be unique.
Display name	Yes	The display name of the user. The name must be 1 to 128 characters in length and can contain letters, digits, hyphens (-), periods (.), and at signs (@).
Role	Yes	The roles to be assigned to the user. **Note** You can enter role names in the field. Fuzzy match is supported.
Organization	Yes	The organization to which the user belongs.
Logon policy	Yes	The logon policy that limits the logon time and IP address of the user. The default policy is automatically associated with new users. **Note** The default policy does not limit the logon time and IP address of the user. To impose limits, you can modify the default logon policy or create a new logon policy for the user. For more information, see Create a logon policy.

Phone	Yes	The phone number of the user. This phone number is used to notify users of resource application and usage. Make sure that this mobile number is valid. **Note** If the phone number of the user changes, update it on the system in a timely manner.
Landline	No	The landline number of the user. The landline number must be 4 to 20 characters in length and can contain only digits and hyphens (-).
Email	Yes	The email address of the user. Emails about the resource application and usage are sent to this email address. Make sure that this email address is valid. **Note** If the email address changes, update it on the system in a timely manner.
DingTalk Key	No	The key of the chatbot for the DingTalk group to which the user belongs. For more information about how to configure the key, see DingTalk development documentation.
Notify User by Email	No	If this option is selected, the Apsara Uni-manager Management Console notifies the user configured as the alert contact by email whenever an alert is generated. **Note** You must configure an email server to notify the user by email. For more information, contact on-site O&M engineers.
Notify User by DingTalk	No	If this option is selected, the Apsara Uni-manager Management Console notifies the user configured as the alert contact by DingTalk whenever an alert is generated.

Create multiple users at a time

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users** > **Users**.
4. Click the **System Users** tab. Click **Batch Create Users**.

> **Note**
>
> - To batch create users, you need to download the template file, fill in the information of the users in the spreadsheet, and then upload the file to the system.
> - A maximum of 100 users can be created at a time.

5. In the **Batch Create Users** dialog box, click **User Creation Template**.
6. In the downloaded template file, fill in the information of the users.

Parameter	Required	Description
Username	Yes	The name of the user. The name must be 2 to 64 characters in length and can contain letters, digits, hyphens (-), underscores (_), and periods (.). **Note** User names must be unique.
Display Name	Yes	The display name of the user. The name must be 1 to 128 characters in length and can contain letters, digits, hyphens (-), periods (.), and at signs (@).
Role	Yes	The role of the user. For a list of roles, click**Permissions** > **Role Permissions** in the left-side navigation pane in the console.
Organization	Yes	The path of the organization to which the user belongs. **Note** Format: root/<Level-1 organization>/<Level-2 organization> If the root organization has been renamed, use the actual name.
Logon Policy	Yes	The name of the logon policy of the user. For a list of logon policies, click **Permissions** > **Access control** in the left-side navigation pane in the console.
Country Code	Yes	The code of the country that the user is in.

Mobile Phone	Yes	The phone number of the user. This phone number is used to notify users of resource application and usage. Make sure that this mobile number is valid. **Note** If the phone number of the user changes, update it on the system in a timely manner.
Landline	No	The landline number of the user. The landline number must be 4 to 20 characters in length and can contain only digits and hyphens (-).
Email	Yes	The email address of the user. Emails about the resource application and usage are sent to this email address. Make sure that this email address is valid. **Note** If the email address changes, update it on the system in a timely manner.
Enable Email Notifications	No	Valid values: **Yes** and **No**. Default value: **No**. If you set the parameter to Yes, the Apsara Uni-manager Management Console notifies the user configured as the alert contact by email whenever an alert is generated. **Note** You must configure an email server to notify the user by email. For more information, contact on-site O&M engineers.
Enable DingTalk Notifications	No	Valid values: **Yes** and **No**. Default value: **No**. If you set the parameter to Yes, the Apsara Uni-manager Management Console notifies the user configured as the alert contact by DingTalk whenever an alert is generated.

DingTalk Key	No	The key of the chatbot for the DingTalk group to which the user belongs. For more information about how to configure the key, see DingTalk development documentation. **Note** This parameter takes effect only when Enable DingTalk Notifications is set to Yes.
Email me the password after the user is created.	No	Valid values: **Yes** and **No**. Default value: **No**. **Note** This parameter takes effect only when Enable Email Notifications is set to Yes.

7. Save the file. Click **Next** in the **Batch Create Users** dialog box. Click **Upload Template**.

> **Note**
> When the template file is uploaded, the system checks the validity of the parameter values. If invalid values exist, the users cannot be created. You need to modify the template file and re-upload it.

8. After the file is uploaded, click **OK**.
9. (Optional) On the **Batch Tasks** tab of the **Users** page, view the status and details of the batch user creation task. If **Succeeded** is displayed in the **Status** column of the task, the task has been completed.

4.3.2.2. Manage system users

You can perform management operations on users such as managing their lifecycle in the console, attaching roles to users, and controlling logons.

An operations administrator can manage the system users of all organizations. An organization administrator can manage only the system users of the corresponding organization and the subordinate organizations.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users > Users**.
4. Click the **System Users** tab and perform the operations described in the following table.

Operation	Procedure

View the information of a user	i. On the Users page, view the organization, role, mobile phone number, email, status, logon policy, and logon time of a user. ii. Click the name of a user, or move the pointer over the More icon in the Actions column and click **View User Information** in the menu that appears. In the panel that appears, view more information of the user.
View the initial password of a user	Move the pointer over the ··· icon in the Actions column corresponding to a user and click **View Initial Password**.
View the resource set that a user is authorized to use	i. On the Users page, move the pointer over the ··· icon in the Actions column corresponding to a user and click **View Resource Set**. ii. Select a role from the drop-down list in the panel and view the resource sets of the role.
Modify a user	i. Click **Edit** in the Actions column corresponding to a user. ii. In the dialog box that appears, modify the information of the user. iii. Click **OK**.
Disable a user	**Important** Disabling a user will forbid the user from logging on to the system and using AccessKey pairs and will immediately log the user off the system. i. Click **Disable** in the Actions column corresponding to a user. ii. In the dialog box that appears, select **I have known the risks and confirm this operation.** iii. Click **Disable**.
Enable a user	**Important** Enabling a user will also enable the associated AccessKey pairs. i. Click **Enable** in the Actions column corresponding to a user. ii. In the message that appears, click **OK**.
Grant permissions to a user	i. Move the pointer over the ··· icon in the Actions column corresponding to a user and click **Authorize**. ii. Modify the **Role** parameter to add, remove, or change the roles of the user. iii. Click **OK**.

Configure the logon policy of a user	i. Move the pointer over the ··· icon in the Actions column corresponding to a user and click **Configure Logon Policy**. ii. In the dialog box that appears, select a logon policy from the Logon policy drop-down list. iii. Click **OK**.
Configure the user groups of a user	i. Move the pointer over the ··· icon in the Actions column corresponding to a user. ii. Add the user to user groups or remove the user from user groups. ▪ Add the user to user groups: Click **Add to User Group**. In the dialog box that appears, select the target user groups from the drop-down list. ▪ Remove the user from user groups: Click **Remove from User Group**. In the dialog box that appears, select the target user groups from the drop-down list. iii. Click **OK**.
Configure multi-factor authentication (MFA) for a user	i. Move the pointer over the ··· icon in the Actions column corresponding to a user and click **MFA Settings**. ii. In the dialog box that appears, turn MFA Settings on or off. ▪ indicates that MFA is enabled. ▪ indicates that MFA is disabled. iii. Click **OK**. **Note** For information about how to associate devices for MFA, see the "User information management" topic.
Delete a RAM user	**Important** ◦ Deleting a user will forbid the user from logging on to the system and using AccessKey pairs, and will immediately log the user off the system. ◦ Deleted user data cannot be restored. Proceed with caution. i. Move the pointer over the ··· icon in the Actions column corresponding to a user and click **Delete**. ii. In the dialog box that appears, select **I have known the risks and confirm this delete operation.** iii. Click **Delete**.

4.3.2.3. Manage historical users

You can restore or permanently delete deleted users.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users > Users**.
4. Click the **Historical Users** tab.

Operation	Procedure
Restore a user	After you restore a deleted user, the user will be added back to the System Users list. i. Click **Restore** in the Actions column corresponding to a user. ii. In the dialog box that appears, select the organization to which the user belongs and the roles of the user. iii. Click **OK**.
Permanently delete a user	After you permanently delete a user, the user will be removed from the Historical Users list and can no longer be restored. i. Click **Delete** in the Actions column corresponding to a user. ii. In the message that appears, click **Delete**.

4.3.3. User group management

4.3.3.1. Create a user group

This topic describes how to create user groups in an organization. You can use user groups to manage permissions of users in batch.

Prerequisites

An organization is created. For more information, see Create an organization.

Background information

Subjects	Description
User groups and organizations	• A user group belongs to only a single organization. • Multiple user groups can be created within an organization.

User groups and users	• A user group can contain multiple users. It can also contain no user. • A user does not necessarily need to belong to a user group. • A user can be added to multiple user groups.
User groups and roles	• A role can be assigned to multiple user groups. • When a role is assigned to a user group, the permissions that the role has are automatically granted to users within the user group.
User groups and resource sets	• A resource set can contain multiple user groups. It can also contain no user group. • A user group can be added to multiple resource sets.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users > User Groups**.
4. Click **Create a user group** in the upper-left corner.
5. In the dialog box that appears, set **User Group Name**, **Organization**, and **Role authorization**.

Parameter	Description
User Group Name	The name of the user group. The name must be 3 to 255 characters in length and can contain letters, digits, underscores (_), hyphens (-), periods (.), and at signs (@).
Organization	The organization to which the user group belongs.
Role authorization	The roles that are assigned to the user group.

6. Click **OK**.

4.3.3.2. Manage user groups

You can modify the name, permissions, and users of a user group.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users > User Groups**.
4. Perform operations on user groups as described in the following table.

Operation	Procedure
View user groups	On the User Groups page, view information of user groups, such as the organization, roles, number of users, and associated resource sets.
Modify a user group	i. Click **Editing** in the Actions column corresponding to a user group. ii. In the dialog box that appears, specify the new name of the user group. iii. Click **OK**.
Add a user	After a user is added to a user group, the user has the permissions of the role that is assigned to the user group. i. Find the user group to which you want to add users and click**Add User** in the Actions column. ii. In the left-side list, select the users to be added and click the > icon to add the users to the right-side list. iii. Click **OK**.
Remove a user	After a user is removed from a user group, the user no longer has the permissions of the role that is assigned to the user group. i. Find the user group from which you want to remove users and move the pointer over the ··· icon in the Actions column. In the menu that appears, select **Remove User**. ii. In the left-side list, select the users you want to remove and click the > icon to add them to the right-side list. iii. Click **OK**.
Assign roles to a user group	You can assign a role to a user group to grant the role permissions to all users within the group. i. Find the target user group and move the pointer over the ··· icon in the Actions column. In the menu that appears, select **Authorized**. ii. In the Role authorization drop-down list, select the roles you want to assign to the user group. iii. Click **OK**.

Delete a group	**Important** When a user group is deleted, resources such as roles, users, and resource sets that are associated with the user group are automatically unbound. i. Find the target user group and move the pointer over the ··· icon in the Actions column. In the menu that appears, select **Delete**. ii. In the message that appears, click **OK**.

4.4. Permission management

4.4.1. Overview

The Apsara Uni-manager Management Console provides permission management features, such as role permissions, data permissions, access control, permission boundaries, and AccessKey logs. You can use the features to manage permissions based on scenarios, users, and cloud services, thereby improving the security of your systems.

Role permissions

You can create roles and assign them to users or user groups, so as to implement fine-grained access control on users. The following operations are supported:

- Create custom roles
- Disable roles
- Enable roles
- Delete roles
- Replicate roles

Data permissions

You can authorize users to access specified cloud instances. You can also view and modify the data permissions granted to users. The following operations are supported:

- Configure the data permissions on cloud instances
- View user data permissions
- Modify user data permissions

Access control

You can configure access policies to control the time and IP addresses for user logon. After a user is associated with an access policy, user logons are limited based on the logon time and IP addresses specified in the policy. The following operations are supported:

- Create access policies
- Modify access policies
- Disable access policies
- Enable access policies
- Delete access policies

Permission boundaries

You can create service boundaries and configure API control policies to block API operations on the core gateway of Apsara Stack. Custom API control policies can be created. The following operations are supported:

- Create service boundaries and API control policies
- View service boundaries and API control policies
- Modify service boundaries and API control policies
- Edit service boundaries and API control policies
- Enable and disable service boundaries and API control policies
- Modify whitelists of service boundaries
- Delete service boundaries and API control policies

AccessKey logs

You can view AccessKey logs to facilitate security management centered around AccessKeys pairs. The following operations are supported:

- View the basic information of AccessKey pairs
- View the cloud services accessed by using an AccessKey pair

4.4.2. Role permissions

4.4.2.1. Overview

A role is a collection of access permissions. In the Apsara Uni-manager Management Console, role types include preset roles, custom roles, RAM role-based authorization, and RAM-based authorization.

Preset role

To implement privilege separation, the Apsara Uni-manager Management Console provides preset roles based on the general use cases in cloud platforms. For more information, see User roles and permissions.

Custom role

You can create custom roles and configure their sharing scopes and permissions on resources, applications, and menus, so as to implement find-grained permission management. Custom roles are commonly used to manage permissions on cloud products.

RAM role-based authorization

You can customize the permission policies, trust policies, and RAM roles for RAM role-based authorization. After you grant access permissions on resources to a RAM role, the trust policy authorizes the cloud service in the level-1 organization to assume the role and access or manage resources in other services. RAM role-based authorization is commonly used to manage the permissions of cloud services or instances.

RAM-based authorization

You can customize the permission policies and user groups for RAM-based authorization. If you want all users in a user group to have all the permissions of the user group on cloud resources, you can use RAM-based authorization. RAM-based authorization is commonly used in cross-account access and temporary authorization scenarios.

Terms

RAM (Resource Access Management)

RAM is a service provided by Apsara Stack. It allows you to manage user identities and resource access permissions.

identity

RAM provides three types of identities: RAM user, RAM user group, and RAM role. RAM users and RAM user groups are physical identities. RAM roles are virtual identities.

RAM user

A physical identity that has a fixed ID and credential information. A RAM user represents a person or an application.

RAM user group

A physical identity that contains a group of RAM users. You can create RAM user groups to classify and authorize RAM users. This simplifies the management of RAM users and permissions.

RAM role

If an entity user assumes a RAM role, the entity user can obtain and use the Security Token Service (STS) token of the role to access the authorized resources.

4.4.2.2. RAM policies

This topic describes the elements, structure, and syntax of RAM policies that are used to define the scope of permissions for custom roles, RAM authorization and RAM role-based authorization.

Structure

RAM policies are in the JSON format, containing two parts:

- Version: the version number of the policy.
- Statements: the statements that specify the permissions. The elements of statements include Effect, Action, Resource, Condition, and Principal.

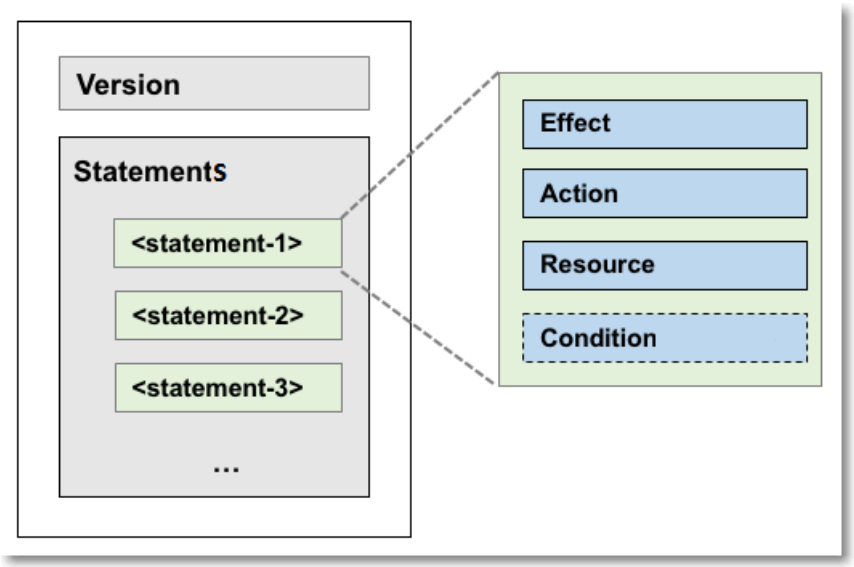

Elements in statements

The elements of statements include Effect, Action, Resource, Condition, and Principal.

Effect (Required)

Specifies whether a statement result is an explicit allow or an explicit deny. Valid values: Allow and Deny.

> **Note**
>
> Precedence of Deny: Deny statements take precedence when a user is assigned multiple policies that contain Allow and Deny statements, and when a single policy contains both Allow and Deny statements.

Example

```
"Effect": "Allow"
```

Action (Required)

The specific API operations involved in the statement.

The Action element is in the following format: `<ram-code>:<action-name>`.

- `<ram-code>`: the product code.
- `<action-name>`: the API operation of the product involved in the policy.

> **Note**
>
> The values of the Action element are case-insensitive. However, to ensure consistency, we recommend that you follow the standard capitalization of the product code and API operation names.

Example

```
"Action": [
  "oss:ListBuckets",
  "ecs:Describe*",
  "rds:Describe*"
]
```

> **Note**
>
> You can use the asterisk (*) as a wildcard character. For example, in the preceding example, `ecs:Describe*` indicates that all the API operations of Elastic Compute Service (ECS) whose names start with Describe are involved in the policy.

Resource (Required)

The resources of cloud services involved in the statement, which are indicated by their Aliyun Resource Names (ARNs)

Resource elements are in the following format: `acs:<ram-code>:<region>:<account-id>:<relative-id>`.

- `acs`: the prefix of ARNs. It stands for Alibaba Cloud Service.
- `<ram-code>`: the product code.

- `<region>`: the region ID of the resource. You can use the asterisk (*) as a wildcard character to indicate regionless resources or resources in all regions.
- `<account-id>`: the user ID of the level-1 organization. This is used to specify the resources that belong to a specific level-1 organization.
- `<relative-id>`: the description of the specific resources. Its syntax is product-specific. For example, the relative-id for an OSS object is a path of that object, such as mybucket/dir1/object1.jpg.

Example

```
"Resource": [
 "acs:ecs:*:*:instance/inst-001",
 "acs:ecs:*:*:instance/inst-002",
 "acs:oss:*:*:mybucket",
 "acs:oss:*:*:mybucket/*"
]
```

Condition (Optional)

The conditions that are required for the policy to take effect. A Condition element, or a condition block, contains one or multiple conditions. Each condition consists of an operator, a condition key, and a condition value.

Condition keys are case-sensitive. The case-sensitivity of the condition values is determined by the operator. For example, assume that a condition key of the string type is specified. if the StringEquals operator is used, the system performs a case-sensitive match. If the StringEqualsIgnoreCase operator is used, the system performs a case-insensitive match.

Description:

- Evaluation logic
  - You can specify one or more values for a condition key. If the value in a request matches one of the values, the condition is met.
  - A condition can have multiple keys that are attached to a single conditional operator. The condition of this type is met only if all requirements for the keys are met.
  - A condition block is met only if all of its conditions are met.
- Operators

  Operators are classified into the following types: string, number, data and time, Boolean, and IP address.

Category	Operator
String	◦ StringEquals ◦ StringNotEquals ◦ StringEqualsIgnoreCase ◦ StringNotEqualsIgnoreCase ◦ StringLike ◦ StringNotLike

Number	◦ NumericEquals ◦ NumericNotEquals ◦ NumericLessThan ◦ NumericLessThanEquals ◦ NumericGreaterThan ◦ NumericGreaterThanEquals
Date and time	◦ DateEquals ◦ DateNotEquals ◦ DateLessThan ◦ DateLessThanEquals ◦ DateGreaterThan ◦ DateGreaterThanEquals
Boolean	Bool
IP address	◦ IpAddress ◦ NotIpAddress

- Condition keys

- General condition keys are in the following format: `acs:<condition-key>`.

General condition key	Type	Description
acs:CurrentTime	Date and time	The time when the server receives the request. Specify the time in the ISO 8601 standard, for example, 2012-11-11T23:59:59Z.
acs:SecureTransport	Boolean	Specifies whether a secure channel is used to send a request. For example, a request can be sent over HTTPS.
acs:SourceIp	IP address	The IP address of the client that sends the requestIt can be a single IP address or a CIDR block. If the value is a single IP address, you must specify the specific IP address rather than a CIDR block. For example, you must specify 10.0.0.1 rather than 10.0.0.1/32.
acs:MFAPresent	Boolean	Specifies whether multi-factor authentication (MFA) is used when the user who sends the request logged on.
acs:PrincipalARN	String	The identity of the requester. This operator is only used in trust policies of RAM roles.

- Service-specific condition keys are in the following format: `<ram-code>:<condition-key>`.

Service-specific condition key	Cloud service	Type	Description
ecs:tag/<tag-key>	Elastic Compute Service (ECS)	String	The condition key for tags of ECS resources, which are customizable.
rds:ResourceTag/<tag-key>	ApsaraDB RDS	String	The condition key for tags of ApsaraDB RDS resources, which are customizable.
oss:Delimiter	Objective Storage Service (OSS)	String	The delimiter used by OSS to categorize object names.
oss:Prefix	OSS	String	The prefix of an OSS object name.

Principal

The entity that is allowed or denied to call the related API operations. This element is available only for resource-based policies. For example, in a trust policy attached to a RAM role, you can specify the resources that are allowed to assume the role in the Principle element. The Principal element is required in resource-based policies.

> **Note**
>
> You cannot specify the Principal element in identity-based policies. Identity-based policies are attached to RAM identities. The RAM identities to which identity-based policies are attached are the principals.

You can specify multiple types of principals for the Principal element in a policy. You can also specify multiple principals of the same type for the Principal element in a policy. If you specify multiple types of principals, separate the types with commas (,). If you specify multiple principals of the same type, separate the principals with commas (,) and enclose the principals with brackets ([]). Multiple principals are associated with each other by the OR operator, meaning that the policy is attached to all principals in the policy.

Example

```
"Principal": {
 "RAM": [
 "acs:ram::123456789012****:root",
 "acs:ram::987654321098****:root"
 ],
 "Service": "ecs.aliyuncs.com"
}
```

Principals are authenticated RAM entities, which are divided into the following types:

- Organization account

  Set the Principal element to the account of the level-1 organization. When you configure the trust policy of a RAM role-based authorization, the account ARN of the selected trust organization will be added to the policy as the principal, such as `acs:ram::165399189392****:root`.

- Cloud service

  Set the Principal element to a cloud service. When you configure the trust policy of a RAM role-based authorization, the identifier of the selected trust service will be added to the policy as the principal, such as `ecs.aliyuncs.com`.

Policy syntax

A RAM policy consists of the version_block and statement_block components.

```
policy = {
 <version_block>,
 <statement_block>
}
```

version_block

The version of the policy. By default, the value is 1.

```
<version_block> = "Version" : ("1")
```

statement_block

The statements. A RAM policy can contain multiple statements.

```
<statement_block> = "Statement" : [ <statement>, <statement>, ... ]
```

- Each statement consists of the effect_block, action_block, resource_block, and condition_block elements. condition_block is an optional element.

  ```
  <statement> = {
   <effect_block>,
   <action_block>,
   <resource_block>,
   <condition_block?>
  }
  ```

  - Valid values of an effect_block element are `Allow` and `Deny`.

    ```
    <effect_block> = "Effect" : ("Allow" | "Deny")
    ```

  - Multiple action_block and resource_block elements can be specified in each statement.

    ```
    <action_block> = "Action" :
     ("*" | [<action_string>, <action_string>, ...])
    <resource_block> = "Resource" :
     ("*" | [<resource_string>, <resource_string>, ...])
    ```

  - Each condition_block element can contain multiple operators and

    ```
    <condition_block> = "Condition" : <condition_map>

    <condition_map> = {
      <condition_type_string> : {
          <condition_key_string> : <condition_value_list>,
          <condition_key_string> : <condition_value_list>,
          ...
      },
      <condition_type_string> : {
          <condition_key_string> : <condition_value_list>,
          <condition_key_string> : <condition_value_list>,
          ...
      }, ...
    }
    ```

Valid values of elements

- If an element value is a string, number, date, time, Boolean value, or an IP address, the value must be enclosed in double quotation marks ("").
- If an element value is a string, you can use asterisks (*) and question marks (?) as wildcard characters.
  - An asterisk (*) indicates 0 or multiple letters. For example, `ecs:Describe*` indicates all the API operations of Elastic Compute Service (ECS) whose names start with Describe.
  - A question mark (?) indicates a single letter.

Policy syntax check

RAM policies follow the JSON format. When you create or update a RAM policy, the system automatically checks the validity of the JSON format. We recommend that you use JSON validators or formatters to help you ensure the validity of the policy. For more information about JSON syntax standards, see RFC 7159.

4.4.2.3. Create a role

You can create custom roles, RAM role-based authorization, and RAM-based authorization to facilitate permission management.

Create a custom role

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions > Role Permissions**.
4. Click **Create Custom Role**.
5. In the **Basic Configuration** step, configure the following parameters and click **Next**.

Parameter	Description
Role name	The name of the role.
Description	The description of the role. This is an optional parameter
Sharing Scope	◦ Global: The permissions of the role are shared by all organizations. ◦ Current Organization: The permissions of the role are shared within the current organization. ◦ Subordinate Organization: The permissions of the role are shared by the current organization and its subordinate organizations.
Management Permissions	◦ All Organizations: The permissions apply to all organizations. ◦ Specified Organization and Subordinate Organizations: The permissions apply to the organization to which the user belongs and its subordinate organizations. ◦ Resource Sets: The permissions apply to the resource sets of the user.
Custom Policy	Turn the switch on or off to enable or disable custom policies. : Custom policies are not supported. : Custom policies are supported.

6. In the **Manage Permissions** step, configure the parameters and click **Next**.

In this step, you can specify the management operations that the role can perform in the Apsara Uni-manager Management Console. You can use the left-side pane to filter the permissions.

7. In the **Application Permissions** step, configure the parameters and click **Next**.

   In this step, you can specify the operations that the role can perform on the cloud products. You can use the left-side pane to filter the permissions.

8. In the **Menu Permissions** step, configure the parameters and click **Next**.

   In this step, you can specify the Menu options that are available to the role. You can use the left-side pane to filter the options.

9. (Optional) In the **Custom Policy** step, configure the parameters and click **Create**.

   In the Set RAM Action for Application section, configure fine-grained permissions of the role on APIs of cloud products. In the Allow and Deny fields, specify the API operations that the role can and cannot call. For more information about the two parameters, see the "Action" section of the Permission policy structure and syntax topic.

   > **Note**
   >
   > The Custom Policy step is available if you turn on Custom Policy in the Basic Configuration step.

0. In the **Create Custom Role** message, click **Complete**.

   You can also click **Authorize now** to assign the role to users or user groups on the **Authorized Personnel** tab.

Create RAM role-based authorization

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions > Role Permissions**.
4. Click **Create Custom Role**.
5. Click **Advanced Settings** at the lower part of the page and click **Crate RAM Role & Attach Policy**.
6. In the **Basic Configuration** step, configure the following parameters and click **Next**.

Parameter	Description
Role name	The name of the RAM role-based authorization.
Description	The description of the RAM role-based authorization. This is an optional parameter
Sharing Scope	◦ Global: The permissions of the RAM role-based authorization are shared by all organizations. ◦ Current Organization: The permissions of the RAM role-based authorization are shared within the current organization. ◦ Subordinate Organization: The permissions of the RAM role-based authorization are shared by the current organization and its subordinate organizations.

7. In the **Configure Custom Policy** step, select the policies and click **Next**.

   If none of the listed policies meets your requirements, you can click **Create Policy** in the upper-right corner of the Select Policy section to create one. For information about how to configure the policy, see Permission policy structure and syntax.

   > **Note**
   > You can select a maximum of five policies for the RAM role-based authorization.

8. In the **Configure Trust Policy** step, configure the **Trust Organization** and **Trust Cloud Service** parameters and click **Next**.

   > **Note**
   > You can select multiple trust organizations and trust cloud services for the RAM role-based authorization.

9. In the **Preview Policy** step, check the selected policies and the trust policy and click **Create**.

Create RAM-based authorization

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions > Role Permissions**.
4. Click **Create Custom Role**.
5. Click **Advanced Settings** at the lower part of the page and click **Create RAM Authorization**.
6. Configure the parameters in the **Basic Information** step.

Parameter	Description
Role name	The name of the RAM-based authorization.
Description	The description of the RAM-based authorization. This is an optional parameter
Sharing Scope	◦ Global: The permissions of the RAM-based authorization are shared by all organizations. ◦ Current Organization: The permissions of the RAM-based authorization are shared within the current organization. ◦ Subordinate Organization: The permissions of the RAM-based authorization are shared by the current organization and its subordinate organizations.

7. Click **Create and Configure RAM Authorization**.
8. In the **Role Permissions** step, configure the authorization.
   - **User Groups**
     a. Click the **User Groups** tab.

b. Click **Add User Group**. In the dialog box that appears, select the target user groups and click **OK**.

- **Permissions**

a. Click the **Permissions** tab.

b. Click **Add Policy**. In the dialog box that appears, select the target policies or create a custom policy, and click **OK**. For information about how to create a policy, see Permission policy structure and syntax.

> **Note**
> You can select a maximum of five policies for the RAM-based authorization.

9. After the configuration is complete, click **OK**.

4.4.2.4. Manage roles

You can change the permissions and status of a role to change the permissions and logon status of the users that are assigned this role.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions > Role Permissions**.
4. (Optional) Select a role type in the drop-down list to filter the roles.
5. Find the target role and perform the operations described in the following table.

Operation	Procedure
View information of a role	i. Click the name of a role. ii. On the page that appears, view detailed information of the role. ▪ **Present role and custom role**: You can click the **Management Permissions**, **Application Permissions**, **Menu Permissions**, and **Authorized Personnel** tabs to view information of the role. ▪ **RAM role-based batch authorization**: You can click the **Custom Policy**, **Trust Policy**, and **RAM Role** tabs to view information of the RAM role-based authorization. ▪ **RAM-based batch authorization**: You can click the **User Groups** and **Permissions** tabs to view information of the RAM-based authorization.

Modify a role	i. Click **Modify** in the Actions column of the role. ii. On the page that appears, click **Modify Basic Information** in the upper-right corner. iii. Modify the permissions of the role. ▪ Custom role: You can click the **Management Permissions**, **Application Permissions**, **Menu Permissions**, and **Authorized Personnel** tabs to modify the permissions of the role. ▪ RAM role-based batch authorization: You can click the **Custom Policy**, **Trust Policy**, and **RAM Role** tabs to modify the permissions of the RAM role-based authorization. ▪ RAM-based batch authorization: You can click the **User Groups** and **Permissions** tabs to modify the permissions of the RAM role-based authorization. **Note** Preset roles cannot be modified.
Disable a role	When a role is disabled, all its permissions are revoked, and no users associated with the role can log on to the console. i. Click **Disable** in the Actions columns corresponding to the role. ii. In the message that appears, click **OK**. **Note** RAM role-based batch authorization and RAM-based batch authorization cannot be disabled.
Enable a role	Click **Enable** in the Actions columns corresponding to a disabled role. **Note** RAM role-based batch authorization and RAM-based batch authorization cannot be enabled.

Copy a role	When you copy a role, a new role with the same sharing scope, management permissions, application permissions, and menu permissions is created. You can adjust the new role based on your needs. i. Move the pointer over the ··· icon in the Actions column corresponding to the role and click **Copy** in the menu that appears. ii. In the message that appears, click **OK**. **Note** RAM role-based batch authorization and RAM-based batch authorization cannot be copied.
Delete a role	**Important** ◦ Preset roles cannot be deleted. ◦ Before your delete a role, make sure that no users or user groups are associated with the role. i. Move the pointer over the ··· icon in the Actions column corresponding to the role and click **Delete** in the menu that appears. ii. In the message that appears, click **OK**.

4.4.3. Data permissions

4.4.3.1. Overview

The data permission management feature allows you to control who can perform what operations on specific instances. You can also view and configure the operations that a user can perform on instances, which is specified in RAM policies.

In Apsara Stack, administrators control users' permissions on the Apsara Uni-manager Management Console, including what operations they can perform, what services they can use, and what modules they can access. However, in actual usage, a substantial amount of operations are performed through calling API operations instead of using the console. For these operations, RAM policies need to be used to control the permissions.

4.4.3.2. Set the data permissions on resource instances

By configuring the data permissions on the Apsara Uni-manager Management Console, organization administrators can control the operations that users can perform on specific cloud resources.

Background information

Data permissions are supported by the following cloud services: Message Queue, Object Storage Service (OSS), Simple Log Service, DataHub, Container Service for Kubernetes (ACK), Key Management Service (KMS), and Tablestore.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions** > **Data Permissions**.
4. Choose the resource set to which the target instance belongs in the navigation tree. On the right side of the page, select a product type and click **Authorize** in the **Actions** column corresponding to the instance.
5. In the **Data Permissions** dialog box, select the user that you want to manage on the left side.

   You can also turn on **Batch Grant Permissions** in the upper-right corner of the dialog box. This way, you can select multiple users and manage the permissions in batch.
6. In the **Instance Permission Settings** section of the dialog box, select or clear the permissions on the instance. Then, click **OK**.

   You can also select **Enable All** in the upper-right corner of the section to grant all permissions to the users.

4.4.3.3. Configure the data permissions of users

You can write policies with JSON statements to configure the permissions of a user on cloud resources.

For information about the JSON syntax and structure of policies, see Permission policy structure and syntax.

> **Note**
>
> Different from **setting the data permissions on resource instances**, the JSON statement method can be used to configure the permissions on all cloud products.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions > Role Permissions**.
4. Click the name of an organization.
5. On the right side of the page, click **Users**.

Operation	Procedure
View user permissions	View the permission policies of a user. i. Click **View Permissions** in the Actions column corresponding to a user. ii. In the Policies section of the dialog box that appears, select a policy and view its content in the right-side editor.

Modify user permissions	You can use JSON statements to specify user permissions. i. Click **Edit Permissions** in the Actions column corresponding to a user. ii. In the Policies section of the dialog box that appears, select a policy and modify its content in the right-side editor. iii. Click **OK**.

4.4.4. Access control

4.4.4.1. Create an access policy

To improve the security of the Apsara Uni-manager Management Console, you can create access policies as an administrator to control logon access based on the logon time and IP addresses of users.

Background information

Access policies are used to control the time periods and IP addresses valid for user logon. After a user is associated with an access policy, user logons are limited based on the logon time and IP addresses specified in the policy.

The default policy generated by the Apsara Uni-manager Management Console does not have limits on the time and IP addresses valid for user logon. This policy cannot be deleted.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions > Access control**.
4. Click **Create an access policy** in the upper-right corner.
5. In the **Create an access policy** dialog box, configure the Name, Sharing Scope, Policy Property, Logon Time, and Logon Address parameters.

   Table 1. Parameters of access policies

Parameter	Description
Name	The name of the access policy. The name can be 2 to 50 characters in length and can contain only letters and digits. The name must be unique in the system.
Description	The description of the access policy.
Sharing Scope	The scope in which the role is visible. Valid values: ◦ **Global**: The role is globally visible. This is the default value. ◦ **Current Organization**: The role is visible only in the current organization and is invisible in subordinate organizations. ◦ **Subordinate Organization**: The role is visible in the current organization and all its subordinate organizations.

Policy Property	The authentication method of the access policy. Valid values: ◦ **Whitelist**: Logon is allowed if the parameter settings are met. ◦ **Blacklist**: Logon is denied if the parameter settings are met.
Logon Time	The permitted logon time period. Users can log on to the Apsara Uni-manager Management Console only within the configured period. Specify the time period in minutes in a 24-hour clock. Example: 16:32. **Note** When the Policy Property parameter is set to Whitelist, the No Time Limit option is available.
Logon Address	The permitted CIDR block. ◦ If the **Policy Property** parameter is set to **Whitelist**, IP addresses within this CIDR block are allowed to log on to the Apsara Uni-manager Management Console. ◦ If the **Policy Property** parameter is set to **Blacklist**, IP addresses within this CIDR block are not allowed to log on to the Apsara Uni-manager Management Console. **Note** When the Policy Property parameter is set to Whitelist, the No CIDR Block Limit option is available.

6. Click **OK**.

4.4.4.2. Manage access policies

You can modify, disable, enable, and delete access policies.

Note

The default access policy cannot be modified, disabled, enabled, or deleted.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions** > **Access control**.
4. Find the target access policy and perform the operations described in the following table.

Operation	Procedure

View an access policy	i. On the Access control page, you can view the information of an access policy, such as its policy properties, status, sharing scope, users, access time, endpoint, and description. ii. You can click the number in the **Users** column to view the users associated with this access policy.
Modify an access policy	i. Click **Modify** in the Actions column corresponding to an access policy. ii. In the dialog box that appears, modify the Name, Description, Policy Property, Logon Time, and Logon Address parameters. iii. Click **OK**.
Disable an access policy	You can disable the policies that are not needed at the moment. **Note** When an access policy is disabled, users that are associated with the policy cannot log on to the Apsara Uni-manager Management Console. i. Click **Disable** in the **Actions** column corresponding to an access policy. ii. In the message that appears, click **OK**.
Enable an access policy	Click **Enable** in the Actions column corresponding to a disabled access policy.
Delete an access policy	**Note** Before you delete an access policy, make sure that it is not associated with any user. i. Click **Delete** in the Actions column corresponding to an access policy. ii. In the message that appears, click **OK**.

4.4.5. Permission boundary

4.4.5.1. Service boundary management

The Apsara Uni-manager Management Console leverages the API control capabilities of the core gateway of Apsara Stack to provide the service boundary feature. You can create service boundaries to implement fine-grained control on API operations based on organizations, usage methods, and time periods.

4.4.5.1.1. Create a service boundary

You can associate API control policies with service boundaries to implement fine-grained control on API operations.

Background information

The Apsara Uni-manager Management Console allows you to control operations by organizations, time, and operation methods through setting service boundaries. You can also use service boundaries to trigger approval processes when cloud operations are performed.

Service boundaries take effect on operations performed through all methods, including cloud accounts, RAM users, and STS tokens. Specially, they also take effect on operations performed on the console. This is because when users perform operations on the console, the system is calling API operations under the hood.

> **Important**
>
> API operations performed across cloud services (such as when Resource Orchestration Service calls an Elastic Compute Service operation to manage the resources) are also controlled by service boundaries. Therefore, inappropriate configurations may interrupt the normal usage of cloud services. You need to add the ARNs of relevant resources to the whitelist of the service boundaries and make sure that you are aware of the impact of the service boundaries before you enable them.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**. In the left-side navigation pane, choose **Permissions** > **Permission Boundary** > **Service Boundary**.
3. On the **Service Boundary Management** page, click **Create Service Boundary**. On the page that appears, configure the parameters.

Section	Parameter	Description
Basic Information	Service Boundary Name	The name of the service boundary.
	Description	The description about the service boundary.
Management Scope	Management Scope	The organizations for which the service boundary takes effect. Valid values: **Global** and **Specified Organization**. When you select Specified Organization, you need to select organizations for the **Authorize Organization** parameter.
	Associate API Control Policy	Select one or more API control policies to be associated with the service boundary. In an API control policy, you can specify conditions based on parameter values of specific API operations. For more information, see Create an API control policy.
Control Mode	Console Calls	The action taken when an operation in the console matches the control policy. ◦ **Reject**: rejects the operation. ◦ **Process Approval**: initiates a preconfigured approval process for the operation. If you select this option, you also need to configure the **Process** parameter by selecting an existing process or create a new process in the **Configure Resource Application Process** dialog box. For more information about processes, see Overview of process management.

	API Calls	The action taken when an API call matches the control policy. The value is set to **Reject** and cannot be changed.
Period Settings	Effective Period	The periods of time during which the service boundary takes effect. ◦ **By Period**: Specify a time period by selecting when the period begins and ends. ◦ **Every Day**: Specify a time period and the effective time in each day during the period. ◦ **Every Week**: Specify a time period, the days of week in which the service boundary takes effect, and the effective time in the days.
Settings	Status	The status the service boundary enters immediately when it is created. Valid values: ◦ **Enable** ◦ **Disable** **Important** We recommend that you disable the service boundary first, and enable it after you have comprehensive knowledge about its influence.
	Whitelist	The resources (RAM authorizations), RAM users, and RAM roles that are exempt from the service boundary. Specify the entities by Alibaba Cloud Resource Names (ARNs). You can fuzzy search and select existing ARNs from the drop-down list, or specify a custom ARN by following the format. **Note** Up to 10 ARNs can be specified in the whitelist of a service boundary.

4. Click **Submit**.

4.4.5.1.2. Manage service boundaries

You can view, modify, enable, disable, and delete service boundaries and manage their whitelists in the Apsara Uni-manager Management Console.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**. In the left-side navigation pane, choose **Permissions** > **Permission Boundary** > **Service Boundary**.
3. On the Service Boundary Management page, manage the service boundaries.

Description	Procedure
View the details of service boundaries	Click the name or ID of the target service boundary. On the page that appears, you can view the details of the service boundary, including its basic information, associated API control policies, approval processes, and whitelists.
Modify service boundaries	◦ Modify the configurations of a service boundary: Click **Edit** in the **Actions** column corresponding to the target service boundary. On the page that appears, you can modify the basic information, management scope, control mode, and effective period of the service boundary. ◦ Modify the API control policies associated with a service boundary: Click **Modify Policy** in the **Actions** column corresponding to a service boundary. On the **Associate Policy** tab of the page that appears, click **Add API Control Policy** or **Remove** to associate or disassociate control policies.
Enable or disable service boundaries	Click **Enable** or **Disable** in the **Actions** column corresponding to the target service boundary.
Manage whitelists	Move your pointer over the More icon in the **Actions** column corresponding to a service boundary and click **Manage Whitelist**. On the **Whitelist** tab of the page that appears, you can click **Add** or **Edit** to add or modify the ARNs that are exempt from the service boundary.
Delete service boundaries	Move your pointer over the More icon in the **Actions** column corresponding to a service boundary and click **Delete**.

4.4.5.2. Control policy management

API control policies are associated with service boundaries to implement fine-grained access control. In an API control policy, you can specify conditions based on parameter values of specific API operations. After it is associated with a service boundary, actions will be taken when the conditions in the API control policy are met. For example, you can define a control policy that enables service boundaries to take actions when users query details of instances whose names contain a specific string.

Supported services

The following table lists the cloud services and their API versions that are supported by API control policies.

> **Note**
>
> The services and API versions listed in the console prevail. We recommend that you refer to relevant API documentation to learn about the operations and their parameters.

Service name	Service code	API version

Elastic Compute Service (ECS)	Ecs	2014-05-26
		2016-03-14
Apsara File Storage NAS	NAS	2017-06-26
Virtual Private Cloud (VPC)	Vpc	2016-04-28
ApsaraDB RDS	Rds	2014-08-15
PolarDB	polardb	2017-08-01
KVStore for Redis	R-kvstore	2015-01-01
ApsaraDB for MongoDB	Dds	2015-12-01
		2022-11-21
API Gateway	CloudAPI	2016-02-01
		2016-07-01
		2016-07-14

Content of control policies

The content of a control policy consists of a coordinate and parameters. The policies are stored as a JSON file that can be modified on the control policy details page.

Element	Description
coordinate	The coordinate of the API operation. Format: `${popCode}:${version}:${action}`. popCode indicates the product code; version indicates the API version; action indicates the API operation. You can use an asterisk (*) as a wildcard character to replace ${version} so that all API versions are included.
parameters	An array that contains details about the rule, including the name of the parameter (key), logical expression (condition), and specified value (value).
┗ condition	The logical relationship between the actual value of the parameter and the specified value.

┕ value	The specified value. **Note** The value can be left empty only if condition is set to NotNull or IsNull.
┕ key	The name of the controlled API operation.

Example

```
[
 {
  "coordinate": "CloudAPI:2016-02-01:AbolishApi",
  "parameters": [
   {
    "condition": "IsNull",
    "value": "",
    "key": "AccessKeyId"
   }
  ]
 }
]
```

Data types of parameter values and logical expressions

The following table lists the logical expressions that can be specified in API control policies.

Type	Logical expression	Description
Numeric (such as INTEGER, LONG, and FLOAT)	NumericEquals	The actual value equals the specified value in the policy.
	NumericNotEquals	The actual value does not equal the specified value in the policy.
	NumericLessThan	The actual value is smaller than the specified value in the policy.
	NumericGreaterThan	The actual value is larger than the specified value in the policy.
	NumericLessThanEquals	The actual value is smaller than or equal to the specified value in the policy.

	NumericGreaterThanEquals	The actual value is larger than or equal to the specified value in the policy.
	IsNull	The parameter is left empty.
	NotNull	The parameter is not empty.
String	StringLike	The actual value includes the specified string in the policy.
	StringNotLike	The actual value does not include the specified string in the policy.
	StringEquals	The actual value matches the specified string in the policy (case-sensitive).
	StringNotEquals	The actual value does not match the specified string in the policy (case-sensitive).
	StringEqualsIgnoreCase	The actual value matches the specified string in the policy (case-insensitive).
	IsNull	The parameter is left empty.
	NotNull	The parameter is not empty.
BOOLEAN	Boolean_Equals	The actual value matches the specified value in the policy.
	IsNull	The parameter is left empty.
	NotNull	The parameter is not empty.

4.4.5.2.1. Create an API control policy

In the Apsara Uni-manager Management Console, you can create API control policies to control API operations on cloud resources based on values of the request parameters.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**. In the left-side navigation pane, choose **Permissions** > **Permission Boundary** > **Control Policies**.
3. On the **API control policy management** page, click **Create an API control policy**.
4. On the **Create Control Policy** page, specify the name and description of the policy and click **Create** in the Control Content section. In the dialog box that appears, configure the following parameters.

Parameter	Description
Cloud Services	The cloud services whose API operations are to be controlled.
Version	The API version. Some cloud services may have multiple API versions. We recommend that you select All Versions for this parameter.
API	The API operation that is controlled. **Note** Each rule is set for a single operation. If you want to specify rules for multiple operations in a single policy, you can click **Add Rule**. The rules in a policy are in **OR** logical relations. This means that when the policy is attached to a service boundary, the policy takes effect as long as one of the rules is met.
Field	The request parameter to be controlled, a specified value, and the logical relation between the values in actual requests and the specified value. The logical relations vary with the types of the selected request parameters. Valid values include **Similar To**, **Equal to String**, **Equal to (case-insensitive)**, **Empty**, **Not Empty**. **Note** You can click **Add Field** to specify multiple conditions. The conditions in a rule are in **AND** logical relations. This means that the rule takes effect only if all conditions are met. **Example**: The following figure shows how to configure the Field parameter for a rule that is created to control the action of creating a VPC named "test".

5. Click **Create**. Then, you can attach the API control policy to service boundaries.

4.4.5.2.2. Manage API control policies

You can view the details of existing API control policies, including their content and association with service boundaries. You can also modify and delete the policies.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**. In the left-side navigation pane, choose **Permissions** > **Permission Boundary** > **Control Policies**.
3. On the page that appears, the existing policies are listed, and you can perform operations on them.

Operation	Procedure
View details of an API control policy	Click the ID of an API control policy. On the**Policy Details** tab of the page that appears, you can view the basic information and control content of the policy. On the **Reference Information** tab, you can view information of the service boundaries with which the policy is associated. You can click **Disassociate** in the Actions column corresponding to a service boundary to cancel the association.
Modify an API control policy	◦ Click **Modify** in the **Actions** column corresponding to a policy. In the dialog box that appears, you can modify its name and description. ◦ Click **Modify Control Content** in the **Actions** column corresponding to a policy. On the policy details page that appears, you can modify or delete the existing rules in the policy, or create new rules. **Note** ▪ When you modify a rule, you cannot change the cloud service, API version, or the API operation involved. ▪ After you delete a rule, service boundaries with which the policy is associated will be affected. Proceed with caution. If you are familiar with the format of control policies, you can click **Custom Edit** to modify the JSON file of the policy. This way, you can modify the cloud service, API version, and the API operation involved. You need to make sure that the modifications comply with the JSON format and that you are aware of the modifications' impact on relevant service boundaries.
Delete an API control policy	Click **Delete** in the **Actions** column corresponding to a policy. **Note** Policies that are associated with service boundaries cannot be deleted.

4.4.6. AccessKey logs

You can learn about the use of AccessKey pairs from AccessKey logs to improve security.

Background information

AccessKey pairs are supported for the following cloud products.

Product category	Product code	Product name
Computing	bms	Bare-metal Management Service (BMS)
	oos	Operation Orchestration Service (OOS)
	ros	Resource Orchestration Service (ROS)
	ess	Auto Scaling
	ehpc	Elastic High Performance Computing (E-HPC)
	ecs	Dedicated Host
	cr/cr-ee	Container Registry (ACR)
	cs	Container Service for Kubernetes (ACK)
Storage	hdr	Hybrid Disaster Recovery (HDR)
	ebs	Elastic Block Storage (EBS)
	pds	Photo and Drive Service
	ots/tablestore	Tablestore **Note** Fat client authentication may be used for Tablestore. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.

	csg	Cloud Storage Gateway (CSG) **Note** Fat client authentication may be used for CSG. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.
	nas	Apsara File Storage NAS
Networking	Cbn	Cloud Enterprise Network (CEN)
	vpc	VPN Gateway
	vpc	Virtual Private Cloud (VPC)
	slb	Server Load Balancer (SLB)
	vpc	NAT gateway
	vpc	IPv6 Gateway
	-	Hybrid Cloud Network
	vpc	High availability virtual IP address (HAVIP)
	vpc	Express Connect
	vpc	Elastic IP address (EIP)
	cloudDns	Apsara Stack DNS
	drds	PolarDB-X 1.0
	rds	ApsaraDB RDS
	r-kvstore	KVStore for Redis
	polardbx	PolarDB-X 2.0

Database		
	hitsdb	Lindorm
	hitsdb	Time Series Database (TSDB)
	gpdb	AnalyticDB for PostgreSQL
	gdb	Graph Database
	drds	PolarDB
	dds	ApsaraDB for MongoDB
	ads	AnalyticDB for MySQL V2.0
	adb	AnalyticDB for MySQL V3.0
	Hbase	ApsaraDB for HBase
	dts	Data Transmission Service (DTS)
	dms-enterprise	Data Management (DMS)
	dbs	Database Backup (DBS)
Middleware	ons-inner	Message Queue (MQ) **Note** Fat client authentication may be used for MQ. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.

	alikafka	Message Queue for Apache Kafka **Note** Fat client authentication may be used for Message Queue for Apache Kafka. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.
	csb	Cloud Service Bus (CSB)
Application Services	cloudapi	API Gateway
	sls	Log Service **Note** Fat client authentication may be used for Log Service. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.
	edas	Enterprise Distributed Application Service (EDAS)
Security	Config	Cloud Config
	kms	Key Management Service (KMS) **Note** Fat client authentication may be used for KMS. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.
	odps	MaxCompute **Note** Fat client authentication may be used for MaxCompute. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.

Big data	datahub	DataHub
	elasticsearch-k8s	Elasticsearch on K8s
	elasticsearch	Elasticsearch
	hologram	Hologres
	dataworks-private-cloud	DataWorks **Note** Fat client authentication may be used for DataWorks. If an AccessKey pair is used for fat client authentication, the AccessKey logs may not record the complete information.
	eas	Machine Learning Platform for AI (PAI)
IoT	linkwan	Apsara Stack Link WAN
	-	Link IoT Edge
	-	IoT Security Operation Center
	iotid	IoT Internet Device ID
	iot	IoT Platform
Monitoring and O&M	cms	CloudMonitor
Disaster recovery	asr	Apsara Stack Resilience for Zone-disaster Recovery

View AccessKey logs

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Permissions** > **AccessKey Logs**.

4. Enter the AccessKey ID to be queried in the search box and click the icon. View the information of the AccessKey pair in the **Basic Information About AccessKey Pair** and **Accessed Services** sections.

   > **Note**
   >
   > AccessKey logs may be generated 1 hour after the AccessKey pair is used. By default, AccessKey logs are stored for seven days. Proceed with caution. If you want the logs to be stored for longer, you need to configure log storage, which is supported by Log Service. For more information about how to configure log storage, see Log storage settings.

   - **Basic Information About AccessKey Pair**
     - **Type**: the type of the AccessKey pair. An AccessKey pair may be a user AccessKey pair or an organization AccessKey pair.
     - **User Logon Name**: the logon name of the user.

       > **Note**
       >
       > This parameter is empty if the queried AccessKey pair is an organization AccessKey pair.

     - **UID**: the ID of the user or organization to which the AccessKey pair belongs.
     - **Organization**: the name of the organization to which the AccessKey pair belongs.

       > **Note**
       >
       > This parameter is empty if the queried AccessKey pair is a user AccessKey pair.

     - **AccessKeyID**: the ID of the AccessKey pair.
     - **Last Used At**: the last time when the AccessKey pair was used. Format: `YYYY:MM:DD:hh:mm:ss`.
     - **Last Used Service**: the cloud service that the AccessKey pair was used to access last time.
   - **Accessed Services**
     - **Events**: Click **Events** in the **Actions** column corresponding to a service to view the API operations called by using the AccessKey pair.
     - **IP Addresses**: Click **IP Addresses** in the **Actions** column corresponding to a service to view the IP addresses of the user of the AccessKey pair.
     - **Resources**: Click **Resources** in the **Actions** column corresponding to a service to view the instances accessed by using the AccessKey pair.

       > **Note**
       >
       > If the AccessKey pair has not been used to access any instance, no record is displayed.

5.Configurations

5.1. Overview

The Apsara Uni-manager Management Console provides a GUI for configuration, which enables simple and centralized management of system configurations.

The configuration center consists of the following modules.

Category	Module	Description
System configuration	Resource Pool Configuration	Manage logical inventory of cloud resources to control the maximum usage of resources.
	Region Settings	Associate regions with geographical locations.
	Security policies	Configure password policies for logons and control logons.
	Authentication Mode	Configure appropriate authentication modes for different scenarios.
	Specifications	Enable or disable available specifications of different cloud resources to control specifications that can be configured when resources are created.
	Menu Settings	Configure the menu bar as required, so as to manage the entry points of service and operation configurations.
	Announcement Management	Manage announcements that inform users of updates, security risks, or emergencies.
	Custom Configurations	Customize elements on the platform such as logon pages and platform logos.
	Authentication by Third Party	Allow the Apsara Uni-manager Management Console (as the identity provider) and third-party systems (as the service providers) to implement single sign-on (SSO) through Security Assertion Markup Language (SAML).
	Message Gateway	Configure email gateways, SMS gateways, DingTalk notifications, so that alerts are sent to alert contacts in the forms of emails, SMS, and DingTalk messages.
Others	Service-linked Roles	Configure permission policies through service-linked roles, so that cloud services in the level-1 organization are authorized to view and manage resources in other organizations.

	Log Storage Settings	Configure the storage methods of logs of cloud services, including operation logs and AccessKey logs.

5.2. System configuration

5.2.1. Resource pool configuration

You can manage the logical inventory of cloud resources to control the maximum usage of resources.

Background information

- If the physical inventory is unlimited, the logical inventory cannot be less than the allocated inventory.

  > **Important**
  >
  > When the logical inventory is less than the used inventory, resources cannot be created.

- If the physical inventory is limited, the logical inventory cannot be less than the used inventory and cannot be greater than the physical inventory.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Resource Pool Configuration**.
4. On the Resource Pools page, click the cloud product that you want to manage.

   You can also use the search box to search for the cloud product by name.
5. In the **Resource Pool Details** pane, click **Modify Inventory**.

   > **Note**
   >
   > If the product is **Elastic Compute Service (ECS)**, you can also configure the quotas of custom instance types. Default quotas cannot be configured.
   >
   > i. In the Resource Pool Details pane, click **Configure Quota**.
   >
   > ii. Add or remove quotas.
   >
   > - Add quotas: In the **Available Quotas** list, click **Add** in the **Actions** column corresponding to an instance type. The instance type will be added to the **Added Quotas** list.
   > - Delete quotas: In the **Added Quotas** list, click **Delete** in the **Actions** column corresponding to an instance type. The default quotas cannot be deleted.
   >
   > iii. Click **OK**.
6. In the **Logical Inventory** column, enter the quotas.
7. Click **Save**.

5.2.2. Region settings

This topic describes how to associate a region with a geographical location. You can modify the location information when the location changes.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Region Settings**.
4. On the Regions page, view the names of regions and their provinces, cities, and districts.
5. Click **Modify Location** in the Actions column corresponding to a region and modify the information of the province, city, and district in the map.

   > **Note**
   > District information is required.

6. Click **OK**.

5.2.3. Security policies

You can configure password policies and logon control to meet various security requirements.

Password policies

1. Log on to the Apsara Uni-manager Management Console as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Security Policies**. On the **Password Policy** tab, configure the following parameters.

Parameter		Description
Basic Settings	Password Length	The minimum length of a password. Valid values: 8 to 32.
	Required Character Types in Password	The character types that must be included in a password. You can select multiple options. Valid values: ◦ **Lowercase Letters** ◦ **Uppercase Letters** ◦ **Digits** ◦ **Special Characters**

	Password Validity (Days)	The validity period of a password. Unit: days. After the validity expires, the password becomes invalid. Valid values: 0 to 1095. **Note** If you set this parameter to 0, passwords will never expire.
Logon Settings	Logon Disabled After Password Expires	Specifies whether to allow logon after the password expires. Valid values: ◦ **Yes** ◦ **No**
	Password Attempts	The maximum number of logon attempts within 1 hour. After the maximum number is reached, the account is locked and cannot be used to log on for a specific period of time. Valid values: 0 to 32. **Note** If you set this parameter to 0, password retries are not limited.
	Account Lockout Duration (Minutes)	The period of time during which logons from an account are not allowed. Unit: minutes. If the maximum number of attempts is reached, the account is locked. During the lockout period, the account cannot be used to log on to the console. Valid values: 0 to 1440. **Note** If you set this parameter to 0, account lockout is disabled.
	Password History Check	The number of previous passwords that cannot be the same with the current password. Valid value: 0 to 24. **Note** If you set this parameter to 0, the system does not check historical passwords. For example, if you set the parameter to 3, the current password cannot be the same as one of the latest three passwords that are used.
	Password Expiration Notification	Specifies whether the system sends notifications about expiration of passwords. Valid values: ◦ **Send Notification** ◦ **No Notification**

Password Modification Settings	Notification Method	The notification method. Valid values: ◦ **Email** ◦ **DingTalk** ◦ **Message** **Note** This parameter is available only when Password Expiration Notification is set to Send Notification
	Notification Template	The template based on which the system sends notifications. For information about how to configure template content, see the "Configure a message template" section of the Message gateway topic. **Note** This parameter is available only when Password Expiration Notification is set to Send Notification.
	Notification Time	The number of days prior to the expiration of a password on which the system sends notifications. Valid values: 1 to 30 **Note** This parameter is available only when Password Expiration Notification is set to Send Notification

4. Click **Save**.

   You can click **Reset** to use the default settings.

Configure logon control

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Security Policies**.
4. Click the **Logon Control** tab.

Parameter	Description
Session Mode	The session mode. Valid values: ◦ Single Session: A user can log on only by using a single client. ◦ Multi-session: A user can log on by using multiple clients at the same time.

Global MFA	Specifies whether to enable global multi-factor authentication (MFA). ◦ When the switch is turned on, MFA is required for all accounts on the platform. ◦ When the switch is turned off, MFA is not required for accounts on the platform.
Effective Scope	The scope of accounts for which global MFA is enabled. Valid values: ◦ Does Not Include System Accounts (Such as Administrator Accounts) ◦ Include System Accounts (Such as Administrator Accounts) **Note** This parameter is available only when Global MFA is turned on.
SSO User Logon	Specifies whether to support single sing-on (SSO) users to log on to the platform. ◦ When the switch is turned on, SSO users can directly log on to this platform. ◦ When the switch is turned off, SSO users cannot directly log on to this platform and can log on only through SSO. **Note** Platform administrators are not subject to this configuration.
Logon Minutes During Which No Actions Occur	Adjust the slider to select a period based on your business requirements. If no actions are performed for the specified number of minutes, the user is forced to log off. The user must log on again.
Logon Hours Before Forced Logoff	Adjust the slider to select a period based on your business requirements. After the specified period, the user is forced to log off. The user must log on again. If you set this parameter to 0, this feature is not enabled.

5. Click **Save**.

5.2.4. Authentication modes

5.2.4.1. Authentication modes

This topic describes the authentication modes. Authentication modes are classified into the token mode and the hybrid mode. You can configure an authentication mode based on your business requirements.

Token mode

In token mode, Resource Access Management (RAM) and Security Token Service (STS) are required. Temporary STS tokens are issued to users and cloud services to assist in permission management.

- In token mode, all cloud services use STS tokens to authenticate users. Cloud services that do not support STS may be affected. For example, only operations administrators and level-1 organization administrators can log on to the console of a cloud service.

  > **Note**
  >
  > On the **Authentication Mode** page, you can view the cloud services that support the token mode.

- After you enable the token mode, authentication of northbound APIs may be affected.

  For example, if you use the AccessKey pair of a RAM user to call an API operation to add, delete, modify, or query specified resources in a specified organization and resource set, and the required RAM policies are not attached to the RAM user, an error may be returned.

  > **Warning**
  >
  > Before you enable the token mode, we recommend that you contact Alibaba Cloud to evaluate the situation. If an error occurs when you call an API operation, disable the token mode.

Hybrid mode

In hybrid mode, cloud services that support RAM and STS can still authenticate users based on STS tokens. For cloud services that do not support STS, permissions are managed based on the authentication logic of Apsara Stack tenant accounts and RAM users. In hybrid mode, if you select one or more cloud services, the cloud services use STS tokens to authenticate users. Cloud services that are not selected still authenticate users in hybrid mode.

5.2.4.2. Configure an authentication mode

You can select an appropriate authentication mode for different scenarios.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Authentication Mode**.
4. Select **Token Mode** or **Hybrid** based on the scenario.

   If you select **Hybrid**, you can specify cloud services to use token-based authentication.

   > **Important**
   >
   > After token-based authentication is enabled, the authentication mode of northbound API operations is affected. We recommend that you contact the Alibaba Cloud service team for evaluation before you enable token-based authentication. If an error occurs when an API operation is called, roll back and disable token-based authentication for the service.

5. In the lower part of the page, click **Save**.

5.2.5. Specification management

5.2.5.1. Specification parameters

As a platform administrator, you can enable and disable specifications of cloud services to control the specifications that are available when users create cloud resources. This topic describes the specification parameters of each configurable cloud service.

Computing

Elastic Compute Service (ECS)

Parameter	Description
Instance Family	The instance family that is divided into different instance types based on the scenarios for which they are suitable.
Specifications Level	The level of the instance type that can be configured for ECS.
vCPUs	The maximum number of vCPUs that can be configured for ECS.
Memory (GB)	The memory size that can be configured for ECS.
Specifications	The code of the instance type that can be configured for ECS.
GPU Type	The GPU type that can be configured for ECS.
GPUs	The number of GPUs that can be configured for an ECS instance.
Supported ENIs	The number of Elastic network interfaces (ENIs) that can be configured for an ECS instance.
Private IP Addresses	The number of private IP addresses that can be configured for an ECS instance.
Disk Type	The disk type that can be configured for an ECS instance.

Storage

Object Storage Service (OSS)

Parameter	Description
Specifications	The specifications that can be configured for OSS.

Specifications Description	The description of the specifications that can be configured for OSS.

Networking

NAT Gateways

Parameter	Description
Specifications	The available sizes of NAT gateways.
Specifications Description	The description of the specifications that can be configured for NAT gateways.

Server Load Balancer (SLB)

Parameter	Description
Specifications	The instance type that can be configured for SLB.
Specification	The category of the SLB specification.
Maximum Number of Connections	The maximum number of connections that can be configured for SLB.
New Connections	The number of new connections that can be configured for SLB.
QPS	The queries per second (QPS) that can be configured for SLB.
Description	The description of the specifications that can be configured for SLB.

IPv6 Gateway

Parameter	Description
Specifications	The available specifications of IPv6 gateways.

Database

ApsaraDB RDS

Parameter	Description
Engine Type	The engine type that can be configured for ApsaraDB RDS.
Minimum Storage (GB)	The minimum amount of storage space that can be configured for ApsaraDB RDS.
Maximum Storage (GB)	The maximum amount of storage space that can be configured for ApsaraDB RDS.
Instance Type Name	The code of the instance type that can be configured for ApsaraDB RDS.
Chip type	The type of chip that can be configured for ApsaraDB RDS.
Version	The engine version that can be configured for ApsaraDB RDS.
CPU Cores	The number of CPU cores that can be configured for ApsaraDB RDS.
Maximum Number of Connections	The maximum number of connections that can be configured for ApsaraDB RDS.
Memory (GB)	The memory size that can be configured for ApsaraDB RDS.
Edition	The instance edition that can be configured for ApsaraDB RDS.
Disk Type	The type of disk that can be configured for ApsaraDB RDS.

KVStore for Redis

Parameter	Description
Version Type	The version type that can be configured for KVStore for Redis.
Engine Version	The engine version that can be configured for KVStore for Redis.

Architecture	The instance architecture that can configured for KVStore for Redis.
Node Type	The type of node that can configured for KVStore for Redis.
Specifications	The instance type that can be configured for KVStore for Redis.
Maximum Number of Connections	The maximum number of connections that can be configured for KVStore for Redis.
Maximum Bandwidth	The maximum bandwidths that can be configured for KVStore for Redis.
CPU Cores	The number of CPU cores that can be configured for KVStore for Redis.
Memory	The memory size that can be configured for KVStore for Redis.
Service Plan	The type of service plan that can be configured for KVStore for Redis.

PolarDB-X 2.0

Parameter	Description
Edition	The edition that can be configured for PolarDB-X 2.0.
Chip Type	The type of chip that can be configured for PolarDB-X 2.0.
Engine Version	The engine version that can be configured for PolarDB-X 2.0.
Instance Type	The instance type that can be configured for PolarDB-X 2.0.
Node Specifications	The node specifications that can be configured for PolarDB-X 2.0, including the specification name, memory size, and CPU cores.

PolarDB-X 1.0

Parameter	Description
Instance Type	The instance type that can be configured for PolarDB-X 1.0.

Instance Type Name	The name of the instance type that can be configured for PolarDB-X 1.0.
Specifications	The specifications that can be configured for PolarDB-X 1.0.
Instance Type Name	The name of the specifications that can be configured for PolarDB-X 1.0.

AnalyticDB for PostgreSQL

Parameter	Description
Specifications	The specifications that can be configured for AnalyticDB for PostgreSQL.
CPU	The number of CPU cores that can be configured for AnalyticDB for PostgreSQL.
Memory	The memory size that can be configured for AnalyticDB for PostgreSQL.
Storage Space	The storage capacity that can be configured for AnalyticDB for PostgreSQL.
Version	The engine version that can be configured for AnalyticDB for PostgreSQL.
Nodes	The nodes that can be configured for AnalyticDB for PostgreSQL.

ApsaraDB for MongoDB

Parameter	Description
Specifications	The specifications that can be configured for ApsaraDB for MongoDB.
Instance Type Name	The name of the instance type that can be configured for ApsaraDB for MongoDB.
Chip Type	The type of chip that can be configured for ApsaraDB for MongoDB.
Version	The engine version that can be configured for ApsaraDB for MongoDB.
Serial Number	The instance architecture that can be configured for ApsaraDB for MongoDB.

Maximum Number of Connections	The maximum number of connections that can be configured for ApsaraDB for MongoDB.
IOPS	The IOPS that can be configured for ApsaraDB for MongoDB.
Storage Space	The storage capacity that can be configured for ApsaraDB for MongoDB.
Minimum Storage	The minimum storage capacity that can be configured for ApsaraDB for MongoDB.
Maximum Storage	The maximum storage capacity that can be configured for ApsaraDB for MongoDB.

Middleware

Cloud-native Gateway

Metric	Description
Specifications	The specifications that can be configured for Cloud-native Gateway.
CPU Cores	The number of CPU cores that can be configured for Cloud-native Gateway.
Memory (GB)	The memory size that can be configured for Cloud-native Gateway.

5.2.5.2. Create or synchronize specifications

You can create or synchronize specifications of cloud resources.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Specifications**.
4. Select a resource type from the **Resource Specification Settings** drop-down list.
5. Click the **Resource Specifications** tab.
   - **Create specifications**

> **Note**
>
> The following resource types support creation of specifications: **NAT Gateway**, **Server Load Balancer (SLB)**, **IPv6 Gateway**, **AnalyticDB for PostgreSQL**, **ApsaraDB RDS**, **PolarDB-X 1.0**, **KVStore for Redis**, **ApsaraDB for MongoDB**, and **Time Series and Spatial-Temporal Database (TSDB)**.

a. Click **Create Specifications**.
b. In the dialog box that appears, configure the parameters. For more information about the specification parameters, see Specification parameters.
c. Click **OK**.

- **Synchronize specifications**

> **Note**
>
> The systems can synchronize the specifications to your current platform by calling API operations. Resource types that support synchronization of specifications are **Elastic Compute Service (ECS)** and **AnalyticDB for MySQL V3.0**.

a. Click **Synchronize Specifications**.
b. In the message that appears, click **OK**.

5.2.5.3. Manage specifications

You can enable or disable specifications of different cloud resources to control specifications that can be configured when resources are created.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Specifications.**

i. In the **Specification Synchronization Settings** section, configure the specifications synchronization mechanism.

> **Important**
>
> After cloud services are updated, new specifications may be introduced. you can configure the specifications synchronization mechanism to make sure the specifications are up-to-date.

a. Click **Edit**. Select **Maintain Current Specifications** or **Synchronize Specifications**.

- **Maintain Current Specifications**: New specifications are not synchronized after version updates.

  > **Note**
  >
  > Select this option when you want to retain the existing configurations after version updates.

- **Synchronize Specifications**: New specifications are synchronized after version updates.

  > **Note**
  >
  > We recommend that you select this option to make sure that specifications are kept up-to-date.

b. Click **Save**.

ii. Select a cloud service from the **Resource Specification Settings** drop-down list and perform the operations described in the following table.

Operation	Procedure
View specifications	▪ Click the **Resource Specifications** tab to view the information of the specifications, including their status. ▪ Click the **Specifications History** tab to view the historical information of the specifications. ▪ Click the **Existing Specifications** tab to view the specifications that have been used and their quantities.
Disable specifications	You can enable a specification so that it is not available when users create the corresponding resources. a. Click the **Resource Specifications** tab. b. Click **Disable** in the **Actions** column corresponding to the specification that you want to disable. c. In the message that appears, click **OK**.
Enable specifications	You can enable disabled specifications. a. Click the **Resource Specifications** tab. b. Click **Enable** in the **Actions** column corresponding to the specification that you want to enable. c. In the message that appears, click **OK**.
Export Instance Types	You can export the specification information to your computer for backup. a. In the message that appears, click **OK**. b. Click **Export**. c. Configure the save path and a file name and click **Download**. The file will be saved as a .csv file.

5.2.6. Manage menu items

You can add links to other platforms to the Apsara Uni-manager Management Console for centralized management and easy access. You can also configure the menu items to manage the access to products and operations.

> **Note**
> The default menu cannot be modified or deleted.

Menu parameters

Parameter	Description

Title	The display name of the menu item.
URL	The URL of the menu item.
Console Type	Different consoles have different domain names. Valid values: oneconsole and other. • oneconsole: You need only to enter the path in the URL field. The domain name is automatically matched. The URL starts with one.console. • other: You must enter the domain name in the URL field.
Icon	The icon displayed in the left-side navigation pane. The icon cannot be changed.
Identifier	The unique identifier of the menu item in the system. This identifier can be used to indicate whether the menu item is selected in the navigation bar. The identifier cannot be changed.
Sort	The display order among the same-level menu items. The larger the value, the lower it is displayed.
Parent Level	The parent level to which the menu item belongs.
API Group	The group to which the menu item belongs. This parameter is available and required when Parent Level is set to **Products**.
Open With	The method to open the page. Valid values: • Default: Opens the page in the current window. • New Window: Opens the page in a new window.
Description	The description of the menu item.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Menu Settings**.
4. On the Main Menu page, perform the operations described in the following table.

Operation	Procedure

Create a menu item	i. Click **Create**. ii. In the dialog box that appears, configure the parameters. For more information about the parameters, see the **Menu parameters** section. iii. Click **OK**.
Modify a menu item	When a menu item is changed, you can modify its information. i. Click **Edit** in the Actions column corresponding to the menu item. ii. In the dialog box that appears, modify the parameters. For more information about the parameters, see the **Menu parameters** section. iii. Click **OK**.
Show a menu item	Select the check box in the **Display** column corresponding to a menu item. This way, the menu item is displayed in the console.
Hide a menu item	Clear the check box in the **Display** column corresponding to a menu item. This way, the menu item is not displayed in the console.
Delete a menu	i. Click **Delete** in the Actions column corresponding to the menu item. ii. Click **OK**.

5.2.7. Announcement management

You can create announcements on the platform to inform users about events such as maintenance, operations, and updates.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, select **Announcement Management**.
4. On the Announcements page, perform operations that are described in the following table.

Operation	Procedure
Create an announcement	i. Click **Create Announcement**. ii. On the Create Announcement page, configure the title, type, and content of the announcement. iii. Click **Create**. After an announcement is created, you need to release it. Released announcements are displayed in the **Announcement** section on the homepage.

Release an announcement	After an announcement in the **To Be Released** state is released, it enters the **Released** state and is displayed in the Announcement section of the homepage of the console. i. Click **Release** in the Actions column corresponding to an announcement. ii. Click **OK**.
Withdraw an announcement	After an announcement in the **Released** state is withdrawn, it enters the **To Be Released** state and is no longer displayed on the homepage. i. Click **Withdrawn** in the Actions column corresponding to an announcement. ii. Click **OK**.
Modify an announcement	You can modify an announcement only if it is in the **To Be Released** state. i. Click **Modify** in the Actions column corresponding to an announcement. ii. Modify the announcement on the page that appears. iii. Click **Modify**.
Delete an announcement	You can delete an announcement when it is no longer needed. Only announcements in the **To Be Released** can be deleted. i. Click **Delete** in the Actions column corresponding to an announcement. ii. Click **OK**.

5.2.8. Custom configurations

You can customize logon pages and platform logos.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Custom Configurations**.

Parameter	Description
Platform language	The languages supported by the platform. You can select multiple languages. Valid values: ◦ Simplified Chinese ◦ Traditional Chinese ◦ English Select the platform default language Select a language as the default language.
Appearance and theme	Select the appearance (light mode or dark mode) and theme color of the console.

Page watermark	Specifies whether to add a watermark on the console. : does not add a watermark. adds the specified watermark.
Switch Settings	◦ Top Navigation Bar Switch: specifies whether to hide the top navigation bar. **Note** If you want to display the hidden top navigation bar after it is hidden, you need to enter the URL of the *Custom configurations of the Apsara Uni-manager Management Console* in your web browser and go to the page as a platform administrator. Then, turn off Top Navigation Bar Switch. ◦ Embedding Switch: specifies whether to enable the embedded features. ◦ You can enter one or multiple URLs to be embedded. Separate multiple URLs with commas (,).
Brand settings	Specify the logos, versions, and copyright notices of the console in different languages. Parameters: ◦ Browser ico icon: Recommended aspect ratio: 1:1 (Example: 32 px × 32 px). Format: ICO, PNG, or JPG. ◦ Platform Logo: Recommended resolution: 160 px × 36 px. Format: PNG or JPG. ◦ Platform name: We recommend that you specify a name that is no longer than eight English letters or four Chinese characters. ◦ Login page: You can configure up to three pages. Each page can have their specific background image and slogan. We recommend that you use an image whose resolution is 1,880 px × 1,600 px in the PNG or JPG format. The slogan must not exceed 120 characters in length. ◦ Copyright Notice: Specify the declaration of the copyright.

4. Click **Determine**.

Related operations

Restore factory settings: Click **Restore factory settings** in the upper-right corner of the page. The configuration takes effect after you log on to the console again.

Important

This operation is irreversible. Proceed with caution.

5.2.9. Third-party authentication

Terms

- **Identity provider (IdP)**: provides management and authentication of identities.

- **Service provider (SP)**: provides services and resources by using the identity management features of the IdP.
- **Security Assertion Markup Language (SAML)**: a protocol for transferring identity and authorization data between an IdP and an SP. The SAML request for SSO can be made by the IdP or the SP.
  - If the request is made by the IdP, the user is redirected to the SP after logging on to the IdP.
  - If the request is made by the SP, the user is redirected to the IdP after visiting the SP, complete the logon, and is redirected back to the SP.
- **OAuth**: an authorization protocol that specifies the following grant types for different use cases: Authorization Code, Client Credentials, Proof Key for Code Exchange (PKCE), and Device Code. The Authorization Code grant type is used by the Apsara Uni-manager Management Console.

5.2.9.1. IdP settings

This topic describes how to configure single sign-on (SSO) when the Apsara Uni-manager Management Console serves as the service provider (SP) and an enterprise system serves as the identity provider (IdP).

Procedure

This section describes the procedures needed when the Apsara Uni-manager Management Console servers as the SP and an enterprise system serves as the IdP.

1. Log on to the Apsara Uni-manager Management Console as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Authentication by Third Party**.
4. On the page that appears, click **Edit** in the **Global authorization jump** and **Global IdP settings** sections and configure the parameters. The following table describes the parameters.

Parameter		Required	Description
Global authorization jump	Enable Redirect	Yes	Specifies whether users are redirected to the configured URL upon logon or logout. Valid values: **Yes** and **No**.
	Authorized URL	No	The URL that users are redirected to upon logon. If you leave this parameter empty, users are redirected to the URL for IdP authentication. **Note** This parameter takes effect if **Enable Redirect** is set to **Yes**.

	Exit URL	No	The URL that users are redirected to upon logout. If you leave this parameter empty, users are redirected to the logon page of the Apsara Uni-manager Management Console. **Note** This parameter takes effect if **Enable Redirect** is set to **Yes**.
	IdP Service Name	Yes	The name of the IdP service. The name can be up to 64 characters in length and can contain letters, digits, and underscores (_). It cannot start with a digit.
	Protocol Type	No	Valid values: **SAML2.0** and **OAUTH2.0**. **Note** If you select **SAML2.0**, you need to use the SAML SP metadata of Apsara Stack. You can click **Download File** to download the metadata.
	IdP Authentication Address	Yes	The URL for IdP authentication. Format: `http://URL` or `https://URL`.
	IdP Certificate	Yes	The Base64-encoded third-party certificate. **Note** This parameter is available if **Protocol Type** is set to **SAML2.0**.
	SLO Endpoint	Yes	The URL for Single Logout (SLO). Format: `http://URL` or `https://URL`. **Note** ◦ The Apsara Uni-manager Management Console uses this URL when it responds to an SLO request from the IdP. ◦ This parameter is available if **Protocol Type** is set to **SAML2.0**.

Global IdP settings	SLO Request Method	Yes	The request method for an SLO. **Note** This parameter is available if **Protocol Type** is set to **SAML2.0**.
	Client ID	Yes	The client ID. **Note** This parameter is available if **Protocol Type** is set to **OAUTH2.0**.
	Client Key	Yes	The client key. **Note** This parameter is available if **Protocol Type** is set to **OAUTH2.0**.
	Service Address for Obtaining Access Token	Yes	The URL for obtaining the access token. Format: `http://URL` or `https://URL`. By default, the POST method is used for this request. **Note** This parameter is available if **Protocol Type** is set to **OAUTH2.0**.

Obtain Verified User Information **Note** ◦ This section is available if **Protocol Type** is set to **OAUTH 2.0**. ◦ Parameters in this section are not part of the standard OAuth 2.0 protocol. Configure the parameters based on your actual needs.	Service Address for Obtaining Verified User	Yes	The URL from which verified users are obtained. Format: `http://URL` and `https://URL`.
	Request Type	No	Valid values: **GET** and **POST**.
	Access Token Request Header Template	Yes	The JSON header template for the requests used to obtain access tokens. **Note** Keep **{{access_token}}** unchanged because it will be replaced by the actual access token. The Apsara Uni-manager Management Console will add the key and value in the template to the header of the requests sent to the service URL.
	Access Token Request Parameter Template	Yes	The JSON parameter template for the requests used to obtain access tokens. **Note** ◦ Keep **{{access_token}}** unchanged because it will be replaced by the actual access token. The Apsara Uni-manager Management Console will add the key and value in the template to the header of the requests sent to the service URL. ◦ This parameter is available if Request Type is set to **GET**.

	Obtain Path for Parsing Verified Username	Yes	The parsing path for the verified username. The objects returned from the service must be JSON objects. **Note** Use periods to indicate the level of the username in the object. For example, `user` indicates that the user is at the root level, and `ascm.user` indicates that the user is at the second level.
	Content-Type	No	Valid values: **application/json; charset=utf-8** and **application/x-www-form-urlencoded**. **Note** This parameter is available if Request Type is set to **POST**.
	Access Token Request Body Template	Yes	The JSON body template for the requests used to obtain access tokens. **Note** Keep **{{access_token}}** unchanged because it will be replaced by the actual access token. The Apsara Uni-manager Management Console will add the key and value in the template to the header of the requests sent to the service URL.

5. Click **Save** for each section.

5.2.10. Message gateway

You can send messages to users through email, SMS, and DingTalk. Before you send notifications, you must configure the methods.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Message Gateway**.
4. Perform the operations described in the following table.

Operation	Procedure

Configure the email gateway	i. Click the **Mail Gateway** tab. ii. Configure the parameters. ▪ **Sender mailbox**: The email address that is used to send messages. ▪ **Sender mailbox password**: The password of the email address. ▪ **Smtp server address**: The address of the SMTP server. Format: smtp.*xxx*.com ▪ **Port**: The port of the SMTP server. ▪ **Encrypted**: Valid values: **No encryption**, **SSL**, and **TLS**. ▪ **Enabled**: Valid values: **Enabled** and **Do not enable**. ▪ **Email Attachment**: Valid values: **Enabled** and **Do not enable**. iii. Click **Submit**. You can click **Reset** to clear the configurations. You can click **Test Mail** and specify a destination mailbox in the dialog box that appears to test the email gateway.
Configure the SMS gateway	i. Click the **SMS Gateway** tab. ii. Configure the parameters. ▪ **SMS gateway address**: The API URL of the SMS gateway. ▪ **Receive Number Parameter Name**: The name of the parameter that specifies the recipient's number in the SMS gateway API. ▪ **SMS content parameter name**: The name of the parameter that specifies the content in the SMS gateway API. ▪ **Request header parameters**: The header parameters for requesting the SMS gateway API. ▪ **Request body parameters**: When the POST method is used, the recipient number parameter name and the content parameter name are included in the request body . When the GET method is used, the two parameter names are included in the query string. ▪ **Content coding**: Valid values: **UTF-8** and **GBK**. ▪ **Transfer mode**: Valid values: **GET** and **POST**. When the parameter is set to GET, the recipient number parameter name and the content parameter name are included in the query string. When the parameter is set to POST, the two parameters are included in the request body. ▪ **Enabled**: Valid values: **Enabled** and **Do not enable**. iii. Click **Submit**. You can click **Reset** to clear the configurations. You can click **Test SMS** and specify a recipient phone number in the dialog box that appears to test the SMS gateway.

Configure the DingTalk method	i. Click the **DingTalk settings** tab. ii. Configure the parameters. **Note** Procedure for obtaining the parameter values: a. Create a DingTalk group chat. b. In the Group Settings panel, click Bot. In the Robot Management dialog box, create a custom robot. c. Go to the robot details page. Set the Security Settings parameter to **Additional Signature** and copy the generated secret. ▪ **Hook address**: The webhook URL of the custom robot. ▪ **Secret key**: The secret key generated in the **Security Settings** parameter of the custom robot. ▪ **Enabled**: Valid values: **Enabled** and **Do not enable**. iii. Click **Submit**. You can click **Reset** to clear the configurations. You can click **Test DingTalk** to test the DingTalk method.

Configure a message template	**Note** ◦ Message templates cannot be created. You can only modify the existing templates. ◦ By default, 17 message templates are provided, including 1 application O&M message template, 1 product message template, 4 operations message templates, and 11 process approval message templates. i. Click **Edit Template** in the **Operation** column corresponding to a template. In the dialog that appears, modify the template. ▪ **Business Type**: The business type of the template. Valid values:**Product Messages**, **Operations Messages**, **Process Approval Message**, **Application O&M Messages**, and **Announcement**. ▪ **Business Name**: The name of the template. The name must be 2 to 50 characters in length. ▪ **Message header template**: The template for titles of messages. You can use the `${}` syntax to specify system built-in variables that are associated with templates. The title template must be 2 to 50 characters in length. **Note** The supported template literals are included in the template content. The system automatically replaces the template literals with the actual values. ▪ **Message content template**: The template for message content. You can use the `${}` syntax to specify system built-in variables that are associated with templates. The content template must be 0 to 1,000 characters in length. **Note** The supported template literals are included in the template content. The system automatically replaces the template literals with the actual values. ▪ **Message default sending method**: Valid values: **Mail**, **SMS**, and **DingTalk**. ii. Click **OK**.

5.3. Other settings

5.3.1. Service-linked role management

When you create a cloud product that needs to access other cloud products, the system automatically creates a service-linked role for the product. The role belongs to the level-1 organization of the organization in which the cloud product is created. A service-linked role is similar to a RAM policy in that it authorizes cloud services in the level-1 organization to view and manage other resources in the organization.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Service-linked Roles**.

Operation	Procedure
Create a service-linked role	i. On the **Service-linked Roles** page, click **Create Service-linked Role**. ii. Select a service and one or multiple authorized organizations. When you select a service, the permissions of the service-linked role are displayed. i. Click **OK**.
View a service-linked role	i. (Optional) On the **Service-linked Roles** page, set **Role Name**, **Service Name**, or **Organization Name**, and click **Search**. ii. Click **View Details** in the **Actions** column corresponding to a role. iii. On the role details page, click the **Role Details** and **Role Policy** tabs to view the information. ▪ On the Role Details tab, you can view the trust policy. ▪ On the Role Policy tab, you can click **View Details** in the **Actions** column corresponding to the policy to view content of the policy.

5.3.2. Log storage settings

You can configure log storage for operation logs from which security administrators can obtain up-to-date information about the resources and functional modules in the Apsara Uni-manager Management Console. You can also export operation logs to your computer.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Log Storage Settings**.
4. Click **Edit** in the lower part of the page.

   The following log storage modes are available: **Default**, **Log Storage for All**, and **Log Storage for Specified Organization**.

Storage Mode	Description
Default	The operations logs of the users of all organizations are stored in the system Logstore for seven days. **Only All Events is available for the Event Type parameter.**

Log Storage for All Organizations	The operation logs of the users of all organizations are stored in the configured Logstore. You need to purchase the Logstore in advance. ◦ You can set the **Storage Duration** parameter to 30, 60, 90, 180, 360, or 720 days. ◦ You can set the **Event Type** parameter to **All Events**, **Write Event**, or **Read Event**.
Log Storage for Specified Organization	The operation logs of the users of specified organizations are stored in the configured Logstore. You need to purchase the Logstore in advance. ◦ You can set the **Storage Duration** parameter to 30, 60, 90, 180, 360, or 720 days. ◦ You can set the **Event Type** parameter to **All Events**, **Write Event**, or **Read Event**. ◦ **In the Settings section, select an organization and configure the parameters for Log Service, including Log Service Region, Project Name, and Logstore Name.** You can click ⊕ to add more organizations. **Important** The organizations must be within the same level-1 organization.

5. Click **Save**.

Related operations

A security administrator can view the operation logs by choosing **Security > Operation Logs** in the top navigation bar of the console.

6.Operations

6.1. Service orders

6.1.1. Overview

In the Apsara Uni-manager Management Console, an order can be viewed as a flow sheet. An order is generated whenever a resource is created or released, or specifications of a resource are changed. Orders record the progress and status of the current resource management process and the basic information of the resources in detail.

The main features of order management include:

- Allow you to view all orders and details of all orders.
- Allow you to view My Orders and details of My Orders.
- Allow you to view My Approval and details of the orders.
- Allow you to query orders by criteria such as approval status and production status.
- Allow you to view the execution progress and execution result of each order.
- Allow you to approve, withdraw, and cancel orders.
- Allow you to manually produce resources for orders that fail to produce.

Orders are generated when resources of Alibaba Cloud services that support service catalogs are created and released, and specifications of the resources are changed. The following Alibaba Cloud services support service catalogs:

- Elastic Compute Service (ECS)
- Elastic Block Storage (EBS)
- ECS snapshots
- Dedicated Host
- Virtual Private Cloud (VPC)
- NAT Gateway
- IPv6 Gateway
- Elastic IP Address (EIP)
- VPN Gateway
- Server Load Balancer (SLB)
- ApsaraDB RDS for PostgreSQL
- ApsaraDB RDS for SQL Server
- ApsaraDB RDS for MySQL
- ApsaraDB RDS for PolarDB
- Message Queue
- High-availability Virtual IP Address (HAVIP)
- API Gateway
- Key Management Service (KMS)
- Express Connect

- AnalyticDB for MySQL V2.0
- AnalyticDB for MySQL V3.0
- PolarDB-X V1.0
- Graph Database (GDB)
- ApsaraDB for MongoDB
- ApsaraDB for Redis
- Time Series Database (TSDB)
- DataHub
- ApsaraDB for Lindorm (Lindorm)
- Object Storage Service (OSS)
- Virtual Border Router (VBR)

6.1.2. Manage orders

In Apsara Uni-manager Management Console, an order is created whenever an instance is created, scaled, or released. On the Orders page, you can view the information about service orders on the My Orders, My Approval, and All Orders tabs. The orders submitted through service catalogs are included.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Service Orders** > **Orders**.
4. Perform the operations described in the following table.

Operation	Procedure	Description

View your orders	i. On the **Orders** page, click the **My Orders** tab. ii. (Optional) Click **Advanced Filter** and filter the orders by approval and production status. Click **Search**. iii. Click the **ID** of an order to view details of the order. ▪ In the **Order Progress** section, view the progress of the current order. ▪ On the **Details** tab, view the information about the involved instance of the order, such as its ID and name, cloud platform, and instance creation time. ▪ On the **Order Requirement Description** tab, view the information about the order, such as the order type, service type, and instance type.	On the My Orders tab of the Orders page, you can view the information of the orders that are created by the current order. The information includes order IDs, business types, organizations, resource sets, cloud services, approval status, production status, update time, submission time, and applicants. You can click the ID of an order to view the progress, results, and details about the requirements.
Manually start resource production for an order	i. On the **Orders** page, click the **My Orders** tab. ii. (Optional) Click **Advanced Filter** and set **Production Status** to **Production Failed**. Click **Search**. iii. Click **Resource Production** in the **Actions** column corresponding to an order. In the message that appears, click **OK**.	After an order is approved, the order enters the production process. A production process involves four order states: Pending Production, Producing, Produced, and Production Failed. An automatic resource production process may fail due to invalid specifications or insufficient quota, inventory, or permissions. If the production fails, the creator of the order can manually start the resource production or contact the administrator to troubleshoot the issue.

Withdraw or cancel an order	i. On the **Orders** page, click the **My Orders** tab. ii. (Optional) Click **Advanced Filter** and set **Approval Status** to **Approving**. Click **Search**. iii. Click **Withdraw** or **Cancel** in the **Actions** column corresponding to an order. In the message that appears, click **OK**.	Orders in the Approving state can be canceled or withdrawn.
View orders that need your approval	i. On the **Orders** page, click the **My Approval** tab. ii. (Optional) Click **Advanced Filter** and filter the orders by **approval status**, **production status**, or **applicants**. Click **Search**. iii. (Optional) Click the **ID** of the order to view the details of the order.	If the you are part of the approval process of the order, you can view all the orders that need your approval on the My Approval tab.
Approve orders that need your approval	i. On the **Orders** page, click the **My Approval** tab. ii. (Optional) Click **Advanced Filter** and filter the orders by **approval status**, **production status**, or **applicants**. Click **Search**. iii. (Optional) Click the **ID** of the order to view the details of the order. iv. Approve an order or approve multiple orders at a time. ▪ Approve a single order: Click **Approve** in the **Actions** column corresponding to an order. In the dialog box that appears, click **Approve** or **Fail** and specify the reason. ▪ Approve multiple orders at a time: Select multiple orders and click **Approval** in the lower-left corner of the page. In the dialog box that appears, click **Approve** or **Fail** and specify the reason.	You can approve an order or approve multiple orders at a time.

View all orders	i. On the **Orders** page, click the **All Orders** tab. **Note** An operations administrator can view all the orders of the platform. An organization administrator can view only the orders within the current organization. ii. (Optional) Click **Advanced Filter** and filter the orders by **approval status**, **production status**, or **applicants**. Click **Search**. iii. (Optional) Click the **ID** of the order to view the details of the order.	You can view a list of orders and the detailed information of the orders. The information includes order IDs, business types, organizations, resource sets, cloud services, approval status, production status, update time, submission time, and applicants. You can click the ID of an order to view the progress, results, and details about the requirements.

6.2. Service management

6.2.1. Service catalog management

6.2.1.1. Overview

A service catalog allows you to apply for, change, and release resources from a unified cloud service portal and provides supporting enterprise-oriented process and quota control.

Service catalogs are used to control products of specific specifications within an enterprise or organization. For example, enterprises or organizations require employees to go through an approval process before the employees activate large-specification Elastic Compute Service (ECS) instances, or enterprises or organizations want to authorize only users in a specific user group or organization to activate large-specification ECS instances. In this case, enterprises or organizations can configure an approval process, specifications, and users in a service catalog. After the service catalog is published, the employees must go through an approval process when they select a large-specification ECS instances. The users specified in the service catalog cannot select the specified specifications when they purchase the service. This way, enterprises or organizations can limit which users to use resources and perform fine-grained control on specifications.

The main features of service catalogs include:

- Allow you to create custom service catalogs.
- Allow you to publish or unpublish custom service catalogs.
- Allow you to modify configurations of service catalogs.
- Allow you to delete service catalogs that are no longer needed.
- Allow you to view all service catalogs and search for specific service catalogs by criteria.
- Allow you to specify the user groups of service catalogs.
- Allow you to bind processes of application, configuration change, and release to service catalogs.

6.2.1.2. Create a service catalog

Enterprises can use service catalogs to manage resources of specific specifications or to authorize only users in a specific user group or organization to activate large-specification resources. After an operations administrator creates and publishes a service catalog, the users specified in the service catalog cannot select the specified specifications when they create the resources, or must go through approval processes to complete the creation.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Services** > **Service Catalog**.
4. Click **Create Custom Service**. Configure the following parameters.

Step in the wizard		Parameter	Description
Basic Information	Basic Information	Service Name	The name of the service catalog. It must be at least 3 characters in length.
		Service Type	The product on which the service catalog imposes limits.
		Description	Optional. Description of the service catalog. It must be at least 3 characters in length.
	Service Authorization	Authorization Type	Valid values: ◦ **Global**: Allows all users of all organizations to access the service. ◦ **Specified Organization**: Allows all users in the specified organizations to access the service. ◦ **Specified User**: Allows specified users to access the service.
		Authorize Organization/Authorize User	When Authorization Type is set to Specified Organization, you need to configure the Authorize Organization parameter. When Authorization Type is set to Specified User, you need to configure the Authorize User parameter.
	Process Configurations	Resource Application	Optional. The approval process for creating resources. For information about how to create a process, see Create a process.

		Change Configuration	Optional. The approval process for changing the configurations of resources. For information about how to create a process, see Create a process.
Field Selection		Fields to be Specified	Select the fields to be specified. The fields for the selected service type are displayed. The fields are organized in a tree chart. For example, if you select Elastic Network Interfaces as the service type, the Area field is a root node under the Create ENI node. Region and Zone are leaf nodes on which no lower-level nodes depend. **Note** Whenever a resource is created, changed, or released in the Apsara Uni-manager Management Console, information about the organizations and resource sets must be provided. Therefore, the Organization and Resource Set fields are not available.
Cloud Service Configurations		Specification Pool	**Add**: Create specification pools that are used to group multiple organizations and resource sets. **Preview Complete Form**: Obtain a preview of all the configurations items when a cloud resource is created.
		Range	Select an organization and a resource set. **Note** The two parameters determine the organization and resource set to which the specifications are available. Therefore, these parameters are required when you specify the configurations of the cloud product.
		Specifications	The parameters that are selected in the Filed Selection step.

6.2.1.3. Manage service catalogs

You can edit, publish, unpublish, and delete service catalogs. You can also grant permissions on service catalogs.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator.
2. In the top navigation bar, choose **Operations**.
3. In the left-side navigation pane, choose **Services** > **Service Catalog**.
4. (Optional) In the search box, enter a **Service Catalog Name** for fuzzy search, or enter a **Service Catalog ID** for precise search.
5. (Optional) Click **Advanced Filter**. You can search for a service catalog by **Service Catalog Name**, **Service Catalog ID**, **Service Type**, **Publishing Status**, **Authorize Organization**, or **Authorize User**.
6. Perform relevant operations to manage the service catalog.

Operation	Procedure	Description
Edit service catalogs	i. Click **Edit** in the **Actions** column corresponding to the service catalog that you want to manage. ii. Follow the instructions of the wizard to modify the configurations of the service catalog, and then click **Submit**.	You cannot edit a service catalog that is in the Published state.
Publish or unpublish a service catalog	Click **Publish** or **Unpublish** in the **Actions** column corresponding to the service catalog that you want to manage. In the message that appears, click **OK**.	After you create a service catalog, you need to publish it. You can unpublish a service catalog that is in the Published state. A service catalog takes effect only after it is published. After a service catalog takes effect, it applies when users in the configured organization create resources. A service catalog becomes invalid when it is unpublished.

Modify the authorization of a service catalog	i. Click **Authorize** in the **Actions** column corresponding to the service catalog that you want to manage to modify the **Authorization Type**. If you set **Authorization Type** to **Specified Organization** or **Specified User**, you must also select **Authorize Organization** or **Authorize User**. ii. After the modification is complete, click **OK**.	After you modify the authorization, all configurations except Service Type apply to the authorized objects specified by Authorization Type. For example, if the original authorization type configured for the service catalog is Global and you modify the authorization type to Specified Organization, all configurations except Authorization Type that are configured for the current service catalog apply to the specified organizations.
Delete a service catalog	Click **Delete** in the **Actions** column corresponding to the service catalog that you want to manage. In the message that appears, click **OK**.	You cannot delete a service catalog that is in the Published state.

6.2.1.4. FAQ

For which products can I create service catalogs?

You can create service catalogs for the following cloud products:

- Elastic Compute Service (ECS)
- Elastic Block Storage (EBS)
- ECS snapshots
- Dedicated Host
- Virtual Private Cloud (VPC)
- NAT Gateway
- IPv6 Gateway
- Elastic IP Address (EIP)
- VPN Gateway
- Server Load Balancer (SLB)
- ApsaraDB RDS for PostgreSQL
- ApsaraDB RDS for SQL Server
- ApsaraDB RDS for MySQL
- ApsaraDB RDS for PolarDB
- Message Queue
- High-availability Virtual IP Address (HAVIP)
- API Gateway

- Key Management Service (KMS)
- Express Connect
- AnalyticDB for MySQL V2.0
- AnalyticDB for MySQL V3.0
- PolarDB-X V1.0
- Graph Database (GDB)
- ApsaraDB for MongoDB
- ApsaraDB for Redis
- Time Series Database (TSDB)
- DataHub
- ApsaraDB for Lindorm (Lindorm)
- Object Storage Service (OSS)
- Virtual Border Router (VBR)

How many product specifications can I specify when I grant permissions in a service catalog?

Alibaba Cloud does not limit the number of product specifications that you can specify when you grant permissions in a service catalog. You can configure multiple product specification sets as needed.

How do I specify the effective scope of a service catalog?

The effective scope of a service catalog can be users or organizations. You can grant permissions to multiple users by user or by organization. Only authorized users can access the service catalog.

What do I do if I was granted permissions in multiple service catalogs?

If you are granted permissions in multiple service catalogs of the same resource type, a selection window appears to allow you to select the service catalog that you use to create instances when you enable the resource type.

6.2.2. Process management

6.2.2.1. Overview

During a process, an operation of a user is reviewed.

In the console, some operations may carry risks and therefore must be approved by a user or role that has higher privilege. Processes are implemented to facilitate refined management of operations on the platform.

Processes can only be associated with the creation, configuration change, and release of resources of the following products:

- Elastic Compute Service (ECS)
- Virtual Private Cloud (VPC)
- NAT Gateway
- IPv6 Gateway

- Elastic IP Address (EIP)
- VPN Gateway
- Server Load Balancer (SLB)
- ApsaraDB RDS
- Message Queue (MQ)
- Apsara File Storage NAS
- Container Registry
- High availability virtual IP address (HAVIP)
- API Gateway
- Key Management Service (KMS)
- Express Connect
- VMware
- AnalyticDB for MySQL V2.0
- AnalyticDB for MySQL V3.0
- AnalyticDB for MySQL V2.0
- PolarDB for Xscale
- Graph Database
- ApsaraDB for HBase
- ApsaraDB for MongoDB
- Polardb-O
- PolarDB-X 2.0
- AnalyticDB for PostgreSQL
- KVStore for Redis
- Time Series Database (TSDB)
- MaxCompute
- DataHub
- Data Lake Analytics (DLA)
- Lindorm
- Object Storage Service (OSS)
- Tablestore

Main features:

- View the information of processes, such as the name, ID, creator, creation time, and status.
- Search for processes by name, ID, and status.
- Create, modify, publish, withdraw, and delete processes.

6.2.2.2. Create a process

You can create different approval processes for different products and organizations.

Procedure

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

2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Services** > **Process Management**.
4. Click **Create Process**.
   i. Specify a process name. The name can be 2 to 64 characters in length. It can contain letters, digits, underscores (_), and hyphens (-).
   ii. Click the ⊕ icon between **Process Initiator** and **Actions** to add one or multiple **optional approvers** or **required approvers**. Click the **Approver** or **Co-approver** node. In the panel that appears, select the users that act as the approver. Click **OK**.

> **Note**
> - Optional Approver: If the node has multiple approvers, the operation passes the node if one approver approves the operation.
> - Required Approver: If the node has multiple approvers, the operation passes the node only if all approvers approve the operation.
> - You can specify the approvers by specifying users, a role, or a user group. You can specify multiple users, but can specify only a single role or user group.

5. Click **Create**.

6.2.2.3. Manage processes

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Services** > **Process Management**.
4. (Optional) You can use process names to fuzzy search processes or use a process ID to exact search a process. You can also filter processes by status.
5. Perform the operations described in the following table.

Operation	Procedure	Description

Modify a process	i. Click **Edit** in the Actions column corresponding to a process. ii. On the page that appears, modify the process and click Save. **Note** If you click **Save As**, a copy of the process will be created. The copy is in the **To Be Released** state.	You can change the name of the process and add, remove, or modify the approvers.
Publish a process	i. Click **Publish** in the Actions column corresponding to a process. ii. In the message that appears, click **OK**. The process enters the **Released** state.	Processes in the To Be Released or Withdrawn state can be published. After a process is published, it can be associated with service catalogs.
Withdraw a process	i. Click **Withdraw** in the Actions column corresponding to a process. ii. In the message that appears, click **OK**. The process enters the **Withdrawn** state.	Processes that are in the Published state and are not associated with any service catalog can be withdrawn. Before you withdraw a process, make sure that no service catalogs are associated with it.
Delete a process	i. Click **Delete** in the Actions column corresponding to a process. ii. In the message that appears, click **OK**.	Processes that are not associated with any service catalog can be deleted. Before you delete a process, make sure that no service catalogs are associated with it.

6.3. Quota management

The Apsara Uni-manager Management Console provides quota management capabilities. The amount of resources on the cloud platform is limited. Therefore, it is important to control the amount of resources available for each organization based on the scope of its responsibilities. This ensures that each organization is allocated with the right amount of resources.

Background information

The platform administrator can configure the quotas for the cloud products in the resource pool. The logical inventory of the cloud platform is equivalent to the quota of the root organization in organization management.

Cloud services that support quota management.

Category	Service
Computing	Elastic Compute Service (ECS) Dedicated Host (DDH) Resource Orchestration Service (ROS) Auto Scaling Container Service for Kubernetes (ACK) Bare-metal Management Service (BMS)
Database	ApsaraDB RDS for MySQL ApsaraDB RDS for PolarDB ApsaraDB RDS for SQL Server ApsaraDB RDS for PostgreSQL PolarDB KVStore for Redis ApsaraDB for MongoDB PolarDB-X 1.0 AnalyticDB for PostgreSQL AnalyticDB for MySQL V3.0 AnalyticDB for MySQL V2.0
Storage	Object Storage Service (OSS) Apsara File Storage NAS Tablestore
Middleware	Message Queue (MQ)
Big data	MaxCompute Machine Learning Platform for AI (PAI) Hologres
Networking	Virtual Private Cloud (VPC) Server Load Balancer (SLB) Elastic IP Address (EIP) High availability virtual IP address (HAVIP) Express Connect Cloud Enterprise Network (CEN)
Application Services	API Gateway

6.3.1. Quota parameters

The organization administrator and users of an organization can create resources within the quota of the cloud products. When the quota of a product in an organization is exhausted, resources of the product can no longer be created. You must increase the quota of the product in the organization before you can create resources again. If no quotas are set, an unlimited amount of resources can be created.

> **Note**
>
> Quotas for resources such as disks, memory, and storage capacity are measured in **GB**, and the quantities are integers. Decimals are rounded to their nearest integer. For example, 0.8 GB is rounded up to 1 GB, and 0.2 GB is rounded down to 0 GB.

Resource Orchestration Service (ROS)

Parameter	Description
Stacks(Units)	The total number of stacks that can be created.

Elastic Compute Service (ECS)

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for ECS.
GPU Quota(GPUs)	The total number of GPUs that can be configured for ECS.
Memory Quota(G)	The total memory size that can be configured for ECS.
SSD(GB)	The total SSD capacity that can be configured for ECS.
Ultra Disk(GB)	The total ultra disk capacity that can be configured for ECS.
Premium Performance Disk(GB)	The total premium performance disk capacity that can be configured for ECS.
Standard Performance Disk(GB)	The total standard perform disk capacity that can be configured for ECS.

Auto Scaling

Parameter	Description

Scaling Group(Units)	The total number of scaling groups that can be created.

Dedicated Host

Parameter	Description
DDH Quota(Units)	The total number of dedicated hosts that can be created.

Bare-metal Management Service

Parameter	Description
Bare Metal Instances(Unit)	The total number of bare metal instances that can be created.
In-band Instance(Sets)	The total number of in-band managed instances that can be created.
Out-of-band Instance(Sets)	The total number of out-of-band managed instances that can be created.

Container Service for Kubernetes (ACK)

Parameter	Description
Cluster(Units)	The total number of clusters that can be created.

Apsara File Storage NAS

Parameter	Description
Disk Quota(GB)	The total storage size that can be configured for Apsara File Storage NAS.

Object Storage Service (OSS)

Parameter	Description
OSS Quota(GB)	The maximum capacity that can be allocated to OSS.

Tablestore

Parameter	Description
Instances(Units)	The maximum number of Tablestore instances that can be created.

Elastic IP Address (EIP)

Parameter	Description
EIP Quota(Units)	The maximum number of EIPs that can be configured.

Server Load Balancer (SLB)

Parameter	Description
Internal IP Quota(Units)	The maximum number of internal IP addresses that can be configured for SLB.
Public IP Address Quota(Units)	The maximum number of public IP addresses that can be configured for SLB.

High availability Virtual IP Address (HAVIP)

Parameter	Description
HAVIP Quota(Units)	The maximum number of HAVIPs that can be created.

Virtual Private Cloud (VPC)

Parameter	Description
VPC Quota(Units)	The maximum number of VPCs that can be created.

Express Connect

Parameter	Description
Connections(Units)	The total number of peering connections and Express Connect circuits that can be created.

Cloud Enterprise Network (CEN)

Parameter	Description
CEN Instance(Units)	The total number of CEN instances that can be created.

ApsaraDB RDS for MySQL

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for ApsaraDB RDS for MySQL.
Memory Quota(G)	The total memory size that can be configured for ApsaraDB RDS for MySQL.
Disk Quota(GB)	The total storage capacity that can be configured for ApsaraDB RDS for MySQL.

ApsaraDB RDS for PolarDB

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for ApsaraDB RDS for PolarDB.
Memory Quota(G)	The total memory size that can be configured for ApsaraDB RDS for PolarDB.
Disk Quota(GB)	The total storage capacity that can be configured for ApsaraDB RDS for PolarDB.

ApsaraDB RDS for PostgreSQL

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for ApsaraDB RDS for PostgreSQL.
Memory Quota(G)	The total memory size that can be configured for ApsaraDB RDS for PostgreSQL.
Disk Quota(GB)	The total storage capacity that can be configured for ApsaraDB RDS for PostgreSQL.

ApsaraDB RDS for SQL Server

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for ApsaraDB RDS for SQL Server.
Memory Quota(G)	The total memory size that can be configured for ApsaraDB RDS for SQL Server.
Disk Quota(GB)	The total storage capacity that can be configured for ApsaraDB RDS for SQL Server.

PolarDB

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for PolarDB.
Memory Quota(G)	The total memory size that can be configured for PolarDB.
Disk Quota(GB)	The total storage capacity that can be configured for PolarDB.

KVStore for Redis

Parameter	Description
Memory Quota(G)	The total memory size that can be configured for KVStore for Redis.

PolarDB-X 1.0

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for PolarDB-X 1.0.

ApsaraDB for MongoDB

Parameter	Description

CPU Quota(Cores)	The total number of CPU cores that can be configured for ApsaraDB for MongoDB.
Memory Quota(G)	The total memory size that can be configured for ApsaraDB for MongoDB.
Disk Quota(GB)	The total storage capacity that can be configured for ApsaraDB for MongoDB.

AnalyticDB for PostgreSQL

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for AnalyticDB for PostgreSQL.
Memory Quota(G)	The total memory size that can be configured for AnalyticDB for PostgreSQL.
Disk Quota(GB)	The total storage capacity that can be configured for AnalyticDB for PostgreSQL.

AnalyticDB for MySQL V2.0

Parameter	Description
CPU Quota(Cores)	The total number of CPU cores that can be configured for AnalyticDB for MySQL V2.0.
Memory Quota(G)	The total memory size that can be configured for AnalyticDB for MySQL V2.0.

AnalyticDB for MySQL V3.0

Parameter	Description
Disk Quota(GB)	The total storage capacity that can be configured for AnalyticDB for MySQL V3.0.

Machine Learning Platform for AI (PAI)

Parameter	Description

Compute Node(Units)	The total number of compute nodes that can be configured for PAI.

Hologres

Parameter	Description
CU Quota(Units)	The total number of capacity units (CUs) that can be configured for Hologres.

MaxCompute

Parameter	Description
CU Quota(Units)	The total number of CUs that can be configured for MaxCompute.
Disk Quota(GB)	The total storage capacity that can be configured for MaxCompute.

Message Queue (MQ)

Parameter	Description
Message Queue Instances(Units)	The total number of MQ instances that can be created.

API Gateway

Parameter	Description
API Groups(Units)	The total number of API groups that can be created.
Extensions(Units)	The total number of extensions that can be created.
Applications(Units)	The total number of applications that can be created.

6.3.2. Manage quotas

The amount of resources on the cloud platform is limited. Therefore, it is important to limit the amount of resources available for each organization based on the scope of its responsibilities. You can create, modify, and clear quotas.

Background information

Quota management in the Apsara Uni-manager Management Console involves the following concepts: physical inventory, quota, allocated resources, used resources, and remaining resources.

- Physical inventory: the physical inventory of a specific type of resource that is available for an organization or resource set. This type of resource can be used to create resource sets for the current organizations, or their subordinate organizations.
- Quota: the total amount of resources available for a target cloud service.
- Allocated resources: the amount of a specific type of resource that an organization allocates to its subordinate organizations. The resources each organization receives are allocated by its parent organization.
- Used resources: the amount of resources that have been used in an organization or a resource set, including the resources used by subordinate organizations and resource sets.
- Remaining resources: the amount of available resources in an organization or resource set, which is physical inventory minus used resources. Resources can be used to create instances only if the amount of required resources is less than the amount of remaining resources.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator or an organization administrator of the current organization.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Quotas** > **Quota Alerts**.
4. In the navigation tree on the **Quotas** page, choose the target organization.
5. Perform the operations described in the following table.

Operation	Procedure
Create or modify a quota	**Note** ◦ When you create a quota for an organization that has a parent organization, the available quota equals the total amount of the parent organization's quota minus the quotas of the other organizations at the same level as the target organization. This rule does not apply when the target organization is a level-1 organization. ◦ When you create a quota for a product of the target organization, make sure the quota for the product has been created in the parent organization. i. On the right side of the page, click the target cloud service and configure the quotas in the panel that appears. For information about the quota parameters of each cloud service, see Quota parameters. ii. Click **Save**.

Clear a quota	**Note** Before you clear a quota for an organization, make sure that no subordinate organizations or resource sets have quotas. i. Click the target product in the product list. In the panel that appears, click **Clear quota**. ii. In the message that appears, click **OK**.

6.3.3. Quota alert management

6.3.3.1. Create a quota alert rule

You can create alert rules for quotas. When the usage of a quota reaches the preset threshold, an alert is generated based on relevant alert rules and an alert notification is sent. In this case, you can submit an application to increase the quota in advance.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Quotas** > **Quota Alerts**.
4. On the **Quota Alerts** page, click **Create Alert**.
5. In the **Crate Quota Alert** dialog box, configure the parameters.

Section	Parameter	Description
Alert Object Information	Organization and Resource Set	The organization or resource set that contains the quota for which you want to create the alert rule.
	Region	The region of the quota.
	Cloud Service Category	The cloud product for which you want to create the alert rule. **Note** Only the products for which quotas have been configured in the organization are available.
	Alert Name	The name of the quota alert rule. The name can contain letters, digits, hyphens (-), and underscores (_).

Alert Information	Alert criteria	◦ **Used**: An alert is generated when the quota used by the organization and its subordinate organizations exceeds the configured threshold. ◦ **Remaining Available Quota**: An alert is generated when the remaining available quota of the organization and its subordinate organizations falls below the configured threshold.
	Alert Threshold	◦ **Resource Type**: The value is set to the cloud product specified in the Cloud Service Category parameter. ◦ **Remaining Available Quota / Quota** or **Used / Quota**: The value is set to the remaining available quota amount or the remaining used amount. ◦ **Threshold**: You can specify a percentage value or a number. ◦ **Alert Threshold**: The threshold for the alert rule. When the used amount reaches the threshold or the remaining available amount falls below the threshold, alerts are generated.
	Alert Callback	Optional. Quota Center sends an HTTP POST request that contains the alert information to the specified URL.

6.3.3.2. Manage quota alerts

You can modify or delete existing quota alert rules. Generated alerts are recorded, so that operations personnel can analyze the alert history to learn about the resource usage. This way, they can optimize the quota allocation and quota alert rules.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Quotas** > **Quota Alerts**.
4. (Optional) On the **Quota Alerts** page, click **Advanced Filter** and enter the name of the alert that you want to modify. You can also filter alerts by specifying the **Organization and Resource Set**, **Region**, and **Cloud Service Category** conditions.
5. Perform the operations described in the following table.

Operation	Procedure
View a quota alert rule	Click the name of an alert rule to view its configuration information.

Modify a quota alert rule	i. Click **Edit** in the **Actions** column corresponding to an alert rule. ii. In the **Edit Quota Alert** dialog box, modify the settings of the **Alert Name**, **Alert criteria**, **Alert Threshold**, and **Alert Callback** parameters. iii. Click **OK**.
View alert history	i. Click **View alert history** in the **Actions** column corresponding to an alert rule. ii. On the **Quota alert history** page, view the alert history. **Note** The quota alert history is updated once a day.
Delete a quota alert rule	i. Click **Delete** in the **Actions** column corresponding to a quota alert rule. ii. In the message that appears, click **OK**.

6.4. Statistical analysis

6.4.1. View reports

The Apsara Uni-manager Management Console provides reports of cloud resources and resource quotas, so that you can view the statistics of cloud resources and quotas in real time. You can also export the reports to your computers.

Background information

Services that support resource reporting

Category	Service	Resource type

Computing	Elastic Compute Service (ECS)	• Automatic snapshot policy • ECS instance • Dedicated host • Deployment set • ECS disk • File sending • ECS recycle bin • ECS command • ECS command execution result • ECS snapshot • Asynchronous task • ECS image • Key pair • Elastic network interface (ENI) • Security group
	Bare-metal Management Service (BMS)	Bare Metal Instance
	Container Registry Standard Edition	• Namespace • Image repository
	Container Registry Advanced Edition	• Namespace (Apsara Stack Advanced Edition) • Image Repository (Apsara Stack Advanced Edition)
	Container Service for Kubernetes (ACK)	ACK cluster
	E-HPC	E-HPC cluster
	Auto Scaling	• Alert task • Scaling group • Scheduled task
	Resource Orchestration Service (ROS)	• Stack • Template
	Cloud Storage Gateway (CSG)	CSG

Storage	Apsara File Storage NAS	• Permission group • File system • Unified namespace
	Object Storage Service (OSS)	• OSS bucket • Single Tunnel
	Tablestore	Tablestore instance
Application services	API Gateway	• API Gateway • API group • API application • API plug-in
Networking	Elastic IP Address (EIP)	EIP
	Express Connect	• Express Connect circuit • Router interface • Virtual border router
	High availability virtual IP address (HAVIP)	HAVIP
	IPv6 Gateway	IPv6 gateway
	NAT Gateway	NAT gateway
	Server Load Balancer (SLB)	• CA certificate • Server Certificate • SLB instance
	DNS within Cloud	• DNS workspace • Private DNS Line • Cross-zone forwarding • VPC forwarding

	Virtual Private Cloud (VPC)	• Network ACL • Router • VPC • vSwitch
	VPN gateway	• Customer gateway • SSL client • SSL server • VPN attachment • VPN gateway
Database	AnalyticDB for MySQL V3.0	AnalyticDB for MySQL cluster
	AnalyticDB for MySQL V2.0	AnalyticDB for MySQL cluster
	ApsaraDB for MongoDB	ApsaraDB for MongoDB Instance
	PolarDB-X 1.0	PolarDB-X instance
	Data Transmission Service (DTS)	• Migration task • Subscription task • Synchronization task
	Graph Database (GDB)	GDB
	AnalyticDB for PostgreSQL	AnalyticDB for PostgreSQL instance
	ApsaraDB for HBase	ApsaraDB an HBase instance
	Time Series Database (TSDB)	TSDB instance
	Lindorm	Lindorm instance
	PolarDB	PolarDB cluster

	PolarDB-X 2.0	PolarDB-X instance
	KVStore for Redis	KVStore for Redis Instance
	ApsaraDB RDS	• ApsaraDB RDS instance • RDS recycle bin
Middleware	Enterprise Distributed Application Service (EDAS)	• EDAS • Cluster list
	Message Queue (MQ)	MQ
Monitoring and O&M	CloudMonitor	Alert rule
	Operation Orchestration Service (OOS)	OOS template
Big data	DataHub	DataHub project
	Elastic Algorithm Service (EAS)	EAS
	Elasticsearch on k8s	Elasticsearch on Kubernetes instance
	E-MapReduce	E-MapReduce cluster
	Realtime Compute for Apache Flink	Project
	MaxCompute	• CU • ODPS engine
Security	Key Management Service (KMS)	Key pair

Services that support quota reporting

Category	Service	Resource type
Computing	ECS	ECS instance
	DDH	DDH
Networking	SLB	SLB instance
	VPC	VPC
	EIP	EIP
Storage	Apsara File Storage NAS	File system
	OSS	OSS bucket
Database	AnalyticDB for PostgreSQL	AnalyticDB for PostgreSQL instance
	ApsaraDB RDS	ApsaraDB RDS for MySQL ApsaraDB RDS for SQL Server ApsaraDB RDS for PostgreSQL PolarDB
	PolarDB-X 1.0	PolarDB-X instance
	AnalyticDB for MySQL V3.0	AnalyticDB for MySQL cluster
Big data	MaxCompute	ODPS engine

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**. In the left-side navigation pane, choose **Statistical Analysis** > **Reports**.
3. Click the **Resource Report** or **Quota Report** tab, and select a resource type from the **Cloud Service** drop-down list. You can filter the reports by specifying the **Organization and Resource Set** and **Region** conditions.
   - A **resource report** contains details of a specific resource type in the specified scope.
   - A **quota report** contains details of the quota of a specific resource type in the specified scope.

4. Select the records in the list and click **Export Selected Reports** in the lower part of the page. The report will be downloaded to your computer as an .XLSX file.

   > **Note**
   > You can customize the fields in the report.

   To export information about multiple resource types, click **Create Report Export Task**. In the dialog box that appears, configure the parameters and create the export task. After the task is executed, choose **Statistical Analysis** > **Download Center** in the left-side navigation pane and download the report on the page that appears.

6.4.2. Report pushing

The Apsara Uni-manager Management Console can automatically send reports to your email address based on the time points, schedules, and desired content configured in report pushing rules. This helps you keep track of the resource usage and facilitates monitoring and alerting.

Background information

A user can configure up to ten report pushing rules for each of their roles. Based on the rules, the system generates reports about enterprise operations and cloud resource usage at the scheduled time and sends them to the email address associated with the user. Only statistics that are within the permissions of the role are included in the reports.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**. In the left-side navigation pane, choose **Statistical Analysis** > **Report Push**.
3. On the Report Pushing page, click **Create Rule** and configure the following parameters.

Step	Parameter	Description
Basic Settings	Rule Name	The name of the report push rule. The name can be 2 to 64 characters in length and can contain letters, digits, hyphens (-), and underscores (_).
	Rule Description	The description of the report push rule.
	Report Name	The name of the generated reports, which will be attached in the emails sent to your email address.
	Push At	The interval at which reports are generated. Valid values: Every Day and Every Week.

	Scheduled Time	The point of time at which reports are generated. **Note** When Push At is set to Every Week, you also need to specify the days of the week on which reports are generated in the **Days** parameter.
Content Settings		Specify the content of the reports. ◦ Enterprise statistics: includes the numbers of and changes in users, roles, organizations, and resource sets. ◦ Operational statistics: includes the usage and alerting of quotas. ◦ Application statistics: the changes in the numbers of resources and alerts. **Note** The reports will only include statistics of the resources that are available to the current role.
Push Preview		A demo report generated based on the configurations for preview.

4. Preview the report and click **Create**. The system will send the reports based on the specified schedule to your email address.
5. You can manage the created rules on the Report Pushing page.
   - Click **Details** in the Actions column corresponding to a rule. On the page that appears, you can view the basic information and push history of the rule. You can also modify the content of the rule.
   - Click **Modify Attributes** in the Actions column corresponding to a rule. In the dialog box that appears, you can modify the name, description, report name, pushing schedules and time points specified in the rule. Content of the reports cannot be modified here. It can be modified only on the rule details page.
   - Click **Enable** or **Disable** in the Actions column corresponding to a rule. The system will not generate reports based on a rule if it has been disabled.

     Note
     However, disabled rules still count toward the ten-rule limit.
   - Move the pointer over the More icon in the Actions column corresponding to a rule and click **Push History**. In the panel that appears, you can view the push tasks that have been created, including their creation time, status, and related exceptions. You can also click Download to download the reports.
   - Move the pointer over the More icon in the Actions column corresponding to a rule and click **Delete**. You can delete the pushing rules that are no longer needed.

6.4.3. Download Center

In Download Center, you can download existing resource reports and quota reports. You can also create download tasks to export specific reports.

Background information

- Exported files are in the **.xlsx** format. You can find the downloaded files from the downloading path of the browser.
- The reports you can download include **metering reports**, **resource reports**, **quota reports**, **user reports**, and **detailed billing reports**.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Statistical Analysis** > **Download Center**.
4. (Optional) Filter the reports by specifying the **Report Name**, **Report Type**, **Cloud Service**, **Status**, and **Start and End Time** conditions.
5. Perform the operations described in the following table.

Operation	Procedure
Download a report	Click **Download Report** in the **Download operation** column corresponding to a report. **Note** You can download reports in the **Completed** state. If the report is in the **Failed** or **Canceled** state, you can click **Retry** to retry the task.
Delete a report	◦ Delete a single report. a. Click **Delete** in the **Download operation** column corresponding to a report. b. In the **Delete Task** message, click **OK**. ◦ Delete multiple reports at a time. a. Select the target reports. b. Click **Delete** below the list.
	i. Click **Create Download Task**.

Create a download task	ii. In the dialog box that appears, configure the parameters. ▪ **Report Name**: The name of the report. ▪ **Report Type**: Valid values include **Metering Report**, **Resource Report**, **Quota Report**, **User Report**, and **Detailed Billing Report**. **Note** When User Report is selected, the parameters below Report Type are not available or required. The user report contains information about all system users of the platform. ▪ **Service**: The target product. **Note** You can click **Select All** to select all the listed products. ▪ **Start and End Time**: The time range of the report. **Note** This parameter is available only when Report Type is set to Metering Report or Detailed Billing Report. ▪ **Organization and Resource Set**: The target organizations and resource sets. **Note** If you want to select all organizations and resource sets, select the **root** organization. ▪ **Region**: The target region. ▪ **Account Period**: The target accounting period. **Note** This parameter is available only when Report Type is set to Detailed Billing Report. iii. Click **OK**.

6.4.4. Carbon emissions report

The Apsara Uni-manager Management Console provides carbon emissions reports to help you improve resource efficiency of the servers in your data centers and minimize environmental impact while ensuring the security and stability of your cloud platform.

On the Apsara Uni-manager Management Console, you can view the amounts and trends of carbon emissions. You can also filter the statistics by cloud services and instances

> **Note**
>
> Only users who are assigned the role of operations administrator, organization administrator, or resource set administrator can view the carbon emissions reports. Custom roles cannot be granted permissions on carbon emissions reports. An operations administrator has carbon emissions statistics permissions on all organizations and resource sets; an organization administrator has permissions on their organization and subordinate organizations; a resource set administrator has permissions on their resource sets.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**. In the left-side navigation pane, choose **Statistical Analysis** > **Carbon Emissions Report**. On the page that appears, you can view the carbon emissions reports.
   - On the **Carbon Emissions Overview** tab, you can have an overview of the carbon emissions of the cloud platform, including the total emissions, emissions over the past 30 days, trends of emissions, and rankings of the services and instances by emissions amounts.
   - On the **Cloud Product Statistics** tab, you can filter statistics about carbon emissions by time periods, organizations, resource sets, regions, and cloud services.
   - On the **Instance Statistics** tab, you can filter statistics about carbon emissions by time periods, organizations, resource sets, regions, cloud services, and instance IDs. Then, you can click **Export Data** to download the statistics to your computer as an .XLSX file.

6.5. Usage statistics

6.5.1. View the usage statistics of cloud resources

The Apsara Uni-manager Management Console shows the number of resource instances that run in the Apsara Stack environment by time, organization, resource set, and region. You can also export statistical reports from the console.

Background information

You can view the usage statistics of the following cloud products or cloud resources:

Category	Service

Computing	• Container Service for Kubernetes (ACK)
	• Bare-metal Management Service (BMS)
	• Container Registry (ACR)
	  ◦ Image repository
	  ◦ Image Repository (Apsara Stack Advanced Edition)
	  ◦ Namespace
	  ◦ Namespace (Apsara Stack Advanced Edition)
	• Elastic Compute Service (ECS)
	  ◦ ECS
	  ◦ Elastic Block Storage (EBS)
	  ◦ Dedicated Host (DDH)
	  ◦ ECS snapshot
	• Intelligent Computing LINGJUN Node
	• Elastic High Performance Computing (E-HPC)
	• Auto Scaling

Networking	• Server Load Balancer (SLB)
	• Virtual Private Cloud (VPC)
	  ◦ VPC
	  ◦ Elastic IP Address (EIP)
	  ◦ High availability virtual IP address (HAVIP)
	  ◦ IPv6 Gateway
	  ◦ NAT Gateway
	• Express Connect
	• VPN Gateway
	• Apsara Stack DNS
	  ◦ DNS workspace

Database	• ApsaraDB RDS
	  ◦ ApsaraDB RDS for MySQL
	  ◦ PolarDB
	  ◦ ApsaraDB RDS for PostgreSQL
	  ◦ ApsaraDB RDS for SQL Server
	• PolarDB
	• AnalyticDB for PostgreSQL
	• ApsaraDB for Redis
	• ApsaraDB for MongoDB
	• Data Transmission Service (DTS)
	• Data Management (DMS)
	• PolarDB-X 2.0
	• PolarDB-X 1.0
	• Graph Database (GDB)
	• ApsaraDB for OceanBase

Storage	• Apsara File Storage NAS • Object Storage Service (OSS) • Tablestore
Middleware	• Cloud Service Bus (CSB) • Message Queue (MQ)

Big data	• Realtime Compute
	• Realtime Compute for Apache Flink
	• DataHub
	• DataV
	• DataWorks
	• Elasticsearch on k8s
	• Hologres
	• MaxCompute
	  ◦ ODPS engine
	• Machine Learning Platform for AI (PAI)
	  ◦ Cloud-native Deep Learning Training
	  ◦ Interactive Modeling (DSW)
	  ◦ Elastic Algorithm Service (EAS)

Application services	• API Gateway • Simple Log Service (SLS)

Security	• Server Guard
	• Security Center
	  ◦ Threat Detection Service
	  ◦ Security Operations Center (SOC)
	  ◦ Cloud Security Scanner
	• Apsara Stack Security - Bastionhost
	  ◦ Bastionhost (Early Version)
	  ◦ Cloud-native Bastionhost
	• DataBase Audit
	• Apsara Stack Security - Data Encryption Service
	• Key Management Service (KMS)
	• Data Security Center (DSC)
	  ◦ Crypto Service
	• Web Application Firewall (WAF)
	• DDoS Traffic Scrubbing
	• Network Detection and Response
	  ◦ Traffic security monitoring
	  ◦ Network detection and response
	• Cloud Firewall
	  ◦ Customer Cloud Firewall
	  ◦ Internet Cloud Firewall

Others	Elastic Desktop Service (EDS)

Procedure

The following procedure uses ECS as an example to describe how to view usage statistics. You can view the usage statistics of other cloud resources in a similar manner.

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**. In the left-side navigation pane, click **Metering Management**.
3. Select **Elastic Compute Service** from the **Cloud Service** drop-down list, and then click **Search**.

   You can configure the following options to filter the statistics: **Organization**, **Resource Set**, **Region**, **Time**, **Instance ID**. You can also click the icon to export the displayed statistics to your computer as an XLS file.

> **Note**
>
> In the console, you can view or export up to 1,000 statistical records to an XLS file. If you want more statistical data, you can call the MeteringQuery operation.

6.6. Bills

6.6.1. View cloud service bills

You can view bills by cloud service type.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Billings** > **Cloud Service Bills**.
4. View the bill overview.
   i. Click the **Bill overview** tab.
   ii. Select the year and the month for which you want to view bills from the **Account period** date picker. You can view different data in the following sections:
   - **Cloud products bill summary**: View the number of cloud services that has billing enabled, the number of cloud services that incur charges, the total fees generated within the selected month, and the total fees generated within the selected year.
   - **Consumption trends of cloud products.Last six months**: View the consumption trend of each cloud service within the last six months. Click **Top 5** or **Top 10** in the upper-right corner to show the top five or top ten cloud services in terms of consumption within each month. Click the or icon to show the consumption trends of cloud services by using a column chart or a line chart.
   - **Top 10Consumption distribution of cloud products**: View the generated fees and their proportions of the top ten cloud services within the selected month.
   - **Consumption details of all cloud products**: View the fees of all cloud services within the selected month. Select Sort by fee from high to low or Sort by fee from low to high in the upper-right corner to sort cloud services by fees.
5. View bill statistics.
   i. Click the **Bill statistics** tab.
   ii. Select the year and the month for which you want to view bills from the date picker.
   iii. (Optional) Select the cloud service and the region that you want to view.
   iv. Click **Search**.
   You can view the billing information that meets the search conditions. The billing information includes the billing period, cloud service, region, fees, and billing cycle.
   v. (Optional) Click the icon to modify fees as an operations administrator.
   vi. (Optional) Click **View details** in the Operation column to view bill details.
   vii. Save the queried data to your computer.
   - Click **Export all bills** to save all queried data to your computer as an XLS file.
   - Select the bills that you want to export and click **Export Selected bills** to save the selected data to your computer as an XLS file.
6. View bill details.
   i. Click the **Bill details** tab.

ii. Select the year and the month for which you want to view bills from the **Account period** date picker.

iii. (Optional) Set Cloud products, Area, and Instance ID.

iv. Click **Search**.

You can view the billing information that meets the search conditions. The billing information includes **Account period**, **Cloud products**, **Tissue**, **Resource set**, **Instance ID**, **Area**, **Billing cycle**, **Free amount**, and **Amount payable**.

v. Save the queried data to your computer.

- Click **Export all bills** to save all queried data to your computer as an XLS file.
- Select the bills that you want to export and click **Export Selected bills** to save the selected data to your computer.

6.6.2. View organization and resource set bills

You can view bills by organization and resource set.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Billings** > **Organizations and Resource Set Bills**.
4. View the bill overview.
   i. Click the **Bill overview** tab.
   ii. Select the year and the month for which you want to view bills from the **Account period** date picker. You can view different data in the following sections:
   - **root(and all subordinate organizations) Bill overview**: View the number of organizations, the number of organizations that incur charges, the total fees generated within the selected month, and the total fees generated within the selected year.
   - **Organizational consumption trend.Last 6 months**: View the consumption trend of each organization within the last six months. Click **Top 5** or **Top 10** in the upper-right corner to show the top five or top ten organizations in terms of consumption within each month. Click the or icon to show the consumption trends of organizations by using a column chart or a line chart.
   - **TOP 10 Resource set consumption distribution**: View the generated fees and their proportions of the top ten organizations within the selected month.
   - **Organization consumption details**: View the fees of all organizations within the selected month. Select Sort by fee from high to low or Sort by fee from low to high in the upper-right corner to sort organizations by fees.
5. View bill statistics.
   i. Click the **Bill statistics** tab.
   ii. Select the year and the month for which you want to view bills from the date picker.
   iii. (Optional) Select the organization and the resource set that you want to view.

iv. Click **Search**.

You can view the billing information that meets the search conditions. The billing information includes **Account period**, **Organization name**, **Resource set name**, **Amount**, and **Billing cycle**.

v. (Optional) Click the icon to modify fees as an operations administrator.

vi. (Optional) Click **View details** in the Operation column to view bill details.

vii. Save the queried data to your computer.

- Click **Export all bills** to save all queried data to your computer as an XLS file.
- Select the bills that you want to export and click **Export Selected bills** to save the selected data to your computer as an XLS file.

6. View bill details.

i. Click the **Bill details** tab.

ii. Select the year and the month for which you want to view bills from the **Account period** date picker.

iii. (Optional) Set **Organization and resource set**, **Cloud products**, **Area**, and Instance ID.

iv. Click **Search**.

You can view the billing information that meets the search conditions. The billing information includes **Account period**, **Cloud products**, **Tissue**, **Resource set**, **Instance ID**, **Area**, **Billing cycle**, **Free amount**, and **Amount payable**.

v. Save the queried data to your computer.

- Click **Export all bills** to save all queried data to your computer as an XLS file.
- Select the bills that you want to export and click **Export Selected bills** to save the selected data to your computer as an XLS file.

6.7. Tag management

You can use tags to classify and manage resources from different dimensions.

6.7.1. Query tags

Each tag consists of a tag key and a tag value. On the **Tag management** page, you can view all tag keys and values, as well as the information about the resources to which the tags are added.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, click **Tags**.
4. In the left-side navigation pane of the **Tag management** page, click the name of the tag whose values that you want to view. You can also view the number of cloud services to which the tag is added.

5. (Optional) In the left-side tag key list, click . In the **Filter Conditions** dialog box,

   configure the following filters based on your business requirements: **The organization of the resource**, **Resource set**, **The region of the resource**, and **Keyword**. You must configure **The organization of the resource**.

6.7.2. Create a tag

This topic describes how to create a tag and add the tag to resources.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, click **Tags**.
4. On the **Tag management** page, click **Create tags**.
   i. In the **Create tags and bind resources** dialog box, select an existing tag key and tag value or enter a new tag key or tag value. Click **Next**.

   > **Note**
   > You can create up to 20 tags and add the tags to resources at a time. A tag key must be unique.

   ii. In the **Create tags and bind resources** dialog box, select the organization and resource set to which the resource belongs, select the region where the resource resides, select the cloud service to which the resource belongs, and then click **Next**.

   > **Note**
   > You can select up to 50 cloud resources at a time.

   iii. Confirm the information about the tag and resource, and click **OK**.

6.8. Cloud resource optimization

6.8.1. Resource utilization rates

Calculation based on all resources

By instance

Except for Object Storage Service (OSS), the utilization rate of a product is calculated by using the following formula: **Utilization rate of a product = 5% × Average utilization rate + 5% × 75th percentile value + 10% × 95th percentile value + 10% × 99th percentile value + 40% × 99.9th percentile value + 30% × 100th percentile value**.

- The CPU, memory, or storage utilization rate of Elastic Compute Service (ECS) is calculated by using the following formula: CPU/Memory/Storage utilization rate of ECS = 5% × Average CPU/memory/storage utilization rate + 5% × 75th percentile value + 10% × 95th percentile value +10% × 99th percentile value + 40% × 99.9th percentile value + 30% × 100th percentile value.

- The CPU, memory, or storage utilization rate of ApsaraDB RDS is calculated by using the following formula: CPU/Memory/Storage utilization rate of ApsaraDB RDS = 5% × Average CPU/memory/storage utilization rate + 5% × 75th percentile value + 10% × 95th percentile value + 10% × 99th percentile value + 40% × 99.9th percentile value + 30% × 100th percentile value.
- The storage utilization rate of OSS is calculated by using the following formula: Storage utilization rate of OSS = Storage usage/Purchased storage capacity.

By organization

The utilization rate of an organization is calculated by using the following formula: **Utilization rate of an organization = (Utilization rate of Instance 1 × Capacity of Instance 1 +... + Utilization rate of Instance N × Capacity of Instance N)/(Capacity of Instance 1 +... + Capacity of Instance N).** The utilization rate can be calculated by week, month, or year. You can also specify a custom calculation cycle.

- The CPU, memory, or storage utilization rate of ECS is calculated by using the following formula: CPU/Memory/Storage utilization rate of ECS = (CPU/Memory/Storage utilization rate of Instance 1 × CPU/Memory/Storage capacity of Instance 1 + CPU/Memory/Storage utilization rate of Instance 2 × CPU/Memory/Storage capacity of Instance 2 +...+ CPU/Memory/Storage utilization rate of Instance N × CPU/Memory/Storage capacity of Instance N)/(CPU/Memory/Storage capacity of Instance 1 + CPU/Memory/Storage capacity of Instance 2 +...+ CPU/Memory/Storage capacity of Instance N).
- The CPU, memory, or storage utilization rate of ApsaraDB RDS is calculated by using the following formula: CPU/Memory/Storage utilization rate of ApsaraDB RDS = (CPU/Memory/Storage utilization rate of Instance 1 × CPU/Memory/Storage capacity of Instance 1 + CPU/Memory/Storage utilization rate of Instance 2 × CPU/Memory/Storage capacity of Instance 2 +...+ CPU/Memory/Storage utilization rate of Instance N × CPU/Memory/Storage capacity of Instance N)/(CPU/Memory/Storage capacity of Instance 1 + CPU/Memory/Storage capacity of Instance 2 +...+ CPU/Memory/Storage capacity of Instance N).
- The storage utilization rate of OSS is calculated by using the following formula: Storage utilization rate of OSS = (Storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Storage utilization rate of Instance 2 × Storage capacity of Instance 2 +...+ Storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +...+ Storage capacity of Instance N).

Calculation based on time periods

By instance

Data provided by CloudMonitor is used to calculate and analyze the utilization rate of ECS, ApsaraDB RDS, or OSS.

The utilization rate is calculated on a daily basis. The instance-specific utilization rate of the previous day is generated every day, and the data is stored for one month. The utilization rates of the previous day are aggregated for each project, and the data is stored for one year. The following data of a project is stored: daily CPU utilization rate, daily memory utilization rate, daily disk utilization rate, and overall utilization rate.

Instance utilization rates

Except for OSS, the daily instance utilization rate of a product is calculated by using the following formula: **Daily instance utilization rate of a product = (Average instance utilization rate during peak hours of a working day) × 70% + (Average instance utilization rate during off-peak hours of a working day) × 20% + Average instance utilization rate during non-working hours × 10%**. The peak hours of working days are from 9:00 to 10:00 and 14:00 to 15:00. The off-peak hours of working days are from 8:30 to 9:00, 10:00 to 12:00, and 15:00 to 17:30.

- The daily CPU, memory, or storage utilization rate of ECS is calculated by using the following formula: Daily CPU/memory/storage utilization rate of ECS = (Average CPU/memory/storage utilization rate during peak hours of a working day) × 70% + (Average CPU/memory/storage utilization rate during off-peak hours of a working day) × 20% + Average CPU/memory/storage utilization rate during non-working hours × 10%. The peak hours of working days are from 9:00 to 10:00 and 14:00 to 15:00. The off-peak hours of working days are from 8:30 to 9:00, 10:00 to 12:00, and 15:00 to 17:30.
- The daily CPU, memory, or storage utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily CPU/memory/storage utilization rate of ApsaraDB RDS = (Average CPU/memory/storage utilization rate during peak hours of a working day) × 70% + (Average CPU/memory/storage utilization rate during off-peak hours of a working day) × 20% + Average CPU/memory/storage utilization rate during non-working hours × 10%. The peak hours of working days are from 9:00 to 10:00 and 14:00 to 15:00. The off-peak hours of working days are from 8:30 to 9:00, 10:00 to 12:00, and 15:00 to 17:30.
- Daily instance utilization rate of OSS = Storage usage/Purchased storage capacity

Overall daily utilization rate of an instance

The overall daily utilization rate of an instance is calculated by using the weighted value of the CPU, memory, and disk utilization rates. The following formula is used: **Overall daily utilization rate of an instance = CPU utilization rate of the instance × 33% + Memory utilization rate of the instance × 33% + Disk utilization rate of the instance × 33%**.

Overall instance utilization rate per unit time

The overall instance utilization rate or average instance utilization rate per day during specific consecutive days, such as 7 days, 14 days, or 1 month, is calculated.

Project-specific utilization rates

The utilization rates of the previous day are aggregated for each project, and the data is stored for one year. The following data of a project is stored: daily CPU utilization rate, daily memory utilization rate, daily disk utilization rate, and overall utilization rate.

Daily CPU utilization rate of a project

The daily CPU utilization rate of a project is calculated by using the following formula: Daily CPU utilization rate of a project = Daily CPU utilization rate of ECS × CPU weight of ECS + Daily CPU utilization rate of ApsaraDB RDS × CPU weight of ApsaraDB RDS.

- The CPU weight of ECS is calculated by using the following formula: CPU weight of ECS = CPU capacity of ECS/(CPU capacity of ECS + CPU capacity of ApsaraDB RDS).
- The CPU weight of ApsaraDB RDS is calculated by using the following formula: CPU weight of ApsaraDB RDS = CPU capacity of ApsaraDB RDS/(CPU capacity of ECS + CPU capacity of ApsaraDB RDS).
- The daily CPU utilization rate of ECS is calculated by using the following formula: Daily utilization rate of ECS = (Daily CPU utilization rate of Instance 1 × CPU capacity of Instance 1 + Daily CPU utilization rate of Instance 2 × CPU capacity of Instance 2 +...+ Daily CPU utilization rate of Instance N × CPU capacity of Instance N)/(CPU capacity of Instance 1 + CPU capacity of Instance 2 +...+ CPU capacity of Instance N).
- The daily CPU utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily CPU utilization rate of ApsaraDB RDS = (Average daily CPU utilization rate of Instance 1 × CPU capacity of Instance 1 + Average daily CPU utilization rate of Instance 2 × CPU capacity of Instance 2 +...+ Average daily CPU utilization rate of Instance N × CPU capacity of Instance N)/(CPU capacity of Instance 1 + CPU capacity of Instance 2 +...+ CPU capacity of Instance N).

Daily memory utilization rate of a project

The daily memory utilization rate of a project is calculated by using the following formula: Daily memory utilization rate of a project = Daily memory utilization rate of ECS × Memory weight of ECS + Daily memory utilization rate of ApsaraDB RDS × Memory weight of ApsaraDB RDS.

- The memory weight of ECS is calculated by using the following formula: Memory weight of ECS = Memory capacity of ECS/(Memory capacity of ECS + Memory capacity of ApsaraDB RDS).
- The memory weight of ApsaraDB RDS is calculated by using the following formula: Memory weight of ApsaraDB RDS = Memory capacity of ApsaraDB RDS/(Memory capacity of ECS + Memory capacity of ApsaraDB RDS).
- The daily memory utilization rate of ECS is calculated by using the following formula: Daily memory utilization rate of ECS = (Average daily memory utilization rate of Instance 1 × Memory capacity of Instance 1 + Average daily memory utilization rate of Instance 2 × Memory capacity of Instance 2 +…+ Average daily memory utilization rate of Instance N × Memory capacity of Instance N)/(Memory capacity of Instance 1 + Memory capacity of Instance 2 +…+ Memory capacity of Instance N).
- The daily memory utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily memory utilization rate of ApsaraDB RDS = (Average daily memory utilization rate of Instance 1 × Memory capacity of Instance 1 + Average daily memory utilization rate of Instance 2 × Memory capacity of Instance 2 +…+ Average daily memory utilization rate of Instance N × Memory capacity of Instance N)/(Memory capacity of Instance 1 + Memory capacity of Instance 2 +…+ Memory capacity of Instance N).

Daily storage utilization rate of a project

The daily storage utilization rate of a project is calculated by using the following formula: Daily storage utilization rate of a project = Daily storage utilization rate of ECS × Storage weight of ECS + Daily storage utilization rate of ApsaraDB RDS × Storage weight of ApsaraDB RDS + Daily storage utilization rate of OSS × Storage weight of OSS.

- The storage weight of ECS is calculated by using the following formula: Storage weight of ECS = (SSD storage capacity of ECS + Ultra disk storage capacity of ECS)/(SSD storage capacity of ECS + Ultra disk storage capacity of ECS + Storage capacity of ApsaraDB RDS + Storage capacity of OSS).
- The storage weight of ApsaraDB RDS is calculated by using the following formula: Storage weight of ApsaraDB RDS = Storage capacity of ApsaraDB RDS/(SSD storage capacity of ECS + Ultra disk storage capacity of ECS + Storage capacity of ApsaraDB RDS + Storage capacity of OSS).
- The storage weight of OSS is calculated by using the following formula: Storage weight of OSS = Storage capacity of OSS/(SSD storage capacity of ECS + Ultra disk storage capacity of ECS + Storage capacity of ApsaraDB RDS + Storage capacity of OSS).
- The daily storage utilization rate of ECS is calculated by using the following formula: Daily storage utilization rate of ECS = (Average daily storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Average daily storage utilization rate of Instance 2 × Storage capacity of Instance 2 +…+ Average daily storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +…+ Storage capacity of Instance N).
- The daily storage utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily storage utilization rate of ApsaraDB RDS = (Average daily storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Average daily storage utilization rate of Instance 2 × Storage capacity of Instance 2 +…+ Average daily storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +…+ Storage capacity of Instance N).

- The daily storage utilization rate of OSS is calculated by using the following formula: Daily storage utilization rate of OSS = (Average daily storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Average daily storage utilization rate of Instance 2 × Storage capacity of Instance 2 +…+ Average daily storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +…+ Storage capacity of Instance N).

Overall utilization rate of a project

The overall utilization rate of a project is calculated by using the weighted value of the CPU, memory, and disk utilization rates. The following formula is used:

Overall daily utilization rate of a project = CPU utilization rate of the project × 33% + Memory utilization rate of the project × 33% + Disk utilization rate of the project × 33%.

Overall project utilization rate per unit time

The overall project utilization rate or average project utilization rate per day during specific consecutive days, such as 7 days, 14 days, or 1 month, is calculated.

Organization-specific utilization rates

All organization-specific utilization rates of the previous day are calculated every day, and the data is stored for one year. The utilization rates of sub-organizations are excluded. The utilization rates are calculated on a daily basis. The following data of the organization is stored: daily CPU utilization rate, daily memory utilization rate, daily disk utilization rate, and overall utilization rate.

Daily CPU utilization rate of an organization

The daily CPU utilization rate of an organization is calculated by using the following formula: Daily CPU utilization rate of an organization = Daily CPU utilization rate of ECS × CPU weight of ECS + Daily CPU utilization rate of ApsaraDB RDS × CPU weight of ApsaraDB RDS.

- The CPU weight of ECS is calculated by using the following formula: CPU weight of ECS = CPU capacity of ECS/(CPU capacity of ECS + CPU capacity of ApsaraDB RDS).
- The CPU weight of ApsaraDB RDS is calculated by using the following formula: CPU weight of ApsaraDB RDS = CPU capacity of ApsaraDB RDS/(CPU capacity of ECS + CPU capacity of ApsaraDB RDS).
- The daily CPU utilization rate of ECS is calculated by using the following formula: Daily utilization rate of ECS = (Average daily CPU utilization rate of Instance 1 × CPU capacity of Instance 1 + Average daily CPU utilization rate of Instance 2 × CPU capacity of Instance 2 + …+ Average daily CPU utilization rate of Instance N × CPU capacity of Instance N)/(CPU capacity of Instance 1 + CPU capacity of Instance 2 +…+ CPU capacity of Instance N).
- The daily CPU utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily CPU utilization rate of ApsaraDB RDS = (Average daily CPU utilization rate of Instance 1 × CPU capacity of Instance 1 + Average daily CPU utilization rate of Instance 2 × CPU capacity of Instance 2 +…+ Average daily CPU utilization rate of Instance N × CPU capacity of Instance N)/(CPU capacity of Instance 1 + CPU capacity of Instance 2 +…+ CPU capacity of Instance N).

Daily memory usage of organization

The daily memory utilization rate of an organization is calculated by using the following formula: Daily memory utilization rate of an organization = Daily memory utilization rate of ECS × Memory weight of ECS + Daily memory utilization rate of ApsaraDB RDS × Memory weight of ApsaraDB RDS.

- The memory weight of ECS is calculated by using the following formula: Memory weight of ECS = Memory capacity of ECS/(Memory capacity of ECS + Memory capacity of ApsaraDB RDS).
- The memory weight of ApsaraDB RDS is calculated by using the following formula: Memory weight of ApsaraDB RDS = Memory capacity of ApsaraDB RDS/(Memory capacity of ECS + Memory capacity of ApsaraDB RDS).
- The daily memory utilization rate of ECS is calculated by using the following formula: Daily memory utilization rate of ECS = (Average daily memory utilization rate of Instance 1 × Memory capacity of Instance 1 + Average daily memory utilization rate of Instance 2 × Memory capacity of Instance 2 +...+ Average daily memory utilization rate of Instance N × Memory capacity of Instance N)/(Memory capacity of Instance 1 + Memory capacity of Instance 2 +...+ Memory capacity of Instance N).
- The daily memory utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily memory utilization rate of ApsaraDB RDS = (Average daily memory utilization rate of Instance 1 × Memory capacity of Instance 1 + Average daily memory utilization rate of Instance 2 × Memory capacity of Instance 2 +...+ Average daily memory utilization rate of Instance N × Memory capacity of Instance N)/(Memory capacity of Instance 1 + Memory capacity of Instance 2 +...+ Memory capacity of Instance N).

Daily storage utilization rate of an organization

The daily storage utilization rate of an organization is calculated by using the following formula: Daily storage utilization rate of an organization = Daily storage utilization rate of ECS × Storage weight of ECS + Daily storage utilization rate of ApsaraDB RDS × Storage weight of ApsaraDB RDS + Daily storage utilization rate of OSS × Storage weight of OSS.

- The storage weight of ECS is calculated by using the following formula: Storage weight of ECS = (SSD storage capacity of ECS + Ultra disk storage capacity of ECS)/(SSD storage capacity of ECS + Ultra disk storage capacity of ECS + Storage capacity of ApsaraDB RDS + Storage capacity of OSS).
- The storage weight of ApsaraDB RDS is calculated by using the following formula: Storage weight of ApsaraDB RDS = Storage capacity of ApsaraDB RDS/(SSD storage capacity of ECS + Ultra disk storage capacity of ECS + Storage capacity of ApsaraDB RDS + Storage capacity of OSS).
- The storage weight of OSS is calculated by using the following formula: Storage weight of OSS = Storage capacity of OSS/(SSD storage capacity of ECS + Ultra disk storage capacity of ECS + Storage capacity of ApsaraDB RDS + Storage capacity of OSS).
- The daily storage utilization rate of ECS is calculated by using the following formula: Daily storage utilization rate of ECS = (Average daily storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Average daily storage utilization rate of Instance 2 × Storage capacity of Instance 2 +...+ Average daily storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +...+ Storage capacity of Instance N).
- The daily storage utilization rate of ApsaraDB RDS is calculated by using the following formula: Daily storage utilization rate of ApsaraDB RDS = (Average daily storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Average daily storage utilization rate of Instance 2 × Storage capacity of Instance 2 +...+ Average daily storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +...+ Storage capacity of Instance N).
- The daily storage utilization rate of OSS is calculated by using the following formula: Daily storage utilization rate of OSS = (Average daily storage utilization rate of Instance 1 × Storage capacity of Instance 1 + Average daily storage utilization rate of Instance 2 × Storage capacity of Instance 2 +...+ Average daily storage utilization rate of Instance N × Storage capacity of Instance N)/(Storage capacity of Instance 1 + Storage capacity of Instance 2 +...+ Storage capacity of Instance N).

Overall daily utilization rate of an organization

The overall daily utilization rate of an organization is calculated by using the weighted value of the CPU, memory, and disk utilization rates. The following formula is used:

Overall daily utilization rate of an organization = CPU utilization rate of the organization × 30% + Memory utilization rate of the organization × 30% + Disk utilization rate of the organization × 30%.

Overall organization utilization rate per unit time

The overall organization utilization rate or average organization utilization rate per day during specific consecutive days, such as 7 days, 14 days, or 1 month, is calculated.

Yearly utilization rate of an organization

- The yearly CPU utilization rate of an organization is calculated by using the following formula: Yearly CPU utilization rate of an organization = (CPU utilization rate of the organization in January + CPU utilization rate of the organization in February +...+ CPU utilization rate of the organization in December)/12.
- The yearly memory utilization rate of an organization is calculated by using the following formula: Yearly memory utilization rate of an organization = (Memory utilization rate of the organization in January + Memory utilization rate of the organization in February +...+ Memory utilization rate of the organization in December)/12.
- The yearly storage utilization rate of an organization is calculated by using the following formula: Yearly storage utilization rate of an organization = (Storage utilization rate of the organization in January + Storage utilization rate of the organization in February +...+ Storage utilization rate of the organization in December)/12.
- The yearly overall utilization rate of an organization is calculated by using the following formula: Yearly overall utilization rate of an organization = (Overall utilization rate of the organization in January + Overall utilization rate of the organization in February +...+ Overall utilization rate of the organization in December)/12.

6.8.2. Resource health scores

Overall health scores of a project

An overall health score of a project is the result that is obtained after the status of idle resources, compliance governance, and resource utilization rate in a project are summarized. Overall health scores of a project are retained for one year.

- Overall health score of a project in a single day = (Total number of resources in the project - Number of idle resources in the project)/Total number of resources in the project × 40 + (Total number of resources in the project - Number of non-compliant resources in the project)/Total number of resources in the project × 30 + Overall resource utilization rate of the project on the day × 30
- Overall health score of a project in recent 30 days = Average of overall health scores of a project in a single day of recent 30 days

Overall health scores of an organization

An overall health score of an organization is the result that is obtained after the status of idle resources, compliance governance, and resource utilization rate in an organization are summarized. Overall health scores of an organization are retained for one year.

- Overall health score of an organization in a single day = (Total number of resources in the organization - Number of idle resources in the organization)/Total number of resources in the organization × 40 + (Total number of resources in the organization - Number of non-compliant resources in the organization)/Total number of resources in the organization × 30 + Overall resource utilization rate of the organization on the day × 30

- Overall health score of an organization in recent 30 days = Average of overall health scores of an organization in a single day of recent 30 days

Overall health scores of a platform

An overall health score of a platform is the result that is obtained after the status of idle resources, compliance governance, and resource utilization rate of all organizations on a platform are summarized. Overall health scores of a platform are retained for one year.

- Overall health score of a platform in a single day = (Total number of resources on the platform - Number of idle resources on the platform)/Total number of resources on the platform × 40 + (Total number of resources on the platform - Number of non-compliant resources on the platform)/Total number of resources on the platform × 30 + Overall resource utilization rate of the platform on the day × 30
- Overall health score of a platform in recent 30 days = Average of overall health scores of a platform in a single day of recent 30 days

6.8.3. Cloud optimization results

The system generates cloud optimization tasks based on algorithms for optimizing resources. You can process the relevant resources and tasks. The data displayed is the analytics result of the previous day.

Prerequisites

Resource optimization settings are completed. For more information, see Resource optimization settings.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Cloud Resource Optimization > Cloud Optimization Results**.
4. Click the **Resources to be Optimized**, **Idle Resources**, or **Non-compliant Resources** tab.
5. (Optional) Click **Advanced Filter** to filter resources by criteria such as **Organization**, **Resource Set**, **Cloud Service**, **Instance ID**, and **Parameter**.
6. The items in the list include **Instance ID/Name**, **Organization**, **Resource Set**, **Cloud Service**, **Parameter**, **Details**, **Metric Data**, **Description**, and **Suggestion**. You can click an instance ID and fix the corresponding issue based on suggestions.

6.8.4. Resource optimization settings

You can configure policies that are used to analyze idle status, optimization plans, and health status of resources. You can modify existing policies.

The following table describes the cloud services for which you can perform resource optimization settings, the type of the resources that can be optimized, and the names of the resource optimization settings.

Cloud service	Type of the resources that can be optimized	Resource optimization setting

Elastic Compute Service (ECS)	Idle resource	• ECS downtime • Idle storage
	Resource to be optimized	• Low disk utilization • High disk utilization • High resource utilization • Low resource utilization
Elastic Block Storage (EBS)	Idle resource	Unbound EBS
Object Storage Service (OSS)	Idle resource	Idle storage
	Resource to be optimized	High disk utilization
Apsara File Storage NAS	Idle resource	No read/write traffic
Elastic IP Address (EIP)	Idle resource	Unbound EIP
Server Load Balancer (SLB)	Idle resource	No access requests
	Non-compliant resources	SLB ports
ApsaraDB RDS	Idle resource	No new connection
	Resource to be optimized	• Low resource utilization • High disk utilization • High resource utilization
KVStore for Redis	Resource to be optimized	• High resource utilization • Low resource utilization
AnalyticDB for PostgreSQL	Resource to be optimized	• Low resource utilization • High resource utilization
ApsaraDB for MongoDB	Resource to be optimized	• Low resource utilization • High resource utilization

AnalyticDB for MySQL V3.0	Resource to be optimized	• Low disk utilization • High disk utilization
PolarDB-X	Resource to be optimized	• High resource utilization • Low resource utilization
Hologres	Resource to be optimized	• High resource utilization • Low resource utilization
Message Queue (MQ)	Idle resource	Congested messages
API Gateway	Idle resource	• No new connection • No read/write traffic

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Cloud Resource Optimization > Optimization Settings**.
4. (Optional) Click the icon in the search box and select **Parameter**, **Category**, or **Cloud Service** from the drop-down list for filtering.
5. Click **Edit** in the **Actions** column corresponding to the configuration item and modify the parameters of the configuration item, such as **Description**, **Time Range**, **Valid value**, **Set Threshold**, and **Suggestions**. The following table describes the parameters of the optimization settings.

> **Note**
> The parameters that can be modified vary based on the type of optimization settings.

Parameter		Description
Basic Information	Description	The description of the setting. The description must be 1 to 300 characters in length.
Time Range	Duration	You can select **1 Day**, **1 Week**, **2 Weeks**, or **30 Days**.
	Hourly Usage	The value of hourly usage of the service. You can select **Average**, **Maximum Value**, or **Minimum Value**.

Valid value	Daily Usage	The value of daily usage of the service. You can select **All** or **Duration**. **Note** If you select **Duration**, you must add new entries to configure the **Time Range** and **Weight** parameters. You can add multiple entries.
Set Threshold	Idle Rules	Configure a threshold for the idle rules of the setting.
Suggestions	-	The suggestions for fixing the optimization settings. The suggestions must be 0 to 300 characters in length.

6. Click **Submit**.

6.9. Billing management

6.9.1. Billing overview

On the Billing Overview page, you can obtain a clear overview of the effective times of the billing configurations on the platform.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Billing Management** > **Billing Overview**.
4. On the **Billing Overview** page, you can view the effective time of each billing configuration in the Gantt chart. By default, the Gantt chart is displayed in days.

   You can perform the following operations to obtain the desired information.

   - **Chang the time unit of the Gantt chart** : Click **Day** in the upper-right corner of the chart and select **Week**, **Month**, **Quarter**, or **Year**.
   - **Search for the target billing configuration**: In the upper-left corner of the Billing Overview page, enter a keyword of the billing configuration to search for it.
   - **View details of a billing configuration**: Click the name of a billing configuration. On the configuration details page, view the detailed information of the configuration.

6.9.2. Billing configurations

6.9.2.1. Create a billing configuration

Billing configurations define the billing plans of products. You can configure the unit prices of cloud products based on the pricing policies or price agreements. The pay-as-you-go and monthly subscription billing methods are supported.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Billing Management** > **Billing Configurations**.
4. On the **Billing Configuration** page, click **Create Billing Configuration**. In the wizard, configure the following parameters.

Step in the wizard	Parameter	Description
Basic Information	Billing Configuration Name	The name of the billing configuration. The name can be 1 to 200 characters in length. **Note** Billing configuration names must be unique.
	Effective Time	The start and end time of the validity period of the billing configuration.
	Billing Method	• Pay-as-you-go • Monthly Subscription
	Plan Priority	A number that indicates the priority of the billing configuration. **Note** A smaller value indicates a higher priority. The value for each billing configuration must be unique. The priority determines the order in which billing configurations are displayed.
	Organization	The organization for which the billing method takes effect.
	Remarks	The remarks of the billing configuration. The remarks must be 1 to 256 characters in length.
	Billing Configurations	Configure the prices and discount rate for each cloud product. The discount rate must be within the range of 0 to 10. **Note** For each product, you can click Advanced Settings in the Actions column to configure their billable items.
		Configure the billable items of each product.

Billing Configurations	Advanced Settings	**Note** Billable items are units of billing for the components of each product. The fees of a product are the sum of the fees generated by all billable items. • Configure Filtering Condition ◦ Exact Match: Bills are generated for all billing data. ◦ Conditional Match: If you select AND, bills are not generated only when all the actual values match the specified values. If you select OR, bills are not generated as long as the actual values match one of the specified values. ◦ Filtering conditions: A condition consists of a field and a value. You can add multiple conditions. When the actual values of the specified fields in a billing record match the specified values, no bills are generated for the billing record. **Note** Configuration of the filtering conditions is available only when Configure Filtering Condition is set to Conditional Match. • Billing Item Management: You can configure the existing billable items or create one or multiple billable items. If you click Create Billable Item, configure the following parameters: ◦ Billable Item Name: The name of the billable item. ◦ Billing Method: The value is set to the billing method of the billing configuration and cannot be changed. ◦ Unit Price Settings: Valid values: Quantity-based Price, Fixed Price, and Stepwise Pricing. ▪ Fixed Price: All specifications of the product are billed at the same price. You can customize the fixed price. The value must be a non-negative number that can be accurate to six decimal places. ▪ Quantity-based Price: The pricing varies based on specifications. You can upload a file from your computer to specify the price of each specification. ▪ Stepwise Pricing: The pricing varies based on tiers of quantities. You can customize the price of each tier.

		◦ Quantity Settings: The quantity that is used in an hour or the quantity that is declared each hour. Valid values: Quantity and Fixed Quantity. You can add multiple quantity settings. ▪ Quantity Type: If you select Quantity, the metering is based on the usage of a specification parameter. If you select Fixed Quantity, the metering is based on a custom quantity number. **Note** When Unit Price Settings is set to Stepwise Pricing, only the Quantity option is available for Quantity Type. ▪ Quantity: When Quantity Type is set to Quantity, select a specification parameter that is used for the metering. When Quantity Type is set to Fixed Quantity, specify a number. ▪ Conversion Factor: The factor that is used to align the quantity and unit price.
Billing Configuration Verification	Instance Type Configuration Ratio	The proportion of specifications in the current cloud product whose prices have been set to a number larger than 0.
	Price Verification	Indicates whether the configuration ratio is 100%. If the configuration ratio of a product is smaller than 100%, Rejected is displayed. **Note** The parameter value is for reference only. You can create the billing configuration even if the configuration ratios of some products are smaller than 100%.

6.9.2.2. Manage billing configurations

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane, choose **Billings Management** > **Billing Configurations**.
4. On the **Billing Configuration** page, perform the following operations.

Operation	Procedure

Modify the basic information of a billing configuration	i. Click **Edit** in the **Actions** column corresponding to a billing configuration. ii. In the **Basic Information** section of the configuration details page, click **Modify Basic Information**. iii. In the dialog box that appears, modify the **Billing Configuration Name**, **Effective Time**, **Plan Priority**, **Organization**, and **Remarks parameters**. iv. Click **Save**.
Modify the pricing configurations of a billing configuration	i. Click **Edit** in the **Actions** column corresponding to a billing configuration. ii. In the pricing list of the page, click **Edit** in the **Actions** column corresponding to a cloud product. Then, change the discount rate and prices of the product. iii. Click **Save** in the **Actions** column. iv. (Optional) Click **Advanced Settings** in the **Actions** column. In the panel that appears, modify the billing items of the product. For information about the parameters about the billing items, see Create a billing configuration. Click **OK**.
Clone a billing configuration	i. Click **Clone** in the **Actions** column of a billing configuration. ii. In the wizard, set the parameters based on the existing configurations. For information about the parameters, see Create a billing configuration. **Note** ▪ Names of billing configurations must be unique. Therefore, you need to modify the name. ▪ The billing method is the same as that of the original billing configuration and cannot be modified. ▪ Priorities of billing configurations must be unique. Therefore, you need to modify the priority. iii. Click **Save**.
Delete a billing configuration	i. Click **Delete** in the **Actions** column corresponding to a billing configuration. You can also click the name of a billing configuration and click **Delete** in the upper-right corner of the billing configuration details page that appears. ii. In the message that appears, click **OK**.

Export billing configurations	i. Click the name of a billing configuration. ii. In the lower part of the billing configuration details page, select a product or **All** from the **Cloud Service** drop-down list. Then, click **Export Price List**. **Note** The exported files are in the **.xlsx** format. The spreadsheet contains the organization, export time, start and end time of the billing configuration, as well as the billing details of each product such as the billing method, discount rate, pricing rule, billable item description, billing fields, specifications, prices, counting type, counting field, unit and conversion factor.

7.Security

7.1. Overview

The Apsara Uni-manger Management Console uses operations logs, AccessKey pairs, and multi-factor authentication (MFA) to enhance the security when users access systems and resources.

Module	Description
Operation Logs	Access and use of resources on the console are recorded in operation logs. • You can filter operation logs by specifying conditions such as the time range, organization, user name, and resource ID. • You can dump logs to Object Storage Service (OSS) and download them to your computer as files.
AccessKey	AccessKey pairs are used by users to access cloud resources. • Users can view their AccessKey pairs. • Users can create AccessKey pairs. Each user can have at most two AccessKey pairs. • Users can enable or disable their AccessKey pairs. At least one AccessKey pair must be enabled for each user. • Users can delete their AccessKey pairs, but at least one AccessKey pairs must be retained for each user.
MFA	MFA is a method for identity authentication in addition to the username and password of an account. It can be used to improve security. • You can check whether the MFA feature is enabled for a user. If not, you can enable it. However, you cannot disable it after it is enabled. • You can bind or unbind virtual MFA devices. • You can reset MFA keys.

7.2. Operation logs

7.2.1. Supported cloud services

Apsara Stack records operation logs for cloud services. The logs contain the access and usage information about the cloud services.

The following table describes the services whose logs are recorded.

Category	Service
	Bare Metal Management Service

Computing	Operation Orchestration Service (OOS)
	Resource Orchestration Service (ROS)
	Auto Scaling
	Elastic High Performance Computing (E-HPC)
	Elastic Compute Service (ECS)
	Container Registry (ACR)
	Container Service for Kubernetes (ACK)
Storage	CDS
	CDS Configuration Manager
	Elastic Block Storage (EBS)
	Object Storage Service (OSS)
	Apsara File Storage NAS
	Tablestore
	Simple Log Service
	Photo and Drive Service
Networking	Server Load Balancer (SLB)
	Virtual Private Cloud (VPC)
	Apsara Stack DNS (in-cloud)
	Apsara Stack DNS

Database	ApsaraDB RDS
	PolarDB
	ApsaraDB for Redis
	ApsaraDB for MongoDB
	PolarDB-X 1.0
	AnalyticDB for PostgreSQL
	AnalyticDB for MySQL V3.0
	Data Transmission Service (DTS)
	Lindorm
	Graph Database (GDB)
	ApsaraDB for HBase
	Time Series Database (TSDB)
	Database Autonomy Service (DAS)
Middleware	Enterprise Distributed Application Service (EDAS)
	Message Queue
	Message Queue for Apache Kafka
	Cloud-native Gateway COP
Application services	API Gateway
	Cloud Config

Security	Key Management Service (KMS)
	Bastionhost (Early Version)
	DataBase Audit
Big data	MaxCompute
	DataHub
	Elasticsearch
	Machine Learning Platform for AI (PAI)
	Hologres
	E-MapReduce
IoT services	IoT Platform
	Alibaba Cloud Link WAN
	IoT Security Operation Center
	IoT Internet Device ID
Monitoring and O&M	CloudMonitor
	ACOS
Disaster Recovery	Apsara Stack Resilience for Zone-disaster Recovery
Others	Apsara Uni-manager
	Hybrid Cloud Networks
	Security Token Service (STS)

	Base Log Service
	Cloud Migration Center
	WUYING Workspace

7.2.2. Manage operation logs

You can view operation logs in which access to and use of resources on the console are recorded. You can export the logs to Object Storage Service (OSS) or download them to your computer. Then, you can perform operations such as behavior analysis, security analysis, resource change tracking, and compliance audit.

Query events

On the Event Query page, you can view all the operations performed on cloud resources through all methods, including the console, API operations, and developer tools.

1. Log on to the Apsara Uni-manager Management Console as a security administrator.
2. In the top navigation bar, choose **Security** > **Global Platform Security** > **Operation Logs**.
3. In the left-side navigation pane of the page that appears, click **Event Query**.
4. You can filter the records by the organization, product type, resource type, and start and end time.

   You can also search for records based on the username, resource ID, source IP address, or keyword.

Parameter	Description
Organization	The organization of the resource on which the operation is performed.
Resource Set	The resource set of the resource on which the operation is performed.
Resource Type	The type of the resource on which the operation is performed.
Resource ID	The ID of the resource on which the operation is performed.
Event Type	The operation that is performed. Operations include creating, modifying, deleting, querying, updating, associating, disassociating, enabling, disabling, applying for, releasing, and changing the ownership of resources.
Status	The status of the operation.
Product Type	The product type of the resource on which the operation is performed.

Call Source	The source of the operation.
Source IP Address	The IP address of the operator.
User Name	The name of the account to which the operator belongs.
Time	The time when the operation is completed.
Operation Name	The name of the operation.
Details	Click to view details of the operation.

Log storage center

You can create asynchronous export tasks to export specific operation logs to OSS buckets for permanent storage.

1. Log on to the Apsara Uni-manager Management Console as a security administrator.
2. In the top navigation bar, choose **Security** > **Global Platform Security** > **Operation Logs**.
3. In the left-side navigation pane, click **Log Storage Center**. On the page that appears, click **Create an OSS export task** and configure the parameters.

Parameter	Description
Task Name	The name of the export task.
Organization	The organization to which the exported logs belong.
Filter Condition Name	The filter conditions for logs. You can select a cloud product or a resource type.
Start and End Time	The time range in which the operations are performed.
User Name	The username of the operator.
Resource ID	The ID of the resource.
Source IP address	The IP address of the operator.

Key Words	Keywords of the target operations.

4. Click **OK**. The OSS export task is created, and the system starts exporting the operation logs to the default OSS bucket as .XLSX files.

   In the task list, you can view the status of the task.

   - In Progress: The task is being executed.
   - Completed: The task has been completed. You can click **Download Report** in the **Download Operation** column corresponding to the task to download the report from the OSS bucket to your computer.
   - Failed: The task failed. You can click **Retry** in the **Download Operation** column to try executing the task again.
   - Expired: The task has expired.

7.3. AccessKey

An AccessKey pair consists of an AccessKey ID and an AccessKey secret. AccessKey pairs are used to authenticate request senders, which can improve cloud security.

> **Important**
>
> A user can have up to two AccessKey pairs, and at least one AccessKey pair must be enabled.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the upper-right corner of the homepage, click the profile picture and click **User information** in the menu that appears.

Operation	Procedure
View AccessKey pairs	In the AccessKey Pair section, view the AccessKey pairs of your account. **Note** The AccessKey IDs and AccessKey secrets provide full permissions on cloud resources within your Apsara Stack tenant account. Make sure that they are confidential.
Create an AccessKey pair	You can delete your old AccessKey pairs and create new ones to implement the rotation of AccessKey pairs. i. In the upper-right corner of the AccessKey Pair section, click **Create AccessKey Pair**. ii. In the message that appears, click **Close**.

Disable an AccessKey pair	**Important** Disabling an AccessKey pair may cause errors in the products or applications that depend on it. Proceed with caution. i. In the AccessKey Pair section, click **Disable** in the **Actions** column corresponding to the target AccessKey pair. ii. In the message that appears, click **OK**.
Enable an AccessKey pair	In the AccessKey Pair section, click **Enable** in the **Actions** column corresponding to a disabled AccessKey pair.
Delete an AccessKey pair	**Important** Deleting an AccessKey pair may cause errors in the products or applications that depend on it. Proceed with caution. i. In the AccessKey Pair section, click **Delete** in the **Actions** column corresponding to the target AccessKey pair. ii. In the message that appears, click **OK**.

7.4. MFA

7.4.1. Overview

Multi-factor authentication (MFA) is an authentication method. In addition to username and password authentication, MFA provides an extra layer of protection.

Introduction to MFA

When MFA is enabled, you must enter your username and password (first security factor) and then a variable verification code (second security factor) from an MFA device when you log on to the Apsara Uni-manager Management Console. Two-factor authentication enhances security for your account.

MFA devices use the Time-based One-time Password (TOTP) algorithm to generate time-dependent 6-digit dynamic verification codes. The Apsara Uni-manager Management Console supports software-based virtual MFA devices. You can install software (such as the Alibaba Cloud app) that supports MFA on your mobile device (such as your mobile phone) to act as a virtual MFA device.

Precautions

The TOTP algorithm requires that the time of the system clock of the Apsara Uni-manager Management Console remain consistent with the standard time on the Internet. Otherwise, discrepancies in time can lead to inconsistent MFA verification codes and you may be unable to log on to the Apsara Uni-manager Management Console.

7.4.2. Manage MFA

Multi-factor authentication (MFA) is an authentication method. In addition to username and password authentication, MFA provides an extra layer of protection. You can enable or disable the MFA feature in the console.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the upper-right corner of the homepage, click the profile picture and click **User information** in the menu that appears.

Operation	Procedure
Bind a virtual MFA device	i. In the upper-right corner of the **MFA** section, click **Bind Virtual MFA Device**. ii. On the **Bind Virtual MFA Device** page, follow the instructions to bind an MFA device. After an MFA device is bound, the MFA status changes to Enabled and the Device Status changes to Bound. **Note** After a virtual MFA device is bound, you must enter a 6-digit MFA verification code in addition to your username and password before you can log on to the Apsara Uni-manager Management Console.
Disable a virtual MFA device	i. In the upper-right corner of the **MFA** section, click **Disable Virtual MFA Device**. ii. In the message that appears, click **Disable Virtual MFA Device**. After you disable MFA, you need only to enter your username and password the next time you log on to the Apsara Uni-manager Management Console.

Related operations

Forcibly enable MFA

Administrators, including the platform administrator, operations administrator, and organization administrator, can check whether their users have MFA enabled. If MFA is disabled for the users, the administrators can forcibly enable MFA for the users.

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users** > **Users**.
4. Click the **System Users** tab.
5. Find the user for whom you want to forcibly enable MFA. Move the pointer over the More icon in the **Actions** column and click **MFA Settings** in the menu that appears.

   You can search for the user by username, organization, or role.
6. In the dialog box that appears, turn on **MFA Settings** and click **OK**.

> **Note**
>
> MFA can be forcibly enabled for users, but cannot be forcibly disabled.
>
> - Before MFA is enabled, the MFA status of the user is **Not Enabled**.
> - After MFA is enabled, the MFA status of the user is **Enabled but Not Bound**. After MFA is forcibly enabled for a user, the user must go to the Bind Virtual MFA Device page to bind a virtual MFA device before the user can log on to the Apsara Uni-manager Management Console.

Reset MFA

If the MFA key of a user is lost, an administrator need to reset the MFA settings.

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users** > **Users**.
4. Click the **System Users** tab.
5. Find the user for whom you want to reset MFA and click the username.

   You can search for the user by username, organization, or role.
6. In the user details panel, click **Reset password**.

8.User information management

On the User Information page, you can manage the information of your account, such as the basic information, password, MFA devices, messages, and AccessKey pairs.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the upper-right corner of the homepage, click the profile picture.
3. Click **User Information** in the menu that appears.

Section	Description
Basic Information	You can view the information of the user. You can click **Edit** in the upper-right corner of the section to modify the parameters, such as the mobile phone number, landline phone number, email, DingTalk token, and remarks.
MFA	You can enable or disable the MFA feature. For more information about MFA, see Manage MFA.
Messages	You can configure the methods for receiving system messages. Make sure that the method is configured. For more information, see Message gateway.
AccessKey	You can manage your AccessKey pairs. For more information, see AccessKey
View Resource Set	You can select different roles to view the associated resource sets.

- In the left part of the page, click **Switch Role** to switch to another role.
- Click **Change Password**. In the dialog box that appears, specify the current password and new password and confirm the password. Click **OK** to change the password.
- Click **View Current Role Policy** to view the information about the policies of the current role.

9.Page integration

Pages in the Apsara Uni-Manager Management Console can be integrated to other platforms as an open console solution. You can integrate the management pages of some services in the Apsara Uni-Manager Management Console to self-built management consoles or consoles provided by third-party ISVs to manage Apsara Stack services in the integrated console. By using the page integration capability, you can manage and customize your services in complex operation, maintenance, and management scenarios in a fast and flexible manner.

The Apsara Uni-Manager Management Console allows you to customize the display language, appearance, theme color, logo, version information of the pages that you integrate to your console. You can also configure to hide the brand information, adjust the layouts, and change the styles for the integrated pages to match the style of your console that is developed by yourself or provided by third-party ISVs. This way, you can provide a unified user experience in your business.

Prerequisites

You must complete the SSO authentication provided by the Apsara Uni-manager Management Console to achieve a unified identity verification mechanism between the Apsara Uni-manager Management Console and the console to which you want to integrate pages. In the mechanism, the Apsara Uni-manager Management Console acts as the service provider and the console to which you want to integrate pages as the identity provider. This way, you can log on to the Apsara Uni-manager Management Console or the console to which you want to integrate pages from the other console without the needs for identity verification.

For more information, see IDP configurations.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an operations administrator.
2. Click the Help icon in the upper-right corner of the page. In the top navigation bar of the page that appears, click **Page Integration**.
3. On the **Page Integration Overview** page, click **Page Integration Configuration**. Configure the parameters described in the following sections:

i. **Destination Application Configuration**: Configure the information about the console to which you want to integrate the pages of the Apsara Uni-manager Management Console.

Section	Parameter	Description
Basic information	**Application Name**	The name of the application to which you want to integrate the pages of the Apsara Uni-manager Management Console. The name must be 3 to 25 characters in length.
	Application Domain Name	The domain name of the application to which you want to integrate the pages of the Apsara Uni-manager Management Console. The domain name must be an HTTPS domain name. Example: https://console.example.com. **Note** You must specify all domain names that are involved in the application, including the domain name of the iFrame page. Otherwise, the integrated pages may not be properly displayed due to the security restrictions of your browsers. Separate multiple domain names with commas (,).
	Application Description	The description of the application to which you want to integrate the pages of the Apsara Uni-manager Management Console.
Switch Settings	**Embedded Switch**	You can turn on this switch if you want that the integrated pages are embedded in your console. If you turn on this switch, the navigation bars of the integrated pages are automatically hidden. If you turn off this switch, the integrated pages cannot be embedded to your console by using iFrame.
	Hide Left-side Navigation Pane	You can turn on this switch to hide the left-side navigation pane of the integrated pages. If you turn on this switch, other pages in the navigation pane are hidden in your console. In this case, you can choose to display only specific pages that are related to the integrated pages by configuring the URLs of the specific pages.

ii. **Page Style Configuration**: Configure the style of the integrated pages to match the style of your console.

Section	Parameter	Description
Platform Language	**Select the languages supported by the platform**	The pages of the Apsara Uni-manager Management Console supports multiple languages. You can select the following languages for the pages: Simplified Chinese, Traditional Chinese, and English.
	Select the default language	Select the default language of the integrated pages.
Appearance & Theme	**Appearance**	Select the page style of the integrated pages. You can select Light Mode or Dark Mode for the pages.
	Theme Color	Select the theme color of the integrated pages.
Page Watermark	**Switch**	You can add watermarks to integrated pages. After you turn on the switch, you can configure the content of the watermark. A watermark can contain up to 10 characters.

iii. **Embedded Page Configuration**: Configure the URLs of the pages that you want to integrate to your console from the Apsara Uni-manager Management Console. Some pages of the following services in the Apsara Uni-manager Management Console can be integrated to other consoles: Elastic Compute Service (ECS), Auto Scaling, ApsaraDB for MongoDB, ApsaraDB RDS, Lindorm, Service Load Balancer (SLB), Virtual Private Cloud (VPC), and Cloud Enterprise Network (CEN).

To create or scale resources in the corresponding pages of your console by clicking buttons in the integrated pages, you must configure operational URLs for the buttons. This way, you can use the integrated pages with your console in a unified manner.

> **Note**
>
> The operational URLs must be a lower-level domain name of the application domain name that you specified in the Destination Application Configuration section and must complete URLs that can be accessed.

4. After the configuration is complete, the pages of the Apsara Uni-manage Management Console are embedded to the page of your console.

   On the **Page Integration Overview** page, you can click **Page URL Embedding** to view the operational URLs of the integrated pages.

10.Cloud Monitor Service

10.1. Overview

CloudMonitor provides an out-of-the-box one-stop enterprise-grade monitoring solution for you to monitor resources on Apsara Stack.

CloudMonitor supports cloud service monitoring, host monitoring, alert service, statistical analysis, event monitoring, application groups, and report download. You can use CloudMonitor to collect data on monitoring metrics of various cloud services, such as Elastic Compute Service (ECS), Server Load Balancer (SLB), Object Storage Service (OSS), and Log Service. The metrics include event metrics and custom metrics.

CloudMonitor helps you implement multidimensional and efficient monitoring. You can improve system service availability and reduce the operations and maintenance (O&M) costs of IT infrastructure.

10.2. Overview

CloudMonitor provides an overview of your cloud resources. You can monitor the resource usage and status of your Alibaba Cloud services. You can also receive and resolve alerts at the earliest opportunity to ensure the availability of your applications. The following sections are displayed on the Overview page in CloudMonitor: Alerting Cloud products, Alert Rules, Event-triggered Alert Rules, Alert History, and Dashboards.

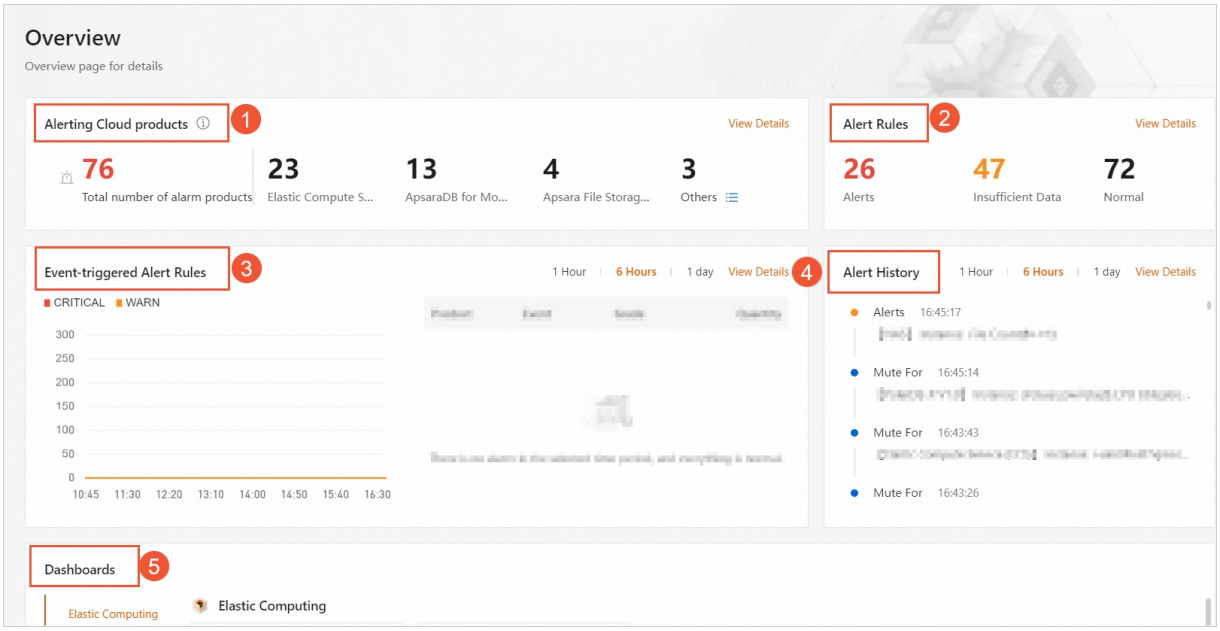

No.	Section	Description
①	**Alerting Cloud products**	This section displays the number of alerts from cloud services that are supported by CloudMonitor and enabled with the alert service. You can click **View Details** to go to the **Alert History List** page to view the details of all alerts.

②	**Alert Rules**	This section displays the numbers of alert rules in the **Alert**, **Insufficient Data**, and **Normal** states. You can click **View Details** to go to the **Alarm Rules** page and view the details of all alert rules. **Note** **An alert rule in the Insufficient Data state is generated in one of the following scenarios:** • The cloud service does not report data of the specified metric. • The cloud service instance is deleted. • The cloud service instance is missing or is not created in the Logstore of Log Service. Further troubleshooting is required.
③	**Event-triggered Alert Rules**	This section displays the status of system events and event-triggered alerts for cloud services. This helps you troubleshoot business exceptions efficiently. The system events include status change events of the CloudMonitor agent, and some system events of Hybrid Disaster Recovery and Elastic Compute Service. CloudMonitor can collect data about system event-triggered alerts, but cannot send alert notifications.
④	**Alert History**	This section displays the detailed records of recent alerts from cloud services that are supported by CloudMonitor and enabled with the alert service. You can click **View Details** to go to the **Alert History List** page and view the detailed alert history.
⑤	**Dashboards**	CloudMonitor automatically collects the monitoring data of resources of cloud services that belong to the current Alibaba Cloud account. You can view the monitoring charts of each cloud service in this section. You can also configure alert rules to monitor the status of resources. If an alert is triggered based on an alert rule, CloudMonitor sends an alert notification to you. This way, you are notified of the status of your resources at the earliest opportunity.

10.3. Cloud service monitoring

CloudMonitor automatically retrieves the resources of cloud services that belong to the current Alibaba Cloud account. You can view the monitoring charts of each cloud service. You can also configure alert rules to monitor resources. If an alert is triggered based on an alert rule, CloudMonitor sends an alert notification to you. This way, you are notified about the status of your resources at the earliest opportunity.

The following table describes the cloud services that can be monitored by using CloudMonitor.

> **Note**
> For information about the metrics of each cloud service, see the "Metrics" topic.

Category	Cloud service
Elastic computing	Elastic Compute Service (ECS) Elastic High Performance Computing (E-HPC)
Application services	API Gateway
Database services	ApsaraDB RDS PolarDB PolarDB-X 2.0 PolarDB-X 1.0 ApsaraDB for MongoDB ApsaraDB for Redis AnalyticDB for PostgreSQL AnalyticDB for MySQL V3.0 Data Transmission Service (DTS)
Storage	Cloud Storage Gateway (CSG) Apsara File Storage NAS (NAS) Object Storage Service (OSS) Tablestore
Middleware	Message Queue (MQ)
Big data	Realtime Compute for Apache Flink DataHub Elasticsearch Hologres
Networking	NAT Gateway Server Load Balancer (SLB) VPN Gateway Elastic IP Address (EIP) Apsara Stack DNS

10.3.1. Metrics

CloudMonitor checks the availability of services based on their metrics. You can configure alert rules and notification policies for these metrics to stay up to date on the running status and performance of monitored service instances.

Elastic computing

Elastic Compute Service (ECS)

Basic metrics

> **Note**
> The following metrics are collected from the virtual host layer and are slightly different from metrics collected inside the operating system.

Console-based metric	Unit	API-based metric	Dimension	Description
(ECS) CPU Utilization	%	CPUUtilization	userId and instanceId	The CPU utilization.
(ECS) InternetInRate	bit/s	InternetInRate	userId and instanceId	The rate of inbound traffic over the Internet.
(ECS) IntranetInRate	bit/s	IntranetInRate	userId and instanceId	The rate of inbound traffic over the intranet.
(ECS) InternetOutRate	bit/s	InternetOutRate	userId and instanceId	The rate of outbound traffic over the Internet.
(ECS) IntranetOutRate	bit/s	IntranetOutRate	userId and instanceId	The rate of outbound traffic over the intranet.
(ECS) DiskReadBPS	byte/s	DiskReadBPS	userId and instanceId	The number of bytes read from all disks per second.
(ECS) DiskWriteBPS	byte/s	DiskWriteBPS	userId and instanceId	The number of bytes written to all disks per second.

(ECS) DiskReadIOPS	count/s	DiskReadIOPS	userId and instanceId	The number of read I/O operations per second (IOPS) of all disks.
(ECS) DiskWriteIOPS	count/s	DiskWriteIOPS	userId and instanceId	The number of write IOPS of all disks.

Operating system metrics

> **Note**
>
> For ECS instances, you must install a monitoring plug-in to collect metrics inside the operating system.
>
> Installation method: In the left-side navigation pane of the CloudMonitor console, click **Host Monitoring**. In the list that appears, select the instances that you want to monitor, and click **Batch Install** in the lower part of the page.
>
> Metric data is displayed in the monitoring chart in 5 to 10 minutes after the monitoring plug-in is installed.

Console-based metric	Unit	API-based metric	Dimension	Description
(Agent)cpu.total	%	cpu_total	userId and instanceId	The CPU utilization.
(Agent)memory.used.utilization	%	memory_usedutilization	userId and instanceId	The memory usage.
(Agent)load.15m	N/A	load_15m	userId and instanceId	The system load in the last 15 minutes. **Note** This metric is not available for Windows.

(Agent)load.1m	N/A	load_1m	userId and instanceId	The system load in the last one minute. **Note** This metric is not available for Windows.
(Agent)load.5m	N/A	load_5m	userId and instanceId	The system load in the last five minutes. **Note** This metric is not available for Windows.
Host.diskusage.utilization	%	diskusage_utilization	userId, instanceId, and device	The disk space usage.
(Agent)fs.inode.utilization_device	%	fs_inodeutilization	userId, instanceId, and device	The inode utilization. **Note** This metric is available only for Linux.
(Agent)disk.read.bps_device	byte/s	disk_readbytes	userId, instanceId, and device	The number of bytes read from the disk per second.
(Agent)disk.read.iops_device	count/s	disk_readiops	userId, instanceId, and device	The number of read IOPS of the disk.
(Agent)disk.write.bps_device	byte/s	disk_writebytes	userId, instanceId, and device	The number of bytes written to the disk per second.

(Agent)disk.write.iops_device	count/s	disk_writeiops	userId, instanceId, and device	The number of write IOPS of the disk.
(Agent)network.tcp.connection_state	count	net_tcpconnection	userId, instanceId, and state	The number of Transmission Control Protocol (TCP) connections in each state.
(Agent) process.expend_pid	N/A	process.expend	userId, instanceId, name, and pid	Detailed statistical metrics for the top 5 processes.
(Agent)process.count_processname	count	process.number	userId, instanceId, and processName	The number of matched processes.
io saturation	%	disk_io_util	userId, instanceId, and device	The I/O saturation.
Agent (zombia_process_number)	count	zombieProcessNumber	userId and instanceId	The total number of zombie processes.
Agent(total_process_number)	count	totalProcessNumber	userId and instanceId	The total number of processes.
Agent(swap_partition_usage)	%	swapUtilization	userId and instanceId	The usage of the swap partition.
Agent(networkout_drop_ratio)	%	networkout_drop_ratio	userId, instanceId, and device	The loss rate of packets that are sent by network interface controllers (NICs)
Agent(networkout_error_ratio)	%	networkout_error_ratio	userId, instanceId, and device	The sending error rate of NICs.
Agent(networkin_drop_ratio)	%	networkin_drop_ratio	userId, instanceId, and device	The loss rate of packets that are received by NICs.

Agent(networkin_error_ratio)	%	networkin_error_ratio	userId, instanceId, and device	The receiving error rate of NICs.

GPU monitoring metrics

Console-based metric	Unit	API-based metric	Dimension	Description
(Agent)GPU-dimension decoder usage	%	gpu_decoder_utilization	userId, instanceId, and gpuId	The decoder usage of a GPU
(Agent)GPU dimension encoder usage	%	gpu_encoder_utilization	userId, instanceId, and gpuId	The encoder usage of a GPU
(Agent)GPU dimension GPU temperature	°C	gpu_gpu_temperature	userId, instanceId, and gpuId	The GPU temperature.
(Agent)GPU dimension GPU usage	%	gpu_gpu_usedutilization	userId, instanceId, and gpuId	The GPU usage.
(Agent) free memory of GPU dimension	byte	gpu_memory_freespace	userId, instanceId, and gpuId	The idle memory of a GPU
(Agent)GPU dimension video memory idle rate	%	gpu_memory_freeutilization	userId, instanceId, and gpuId	The memory idle rate of a GPU
(Agent)GPU dimension memory usage	byte	gpu_memory_usedspace	userId, instanceId, and gpuId	The used memory of a GPU
(Agent)GPU dimension memory usage	%	gpu_memory_usedutilization	userId, instanceId, and gpuId	The memory usage of a GPU
(Agent)GPU dimension GPU power	W	gpu_power_readings_power_draw	userId, instanceId, and gpuId	The GPU power.

Elastic High Performance Computing (E-HPC)

Console-based metric	Unit	API-based metric	Dimension	Description
created job num	count	createdJobs	Uid, ClusterId, and User	The number of submitted jobs.
finished job num	count	finishedJobs	Uid, ClusterId, and User	The number of completed jobs.
queued job num	count	queuedJobs	Uid, ClusterId, and User	The number of queued jobs.
running job num	count	runningJobs	Uid, ClusterId, and User	The number of running jobs.

Application services

API Gateway

Console-based metric	Unit	API-based metric	Dimension	Description
Latency	ms	Latency	userId and apiUid	The response time.
SumQPS	count	SumQPS	userId and apiUid	The total QPS.
TrafficRX	KB	TrafficRX	userId and apiUid	The inbound traffic.
TrafficTX	KB	TrafficTX	userId and apiUid	The outbound traffic.
code2XX Number	count	code2XX	userId and apiUid	The number of times a 2XX HTTP status code is returned.
code4XX Number	count	code4XX	userId and apiUid	The number of times a 4XX HTTP status code is returned.

code5XX Number	count	code5XX	userId and apiUid	The number of times a 5XX HTTP status code is returned.

Database services

ApsaraDB RDS

By default, an old link is used for ApsaraDB RDS. The period of data collection is 60 to 300 seconds.

ApsaraDB RDS supports the following engines: MySQL, SQL Server, PolarDB O, and PostgreSQL.

RDS MySQL

Console-based metric	Unit	API-based metric	Dimension	Description
ConnectionsUtilization	%	ConnectionUsage	userId and instanceId	The connection utilization.
CPU Utilization	%	CpuUsage	userId and instanceId	The CPU utilization.
DiskUtilization	%	DiskUsage	userId and instanceId	The disk space usage.
IOPSUtilization	%	IOPSUsage	userId and instanceId	The number of I/O requests per second.
MemoryUtilization	%	MemoryUsage	userId and instanceId	The memory usage.
NetworkInRate	bit/s	MySQL_NetworkInNew	userId and instanceId	The inbound traffic to the ApsaraDB RDS for MySQL instance per second.
NetworkOutRate	bit/s	MySQL_NetworkOutNew	userId and instanceId	The outbound traffic from the ApsaraDB RDS for MySQL instance per second.

MySQL_QPS	count/s	MySQL_QPS	userId and instanceId	The number of queries per second (QPS) of the ApsaraDB RDS for MySQL instance.
Docker Cpu Usage	%	DockerCpuUsage	userId and instanceId	The CPU utilization of the Docker container.

RDS SQLServer

Console-based metric	Unit	API-based metric	Dimension	Description
ConnectionsUtilization	%	ConnectionUsage	userId and instanceId	The connection utilization.
CPU Utilization	%	CpuUsage	userId and instanceId	The CPU utilization.
DiskUtilization	%	DiskUsage	userId and instanceId	The disk utilization.
IOPSUtilization	%	IOPSUsage	userId and instanceId	The number of I/O requests per second.
MemoryUtilization	%	MemoryUsage	userId and instanceId	The memory usage.
NetworkInRate	bit/s	SQLServer_NetworkInNew	userId and instanceId	The inbound traffic to the ApsaraDB RDS for SQL Server instance per second.
NetworkOutRate	bit/s	SQLServer_NetworkOutNew	userId and instanceId	The outbound traffic from the ApsaraDB RDS for SQL Server instance per second.

RDS PolarDB O

Console-based metric	Unit	API-based metric	Dimension	Description
ConnectionsUtilization	%	ConnectionUsage	userId and instanceId	The connection utilization.
CPU Utilization	%	CpuUsage	userId and instanceId	The CPU utilization.
DiskUtilization	%	DiskUsage	userId and instanceId	The disk utilization.
IOPSUtilization	%	IOPSUsage	userId and instanceId	The number of I/O requests per second.
MemoryUtilization	%	MemoryUsage	userId and instanceId	The memory usage.

RDS PostgreSQL

- **ApsaraDB RDS for PostgreSQL instances with local SSDs**

Console-based metric	Unit	API-based metric	Dimension	Description
ConnectionsUtilization	%	ConnectionUsage	userId and instanceId	The connection utilization.
CPU Utilization	%	CpuUsage	userId and instanceId	The CPU utilization.
DiskUtilization	%	DiskUsage	userId and instanceId	The disk utilization.
IOPSUtilization	%	IOPSUsage	userId and instanceId	The number of I/O requests per second.
MemoryUtilization	%	MemoryUsage	userId and instanceId	The memory usage.

- **ApsaraDB RDS for PostgreSQL with ESSDs**

Console-based metric	Unit	API-based metric	Dimension	Description
Average active connections per CPU_PG cloud disk	count	ActiveConnections	userId and instanceId	The average number of active connections per CPU.
Inode Usage_PG Disk	%	InodeUsage	userId and instanceId	The Inode usage.
Read-only synchronization delay_PG cloud disk	Seconds	ReadLag	userId and instanceId	The synchronization latency of the read-only instance.
Database Age_PG Disk	xids	DBAge	userId and instanceId	The age of the database.
Read-only Stream replication state_pg cloud disk	value	StreamingStatus	userId and instanceId	The stream replication status of the read-only instance. ◦ 0: disconnected. ◦ 1: normal
Pg_connection usage	%	conn_usage	userId and instanceId	The connection utilization.
PG_CPU usage	%	cpu_usage	userId and instanceId	The CPU utilization.
PG_IOPS usage	%	iops_usage	userId and instanceId	The IOPS usage.
Pg_disk space usage	%	local_fs_size_usage	userId and instanceId	The disk space usage.
PG_Memory Usage	%	mem_usage	userId and instanceId	The memory usage.

ApsaraDB RDS (New)

ApsaraDB RDS also supports a new link. The data collection interval is 60 seconds.

ApsaraDB RDS (New) supports the following engines: MySQL, PolarDB O, and PostgreSQL.

RDS MySQL

Console-based metric	Unit	API-based metric	Dimension	Description
Connection Usage (New)	%	ConnectionUsage_new	userId and instanceId	The connection utilization.
CPU usage (new)	%	CpuUsage_new	userId and instanceId	The CPU utilization.
Disk usage (new)	%	DiskUsage_new	userId and instanceId	The disk space usage.
IOPS usage (new)	%	IOPSUsage_new	userId and instanceId	The number of I/O requests per second.
Memory usage (new)	%	MemoryUsage_new	userId and instanceId	The memory usage.
MySQL Network inflow bandwidth (new)	bit/s	MySQL_NetworkInNew_new	userId and instanceId	The inbound traffic to the ApsaraDB RDS for MySQL instance per second.
MySQL Network outflow bandwidth (new)	bit/s	MySQL_NetworkOutNew_new	userId and instanceId	The outbound traffic from the ApsaraDB RDS for MySQL instance per second.
MySQL queries per second (new)	count/s	MySQL_QPS_new	userId and instanceId	The number of queries per second for the ApsaraDB RDS for MySQL instance.
DataDelay	Seconds	DataDelay	userId and instanceId	The latency between primary and secondary instances.

RWDelay (New)	Seconds	RWDelay	userId, instanceId, and slave	The read/write latency.
ThreadsConnected	count	ThreadsConnected	userId and instanceId	The number of thread connections.
ThreadsRunning	count	ThreadsRunning	userId and instanceId	The number of active threads.
SlowQueriesPs	count	SlowQueriesPs	userId and instanceId	The number of slow queries per second.
Total Wait Time for Row Locks (New)	ms	InnodbRowLockWaits	userId and instanceId	The number of wait times for row locks.
Average Wait Time for Row Locks (New)	ms	InnodbRowLockTimeAvg	userId and instanceId	The average wait time for row locks.
Row Lock Waits (New)	count	InnodbRowLockWaits	userId and instanceId	The number of wait times for row locks.
Current Total Connections (New)	count	TotalConnections	userId and instanceId	The total number of connections.

RDS PolarDB O

Console-based metric	Unit	API-based metric	Dimension	Description
Connection Usage (New)	%	ConnectionUsage_new	userId and instanceId	The connection utilization.
CPU usage (new)	%	CpuUsage_new	userId and instanceId	The CPU utilization.
Disk usage (new)	%	DiskUsage_new	userId and instanceId	The disk utilization.

IOPS usage (new)	%	IOPSUsage_new	userId and instanceId	The number of I/O requests per second.
Memory usage (new)	%	MemoryUsage_new	userId and instanceId	The memory usage.
PPAS Primary/Standby Disconnect	count	PPAS_ReplicationNum	userId and instanceId	The number of disconnections between primary and secondary ApsaraDB RDS for PPAS instances.
PPAS primary/standby delay	Seconds	PPAS_ReplicationLag	userId and instanceId	The latency between primary and secondary ApsaraDB RDS for PPAS instances.

RDS PostgreSQL

- **ApsaraDB RDS for PostgreSQL instances with local SSDs**

Console-based metric	Unit	API-based metric	Dimension	Description
Connection Usage (New)	%	ConnectionUsage_new	userId and instanceId	The connection utilization.
CPU usage (new)	%	CpuUsage_new	userId and instanceId	The CPU utilization.
Disk usage (new)	%	DiskUsage_new	userId and instanceId	The disk utilization.
IOPS usage (new)	%	IOPSUsage_new	userId and instanceId	The number of I/O requests per second.

Memory usage (new)	%	MemoryUsage_new	userId and instanceId	The memory usage.
Number of Primary/Secondary Connections (Value 0 indicates disconnected. Value 1 indicates connected.)	count	PG_ReplicationNum	userId and instanceId	The number of primary/secondary connections for the ApsaraDB RDS for PPAS instance.
PG primary/standby delay	Seconds	PG_ReplicationLag	userId and instanceId	The latency between primary and secondary ApsaraDB RDS for PostgreSQL instances.

- **ApsaraDB RDS for PostgreSQL instances with standard SSDs and ESSDs**

Console-based metric	Unit	API-based metric	Dimension	Description
Average number of active connections per CPU_PG cloud disk (new)	count	ActiveConnections_new	userId and instanceId	The average number of active connections per CPU.
Pg_connection usage	%	conn_usage	userId and instanceId	The connection utilization.
PG_CPU usage	%	cpu_usage	userId and instanceId	The CPU utilization.
PG_IOPS usage	%	iops_usage	userId and instanceId	The IOPS usage.
INODE usage_PG cloud disk (new)	%	InodeUsage_new	userId and instanceId	The Inode usage.
Pg_disk space usage	%	local_fs_size_usage	userId and instanceId	The disk space usage.
PG_Memory Usage	%	mem_usage	userId and instanceId	The memory usage.

Read-only synchronization delay_PG cloud disk (new)	Seconds	ReadLag_new	userId and instanceId	The synchronization latency of the read-only instance.
Database age_PG cloud disk (new)	xids	DBAge_new	userId and instanceId	The age of the database.
Read-only Stream Copy Status_PG Cloud Disk (New)	value	StreamingStatus_new	userId and instanceId	The streaming replication status of read-only instances
PG Disk_Number of Primary/Secondary Connections	count	PG_ESSD_ReplicationNum	userId and instanceId	The number of primary/secondary connections for the ApsaraDB RDS for PPAS instance.
PG cloud disk_primary/standby delay	Seconds	PG_ESSD_ReplicationLag	userId and instanceId	The latency between primary and secondary ApsaraDB RDS for PostgreSQL instances.

PolarDB-X V2.0

Compute node metrics

Console-based metric	Unit	API-based metric	Dimension	Description
CPU usage	%	CPUUtilization	userId and instanceId	The CPU utilization.
Total Sessions	count	Connectioncount	userId and instanceId	The total number of sessions.
Requests per second	count	LogicQPS	userId and instanceId	The number of requests per second.

Average Logical Response Time	ms	LogicRT	userId and instanceId	The average response time of logical requests.
Memory Utilization	%	MemoryUtilization	userId and instanceId	The memory usage.
DRDS network input bandwidth	bit/s	NetworkInputTraffic	userId and instanceId	The inbound bandwidth.
DRDS_network output bandwidth	bit/s	NetworkOutputTraffic	userId and instanceId	The outbound bandwidth.
Physical QPS	count	PhysicsQPS	userId and instanceId	The physical QPS.
Physical RT	ms	PhysicsRT	userId and instanceId	The physical response time.
Number of active threads	count	Threadcount	userId and instanceId	The number of active threads.
Heap Memory Usage	%	HeapMemoryUsage	userId and instanceId	The heap memory usage.
Log directory size	byte	LogSize	userId and instanceId	The size of the log directory.
Errors per Second	count	SecondErrorcount	userId and instanceId	The number of errors per second.
Physical slow SQL per second	count	SecondPhysicalSqlcount	userId and instanceId	The number of physical slow SQL queries per second.
Slow requests per second	count	SecondRequestcount	userId and instanceId	The number of slow requests per second.
spill Directory Size	byte	SpillSize	userId and instanceId	The size of the spill directory.
Temporary Directory Size	byte	TmpSize	userId and instanceId	The size of the temporary directory.

Data node metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Connection Usage (New)	%	ConnectionUsage_new	userId and instanceId	The connection utilization.
CPU usage (new)	%	CpuUsage_new	userId and instanceId	The CPU utilization.
Disk usage (new)	%	DiskUsage_new	userId and instanceId	The disk space usage.
IOPS usage (new)	%	IOPSUsage_new	userId and instanceId	The number of I/O requests per second.
Memory usage (new)	%	MemoryUsage_new	userId and instanceId	The memory usage.
MySQL Network inflow bandwidth (new)	bit/s	MySQL_NetworkInNew_new	userId and instanceId	The inbound traffic to the MySQL data node per second.
MySQL Network outflow bandwidth (new)	bit/s	MySQL_NetworkOutNew_new	userId and instanceId	The outbound traffic from the MySQL data node per second.
MySQL queries per second (new)	count/s	MySQL_QPS_new	userId and instanceId	The number of queries per second on the MySQL data node.

PolarDB-X 1.0

Console-based metric	Unit	API-based metric	Dimension	Description
CPU Utilization	%	CPUUtilization	useId and InstanceId	The CPU utilization.
Memory Usage	%	MemoryUtilization	useId and InstanceId	The memory usage of the RDS instance.

Network input	byte/s	NetworkInputTraffic	useId and InstanceId	The number of network inbound traffic.
Network output	byte/s	NetworkOutputTraffic	useId and InstanceId	The number of network outbound traffic.
Average number of logical requests per second	count	LogicQPS	useId and InstanceId	The average number of logical requests per second.
Average Logical Response Time	ms	LogicRT	useId and InstanceId	The average response time of logical requests.
Average number of physical requests per second	count	PhysicsQPS	useId and InstanceId	The average number of physical requests per second.
Average Physical Response Time	ms	PhysicsRT	useId and InstanceId	The average response time of physical requests.
Total sessions	count	Connectioncount	useId and InstanceId	The total number of sessions.
Number of active sessions	count	Threadcount	useId and InstanceId	The number of active sessions.

ApsaraDB for MongoDB

Basic metrics

Console-based metric	Unit	API-based metric	Dimension	Description
CPU Utilization	%	CPUUtilization	userId, instanceId, and role	The CPU utilization.
Memory Usage	%	MemoryUtilization	userId, instanceId, and role	The memory usage.
Disk Usage	%	DiskUtilization	userId, instanceId, and role	The disk space utilization.

IOPS Usage	%	IOPSUtilization	userId, instanceId, and role	The IOPS usage.
Connection Usage	%	ConnectionUtilization	userId, instanceId, and role	The connection usage.
QPS	count	QPS	userId, instanceId, and role	The number of requests per second.
Connection Usage	count	ConnectionAmount	userId, instanceId, and role	The number of used connections.

Disk capacity metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Disk Size Occupied by Instances	byte	InstanceDiskAmount	userId, instanceId, and role	The disk space occupied by the instance.
Disk Size Occupied by Data	byte	DataDiskAmount	userId, instanceId, and role	The disk space occupied by data.
Disk Size Occupied by Logs	byte	LogDiskAmount	userId, instanceId, and role	The disk space occupied by logs.

Network request metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Cache Inbound Bandwidth	byte	IntranetIn	userId, instanceId, and role	The inbound traffic over the intranet.
Cache Outbound Bandwidth	byte	IntranetOut	userId, instanceId, and role	The outbound traffic over the intranet.

Number of Requests	count	NumberRequests	userId, instanceId, and role	The number of requests.

Operation count metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Number of Command Operations	count	OpCommand	userId, instanceId, and role	The number of Command operations.
Number of Delete Operations	count	OpDelete	userId, instanceId, and role	The number of Delete operations.
Number of Getmore Operations	count	OpGetmore	userId, instanceId, and role	The number of Getmore operations.
Number of Insert Operations	count	OpInsert	userId, instanceId, and role	The number of Insert operations.
Number of Query Operations	count	OpQuery	userId, instanceId, and role	The number of Query operations.
Number of Update Operations	count	OpUpdate	userId, instanceId, and role	The number of Update operations.

ApsaraDB for MongoDB (New)

Basic metrics

Console-based metric	Unit	API-based metric	Dimension	Description
CPUUtilization_new	%	CPUUtilization_new	userId, instanceId, and role	The CPU utilization.

Memory usage (new)	%	MemoryUtilization_new	userId, instanceId, and role	The memory usage.
Disk Usage (New)	%	DiskUtilization_new	userId, instanceId, and role	The disk space utilization.
IOPS usage (new)	%	IOPSUtilization_new	userId, instanceId, and role	The IOPS usage.
Connection Usage (New)	%	ConnectionUtilization_new	userId, instanceId, and role	The connection usage.
Visits per second (new)	count	QPS_new	userId, instanceId, and role	The number of requests per second.
Connection Usage (New)	count	ConnectionAmount_new	userId, instanceId, and role	The number of used connections.
Current Total Connections (New)	count	TotalConnections	userId, instanceId, and role	The total number of connections.

Disk capacity metrics

Console-based metric	Unit	API-based metric	Dimension	Description
The amount of disk space occupied by the instance (new)	byte	InstanceDiskAmount_new	userId, instanceId, and role	The disk space occupied by the instance.
Amount of disk space occupied data (new)	byte	DataDiskAmount_new	userId, instanceId, and role	The disk space occupied by data.
The amount of disk space occupied by the log (new)	byte	LogDiskAmount_new	userId, instanceId, and role	The disk space occupied by logs.

Network request metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Intranet Inflow (New)	byte	IntranetIn_new	userId, instanceId, and role	The inbound traffic over the intranet.
Intranet outbound traffic (new)	byte	IntranetOut_new	userId, instanceId, and role	The outbound traffic over the intranet.
Number of requests (new)	count	NumberRequests_new	userId, instanceId, and role	The number of requests.

Operation count metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Number of Command Operations (new)	count	OpCommand_new	userId, instanceId, and role	The number of Command operations.
Number of Delete operations (new)	count	OpDelete_new	userId, instanceId, and role	The number of Delete operations.
Number of Getmore operations (new)	count	OpGetmore_new	userId, instanceId, and role	The number of Getmore operations.
Number of Insert operations (new)	count	OpInsert_new	userId, instanceId, and role	The number of Insert operations.
Number of Query operations (new)	count	OpQuery_new	userId, instanceId, and role	The number of Query operations.

Number of Update operations (new)	count	OpUpdate_new	userId, instanceId, and role	The number of Update operations.

Basic metrics (Cluster Edition)

Console-based metric	Unit	API-based metric	Dimension	Description
(Cluster) CPU usage	%	CPUUtilization_new_cluster	userId, instanceId, and role	The CPU utilization.
(Cluster Version) Memory Usage	%	MemoryUtilization_new_cluster	userId, instanceId, and role	The memory usage.
(Cluster Edition) Disk Usage	%	DiskUtilization_new_cluster	userId, instanceId, and role	The disk space usage.
(Cluster version) IOPS usage	%	IOPSUtilization_new_cluster	userId, instanceId, and role	The IOPS usage.
(Cluster Edition) Connection Usage	%	ConnectionUtilization_new_cluster	userId, instanceId, and role	The connection utilization.
(Cluster Edition) Visits per Second	count	QPS_new_cluster	userId, instanceId, and role	The number of visits per second.
(Cluster Edition) Connection Usage	count	ConnectionAmount_new_cluster	userId, instanceId, and role	The number of used connections.
(Cluster Edition) Total Number of Connections	count	TotalConnections_cluster	userId, instanceId, and role	The total number of connections.

Disk capacity metrics (Cluster Edition)

Console-based metric	Unit	API-based metric	Dimension	Description
(Cluster Edition) The amount of disk space occupied by the instance	bytes	InstanceDiskAmount_new_cluster	userId, instanceId, and role	The disk space occupied by the instance.

(Cluster Edition) The amount of disk space occupied by data	bytes	DataDiskAmount_new_cluster	userId, instanceId, and role	The disk space occupied by data.
(Cluster Edition) The amount of disk space occupied by logs	bytes	LogDiskAmount_new_cluster	userId, instanceId, and role	The disk space occupied by logs.

Network request metrics (Cluster Edition)

Console-based metric	Unit	API-based metric	Dimension	Description
(Cluster Version) Intranet Inflow	bytes	IntranetIn_new_cluster	userId, instanceId, and role	The inbound traffic over the intranet.
(Cluster Version) Internal outbound traffic	bytes	IntranetOut_new_cluster	userId, instanceId, and role	The outbound traffic over the intranet.
(Cluster Version) Number of Requests	count	NumberRequests_new_cluster	userId, instanceId, and role	The number of requests.

Operation count metrics (Cluster Edition)

Console-based metric	Unit	API-based metric	Dimension	Description
(Cluster Edition) Command Operation Times	count	OpCommand_new_cluster	userId, instanceId, and role	The number of Command operations.
(Cluster version) Delete operations	count	OpDelete_new_cluster	userId, instanceId, and role	The number of Delete operations.
(Cluster version) Getmore Operations	count	OpGetmore_new_cluster	userId, instanceId, and role	The number of Getmore operations.
(Cluster version) Insert operations	count	OpInsert_new_cluster	userId, instanceId, and role	The number of Insert operations.

(Cluster Version) Query Operation Times	count	OpQuery_new_cluster	userId, instanceId, and role	The number of Query operations.
(Cluster Edition) Update Operation Times	count	OpUpdate_new_cluster	userId, instanceId, and role	The number of Update operations.

ApsaraDB for Redis

Console-based metric	Unit	API-based metric	Dimension	Description
Connections Usage	%	ConnectionUsage	userId and instanceId	The connection usage.
CPU Utilization	%	CpuUsage	userId and instanceId	The CPU utilization.
Failed Operations	count/s	Failedcount	userId and instanceId	The number of failed operations.
Cache Inbound Bandwidth	byte/s	IntranetIn	userId and instanceId	The inbound traffic over the intranet.
Inbound Bandwidth Usage	%	IntranetInRatio	userId and instanceId	The inbound bandwidth usage.
Cache Outbound Bandwidth	byte/s	IntranetOut	userId and instanceId	The outbound traffic over the intranet.
Outbound Bandwidth Usage	%	IntranetOutRatio	userId and instanceId	The outbound bandwidth usage.
Memory Usage	%	MemoryUsage	userId and instanceId	The memory usage.
Used QPS	count/S	UsedQPS	userId and instanceId	The number of used QPS.

ApsaraDB for Redis (New)

Standard Edition (read/write splitting)

Console-based metric	Unit	API-based metric	Dimension	Description
Usage of Connections (Standard Edition)	%	StandardConnectionUsage	userId and instanceId	The connection utilization for the ApsaraDB for Redis instance of Standard Edition.
Inflow Bandwidth (Standard Edition)	KB/s	StandardIntranetIn	userId and instanceId	The inbound bandwidth to the ApsaraDB for Redis instance of Standard Edition.
Inflow Bandwidth Usage (Standard Edition)	%	StandardIntranetInRatio	userId and instanceId	The usage of the inbound traffic to the ApsaraDB for Redis instance of Standard Edition.
Outflow Bandwidth (Standard Edition)	KB/s	StandardIntranetOut	userId and instanceId	The outbound bandwidth from the ApsaraDB for Redis instance of Standard Edition.
Outflow Bandwidth Usage (Standard Edition)	%	StandardIntranetOutRatio	userId and instanceId	The usage of the outbound traffic from the ApsaraDB for Redis instance of Standard Edition.
Memory Usage (Standard Edition)	%	StandardMemoryUsage	userId and instanceId	The memory usage of the ApsaraDB for Redis instance of Standard Edition.
QPS usage rate (Standard version)	%	StandardQPSUsage	userId and instanceId	The QPS usage of the ApsaraDB for Redis instance of Standard Edition.
CPU usage (standard version)	%	StandardCpuUsage	userId and instanceId	The CPU utilization of the ApsaraDB for Redis instance of Standard Edition.

Average visits per second (standard version)	count/S	StandardUsedQPS	userId and instanceId	The average number of visits per second on the ApsaraDB for Redis instance of Standard Edition.
Total QPS (Standard Edition)	count/S	StandardTotalQps	userId and instanceId	The total QPS of the ApsaraDB for Redis instance of Standard Edition.
Current Total Connections (Standard Edition)	count	StandardTotalConnections	userId and instanceId	The total number of connections of the ApsaraDB for Redis instance of Standard Edition.

Cluster Edition

Console-based metric	Unit	API-based metric	Dimension	Description
Usage of Connections (Cluster Edition)	%	ShardingConnectionUsage	userId and instanceId	The number of connections used by the ApsaraDB for Redis instance of Cluster Edition.
Inflow Bandwidth (Cluster Version)	KB/s	ShardingIntranetIn	userId and instanceId	The inbound bandwidth to the ApsaraDB for Redis instance of Cluster Edition.
Outflow Bandwidth (Cluster Version)	KB/s	ShardingIntranetOut	userId and instanceId	The outbound bandwidth from the ApsaraDB for Redis instance of Cluster Edition.
Inflow Bandwidth Usage (Cluster Version)	%	ShardingIntranetInRatio	userId and instanceId	The usage of the inbound bandwidth to the ApsaraDB for Redis instance of Cluster Edition.
Outflow Bandwidth Usage (Cluster Version)	%	ShardingIntranetOutRatio	userId and instanceId	The usage of the outbound bandwidth from the ApsaraDB for Redis instance of Cluster Edition.

Memory usage (cluster version)	%	ShardingMemoryUsage	userId and instanceId	The memory usage of the ApsaraDB for Redis instance of Cluster Edition.
QPS usage (cluster version)	%	ShardingQPSUsage	userId and instanceId	The QPS usage of the ApsaraDB for Redis instance of Cluster Edition.
CPU usage (cluster version)	%	ShardingCpuUsage	userId and instanceId	The CPU utilization of the ApsaraDB for Redis instance of Cluster Edition.
Average accesses per second (cluster version)	count/S	ShardingUsedQPS	userId and instanceId	The average number of visits per second on the ApsaraDB for Redis instance of Cluster Edition.
Total QPS (Cluster Edition)	count/S	ShardingTotalQps	userId and instanceId	The total QPS of the ApsaraDB for Redis instance of Cluster Edition.
Total Current Connections (Cluster Edition)	count	ShardingTotalConnections	userId and instanceId	The total number of connections used by the ApsaraDB for Redis instance of Cluster Edition.

Cluster Edition - Node

Console-based metric	Unit	API-based metric	Dimension	Description
Connections Usage (Node)	%	NodeConnectionUsage	userId, instanceId, and nodeId	The connection usage.
Inflow bandwidth (node)	KB/s	NodeIntranetIn	userId, instanceId, and nodeId	The inbound bandwidth.
Outflow bandwidth (node)	KB/s	NodeIntranetOut	userId, instanceId, and nodeId	The outbound bandwidth.

Inflow Bandwidth Usage (Node)	%	NodeIntranetInRatio	userId, instanceId, and nodeId	The usage of the inbound bandwidth.
Outflow bandwidth usage (nNode)	%	NodeIntranetOutRatio	userId, instanceId, and nodeId	The usage of the outbound bandwidth.
Memory usage (node)	%	NodeMemoryUsage	userId, instanceId, and nodeId	The memory usage.
QPS usage (node)	%	NodeQPSUsage	userId, instanceId, and nodeId	The QPS usage.
CPU usage (node)	%	NodeCpuUsage	userId, instanceId, and nodeId	The CPU utilization.
Average access per second (node)	count	NodeUsedQPS	userId, instanceId, and nodeId	The average number of visits per second.
Total QPS (node)	count/S	NodeTotalQps	userId, instanceId, and nodeId	The total QPS.
Total Current Connections (Nodes)	count	NodeTotalConnections	userId, instanceId, and nodeId	The total number of connections.

AnalyticDB for PostgreSQL

Computing-storage integration metrics

Console-based metric	Unit	API-based metric	Dimension	Description
[Computing and Storage Integration] Connection Usage Rate	%	ConnectionUsage	userId, instanceId, and role	The connection usage.
[Integration of Computing and Storage] CPU Usage	%	CpuUsage	userId, instanceId, and role	The CPU utilization.
[Computing and Storage Integration] Disk Usage	%	DiskUsage	userId, instanceId, and role	The disk space usage.

[Integration of computing and storage] IO throughput usage	%	IOPSUsage	userId, instanceId, and role	The IOPS utilization.
[Computing and Storage Integration] Memory Usage	%	MemoryUsage	userId, instanceId, and role	The memory usage.

Storage reservation metrics

Console-based metric	Unit	API-based metric	Dimension	Description
[Storage Reservation] Coordinate the number of node connections	count	adbpg_conn_count	instance_id, instance_component, and hostname	The number of connections of the coordinator node.
[Storage Reservation] Node Disk Usage	MB	node_disk_used_mb	instance_id, instance_component, and hostname	The disk usage of the node.
node_cpu_used_percent	%	node_cpu_used_percent	instance_id, instance_component, and hostname	The CPU utilization of the node.
node_mem_used_percent	%	node_mem_used_percent	instance_id, instance_component, and hostname	The memory usage of the node.
node_disk_iops	count	node_disk_iops	instance_id, instance_component, and hostname	The IOPS of the node.
node_disk_used_percent	%	node_disk_used_percent	instance_id, instance_component, and hostname	The disk usage of the node.

[Storage Reservation] Instance Disk Usage	MB	adbpg_instance_disk_used_mb	instance_id	The disk usage of the instance.
[Storage Reservation] Instance Disk Usage	%	adbpg_instance_disk_usage_percent	instance_id	The disk usage of the instance.
[Storage Reservation] Number of Abnormal Coordinating Nodes	count	adbpg_master_cnt_unhealth	instance_id	The number of abnormal coordinator nodes.
adbpg_segment_cnt_unhealth	count	adbpg_segment_cnt_unhealth	instance_id	The number of abnormal compute nodes.
[Storage Reservation] Number of nodes that coordinate nodes that are out-of-sync	count	adbpg_master_cnt_repl_not_sync	instance_id	The number of coordinator nodes that are in the out-of-sync state.
[Storage Reservation] Number of nodes in an out-of-sync state for computing nodes	count	adbpg_segment_cnt_repl_not_sync	instance_id	The number of compute nodes that are in the out-of-sync state.
adbpg_query_blocked	count	adbpg_query_blocked	instance_id	The number of blocked queries on the instance.
adbpg_query_queued	count	adbpg_query_queued	instance_id	The number of queued queries on the instance.
Storage Reservation Instance Active Connections	count	adbpg_conn_active	instance_id	The number of active connections on the instance.
[Storage Reservation] Number of active queries for instances	count	adbpg_query_active	instance_id	The number of active queries on the instance.

adbpg_conn_idle	count	adbpg_conn_idle	instance_id	The number of idle connections on the instance.

AnalyticDB for PostgreSQL (new)

Computing and storage integration metrics (master node)

Console-based metric	Unit	API-based metric	Dimension	Description
[Computing and Storage Integration] Primary Node Connection Number Usage (New)	%	ConnectionUsageMaster	userId, instanceId, and role	The usage of primary node connections.
[Integration of computing and storage] master node CPU usage (new)	%	CpuUsageMaster	userId, instanceId, and role	The CPU utilization of the primary node.
[Computing and Storage Integration] Primary Node Disk Usage (New)	%	DiskUsageMaster	userId, instanceId, and role	The disk space usage of the primary node.
[Integration of computing and storage] master node IO throughput usage (new)	%	IOPSUsageMaster	userId, instanceId, and role	The IOPS usage of the primary node.
[Computing and Storage Integration] Master Node Memory Usage (New)	%	MemoryUsageMaster	userId, instanceId, and role	The memory usage of the primary node.

Computing and storage integration metrics (compute node group)

Console-based metric	Unit	API-based metric	Dimension	Description

[Integration of computing and storage] CPU usage of computing group nodes (new)	%	CpuUsageSegment	userId, instanceId, and role	The CPU utilization of the nodes in the compute group.
[Computing and Storage Integration] Computing Group Node Disk Usage (New)	%	DiskUsageSegment	userId, instanceId, and role	The disk space usage of the nodes in the compute group.
[Integration of computing and storage] I/O throughput usage of computing group nodes (new)	%	IOPSUsageSegment	userId, instanceId, and role	The IOPS usage of the nodes in the compute group.
[Computing and Sstorage Integration] Computing Group Node Memory Usage (New)	%	MemoryUsageSegment	userId, instanceId, and role	The memory usage of the nodes in the compute group.

AnalyticDB for MySQL V3.0

Console-based metric	Unit	API-based metric	Dimension	Description
Connections	count	connections	userId and instanceId	The rate of inbound packets of the connection.
AvgCpuUsed	%	worker_avg_cpu_used	userId and instanceId	The CPU utilization of the node.
executorAvgCpuUsed	%	executor_avg_cpu_used	userId and instanceId	The CPU utilization of the executor.
MaxNodeDiskUsedPercent	%	worker_max_node_disk_used_percent	userId and instanceId	The maximum disk utilization of the node in the reserved mode.

DiskUsedPercent	%	disk_used_percent	userId and instanceId	The disk utilization in the reserved mode.
MaxNodeDiskUsage	MB	worker_max_node_disk_used_MB	userId and instanceId	The maximum disk usage of the node in the elastic mode.
ColdDataDiskUsage	MB	ColdDataDiskUsage	userId and instanceId	The disk space occupied by cold data in the elastic mode.
HotDataDiskUsage	MB	HotDataDiskUsage	userId and instanceId	The disk space occupied by hot data in the elastic mode.
workerReadbytes	MB	worker_avg_read_bytes	userId and instanceId	The read throughput of the worker.
workerWritebytes	MB	worker_avg_write_bytes	userId and instanceId	The write throughput of the worker.
executorReadbytes	MB	executor_avg_read_bytes_MB	userId and instanceId	The read throughput of the executor.
executorWritebytes	MB	executor_avg_write_bytes_MB	userId and instanceId	The write throughput of the executor.
workerReadIOPS	count/s	worker_avg_read_iops	userId and instanceId	The read IOPS of the worker.
workerWriteIOPS	count/s	worker_avg_write_iops	userId and instanceId	The write IOPS of the worker.
executorReadIOPS	count/s	executor_avg_read_iops	userId and instanceId	The read IOPS of the executor.
executorWriteIOPS	count/s	executor_avg_write_iops	userId and instanceId	The write IOPS of the executor.

workerIOUtil	%	worker_avg_io_util	userId and instanceId	The disk I/O usage.
QPS	count/s	qps	userId and instanceId	The number of requests per second.
QueryAvgRt	ms	query_avg_rt	userId and instanceId	The average query time.
QueryMaRt	ms	query_max_rt	userId and instanceId	The maximum query time.
QueryMaxWaitTime	ms	query_total_wait_time	userId and instanceId	The average query waiting time.
tps	count/s	tps	userId and instanceId	The number of transactions per second.
InsertAvgRt	ms	insert_avg_rt	userId and instanceId	The average write response time.
InsertMaxRt	ms	insert_max_rt	userId and instanceId	The maximum write response time.
InsertInbytes	MB	insert_in_bytes	userId and instanceId	The write throughput.

Data Transmission Service (DTS)

Console-based metric	Unit	API-based metric	Dimension	Description
ChangeTrackingLatency	Seconds	ChangeTrackingLatency	userId and instanceId	The latency of the subscription task.
SynchronizationLatency	Seconds	SynchronizationLatency	userId and instanceId	The latency of the synchronization task.

MigrationLatency	Seconds	MigrationLatency	userId and instanceId	The latency of the migration task.

Storage

Cloud Storage Gateway (CSG)

Console-based metric	Unit	API-based metric	Dimension	Description
Gateway CPU usage percentage	%	csg_gateway_cpu_usage_percentage	userId and instanceId	The CPU utilization of the gateway.
Gateway memory usage percentage	%	csg_gateway_memory_usage_percentage	userId and instanceId	The memory usage of the gateway.
Gateway cache usage percentage	%	csg_gateway_cache_usage_percentage	userId, instanceId, and shareName	The cache space usage of the gateway.
Gateway running state	count	csg_gateway_running_state	userId, instanceId, and shareName	The status of the gateway.
Gateway share Upload bucket queue length	count	csg_gateway_file_share_upload_queue_num	userId, instanceId, and shareName	The queue length for upload from share to bucket.
Gateway file share remaining meta space	byte	csg_gateway_file_share_remaining_meta_space_bytes	userId, instanceId, and shareName	The remaining metadata space of the gateway share.
Gateway file share throttling state	count	csg_gateway_file_share_throttling_state	userId, instanceId, and shareName	The throttling status of the gateway share.
Gateway shared read/write rate	KBps	csg_gateway_file_share_read_write_rate_KBps	userId, instanceId, and shareName	The read and write rate of the gateway share.
Gateway file share read iops	KBps	csg_gateway_file_share_read_IOPS	userId, instanceId, and shareName	The read IOPS of the gateway share.
Gateway file share write iops	KBps	csg_gateway_file_share_write_IOPS	userId, instanceId, and shareName	The write IOPS of the gateway share.

Gateway file share upload OSS rate	KBps	csg_gateway_file_share_upload_rate_KBps	userId, instanceId, and shareName	The rate at which the gateway share uploads data to an OSS bucket.

Apsara File Storage NAS (NAS)

File system metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Aligned Size	byte	AlignedSize	userId and fileSystemId	The data volume of general-purpose NAS file systems (excluding IA storage media).
Storage Usage of IA Storage Medium	byte	SecondaryAlignedSize	userId and fileSystemId	The data volume of IA storage media.
File count	byte	Filecount	userId and fileSystemId	The number of files.
Read IOPS	op/s	IopsRead	userId and fileSystemId	The average read IOPS per minute.
Write IOPS	op/s	IopsWrite	userId and fileSystemId	The average write IOPS per minute.
Read Latency	ms	LatencyRead	userId and fileSystemId	The average read latency per minute.
Write Latency	ms	LatencyWrite	userId and fileSystemId	The average write latency per minute.
Metadata QPS	op/s	QpsMeta	userId and fileSystemId	The average QPS of metadata per minute.
Metadata delay	ms	LatencyMeta	userId and fileSystemId	The average latency of metadata.
Read Throughput	byte/s	ThruputRead	userId and fileSystemId	The average read throughput per minute.

Write Throughput	byte/s	ThruputWrite	userId and fileSystemId	The average write throughput per minute.
Maximum usable capacity	byte	Quota	userId and fileSystemId	The maximum available capacity.
Maximum number of files	count	MaxFilecount	userId and fileSystemId	The maximum number of files.

Object Storage Service (OSS)

Besic metrics

Console-based metric	Unit	API-based metric	Dimension	Description
RequestValidRate	%	RequestValidRate	userId and BucketName	The rate of valid requests.
ValidRequestcount	count	ValidRequestcount	userId and BucketName	The number of valid requests.
TotalRequestcount	count	TotalRequestcount	userId and BucketName	The total number of requests.
Availability	%	Availability	userId and BucketName	The availability. Stay alert if this metric is lower than 100%.
InternetSend	byte	InternetSend	userId and BucketName	The outbound traffic over the Internet.
InternetRecv	byte	InternetRecv	userId and BucketName	The inbound traffic over the Internet.
IntranetSend	byte	IntranetSend	userId and BucketName	The outbound traffic over the intranet.
IntranetRecv	byte	IntranetRecv	userId and BucketName	The inbound traffic over the intranet.

CDNSend	byte	CdnSend	userId and BucketName	The CDN outbound traffic.
CDNRecv	byte	CdnRecv	userId and BucketName	The CDN inbound traffic.
SyncSend	byte	SyncSend	userId and BucketName	The cross-region replication outbound traffic.
SyncRecv	byte	SyncRecv	userId and BucketName	The cross-region replication inbound traffic.
MeteringStorageUtilization	byte	MeteringStorageUtilization	userId and BucketName	The currently used storage.
Storage Usage	%	MeteringStorageUtilizationPercent	userId and BucketName	The storage usage of the bucket.

Request status metrics

Console-based metric	Unit	API-based metric	Dimension	Description
ServerErrorcount	count	ServerErrorcount	userId and BucketName	The number of requests with a server error.
ServerErrorRate	%	ServerErrorRate	userId and BucketName	The ratio of requests with a server error.
NetworkErrorcount	count	NetworkErrorcount	userId and BucketName	The number of requests with a network error.
NetworkErrorRate	%	NetworkErrorRate	userId and BucketName	The ratio of requests with a network error.

AuthorizationErrorcount	count	AuthorizationErrorcount	userId and BucketName	The number of requests with an authorization error.
ResourceNotFoundErrorcount	count	ResourceNotFoundErrorcount	userId and BucketName	The number of requests with an error that indicates resources cannot be found.
Client Timeout Error Request Proportion	count	ClientTimeoutErrorcount	userId and BucketName	The number of requests with a client timeout error.
ClientOtherErrorcount	count	ClientOtherErrorcount	userId and BucketName	The number of requests with other client errors.
AuthorizationErrorRate	%	AuthorizationErrorRate	userId and BucketName	The ratio of requests with an authorization error.
ResourceNotFoundErrorRate	%	ResourceNotFoundErrorRate	userId and BucketName	The ratio of requests with an error that indicates resources cannot be found.
Client Timeout Error Request Proportion by User	%	ClientTimeoutErrorRate	userId and BucketName	The ratio of requests with a client timeout error.
Client Other Error Request Proportion	%	ClientOtherErrorRate	userId and BucketName	The ratio of requests with other client errors.
Successcount	count	Successcount	userId and BucketName	The number of successful requests.
Redirectcount	count	Redirectcount	userId and BucketName	The number of redirected requests.

SuccessRate	%	SuccessRate	userId and BucketName	The ratio of successful requests.
RedirectRate	%	RedirectRate	userId and BucketName	The ratio of redirected requests.

Maximum latency metrics

Console-based metric	Unit	API-based metric	Dimension	Description
MaxGetObjectE2eLatency	ms	MaxGetObjectE2eLatency	userId and BucketName	The maximum E2E latency of GetObject requests.
MaxGetObjectServerLatency	ms	MaxGetObjectServerLatency	userId and BucketName	The maximum server latency of GetObject requests.
MaxHeadObjectE2eLatency	ms	MaxHeadObjectE2eLatency	userId and BucketName	The maximum E2E latency of HeadObject requests.
MaxHeadObjectServerLatency	ms	MaxHeadObjectServerLatency	userId and BucketName	The maximum server latency of HeadObject requests.
MaxPutObjectE2eLatency	ms	MaxPutObjectE2eLatency	userId and BucketName	The maximum E2E latency of PutObject requests.
MaxPutObjectServerLatency	ms	MaxPutObjectServerLatency	userId and BucketName	The maximum server latency of PutObject requests.
MaxAppendObjectE2eLatency	ms	MaxAppendObjectE2eLatency	userId and BucketName	The maximum E2E latency of AppendObject requests.

MaxAppendObjectServerLatency	ms	MaxAppendObjectServerLatency	userId and BucketName	The maximum server latency of AppendObject requests.
MaxUploadPartE2eLatency	ms	MaxUploadPartE2eLatency	userId and BucketName	The maximum E2E latency of UploadPart requests.
MaxUploadPartServerLatency	ms	MaxUploadPartServerLatency	userId and BucketName	The maximum server latency of UploadPart requests.
MaxUploadPartCopyE2eLatency	ms	MaxUploadPartCopyE2eLatency	userId and BucketName	The maximum E2E latency of UploadPartCopy requests.
MaxUploadPartCopyServerLatency	ms	MaxUploadPartCopyServerLatency	userId and BucketName	The maximum server latency of UploadPartCopy requests.

Metrics of successful requests

Console-based metric	Unit	API-based metric	Dimension	Description
GetObjectcount	Frequency	GetObjectcount	userId and BucketName	The number of successful GetObject requests.
HeadObjectcount	count	HeadObjectcount	userId and BucketName	The number of successful HeadObject requests.
PostObjectcount	count	PostObjectcount	userId and BucketName	The number of successful PostObject requests.

AppendObjectcount	count	AppendObjectcount	userId and BucketName	The number of successful AppendObject requests.
UploadPartcount	count	UploadPartcount	userId and BucketName	The number of successful UploadPart requests.
UploadPartCopycount	count	UploadPartCopycount	userId and BucketName	The number of successful UploadPartCopy requests.
DeleteObjectcount	count	DeleteObjectcount	userId and BucketName	The number of successful DeleteObject requests.
DeleteObjectscount	count	DeleteObjectscount	userId and BucketName	The number of successful DeleteObjects requests.

Tablestore (OTS)

Console-based metric	Unit	API-based metric	Dimension	Description
Instance VCU	count/s	InstanceVCU	userId and instanceName	The VCUs of the instance.
Reserved VCU	count/s	InstanceReservedVCU	userId and instanceName	The reserved VCUs of the instance.
Internal Read CU	count/s	InstanceFreeReadCU_NoOp	userId and instanceName	The internal read CUs.
Internal Write CU	count/s	InstanceFreeWriteCU_NoOp	userId and instanceName	The internal write CUs.
Over ReadCU	count/s	InstanceOverReadCU_NoOp	userId and instanceName	The excess read CUs of the instance.

Over WriteCU	count/s	InstanceOverWriteCU_NoOp	userId and instanceName	The excess write CUs of the instance.
Read CU	count/s	InstanceReadCU_NoOp	userId and instanceName	The operation read CUs.
Internal Write CU	count/s	InstanceWriteCU_NoOp	userId and instanceName	The operation write CUs.
Total Request	count/s	InstanceRequestNumber_NoOp	userId and instanceName	The total number of requests.
2xx Request	count/s	Instancecount2xxNumber_NoOp	userId and instanceName	The number of requests that return a 2XX HTTP status code.
429 Request	count/s	Instancecount429Number_NoOp	userId and instanceName	The number of requests that return HTTP status code 429.
4xx Request	count/s	Instancecount4xxNumber_NoOp	userId and instanceName	The number of requests that return a 4XX HTTP status code.
5xx Request	count/s	Instancecount5xxNumber_NoOp	userId and instanceName	The number of requests that return a 5XX HTTP status code.
Internet Inbound	byte/s	InstanceExtranetIn_NoOp	userId and instanceName	The inbound traffic over the Internet.
Internet Outbound	byte/s	InstanceExtranetOut_NoOp	userId and instanceName	The outbound traffic over the Internet.
Free Outbound Internet	byte/s	InstanceFreeExtranetOut_NoOp	userId and instanceName	The free outbound traffic over the Internet.

Total Rows	count/s	InstanceRowcount_NoOp	userId and instanceName	The total number of rows.
Failed Rows	count/s	InstanceFailRowcount_NoOp	userId and instanceName	The number of failed rows.
Storage Capacity	byte	StorageSize	userId and instanceName	The storage size.

Middleware

Message Queue for Apache RocketMQ

Instance metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Message retention duration	h	MessageRetentionPeriod	userId and instanceId	The retention period of the message.
ReceiveMessageCountPerInstance	count/min	ReceiveMessageCountPerInstance	userId and instanceId	The number of messages that the consumer receives per minute.
SendMessageCountPerInstance	count/min	SendMessageCountPerInstance	userId and instanceId	The number of messages that the producer sends per minute.
Consumer(Instance) recevies messages TPS	count/min	ReceiveMessageTPSPerInstance	userId and instanceId	The TPS for receiving messages by the consumer per minute.
Producer(Instance) sends messages TPS	count/min	SendMessageTPSPerInstance	userId and instanceId	The TPS for sending messages by the producer per minute.

Group metrics

Console-based metric	Unit	API-based metric	Dimension	Description
ConsumerLagPerGid	count	ConsumerLag	userId, instanceId, and groupId	The number of accumulated messages.
ReceiveMessagecountPerGid	count/min	ReceiveMessagecountPerGid	userId, instanceId, and groupId	The number of messages that the consumer receives per minute.
SendDLQMessagecountPerGid	count/min	SendDLQMessagecountPerGid	userId, instanceId, and groupId	The number of dead-letter messages per minute.
Producer(GroupId) Number of messages sent per minute	count/min	SendMessagecountPerGid	userId, instanceId, and groupId	The number of messages that the producer sends per minute.
Consumer(GroupId) receives message TPS	count/min	ReceiveMessageTPSPerGid	userId, instanceId, and groupId	The number of TPS received by the consumer for messaging per minute.
Producer(GroupId) sends messages TPS	count/min	SendMessageTPSPerGid	userId, instanceId, and groupId	The number of TPS sent by the producer for messaging per minute.

Topic metrics

Console-based metric	Unit	API-based metric	Dimension	Description
ReceiveMessagecountPerTopic	count/min	ReceiveMessagecountPerTopic	userId, instanceId, and topic	The number of messages that the consumer receives per minute.

SendMessagecountPerTopic	count/min	SendMessagecountPerTopic	userId, instanceId, and topic	The number of messages that the producer sends per minute.
Consumer(Topic) receives message TPS	count/min	ReceiveMessageTPSPerTopic	userId, instanceId, and topic	The number of TPS received by the consumer for messaging per minute.
Producer(Topic) sends messages TPS per minute	count/min	SendMessageTPSPerTopic	userId, instanceId, and topic	The number of TPS sent by the producer for messaging per minute.

Group and topic metrics

Console-based metric	Unit	API-based metric	Dimension	Description
ConsumerLagPerGidTopic	count	ConsumerLagPerGidTopic	userId, instanceId, groupId, and topic	The number of accumulated messages.
ReceiveMessagecountPerGidTopic	count/min	ReceiveMessagecountPerGidTopic	userId, instanceId, groupId, and topic	The number of messages that the consumer receives per minute.
SendDLQMessagecountPerGidTopic	count/min	SendDLQMessagecountPerGidTopic	userId, instanceId, groupId, and topic	The number of dead-letter messages per minute.
Producer(GroupId & Topic) Number of messages sent per minute	count/min	SendMessagecountPerGidTopic	userId, instanceId, groupId, and topic	The number of messages that the producer sends per minute.

Big data

Realtime Compute (Blink)

Console-based metric	Unit	API-based metric	Dimension	Description
FetchedDelay	Seconds	FetchedDelay	userId, regionId, projectName, and jobName	The processing latency.
Response Delay	%	inputDelay	userId, regionId, projectName, and jobName	The service latency.
Read RPS	rps	ParserTpsRate	userId, regionId, projectName, and jobName	The read records per second (RPS).
Write RPS	rps	SinkOutTpsRate	userId, regionId, projectName, and jobName	The write RPS.
FailoverRate	%	TaskFailoverRate	userId, regionId, projectName, and jobName	The failover rate.

Realtime Compute for Apache Flink

Basic metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Number of Times a Job Is Restarted upon a Failure	count	flink_jobmanager_job_numRestarts	userId and instanceId	The number of times that the Flink job is restarted upon failure.
Job Duration	Value	flink_jobmanager_job_uptime	userId and instanceId	The period of time during which the Flink job has been run.
Running Jobs	count	flink_jobmanager_numRunningJobs	userId and instanceId	The number of running Flink jobs.
Available Slots	count	flink_jobmanager_taskSlotsAvailable	userId and instanceId	The number of available Flink slots.

Total Number of Slots for a Task	count	flink_jobmanager_taskSlotsTotal	userId and instanceId	The total number of slots in a Flink job.
TM Quantity	count	flink_jobmanager_numRegisteredTaskManagers	userId and instanceId	The number of Flink task managers.
Processing Latency	ms	flink_taskmanager_job_task_operator_currentEmitEventTimeLag	userId and instanceId	The latency of the Flink service.
Transmission Latency	Seconds	flink_taskmanager_job_task_operator_currentFetchEventTimeLag	userId and instanceId	The transmission latency of the Flink service.
Total Number of Input Data Records of All Operators	count	flink_taskmanager_job_task_operator_numRecordsIn	userId and instanceId	The total number of input records of all Flink operators.
Total Number of Output Data Records	count	flink_taskmanager_job_task_operator_numRecordsOut	userId and instanceId	The total number of Flink output records.
Number of Input Data Records per Second for All Data Streams	count/s	flink_taskmanager_job_task_operator_numRecordsInPerSecond	userId and instanceId	The number of input records per second for all Flink data streams.
Number of Output Data Records per Second for All Data Streams	count/s	flink_taskmanager_job_task_operator_numRecordsOutPerSecond	userId and instanceId	The number of output records per second for all Flink data streams.
Number of Data Records Not Read by the Source	count	flink_taskmanager_job_task_operator_pendingRecords	userId and instanceId	The number of unread data records in the Flink source.

Duration for Which Data Is Not Processed in the Source	ms	flink_taskmanager_job_task_operator_sourceIdleTime	userId and instanceId	The amount of time during which the Flink source does not process data.

Checkpoint metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Total Number of Checkpoints	count	flink_jobmanager_job_totalNumberOfCheckpoints	userId and instanceId	The total number of Flink checkpoints.
Number of Failed Checkpoints	count	flink_jobmanager_job_numberOfFailedCheckpoints	userId and instanceId	The number of failed Flink checkpoints.
Number of Completed Checkpoints	count	flink_jobmanager_job_numberOfCompletedCheckpoints	userId and instanceId	The number of completed Flink checkpoints.
Number of In-progress Checkpoints	count	flink_jobmanager_job_numberOfInProgressCheckpoints	userId and instanceId	The number of ongoing Flink checkpoints.
Duration of the Last Checkpoint	ms	flink_jobmanager_job_lastCheckpointDuration	userId and instanceId	The duration of the last Flink checkpoint.
Size of the Last Checkpoint	bytes	flink_jobmanager_job_lastCheckpointSize	userId and instanceId	The size of the last Flink checkpoint.

State metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Maximum Latency in a State Data Cleanup	Nanosecond	flink_taskmanager_job_task_operator_state_name_stateClearLatency	userId and instanceId	The maximum latency of a Flink cleanup in a single state.

Maximum Latency in Single ValueState Access (Get)	Nanosecond	flink_taskmanager_job_task_operator_state_name_valueStateGetLatency	userId and instanceId	The maximum access latency of the Flink job in the Value state at one time (Get).
Maximum Latency in Single ValueState Access (Update)	Nanosecond	flink_taskmanager_job_task_operator_state_name_valueStateUpdateLatency	userId and instanceId	The maximum access latency of the Flink job in the Value state at one time (Update).
Maximum Latency in Single AggregatingState Access (Get)	Nanosecond	flink_taskmanager_job_task_operator_state_name_aggregatingStateGetLatency	userId and instanceId	The maximum latency of the Flink job in the Aggregating state at one time (Get).
Maximum Latency in Single AggregatingState Access (Add)	Nanosecond	flink_taskmanager_job_task_operator_state_name_aggregatingStateAddLatency	userId and instanceId	The maximum latency of the Flink job in the Aggregating state at one time (Add).
Maximum Latency in Single AggregatingState Access (MerageNamespaces)	Nanosecond	flink_taskmanager_job_task_operator_state_name_aggregatingStateMergeNamespacesLatency	userId and instanceId	The maximum latency of the Flink job in the Aggregating state at one time (MerageNamespaces).
Maximum Latency in Single ReducingState Access (Get)	Nanosecond	flink_taskmanager_job_task_operator_state_name_reducingStateGetLatency	userId and instanceId	The maximum latency of the Flink job in the Reducing state at one time (Get).
Maximum Latency in Single ReducingState Access (Add)	Nanosecond	flink_taskmanager_job_task_operator_state_name_reducingStateAddLatency	userId and instanceId	The maximum latency of the Flink job in the Reducing state at one time (Add).

Maximum Latency in Single ReducingState Access (MerageNamespaces)	Nanosecond	flink_taskmanager_job_task_operator_state_name_reducingStateMergeNamespacesLatency	userId and instanceId	The maximum latency of the Flink job in the Reducing state at one time (MerageNamespaces).
Maximum Latency in Single MapState Access (IsEmpty)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateIsEmptyLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (IsEmpty).
Maximum Latency in Single MapState Access (IteratorHasNext)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateIteratorHasNextLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (IteratorHasNext).
Maximum Latency in Single MapState Access (IteratorNext)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateIteratorNextLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (IteratorNext).
Maximum Latency in Single MapState Access (IteratorRemove)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateIteratorRemoveLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (IteratorRemove).
Maximum Latency in Single MapState Access (Get)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateGetLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (Get).
Maximum Latency in Single MapState Access (Put)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStatePutLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (Put).

Maximum Latency in Single MapState Access (PutAll)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStatePutAllLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (PutAll).
Maximum Latency in Single MapState Access (Remove)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateRemoveLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (Remove).
Maximum Latency in Single MapState Access (Contains)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateContainsLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (Contains).
Maximum Latency in Single MapState Access (EntriesInit)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateEntriesInitLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (EntriesInit).
Maximum Latency in Single MapState Access (KeysInit)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateKeysInitLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (KeysInit).
Maximum Latency in Single MapState Access (ValuesInit)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateValuesInitLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (ValuesInit).

Maximum Latency in Single MapState Access (IteratorInit)	Nanosecond	flink_taskmanager_job_task_operator_state_name_mapStateIteratorInitLatency	userId and instanceId	The maximum latency of the Flink job in the Map state at one time (IteratorInit).
Maximum Latency in Single ListState Access (Get)	Nanosecond	flink_taskmanager_job_task_operator_state_name_listStateGetLatency	userId and instanceId	The maximum latency of the Flink job in the List state at one time (Get).
Maximum Latency in Single ListState Access (Add)	Nanosecond	flink_taskmanager_job_task_operator_state_name_listStateAddLatency	userId and instanceId	The maximum latency of the Flink job in the List state at one time (Add).
Maximum Latency in Single ListState Access (AddAll)	Nanosecond	flink_taskmanager_job_task_operator_state_name_listStateAddAllLatency	userId and instanceId	The maximum latency of the Flink job in the List state at one time (AddAll).
Maximum Latency in Single ListState Access (Update)	Nanosecond	flink_taskmanager_job_task_operator_state_name_listStateUpdateLatency	userId and instanceId	The maximum latency of the Flink job in the List state at one time (Update).
Maximum Latency in Single ListState Access (MergeNamespaces)	Nanosecond	flink_taskmanager_job_task_operator_state_name_listStateMergeNamespacesLatency	userId and instanceId	The maximum latency of the Flink job in the List state at one time (MergeNamespaces).
Maximum Latency in Single SortedMapState Access (FirstEntry)	Nanosecond	flink_taskmanager_job_task_operator_state_name_sortedMapStateFirstEntryLatency	userId and instanceId	The maximum latency of the Flink job in the Sorted Map state at one time (FirstEntry).

Maximum Latency in Single SortedMapState Access (LastEntry)	Nanosecond	flink_taskmanager_job_task_operator_state_name_sortedMapStateLastEntryLatency	userId and instanceId	The maximum latency of the Flink job in the Sorted Map state at one time (LastEntry).

Task metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Input Pool Usage	Value	flink_taskmanager_job_task_buffers_inPoolUsage	userId and instanceId	The usage of the Flink input pool.
Output Pool Usage	Value	flink_taskmanager_job_task_buffers_outPoolUsage	userId and instanceId	The usage of the Flink output pool.
Input Queue Length	Value	flink_taskmanager_job_task_buffers_inputQueueLength	userId and instanceId	The length of the Flink input queue.
Output Queue Length	Value	flink_taskmanager_job_task_buffers_outputQueueLength	userId and instanceId	The length of the Flink output queue.
Total Number of Input bytes	bytes	flink_taskmanager_job_task_operator_numbytesIn	userId and instanceId	The total number of Flink input bytes.
Total Number of Input bytes per Second (Local)	bytes/s	flink_taskmanager_job_task_numbytesInLocalPerSecond	userId and instanceId	The total number of Flink input bytes per second (Local).
Total Number of Input bytes per Second (Remote)	bytes/s	flink_taskmanager_job_task_numbytesInRemotePerSecond	userId and instanceId	The total number of Flink input bytes per second (Remote).
Total Number of Input Buffers per Second (Local)	bytes/s	flink_taskmanager_job_task_numBuffersInLocalPerSecond	userId and instanceId	The total number of Flink input buffers per second (Local).

Total Number of Input Buffers per Second (Remote)	bytes/s	flink_taskmanager_job_task_numBuffersInRemotePerSecond	userId and instanceId	The total number of Flink input buffers per second (Remote).
Total Number of Input bytes per Second (Operator)	bytes/s	flink_taskmanager_job_task_operator_numbytesInPerSecond	userId and instanceId	The total number of Flink input bytes per second (Operator).
Total Number of Output bytes	bytes	flink_taskmanager_job_task_operator_numbytesOut	userId and instanceId	The total number of Flink output bytes.
Total Number of Output bytes per Second	bytes/s	flink_taskmanager_job_task_numbytesOutPerSecond	userId and instanceId	The total number of Flink output bytes per second.
Total Number of Output Buffers per Second	bytes/s	flink_taskmanager_job_task_numBuffersOutPerSecond	userId and instanceId	The total number of Flink output buffers per second.
Total Number of Output bytes per Second (Operator)	bytes/s	flink_taskmanager_job_task_operator_numbytesOutPerSecond	userId and instanceId	The total number of Flink output bytes per second (Operator).
Total Number of Data Records That Flow into Each Subtask	count	flink_taskmanager_job_task_numRecordsIn	userId and instanceId	The total amount of data received by each Flink subtask.
Total Number of Data Records That Are Exported by Each Subtask	count	flink_taskmanager_job_task_numRecordsOut	userId and instanceId	The total amount of data output by each Flink subtask.
Total Number of Data Records That Flow into Each Subtask per Second	count/s	flink_taskmanager_job_task_numRecordsInPerSecond	userId and instanceId	The total amount of data received by each Flink subtask per second.
Total Number of Data Records That Are Exported by Each Subtask per Second	count/s	flink_taskmanager_job_task_numRecordsOutPerSecond	userId and instanceId	The total amount of data output by each Flink subtask per second.

Time Required by Each Subtask to Export the Last Data Record to the Sink	Seconds	flink_taskmanager_job_task_operator_currentSendTime	userId and instanceId	The amount of time that is used by each Flink subtask to send the most recent data to the downstream system.
Time When Each Task Receives the Latest Watermark	Value	flink_taskmanager_job_task_currentInputWatermark	userId and instanceId	The time when each Flink task receives the most recent watermark.
Time When Each Task Receives the Latest Watermark (Operator)	Value	flink_taskmanager_job_task_operator_currentInputWatermark	userId and instanceId	The time when each Flink task receives the most recent watermark (Operator).
Time When Each Task Transmits the Latest Watermark (Operator)	Value	flink_taskmanager_job_task_operator_currentOutputWatermark	userId and instanceId	The time when the most recent watermark is output by each Flink task (Operator).
Watermark Latency	ms	flink_taskmanager_job_task_operator_watermarkLag	userId and instanceId	The latency of the Flink watermark.

Status metrics

Console-based metric	Unit	API-based metric	Dimension	Description
CPU Load of a Single JM	%	flink_jobmanager_Status_JVM_CPU_Load	userId and instanceId	The CPU load of a single Flink JM.
CPU Load of a Single TM	%	flink_taskmanager_Status_JVM_CPU_Load	userId and instanceId	The CPU load of a single Flink TM.
CPU Utilization of a Single TM	%	flink_taskmanager_Status_ProcessTree_CPU_Usage	userId and instanceId	The CPU utilization of a single Flink TM.
Heap Memory of JM (Used)	bytes	flink_jobmanager_Status_JVM_Memory_Heap_Used	userId and instanceId	The heap memory of the FlinkJM (Used).

Heap Memory of JM (Committed)	bytes	flink_jobmanager_Status_JVM_Memory_Heap_Committed	userId and instanceId	The heap memory of the FlinkJM (Committed).
Heap Memory of JM (Max)	bytes	flink_jobmanager_Status_JVM_Memory_Heap_Max	userId and instanceId	The heap memory of the FlinkJM (Max).
Non-Heap Memory of JM (Used)	bytes	flink_jobmanager_Status_JVM_Memory_NonHeap_Used	userId and instanceId	The non-heap memory of the FlinkJM (Used).
Non-Heap Memory of JM (Committed)	bytes	flink_jobmanager_Status_JVM_Memory_NonHeap_Committed	userId and instanceId	The non-heap memory of the FlinkJM ((Committed).
Non-Heap Memory of JM (Max)	bytes	flink_jobmanager_Status_JVM_Memory_NonHeap_Max	userId and instanceId	The non-heap memory of the FlinkJM (Max).
Used Heap Memory of TM	bytes	flink_taskmanager_Status_JVM_Memory_Heap_Used	userId and instanceId	The heap memory of the FlinkTM (Used).
Committed Heap Memory of TM	bytes	flink_taskmanager_Status_JVM_Memory_Heap_Committed	userId and instanceId	The heap memory of the FlinkTM (Committed).
Max. Heap Memory of TM	bytes	flink_taskmanager_Status_JVM_Memory_Heap_Max	userId and instanceId	The heap memory of the FlinkTM (Max).
Used Non-Heap Memory of TM	bytes	flink_taskmanager_Status_JVM_Memory_NonHeap_Used	userId and instanceId	The non-heap memory of the FlinkTM (Used).
Committed Non-Heap Memory of TM	bytes	flink_taskmanager_Status_JVM_Memory_NonHeap_Committed	userId and instanceId	The non-heap memory of the FlinkTM (Committed).

Max. Non-Heap Memory of TM	bytes	flink_taskmanager_Status_JVM_Memory_NonHeap_Max	userId and instanceId	The non-heap memory of the FlinkTM (Max).
Memory of the Entire Process by Using Linux	bytes	flink_taskmanager_Status_ProcessTree_Memory_RSS	userId and instanceId	The memory of the entire process that Flink obtains through Linux.
Number of Threads of JM	count	flink_jobmanager_Status_JVM_Threads_count	userId and instanceId	The number of FlinkJM threads.
Number of Threads of TM	count	flink_taskmanager_Status_JVM_Threads_count	userId and instanceId	The number of FlinkTM threads.
Number of GCs on JM (ParNew)	count	flink_jobmanager_Status_JVM_GarbageCollector_ParNew_count	userId and instanceId	The number of FlinkJM GCs (ParNew).
Number of GCs on JM (ConcurrentMarkSweep)	count	flink_jobmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_count	userId and instanceId	The number of FlinkJM GCs (ConcurrentMarkSweep).
Duration of Each GC on JM (ParNew)	ms	flink_jobmanager_Status_JVM_GarbageCollector_ParNew_Time	userId and instanceId	The time of each JM GC in Flink (ParNew).
Duration of Each GC on JM (ConcurrentMarkSweep)	ms	flink_jobmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_Time	userId and instanceId	The time of each JM GC in Flink (ConcurrentMarkSweep).
Number of GCs on TM (ParNew)	count	flink_taskmanager_Status_JVM_GarbageCollector_ParNew_count	userId and instanceId	The number of FlinkTM GCs (ParNew).
Number of GCs on TM (ConcurrentMarkSweep)	count	flink_taskmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_count	userId and instanceId	The number of FlinkTM GCs (ConcurrentMarkSweep).

Duration of Each GC on TM (ParNew)	ms	flink_taskmanager_Status_JVM_GarbageCollector_ParNew_Time	userId and instanceId	The time of each TM GC in Flink (ParNew).
Duration of Each GC on JM (ConcurrentMarkSweep)	ms	flink_taskmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_Time	userId and instanceId	The time of each TM GC in Flink (ConcurrentMarkSweep).
Total Classes Loaded After the JVM Where the JM Resides Is Created	Value	flink_jobmanager_Status_JVM_ClassLoader_ClassesLoaded	userId and instanceId	The total number of classes loaded after the JVM where the FlinkJM resides is created.
Total Classes Unloaded After the JVM Where the JM Resides Is Created	Value	flink_jobmanager_Status_JVM_ClassLoader_ClassesUnloaded	userId and instanceId	The total number of unloaded classes after the JVM where the FlinkJM resides is created.
Total Classes Loaded After the JVM Where the TM Resides Is Created	Value	flink_taskmanager_Status_JVM_ClassLoader_ClassesLoaded	userId and instanceId	The total number of classes loaded after the JVM where the FlinkTM resides is created.
Total Classes Unloaded After the JVM Where the TM Resides Is Created	Value	flink_taskmanager_Status_JVM_ClassLoader_ClassesUnloaded	userId and instanceId	The total number of unloaded classes after the JVM where the FlinkTM resides is created.

DataHub

Console-based metric	Unit	API-based metric	Dimension	Description
subscription delay time	Seconds	sub_delay	userId and subId	The latency between the time when a user subscribes to data and the current system time.

Elasticsearch (ECS)

Instance metrics

Console-based metric	Unit	API-based metric	Dimension	Description
ClusterStatus	value	ClusterStatus	userId and instanceId	The status of the cluster. Valid values: • 0: green, which indicates normal. • 1: yellow, which indicates warning. • 2: red, which indicates exception.
InstanceIdCPUUtilization	%	InstanceIdCPUUtilization	userId and instanceId	The CPU utilization of the instance.
NodeDiskUtilization	%	InstanceIdDiskUtilization	userId and instanceId	The disk space usage of the instance.
NodeHeapMemoryUtilization	%	InstanceIdHeapMemoryUtilization	userId and instanceId	The HeapMemory usage of the instance.
Number of Times Garbage Collection Is Performed in Old Generation Space on ES Node (ECS)	count	InstanceIdOldGCCollectioncount	userId and instanceId	The number of times old garbage collection is performed.
Time Consumed by Garbage Collection in Old Generation Space on ES Node (ECS)	ms	InstanceIdOldGCCollectionTimeInMillis	userId and instanceId	The time consumed by old garbage collection.
Number of Times Garbage Collection Is Performed in Yound Generation Space on ES Node (ECS)	count	InstanceIdYoungGCCollectioncount	userId and instanceId	The number of times young garbage collection is performed.
Time Consumed by Garbage Collection in Young Generation Space on ES Node (ECS)	ms	InstanceIdYoungGCCollectionTimeInMillis	userId and instanceId	The time consumed by young garbage collection.

Node metrics

Console-based metric	Unit	API-based metric	Dimension	Description
NodeCPUUtilization	%	NodeCPUUtilization	userId, instanceId, and nodeName	The CPU utilization of the node.
NodeDiskUtilization	%	NodeDiskUtilization	userId, instanceId, and nodeName	The disk usage of the node.
NodeHeapMemoryUtilization	%	NodeHeapMemoryUtilization	userId, instanceId, and nodeName	The HeapMemory usage of the node.
ES node dimension old age garbage collection times	count	OldGCCollectioncount	userId, instanceId, and nodeName	The number of times old garbage collection is performed.
The ES node dimension takes total time for garbage collection in the old age.	ms	OldGCCollectionTimelnMillis	userId, instanceId, and nodeName	The time consumed by old garbage collection.
Number of Young Generation Garbage Collection in ES Node Dimension	count	YoungGCCollectioncount	userId, instanceId, and nodeName	The number of times young garbage collection is performed.
Total Time-consuming for Young Generation Garbage Collection in ES Node Dimension	ms	YoungGCCollectionTimeInMillis	userId, instanceId, and nodeName	The time consumed by young garbage collection.

Elasticsearch (Kubernetes)

Instance metrics

Console-based metric	Unit	API-based metric	Dimension	Description

ClusterStatus	value	ClusterStatus	userId and instanceId	The status of the cluster. Valid values: • 0: green, which indicates normal. • 1: yellow, which indicates warning. • 2: red, which indicates exception.
Number of tasks waiting to be executed	count	PendingTaskcount	userId and instanceId	The number of tasks that are waiting to be executed.
The maximum waiting time for a task to be executed	ms	PendingTaskMaxWaitTime	userId and instanceId	The maximum waiting time for a task to be executed.
Number of data nodes	count	DataNodes	userId and instanceId	The number of data nodes.
Active Shards	count	ActiveShards	userId and instanceId	The number of active shards.
Unallocated Shards	count	UnassingnedShards	userId and instanceId	The number of unallocated shards.

Node metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Node CPU usage (elasticsearch-k8s)	%	NodeCPUUtil	userId, instanceId, and nodeName	The CPU utilization of the node.
Node Disk Usage (elasticsearch-k8s)	%	NodeDiskUtil	userId, instanceId, and nodeName	The disk utilization of the node.
NodeHeapMemoryUtilization	%	NodeHeapMemoryUtil	userId, instanceId, and nodeName	The heap memory usage of the node.

Node cpu 1min load	value	NodeLoad1	userId, instanceId, and nodeName	The average load of the node over the past 1 minute.
Node cpu 5min load'	value	NodeLoad5	userId, instanceId, and nodeName	The average load of the node over the past 5 minutes.
Node cpu 15min load	value	NodeLoad15	userId, instanceId, and nodeName	The average load of the node over the past 15 minutes.
Node OldGC times-Cumulative Value	count	OldGCTotalcount	userId, instanceId, and nodeName	The total number of times old garbage collection is performed.
Node OldGC Time-Accumulation Value	ms	OldGCTotalTime	userId, instanceId, and nodeName	The time consumed by old garbage collection.
Node OldGC average time consumption (millinseconds)	ms	OldGCCostTime	userId, instanceId, and nodeName	The average time consumed by old garbage collection.
Node OldGC rate	count/s	OldGCRate	userId, instanceId, and nodeName	The rate of old garbage collection.
Node YoungGC times-Cumulative Value	count	YoungGCTotalcount	userId, instanceId, and nodeName	The total number of times young garbage collection is performed.
Node YoungGC Time-Accumulation Value	count	YoungGCTotalTime	userId, instanceId, and nodeName	The total time consumed by young garbage collection.

Average Node YoungGC Time Consumption (MS)	ms	YoungGCCostTime	userId, instanceId, and nodeName	The average time consumed by young garbage collection.
Node YoungGC rate	count/s	YoungGCRate	userId, instanceId, and nodeName	The rate of young garbage collection.
Node free disk space	GB	AvailableDiskSize	userId, instanceId, and nodeName	The available disk space of the node.
Node heap memory usage	GB	NodeHeapMemory	userId, instanceId, and nodeName	The heap memory usage of the node.
Number of node documents	count	NodeDocumentcount	userId, instanceId, and nodeName	The number of node documents.
Number of node segments	count	NodeSegmentcount	userId, instanceId, and nodeName	The number of segments on the node.
Disk IO Rate	count/s	NodeDiskIOOperationsRate	userId, instanceId, and nodeName	The disk I/O rate.
Disk write rate	count/s	NodeWriteIOOperationsRate	userId, instanceId, and nodeName	The disk write rate.
Disk read rate	count/s	NodeReadIOOperationsRate	userId, instanceId, and nodeName	The disk read rate.
Write thread	count	IndexThread	userId, instanceId, and nodeName	The number of write threads.
Query thread	count	WriteThread	userId, instanceId, and nodeName	The number of query threads.

Node data size	MB	NodeData	userId, instanceId, and nodeName	The size of data on the node.
Node lucene memory usage	KB	NodeIndexMemoryLucene	userId, instanceId, and nodeName	The memory usage of lucene on the node.
Node terms memory usage	KB	NodeIndexMemoryTerms	userId, instanceId, and nodeName	The memory usage of terms on the node.
Node points memory usage	KB	NodeIndexMemoryPoints	userId, instanceId, and nodeName	The memory usage of points on the node.
Average time consumption of a single query (ms)	ms	NodeSearchLatency	userId, instanceId, and nodeName	The average amount of time that is consumed by a single query.
Average write time (ms)	ms	NodeIndexingLatency	userId, instanceId, and nodeName	The average duration of a single write operation.

Index metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Write rate	count/s	IndexRate	userId, instanceId, and indexName	The write rate.
Average write time (ms)	ms	AvgIndexTime	userId, instanceId, and indexName	The average amount of time that is consumed by a single write operation.
Query rate	count/s	SearchRate	userId, instanceId, and indexName	The query rate.

Average time consumption of a single query (ms)	ms	AvgSearchTime	userId, instanceId, and indexName	The average amount of time that is consumed by a single query.
Number of indexed documents	count	IndexDocummentcount	userId, instanceId, and indexName	The number of indexed documents.
Total Number of Segments	count	IndexSegmentAllShard	userId, instanceId, and indexName	The total number of segments.
Number of main segmentation	count	IndexSegmentPrimaries	userId, instanceId, and indexName	The number of primary shard segments.
Field memory	byte	FielddataMemorybytesTotal	userId, instanceId, and indexName	The field memory.
Index terms memory	byte	IndexMemoryTerms	userId, instanceId, and indexName	The memory of index terms.
Node points memory usage	byte	IndexMemoryPoints	userId, instanceId, and indexName	The memory of index points.
Index data size	byte	IndexDisk	userId, instanceId, and indexName	The size of the index data.
Index master slice data size	byte	IndexDiskPrimaries	userId, instanceId, and indexName	The data size of the index primary shard.

Hologres

Instance status metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Total Connections Used	count	connection_num	userId and instanceId	The number of SQL connections.

Connection Usage	%	connection_usage	userId and instanceId	The connection usage.
Export QPS	count/s	record_output_qps	userId and instanceId	The QPS.
Micro batch size	count/s	scan_batchsize	userId and instanceId	The size of the micro batch.
CPU Utilization	%	cpu_usage	userId and instanceId	The CPU utilization.
memory usage	%	memory_usage	userId and instanceId	The memory usage.
storage_usage	byte	storage_usage	userId and instanceId	The storage usage.
I/O Write Throughput	byte/s	io_write_throughput	userId and instanceId	The write I/O throughput per second.
I/O Read Throughput	byte/s	io_read_throughput	userId and instanceId	The read I/O throughput per second.

Running status metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Failed Query Qps	count/s	failed_query_qps	userId and instanceId	The number of failed QPS.
Query Latency Detail	ms	query_latency	instanceId, userId, and cmdType	The query latency.
Query QPS Detail	count/s	query_qps	instanceId, userId, and cmdType	The QPS details.

Network

NAT Gateway (NAT)

Instance metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Total number of sessions	count/s	curr_num	userId and instanceId	The total number of sessions.
Capacity Discard Connections (Cumulative Value)	count	capacity_drop	userId and instanceId	The number of capacity-dropped connections.
Number of new discarded connections (cumulative value)	count	ratelimit_drop	userId and instanceId	The number of new dropped connections.
Capacity Discard Connections Per Second	count/s	capacity_drop_inc	userId and instanceId	The number of capacity-dropped connections per second.
Number of new and discarded connections per second.	count/s	ratelimit_drop_inc	userId and instanceId	The number of new dropped connections per second.
Number of new connections per second	count/s	establish_rate	userId and instanceId	The number of new connections per second.

SNAT metrics

Console-based metric	Unit	API-based metric	Dimension	Description
Network inflow bandwidth	bit/s	snet_rx.rate	userId, instanceId, and snatentryId	The inbound bandwidth.
snet_tx.rate	bit/s	snet_tx.rate	userId, instanceId, and snatentryId	The outbound bandwidth.

Server Load Balancer (SLB)

Instance metrics

Console-based metric	Unit	API-based metric	Dimension	Description
TCPActiveConnection	count	InstanceActiveConnection	userId and instanceId	The number of active TCP connections.
TCPInactiveConnection	count	InstanceInactiveConnection	userId and instanceId	The number of inactive connections on the port.
InstanceMaxConnection	count/s	InstanceMaxConnection	userId and instanceId	The number of concurrent connections per second on the port.
TCPNewConnection	count	InstanceNewConnection	userId and instanceId	The number of new TCP connections.
PackageInPerSecond	count/s	InstancePacketRX	userId and instanceId	The number of inbound packets per second.
PackageOutPerSecond	count/s	InstancePacketTX	userId and instanceId	The number of outbound packets per second.
InternetInRate	bit/s	InstanceTrafficRXNew	userId and instanceId	The inbound bandwidth over the Internet.
InternetOutRate	bit/s	InstanceTrafficTXNew	userId and instanceId	The outbound bandwidth over the Internet.
Instance health check the number of backend abnormal ECS instances	count	InstanceUnhealthyServercount	userId and instanceId	The number of unhealthy ECS instances.
InstanceDropConnection	count/s	InstanceDropConnection	userId and instanceId	The number of connections dropped per second.
InstanceDropTrafficRX	bit/s	InstanceDropTrafficRX	userId and instanceId	The number of inbound bits dropped from the listener per second.

InstanceDropTrafficTX	bit/s	InstanceDropTrafficTX	userId and instanceId	The number of outbound bits dropped from the listener per second.
InstanceDropPacketRX	count/s	InstanceDropPacketRX	userId and instanceId	The number of inbound packets dropped from the listener per second.
InstanceDropPacketTX	count/s	InstanceDropPacketTX	userId and instanceId	The number of outbound packets dropped from the listener per second.
Seven-layer instance Qps	req/s	InstanceSevenFloorQps	userId and instanceId	The QPS per second at Layer 7.
InstanceFailConnection	count	InstanceFailConnection	userId and instanceId	The number of failed connections on the instance.

Port metrics

Console-based metric	Unit	API-based metric	Dimension	Description
TCPActiveConnection	count	ActiveConnection	userId, instanceId, and port	The number of active TCP connections per second.
TCPInactiveConnection	count	InactiveConnection	userId, instanceId, and port	The number of inactive connections per second on the port.
MaxConnection	count/s	MaxConnection	userId, instanceId, and port	The number of concurrent connections per second on the port.
TCPNewConnection	count	NewConnection	userId, instanceId, and port	The number of new TCP connections.
PackageInPerSecond	count/s	PacketRX	userId, instanceId, and port	The number of inbound packets per second.

PackageOutPerSecond	count/s	PacketTX	userId, instanceId, and port	The number of outbound packets per second.
InternetInRate	bit/s	TrafficRXNew	userId, instanceId, and port	The inbound bandwidth over the Internet.
InternetOutRate	bit/s	TrafficTXNew	userId, instanceId, and port	The outbound bandwidth over the Internet.
UnhealthyServercount	count	UnhealthyServercount	userId, instanceId, and port	The number of unhealthy ECS instances.
DropConnection	count/s	DropConnection	userId, instanceId, and port	The number of connections dropped from the listener per second.
DropTrafficRX	bit/s	DropTrafficRX	userId, instanceId, and port	The number of inbound bits dropped from the listener per second.
DropTrafficTX	bit/s	DropTrafficTX	userId, instanceId, and port	The number of outbound bits dropped from the listener per second.
DropPacketRX	count/s	DropPacketRX	userId, instanceId, and port	The number of inbound packets dropped from the listener per second.
DropPacketTX	count/s	DropPacketTX	userId, instanceId, and port	The number of outbound packets dropped per second.
Seven layers Qps	req/s	SevenFloorQps	userId, instanceId, and port	The QPS per second at Layer 7.
Number of connection failures per second	count	FailConnection	userId, instanceId, and port	The number of failed connections per second.

Elastic IP Address (EIP)

Console-based metric	Unit	API-based metric	Dimension	Description
net_rx.rate	bit/s	net_rx.rate	userId and instanceId	The inbound bandwidth.
net_rxPkgs.rate	count/s	net_rxPkgs.rate	userId and instanceId	The rate of inbound packets.
net_tx.rate	bit/s	net_tx.rate	userId and instanceId	The outbound bandwidth.
net_txPkgs.rate	count/s	net_txPkgs.rate	userId and instanceId	The rate of outbound packets.
Out Ratelimit Drop Speed	Packets/s	out_ratelimit_drop_speed	userId and instanceId	The outbound packet loss rate.
In Ratelimit Drop Speed	Packets/s	in_ratelimit_drop_speed	userId and instanceId	The inbound packet loss rate.

VPN Gateway

Console-based metric	Unit	API-based metric	Dimension	Description
IPSec.connection.rxPkgs	Packets/s	ipsec.rxPkgs	userId, instanceId, and ipsecId	The rate of inbound packets of the connection.
IPSec.connection.txpkgs	Packets/s	ipsec.txPkgs	userId, instanceId, and ipsecId	The rate of outbound packets of the connection.
IPSec.connection.rx.rate	bit/s	ipsec_rx.rate	userId, instanceId, and ipsecId	The number of inbound bits of the connection per second.

IPSec.connection.tx.rate	bit/s	ipsec_tx.rate	userId, instanceId, and ipsecId	The number of outbound bits of the connection per second.
SSL Client count	count	ssl_client.count	userId and instanceId	The number of connections.

Apsara Stack DNS

Console-based metric	Unit	API-based metric	Dimension	Description
Latency	ms	Latency	userId and instanceId	The parsing latency.
ParsingSucceeded	count	ParsingSucceeded	userId and instanceId	The number of parsed records
Total number of analyses	count	ParsingTotal	userId and instanceId	The total number of parsed records.

10.3.2. View the monitoring data

On the monitoring page of a cloud service, you can view the status of resources and the metrics.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Cloud Service Monitoring**.
4. On the **Cloud Service Monitoring** page, select a cloud service whose monitoring data you want to view.
5. On the monitoring page of the cloud service, click **Monitoring Charts** in the **Actions** column of a resource. On the page that appears, you can view the monitoring charts of the resource.

> **Note**
> You can view the monitoring data of the last 30 days at most.

10.3.3. View alert rules

On the monitoring page of a cloud service, you can view all alert rules of the cloud service in the current resource set and organization.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Cloud Service Monitoring**.
4. On the **Cloud Service Monitoring** page, select a cloud service whose alert rules you want to view.
5. On the monitoring page of the cloud service, click **View Alert Rules**.
6. Optional. Find the alert rule whose details you want to view. Click **Rule Details** in the **Actions** column of the alert rule. In the **Alert Rules/Alert Logs/Faulty Instances** panel, view the rule details, alert history, and faulty resources.

10.4. Host monitoring

After the CloudMonitor agent is installed on your Elastic Compute Service (ECS) instances, CloudMonitor can collect various monitoring data of metrics about the operating system that the ECS instances run. You can configure alert rules for the metrics to reduce monitoring costs.

- **Easy-to-use agent**: After the CloudMonitor agent is installed on your ECS instances, CloudMonitor can collect various monitoring data of metrics about the operating system that the ECS instances run.
- **Various metrics**: Host monitoring supports metrics about the CPU, memory, disks, and network to meet your basic O&M requirements.
- **Business-specific process monitoring**: Host monitoring collects the CPU utilization, memory usage, and the number of opened files for active processes. This helps you monitor the resource usage of your ECS instances.
- **Flexible alerting**: Host monitoring allows you to manage ECS instances across regions and configure alert rules by application group. This reduces the costs of monitoring management.

10.4.1. Install a CloudMonitor agent

You can install the CloudMonitor agent on Elastic Compute Service (ECS) instances to collect the monitoring data of operating systems.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Host Monitoring**.
4. On the **Host Monitoring** page, select the check boxes of the ECS instances for which you want to install or upgrade the agent, and then click **Batch Install** in the lower-left corner of the instance list.

   > **Note**
   >
   > You can only install the plug-in of CloudMonitor for ECS instances that run the **Linux** and **Windows** operating systems.

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

> **Note**
>
> The installation or upgrade process takes approximately five minutes. After the agent status changes from **Installing** to **Running**, the CloudMonitor agent is installed or upgraded.

10.4.2. Process monitoring

CloudMonitor provides the process monitoring feature to help you collect information about the active processes of Elastic Compute Service (ECS) instances during a specific period of time. You can collect the following information: CPU utilization, memory usage, and the number of opened files. If you add a process keyword, you can obtain the number of processes that match the keyword.

Prerequisites

The CloudMonitor agent is installed on the ECS instances. For more information, see Install a CloudMonitor agent.

Background information

CloudMonitor identifies the top five processes with the highest CPU utilization every minute and records the CPU utilization, memory usage, and the number of opened files in these processes.

- CPU utilization and memory usage of a process: You can check the output of the **top** command in the Linux operating system to understand the two metrics.
- Number of opened files in a process: You can check the output of the **lsof** command in the Linux operating system to understand this metric.

CloudMonitor collects the CPU utilization of only the top five CPU-intensive processes. The following issues may occur:

- If a process occupies multiple CPU cores, the collected CPU utilization may exceed 100%. This is because CloudMonitor collects the total CPU utilization of the process on all the occupied CPU cores.
- If the top five CPU-intensive processes change during the time period specified in your query, the process list displays all the processes that have ranked top five over the specified time period. The time displayed in the list indicates when the processes last ranked top five.
- CloudMonitor collects the CPU utilization, memory usage, and number of opened files for a process only when the process is among the top five CPU-intensive processes. If a process has not stayed in the top five within the time period specified in your query, its monitoring data is displayed as dispersed data points in the monitoring charts. The density of the data points indicates the activity level of the process on the host. Example:
  - In the following figures, the data points for the wrapper process are sparse and discontinuous in the monitoring charts. This is because the wrapper process did not continuously rank top five with the highest CPU utilization during the specified time period.
  - In the following figures, the data points for the java process are dense and continuous in the monitoring charts. This is because the java process continuously ranked top five with the highest CPU utilization during the specified time period.

Add or remove a process for monitoring

You can monitor the number of processes on ECS instances and collect the number of key processes to obtain the availability of the key processes.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Host Monitoring**.
4. On the **Host Monitoring** page, click the name of the ECS instance that you want to manage.
5. Click the **Process Monitoring** tab. In the **Host Process Number(value)** section, click **Add Process Monitor** in the upper-right corner.
   - **Add a process**: In the **Add Process Monitor** panel, enter a process name and click **Add**.
   - **Remove a process**: In the **Add Process Monitor** panel, find the process that you want to remove, and click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

Examples of adding processes for monitoring

The following processes are running on an ECS instance:

- `/usr/bin/java -Xmx2300m -Xms2300m org.apache.catalina.startup.Bootstrap`
- `/usr/bin/ruby`
- `nginx -c /etc/nginx/nginx.conf`

You can add the following process keywords to obtain the monitoring results:

- ruby: One process matches this keyword by name. Therefore, the number of processes that match this keyword is 1.
- nginx: One process matches this keyword by name and parameter. Therefore, the number of processes that match this keyword is 1.
- /usr/bin: Two processes match this keyword by path. Therefore, the number of processes that match this keyword is 2.
- apache.catalina: One process matches this keyword by parameter. Therefore, the number of processes that match this keyword is 1.
- nginx.conf: One process matches this keyword by parameter. Therefore, the number of processes that match this keyword is 1.
- -c: One process matches this keyword by parameter. Therefore, the number of processes that match this keyword is 1.

Create an alert rule for process monitoring

After you add a process for monitoring, you can configure alert rules for processes under monitoring. When the number of processes exceeds a specified threshold, the alert contacts can receive alert notifications.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Host Monitoring**.
4. Click **Create Alert Rule**.

5. In the **Add Rule Description** panel that appears, select the **(Agent)process.count_processname** metric from the **Metric Name** drop-down list, and configure other parameters based on your business requirements.

   If multiple processes run on the host and the number of processes that match each process keyword varies, you can go back to the Create Alert Rule panel and click **Add Rule Description** to configure alert rules for the processes one by one. For more information about how to configure parameters in an alert rule, see Create an alert rule.

6. Click **OK**.

10.5. Alert service

You can manage alert rules, alert history, and alert templates in the Alert Service module of CloudMonitor. You can configure alert rules to specify how CloudMonitor checks monitoring data. If monitoring data triggers the alert rules, CloudMonitor sends alert notifications. When an alert is generated, CloudMonitor records the alert information in the alert history. You can review and analyze the alert history to troubleshoot issues and make predictions. Alert templates simplify the process of creating and maintaining alert rules when you need to manage a large number of cloud resources. This helps you monitor cloud services more efficiently.

10.5.1. Alert history

You can specify the cloud service, instance, alert level, and other filter conditions to view the alert history. You can also export and download the alert history to your on-premises machine for data analysis.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Alert Service** > **Alert History**.
4. Optional. On the **Alert History List** page, specify a rule ID, rule name, cloud service, alert level, or instance ID, and a time range to filter the alert history that you want to view.
5. On the Alert History List page, perform the required operations. The following table describes the supported operations.

Operation	Procedure
View the alert history	On the **Alert History** List page, view the details of the alert history that you filtered out. The details include the product, faulty instance, point in time when the alert occurred, alert rule name, status, alert level, and notification contact. You can click the alert rule name to view the details of the alert rule. **Note** By default, you can query the alert history of the last two days.

Export and download a report on the alert history	i. On the **Alert History List** page, click **Export report**. ii. In the **Create an alarm export task** dialog box, select a time range, enter a task name in the **Task Name** field, and then click **OK**. iii. In the left-side navigation pane, click **Export Report**. On the **Report download** page, click the **Alarm history** tab. iv. Find the report that you exported and click **Download report** in the **Actions** column. v. In the **Download report task** message, click **OK**. **Note** The report is in the **.xlsx** format. You can find the downloaded report in the default download directory of your browser.

10.5.2. Alert rules

You can configure alert rules for host monitoring metrics and instances in cloud service monitoring. You can apply an alert rule to all resources, an application group, or an instance. Alert Service sends alert notifications to contacts by using emails, DingTalk messages, and webhook URLs.

10.5.2.1. Create an alert rule

You can create alert rules for cloud services that are connected to CloudMonitor, such as Elastic Compute Service (ECS) instances.

During the validity period of an alert rule, CloudMonitor checks whether the metric values of cloud services exceed the specified thresholds at the default interval which is at least 60 seconds. If a metric value exceeds the threshold, CloudMonitor sends alerts to the specified contact or alert processing platforms.

Note

ApsaraDB RDS is an exception when the old data transmission link is used for monitoring. In this case, CloudMonitor checks the metrics of ApsaraDB RDS once in a period that ranges from 60s to 300s.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Alert Service** > **Alert Rule**. On the **Alert Rules** page, click **Create Alert Rule**.

4. In the **Create Alert Rule** panel, configure the required parameters. The following table describes the parameters in an alert rule.

Parameter	Description
Product	Select the cloud service that you want to monitor.
Resource Range	The resources that are associated with the alert rule. ◦ **Instances**: Select one or more instances in the resource set of the current organization. ◦ **Resource Group**: Select a resource set in the current organization and region. ◦ **Application grouping**: Select an application group in the current organization and region.
Rule Description	Specify the conditions of the alert rule. If the metrics meet the specified conditions, the alert rule is triggered. To configure the trigger conditions, perform the following steps: i. Click **Add Rule Description**. ii. In the **Add Rule Description** panel, configure the **Rule Name**, **Metric Name**, **Comparison**, and **Threshold** and **Alert Level** parameters. The Comparison parameter is optional. iii. Click **OK**.
Mute For	Specify the interval at which you want CloudMonitor to send alert notifications again in the cases that the alert is not cleared. Valid values: **5 Minutes**, **15 Minutes**, **30 Minutes**, **60 Minutes**, **3 Hours**, **6 Hours**, **12 Hours**, and **24 Hours**. An alert is triggered if the conditions of an alert rule are met. CloudMonitor does not resend an alert notification when the alert is triggered within the mute period. If the alert is not cleared after the mute period ends, CloudMonitor resends alert notifications. For example, if you set the **Mute For** parameter to **12 Hours**, CloudMonitor resends alert notifications only if the alert still persists after 12 hours.
Effective Period	Specify the period during which the alert rule is effective. CloudMonitor checks whether the monitoring data can trigger the alert rule only during the effective period.
Alert Callback	Specify a URL that can be accessed by CloudMonitor. **Note** CloudMonitor sends HTTP POST requests to push alert notifications to the specified URL. Only a URL of the HTTP protocol is supported. IP addresses are not supported.

Alert Contact Group	Select user groups or users to which you want to send alert notifications. By default, the value of this parameter is a dummy user used to test whether the alert rule is effective.

Parameters in the rule description

Parameter	Description
Rule name	Specify a name for the alert rule.
Metric Name	Select a metric of the cloud service that you want to monitor. For more information about the metrics of cloud services, see the "Metrics" topic.
Comparison	Specify the comparison condition between the monitoring data and alert threshold. Valid values: **>**, **>=**, **<**, **<=**, **=**, **!=**, **Compared With Yesterday Rise**, and **Compared With Yesterday Decline**.
Threshold and Alert Level	• **Threshold**: Specify the type of the threshold, the threshold, the monitoring period, and the number of times for which the monitoring data exceeds the threshold. The threshold is a percentage. Valid types of the threshold value: **Average**, **Minimum**, **Maximum**. Valid monitoring periods: **continuous**, **Within 5 minutes**, **Within 10 minutes**, **Within 15 minutes**, **Within 30 minutes**, and **Within 60 minutes**. • **Alert Level**: The valid values include **Critical**, **Warn**, and **Info**.

Sample alert rule

Create an alert rule named example to monitor the CPU utilization of ECS instances. When the CPU utilization is equal to or greater than 70% three times within five intervals, an alert of the Critical level is generated.

The following figure shows the parameter configurations in the rule description.

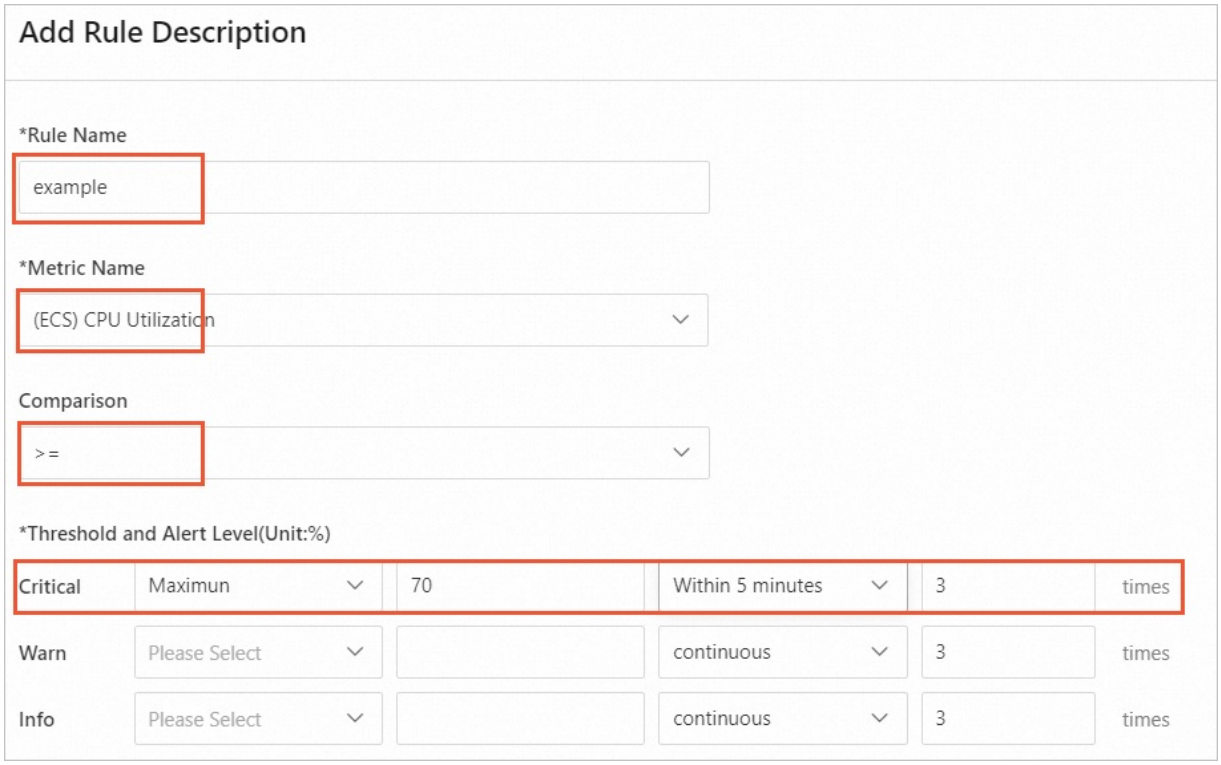

10.5.2.2. Manage alert rules

You can modify, disable, enable, and delete alert rules on the Alert Rule page in CloudMonitor.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Alert Service** > **Alert Rule**.
4. Find the alert rule that you want to manage and perform required operations. The following table describes supported operations.

Operation	Procedure
View the details of an alert rule	Find the alert rule whose details you want to view and click**Rule Details** in the **Actions** column. In the **Alert Rules/Alert Logs/Faulty Instances** panel, view the details of the alert rule. The rule details include the **product name**, **alert status**, **metric name**, **associated resources**, and others.
View the alert history of an alert rule	Find the alert rule whose alert history you want to view and click **Alert Logs** in the **Actions** column. In the **Alert Rules/Alert Logs/Faulty Instances** panel, view the alert history of the alert rule.

Modify an alert rule	i. On the Alert Rules page, find the alert rule that you want to modify and click **Modify** in the **Actions** column. ii. In the **Modify Alert Rule** panel, modify the **Rule Description**, **Mute For**, **Effective Period**, **Alert Callback**, and **Alert Contact Group** parameters. For more information about the parameters, see Create an alert rule. iii. After the modification is complete, click **OK**.
Disable one or more alert rules	If you stop a cloud service for maintenance or upgrade purposes, you can disable the alert rules that are associated with the cloud service. This way, manual operations on the service do not trigger CloudMonitor alert notifications. **Note** Make sure that the alert rule is in the **Normal** state before you disable the alert rule. ◦ Disable a single alert rule a. On the Alert Rules page, find the alert rule that you want to disable, click the ··· icon in the **Actions** column, and then click **Disable**. b. In the **Disable Alert Rule** message, click **OK**. ◦ Disable multiple alert rules at a time a. Select the check boxes of the alert rules that you want to disable and click **Disable** in the lower-left corner of the alert rule list. b. In the **Disable Alert Rule** message, click **OK**.
Enable one or more alert rules	After you maintain or upgrade a cloud service, you can enable all alert rules in the Disabled state. ◦ Enable a single alert rule a. Find the alert rule that you want to enable, click the ··· icon in the **Actions** column, and then click **Enable**. b. In the **Enable Alert Rule** message, click **OK**. ◦ Enable multiple alert rules at a time a. Select the check boxes of the alert rules that you want to enable and click **Enable** in the lower-left corner of the alert rule list. b. In the **Enable Alert Rule** message, click **OK**.

Delete one or more alert rules	◦ Delete a single alert rule a. Find the alert rule that you want to delete, click the ··· icon in the Actions column, and then click **Delete**. b. In the **Delete Alert** message, click **OK**. ◦ Delete multiple alert rules at a time a. Select the check boxes of the alert rules that you want to delete and click **Delete** in the lower-left corner of the alert rule list. b. In the **Delete Alert** message, click **OK**.
Export and download an alert rule report	i. On the **Alert Rules** page, click **Export report**. ii. In the **Create an alarm export task** dialog box, enter a name in the **Task name** field. iii. Click **OK**. iv. In the left-side navigation pane, click **Export Report**. On the **Report download** page, click the **Alarm rules** tab. v. Find the report that you exported and click **Download Report** in the **Actions** column. vi. In the **Download report task** message, click **OK**.

10.5.3. Alert templates

10.5.3.1. Common alerting templates

CloudMonitor provides preset alerting templates for Elastic Computer Service (ECS), ApsaraDB RDS, Object Storage Service (OSS), Server Load Balancer (SLB), Elastic IP Address, Apsara Stack DNS, Hologres, and KVStore for Redis. You can directly apply the preset alerting templates on resource sets or applications groups to implement alerting based on common metrics. You can also create custom templates on top of the preset templates to suit your business needs.

CloudMonitor provides preset alerting templates that help you quickly implement alerting on cloud resources based on common metrics.

ECS

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
cpu_total	>=85.00%	5	INFO	The CPU utilization of an ECS instance reaches 85% for five consecutive times.
	>=95.00%	5	WARN	The CPU utilization of an ECS instance reaches 95% for five consecutive times.

memory_usedutilization	>=80.00%	5	INFO	The memory usage of an ECS instance reaches 80% for five consecutive times.
diskusage_utilization	>=85.00%	5	INFO	The disk usage of an ECS instance reaches 85% for five consecutive times.
	>=95.00%	5	WARN	The disk usage of an ECS instance reaches 95% for five consecutive times.

ApsaraDB RDS

ApsaraDB RDS (New)

The templates are applicable to ApsaraDB RDS instances that run the MySQL, PolarDB for Oracle, and SQL Server engines.

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
ConnectionUsage_new	>=80.00%	5	INFO	The connection usage of an ApsaraDB RDS instance reaches 80% for five consecutive times.
CpuUsage_new	>=65.00%	5	INFO	The CPU utilization of an ApsaraDB RDS instance reaches 65% for five consecutive times.
DataDelay	>=10	5	WARN	The latency of a read-only ApsaraDB RDS instance reaches 10 seconds for five consecutive times.
DiskUsage_new	>=70.00%	5	INFO	The disk usage of an ApsaraDB RDS instance reaches 70% for five consecutive times.
IOPSUsage_new	>=80.00%	5	INFO	The IOPS usage of an ApsaraDB RDS instance reaches 80% for five consecutive times.
MemoryUsage_new	>=80.00%	5	INFO	The memory usage of an ApsaraDB RDS instance reaches 80% for five consecutive times.

ApsaraDB RDS for PostgreSQL

The templates are applicable only to ApsaraDB RDS for PostgreSQL instances.

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
cpu_usage	>=80.00%	5	INFO	The CPU utilization of an ApsaraDB RDS for PostgreSQL instance reaches 80% for five consecutive times.
mem_usage	>=80.00%	5	INFO	The memory usage of an ApsaraDB RDS for PostgreSQL instance reaches 80% for five consecutive times.
local_fs_size_usage	>=85.00%	5	WARN	The disk usage of an ApsaraDB RDS for PostgreSQL instance reaches 85% for five consecutive times.
ReadLag_new	>=7200	5	INFO	The latency of a read-only ApsaraDB RDS for PostgreSQL instance reaches 7,200 seconds for five consecutive times.

OSS

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
Availability	<100%	3	INFO	The availability of an OSS bucket goes below 100% for three consecutive times.

SLB

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
InstanceTrafficRXNew_SUM	>=4000000000	1	WARN	The inbound bandwidth of an SLB instance reaches 4 Gbit/s.
	>=4500000000	1	Critical	The inbound bandwidth of an SLB instance reaches 4.5 Gbit/s.
	>=4000000000	1	WARN	The outbound bandwidth of an SLB instance reaches 4 Gbit/s.

InstanceTrafficTXNew_SUM	>=4500000000	1	Critical	The outbound bandwidth of an SLB instance reaches 4.5 Gbit/s.
InstanceNewConnection	>=16000	1	WARN	The number of new connections to an SLB instance per second reaches 16,000.
	>=18000	1	Critical	The number of new connections to an SLB instance per second reaches 18,000.
InstanceMaxConnection	>=160000	1	WARN	The maximum number of concurrent connections per second on the port of an SLB instance reaches 160,000.
	>=180000	1	Critical	The maximum number of concurrent connections per second on the port of an SLB instance reaches 180,000.
InstanceDropConnection	>=10	3	WARN	The number of connections dropped by an SLB instance per second reaches 10 for three consecutive times.
	>=200	3	Critical	The number of connections dropped by an SLB instance per second reaches 200 for three consecutive times.
InstanceDropTrafficRX	>=100000	3	WARN	An SLB instance drops 100,000 bits of inbound traffic or more per second for three consecutive times.
	>=2000000	3	Critical	An SLB instance drops 2,000,000 bits of inbound traffic or more per second for three consecutive times.
InstanceDropTrafficTX	>=100000	3	WARN	An SLB instance drops 100,000 bits of outbound traffic or more per second for three consecutive times.
	>=2000000	3	Critical	An SLB instance drops 2,000,000 bits of outbound traffic or more per second for three consecutive times.
InstanceUnhealthyServerCount	One instance more than the previous day	3	WARN	The number of abnormal ECS instances detected in the health check of an SLB instance grows by 1 for three consecutive days.

EIP

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
out_ratelimit_drop_speed	>0	3	INFO	The maximum outbound packet loss rate of an EIP is greater than 0 for three consecutive times.
in_ratelimit_drop_speed	>0	3	INFO	The maximum inbound packet loss rate of an EIP is greater than 0 for three consecutive times.

Apsara Stack DNS

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
ParsingTotal	>40000	3	WARN	The total number of DNS domain names resolved is greater than 40,000 for three consecutive times.
	>45000	3	Cirtical	The total number of DNS domain names resolved is greater than 45,000 for three consecutive times.

KVStore for Redis

KVStore for Redis (Standard Architecture)

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
StandardCpuUsage	>=65.00%	5	WARN	The CPU utilization of a KVStore for Redis instance reaches 65% for five consecutive times.
StandardIntranetInRatio	>=80.00%	5	WARN	The inbound bandwidth usage of a KVStore for Redis instance reaches 80% for five consecutive times.

StandardIntranetOutRatio	>=80.00%	5	WARN	The outbound bandwidth usage of a KVStore for Redis instance reaches 80% for five consecutive times.
StandardMemoryUsage	>=80.00%	5	WARN	The memory usage of a KVStore for Redis instance reaches 80% for five consecutive times.
StandardQPSUsage	>=80.00%	5	WARN	The QPS usage of a KVStore for Redis instance reaches 80% for five consecutive times.
StandardConnectionUsage	>=80.00%	5	WARN	The connection usage of a KVStore for Redis instance reaches 80% for five consecutive times.

KVStore for Redis (Cluster Architecture)

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
ShardingCpuUsage	>=65.00%	5	WARN	The CPU utilization of a KVStore for Redis instance reaches 65% for five consecutive times.
ShardingIntranetIn	>=80.00%	5	WARN	The inbound bandwidth usage of a KVStore for Redis instance reaches 80% for five consecutive times.
ShardingIntranetOut	>=80.00%	5	WARN	The outbound bandwidth usage of a KVStore for Redis instance reaches 80% for five consecutive times.
ShardingMemoryUsage	>=80.00%	5	WARN	The memory usage of a KVStore for Redis instance reaches 80% for five consecutive times.
ShardingQPSUsage	>=80.00%	5	WARN	The QPS usage of a KVStore for Redis instance reaches 80% for five consecutive times.
ShardingConnectionUsage	>=80.00%	5	WARN	The connection usage of a KVStore for Redis instance reaches 80% for five consecutive times.

Hologres

Monitoring item	Threshold	Consecutive occurrences	Alert level	Rule description
cpu_usage	>=99.00 %	3	WARN	The CPU utilization of a Hologres instance reaches 99% for three consecutive times.
memory_usage	>=90.00 %	3	WARN	The memory usage of a Hologres instance reaches 90% for three consecutive times.

10.5.3.2. Create an alert template

If you want to monitor a large number of resources, you can use the alert template feature. This feature helps you save multiple alert rules on the metrics of various cloud services as a template. You can create or modify alert rules by using alert templates. This saves repeated workloads of configuring each alert rule.

Background information

You must use alert templates together with application groups. You can create an application group first. Then, you can create an alert template and apply the alert template to the application group to create alert rules for the application group. This simplifies the creation and maintenance of alert rules. For more information about how to create an application group, see Create an application group.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Alert Service** > **Alert Templates**.
4. On the **Alert Templates** page, click **Create Alert Template**.
5. In the **Create alarm Template** panel, configure the following parameters.

Parameter	Description
Template Name	Specify a name for the alert template. The name can contain letters, digits, and underscores (_). We recommend that the name be 1 to 30 characters in length.
Description	Optional. Specify the description of the alert template. We recommend that the description be 1 to 100 characters in length.
Clone	Optional. Select an existing alert template and modify the configurations to create an alert template.
Select Alibaba Cloud Service	Select one or more cloud services.

6. Click **Add Rules** under the cloud service that you specify to configure an alert rule. The following table describes the parameters of the alert rule.

Parameter	Description
Rule Name	The name of the alert rule.
Metric Name	The monitoring metric that may trigger alerts. For more information about the metrics of cloud services, see the "Metrics" topic.
Comparison	The comparison condition between metric data and thresholds. If the metric data exceeds the specified thresholds, the alert rule is triggered. Valid comparison conditions: >, >=, <, <=, =, !=, Compared With Yesterday Rise, and Compared With Yesterday Decline.
Threshold and Alert Level	The conditions, alert thresholds, and alert levels of the alert rule. You can specify multiple alert levels with different conditions. ◦ **Alert levels:** ▪ **Critical** ▪ **Warn** ▪ **Info** ◦ **Number of times for which the metric data is equal to or greater than the threshold: The number of times determines whether the alert notification is sent.** Valid values: ▪ **1 Consecutive Cycles** ▪ **2 Consecutive Cycles** ▪ **3 Consecutive Cycles** ▪ **4 Consecutive Cycles** ▪ **5 Consecutive Cycles** ▪ **10 Consecutive Cycles** ▪ **15 Consecutive Cycles** ▪ **20 Consecutive Cycles**

7. After you add rules for each cloud service, click **OK**.

10.5.3.3. Manage alert templates

You can modify, clone, export, and delete alert templates on the Alert Templates page in CloudMonitor.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Alert Service** > **Alert Templates**.

4. Select an alert template, and perform required operations. The following table describes supported operations.

Operation	Procedure
Modify an alert template	i. Find the alert template that you want to modify and click **Modify** in the **Actions** column. ii. In the **Modify alarm Template** panel, modify the description, Alibaba Cloud services, and rules of the alert template. iii. Click **OK**. **Note** You cannot modify default templates.
Clone an alert template	You can clone an alert template to copy the template configurations to a new alert template. i. Find the alert template that you want to clone and click **Clone** in the **Actions** column. ii. In the **Clone alarm Template** panel, specify a name for the new template, select Alibaba Cloud services, and configure alert rules based on your business requirements. iii. Click **OK**. iv. In the **Alert Template Created/Modified** message, click **Cancel**. **Note** If you click Cancel, the alert template is created but not applied to the application group.
Delete one or more alert templates	**Warning** Deleted alert templates cannot be restored. ◦ Delete a single alert template a. Find the alert template that you want to delete and click **Delete** in the **Actions** column. b. In the **Delete Alert Template** message, click **OK**. ◦ Delete multiple alert templates at a time a. Select the check boxes of the alert templates that you want to delete and click **Delete** in the lower-left corner of the page. b. In the **Delete Alert Template** message, click **OK**.

Apply one or more alert templates to application groups or resource sets	Alert templates must be used together with application groups. You can apply alert templates to application groups to create alert rules for resources from different cloud services. Make sure that application groups or resource sets, and alert templates are available. For information about how to create the resources, see Create an alert template, Create an application group, and Create a resource set. **Important** After you select and apply an alert template to a specified application group or resource set, CloudMonitor deletes the existing alert rules that are created based on the original alert template except for the process monitoring and availability monitoring rules. Then, CloudMonitor creates new alert rules based on the new alert template. We recommend that you apply an alert template to up to five application groups or resource sets. Otherwise, the alert template may fail to take effect for some application groups or resource sets. ◦ Apply an alert template to application groups or resource sets a. Find the alert template that you want to apply, click the icon in the **Actions** column, and then click **Apply alarm Template**. b. In the **Apply Templates to Groups** dialog box, configure the **Resource Group**, **Resource selection**, **Mute For**, **Effective Period**, **Alert Callback**, and **Alert Contact Group** parameters. You can set the **Resource Group** parameter to **Resource set** or **Application grouping**. c. Click **OK**. d. In the **Apply Template to Group** message, click **OK**. ◦ Apply multiple alert templates to application groups or resource sets a. Select the check boxes of the alert templates that you want to apply, and then click **Apply Templates** in the lower-left corner. b. In the **Apply Templates to Groups** dialog box, configure the **Resource Group**, **Resource selection**, **Mute For**, **Effective Period**, **Alert Callback**, and **Alert Contact Group** parameters. You can set the **Resource Group** parameter to **Resource set** or **Application grouping**. c. Click **OK**. d. In the **Apply Template to Group** message, click **OK**.

Export and download a report on alert templates	i. On the **Alert Templates** page, click **Export report**. ii. In the **Create an alarm export task** dialog box, enter a name in the **Task name** field. iii. Click **OK**. iv. In the left-side navigation pane, click **Export Report**. On the **Report download** page, click the **Alarm Template** tab. v. Find the report that you exported and click **Download** report in the **Actions** column. vi. In the **Download report task** message, click **OK**. **Note** The report is in the **.xlsx** format. You can find the downloaded report in the default download directory of your browser.

10.6. Statistical analysis

You can use CloudMonitor to analyze the monitoring data of all supported cloud services in a specified organization.

10.6.1. Sorting

You can use CloudMonitor to sort and analyze the monitoring data of supported cloud services in a specific organization.

Limits

- CloudMonitor allows you to sort the data of up to 1,000 instances of a cloud service. Data computing fails if the number of instances exceeds this limit.
- You can sort only data of a specific day. You must select a day out of the last 30 days.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Statistical Analysis > Sorting**.
4. Select a time range (the last 1 hour, 3 hours, 6 hours, 12 hours, or 1 day) or specify a custom time range. Then, select a cloud service and a metric.
5. Configure the following parameters in a sorting rule:

- **Ascending order** or **Descending order**: If you select Ascending order, the monitoring data is sorted in ascending order top down. If you select Descending order, the monitoring data is sorted in descending order top down.
- **Average**, **Minimum**, or **Maximum**: The average, minimum, or maximum values of the specified metric of instances are sorted.
- **10**, **50**, **100**, or **200**: The top 10, 50, 100, or 200 monitoring data entries are displayed.

6. Click **Order**.

10.7. Event monitoring

The Event Monitoring module provides a unified portal for you to query and manage system events of CloudMonitor, Elastic Compute Service (ECS), Hybrid Disaster Recovery (HDR), Cloud Config, ASR for Backup and Recovery, Elastic Desktop Service (EDS), Open Network Services Platform (ONSP), and Resource Orchestration Service (ROS). You can view the status of the cloud services and receive alert notifications. This helps you troubleshoot business exceptions at the earliest opportunity.

10.7.1. System events

10.7.1.1. View system events

Background information

The following table describes the cloud services supported by CloudMonitor and their system events.

Cloud service	System event
CloudMonitor	• Agent_Status_Stopped • Agent_Status_Running
	• Disk:ConvertToPostpaidCompleted • Disk:DiskOperationCompleted • Disk:ErrorDetected:Executed • Disk:ErrorDetected:Executing • Disk:OverduePaymentRelease • Disk:Stalled:Executed • Disk:Stalled:Executing • Instance:AutoReactivateCompleted • Instance:BurstablePerformanceRestricted • Instance:InstanceFailure.PerformanceImpact:Executed • Instance:InstanceFailure.PerformanceImpact:Executing • Instance:InstanceFailure.Reboot:Executed • Instance:InstanceFailure.Reboot:Executing • Instance:LiveMigrationAcrossDDH • Instance:PerformanceModeChange • Instance:PreemptibleInstanceInterruption • Instance:Security.TpmAlert:Executed • Instance:Security.TpmAlert:Executing • Instance:StateChange • Instance:SystemFailure.Delete:Avoided • Instance:SystemFailure.Delete:Executed • Instance:SystemFailure.Delete:Executing • Instance:SystemFailure.PerformanceImpact:Executed • Instance:SystemFailure.PerformanceImpact:Executing

Elastic Compute Service (ECS)	• Instance:SystemFailure.Reboot:Executed • Instance:SystemFailure.Reboot:Executing • Instance:SystemFailure.Reboot:Failed • Instance:SystemFailure.Redeploy:Avoided • Instance:SystemFailure.Redeploy:Canceled • Instance:SystemFailure.Redeploy:Executed • Instance:SystemFailure.Redeploy:Executing • Instance:SystemFailure.Redeploy:Scheduled • Instance:SystemFailure.Stop:Executed • Instance:SystemFailure.Stop:Executing • Instance:SystemMaintenance.IsolateErrorDisk:Avoided • Instance:SystemMaintenance.IsolateErrorDisk:Canceled • Instance:SystemMaintenance.IsolateErrorDisk:Executed • Instance:SystemMaintenance.IsolateErrorDisk:Executing • Instance:SystemMaintenance.IsolateErrorDisk:Failed • Instance:SystemMaintenance.IsolateErrorDisk:Inquiring • Instance:SystemMaintenance.PerformanceImpact:Executed • Instance:SystemMaintenance.PerformanceImpact:Executing • Instance:SystemMaintenance.ReInitErrorDisk:Avoided • Instance:SystemMaintenance.ReInitErrorDisk:Canceled • Instance:SystemMaintenance.ReInitErrorDisk:Executed • Instance:SystemMaintenance.ReInitErrorDisk:Executing • Instance:SystemMaintenance.ReInitErrorDisk:Failed • Instance:SystemMaintenance.ReInitErrorDisk:Inquiring • Instance:SystemMaintenance.Reboot:Avoided • Instance:SystemMaintenance.Reboot:Canceled • Instance:SystemMaintenance.Reboot:Executed • Instance:SystemMaintenance.Reboot:Executing • Instance:SystemMaintenance.Reboot:Failed • Instance:SystemMaintenance.Reboot:Inquiring • Instance:SystemMaintenance.Reboot:Scheduled • Instance:SystemMaintenance.RebootAndIsolateErrorDisk:Avoided • Instance:SystemMaintenance.RebootAndIsolateErrorDisk:Canceled • Instance:SystemMaintenance.RebootAndIsolateErrorDisk:Executed • Instance:SystemMaintenance.RebootAndIsolateErrorDisk:Executing • Instance:SystemMaintenance.RebootAndIsolateErrorDisk:Inquiring • Instance:SystemMaintenance.RebootAndReInitErrorDisk:Avoided • Instance:SystemMaintenance.RebootAndReInitErrorDisk:Canceled • Instance:SystemMaintenance.RebootAndReInitErrorDisk:Executed • Instance:SystemMaintenance.RebootAndReInitErrorDisk:Executing • Instance:SystemMaintenance.RebootAndReInitErrorDisk:Inquiring • Instance:SystemMaintenance.Redeploy:Avoided

	• Instance:SystemMaintenance.Redeploy:Canceled • Instance:SystemMaintenance.Redeploy:Executed • Instance:SystemMaintenance.Redeploy:Executing • Instance:SystemMaintenance.Redeploy:Inquiring • Instance:SystemMaintenance.Redeploy:Scheduled • Instance:SystemMaintenance.Stop:Avoided • Instance:SystemMaintenance.Stop:Canceled • Instance:SystemMaintenance.Stop:Executed • Instance:SystemMaintenance.Stop:Executing • Instance:SystemMaintenance.Stop:Failed • Instance:SystemMaintenance.Stop:Scheduled • NetworkInterface:NetworkInterfaceOperateCompleted • Snapshot:CreateSnapshotCompleted • Instance:VmLivemigrateException • Instance:VmDownCausedByNcDown • Instance:VmDownEnd • Instance:VmPanicEvent • Instance:VmCrashEvent
Hybrid Disaster Recovery (HDR)	• PrimarySiteVpc:Deleted • SecondarySiteVpc:Deleted • PrimarySiteNetwork:Deleted • SecondarySiteNetwork:Deleted • PrimarySecurityGroup:Deleted • SecondarySecurityGroup:Deleted • SecurityGroup:Inconsistency • ProductSiteNetwork:Inconsistency • DrSiteNetwork:Inconsistency • ProductSiteStorage:Inconsistency • DrSiteStorage:Inconsistency • PrimarySiteComputation:Inconsistency • DrSiteComputation:Inconsistency • ReplicationGroup:Inconsistency • ReplicationPair:Inconsistency • Gateway:Inconsistency • ShadowEcs:Inconsistency
Cloud Config	• Resource:RuleMismatch • Resource:RuleMatch

ASR for Backup and Recovery	• BackupPlan • RecoveryTask • MigratePlan • LICENSE
Elastic Desktop Service (EDS)	• ServiceStateChange • ServiceError • SeviceTimeout
Open Network Services Platform (ONSP)	• onsp:vnfMemberError • onsp:vnfError • onsp:endpointError • onsp:clusterCreateError • onsp:clusterDeleteError • onsp:endpointCreateError • onsp:endpointDeleteError • onsp:RouteEntryCreateError • onsp:RouteEntryDeleteError • onsp:EipCreateError • onsp:EipDeleteError
Resource Orchestration Service (ROS)	• Stack:StatusChange • StackResource:StatusChange • StackGroup:OperationStatusChange

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over **Products** and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Event Monitoring** > **System Events**.
4. On the **Event Monitoring** tab, specify the product, event type, event name, and time range and click **Search**.
5. In the event list, click **Details** corresponding to a record. In the panel that appears, you can view the details presented in the **JSON** format.

10.7.1.2. Create a system event-triggered alert rule

You can create alert rules for system events of cloud service. If an exception occurs, an alert is generated based on the rule to help you analyze and troubleshoot the issue.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Event Monitoring** > **System Events**.
4. On the **System Event** page, click the **Event-triggered Alert Rules** tab, and then click **Create Event-triggered Alert Rule**.
5. In the **Create/Modify Event-triggered Alert Rule** panel, configure the parameters. The following table describes the parameters.

Parameter	Description
Alert Rule Name	The name of the system event-triggered alert rule. The name must be 1 to 30 characters in length and can contain letters, digits, and underscores (_).
Product Type	The Alibaba Cloud service to which the system event-triggered alert rule is applied. Valid values: **CloudMonitor**, **Hybrid Disaster Recovery**, and **Elastic Compute Service (ECS)**.
Event Level	The level of the event that triggers alerts. Valid values: ◦ If you set the Product Type parameter to CloudMonitor, the valid value is fixed to **CRITICAL**. ◦ If you set the Product Type parameter to Hybrid Disaster Recovery, the valid value is fixed to **CRITICAL**. ◦ If you set the Product Type parameter to Elastic Compute Service (ECS), the valid values are **CRITICAL** and **WARN**.
Event Name	The name of the event that triggers alerts.

10.7.1.3. Manage system event-triggered alert rules

You can modify, disable, enable, and delete system event-triggered alert rules on the System Event page in CloudMonitor.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, choose **Event Monitoring** > **System Events**.
4. Click the **Event-triggered Alert Rules** tab, find the system event-triggered alert rule that you want to manage, and then perform required operations. The following table describes supported operations.

Operation	Procedure

Modify a system event-triggered alert rule	i. Find the system event-triggered alert rule that you want to modify, and click **Modify** in the **Actions** column. In the panel that appears, modify the **Product Type**, **Event Level**, and **Event Name** parameters based on your business requirements. For more information about the parameters, see Create a system event-triggered alert rule. ii. Click **OK**.
Disable one or more system event-triggered alert rules	**Note** You can disable only event-triggered alert rules in the Enabled state. After you disable an alert rule, the rule no longer takes effect. ◦ Disable a single system event-triggered alert rule a. Find the system event-triggered alert rule that you want to disable and click **Disable** in the **Actions** column. b. In the message that appears, click **OK**. ◦ Disable multiple system event-triggered alert rules at a time a. Select the check boxes of system event-triggered alert rules that you want to disable and click **Disable** in the lower-left corner of the alert rule list. b. In the message that appears, click **OK**.
Enable one or more system event-triggered alert rules	You can enable system event-triggered alert rules in the Disabled state. ◦ Enable a single system event-triggered alert rule a. Find the system event-triggered alert rule that you want to enable and click **Enable** in the **Actions** column. b. In the message that appears, click **OK**. ◦ Enable multiple system event-triggered alert rules at a time a. Select the check boxes of the system event-triggered alert rules that you want to enable, and click **Enable** in the lower-left corner of the alert rule list. b. In the message that appears, click **OK**.
Delete one or more system event-triggered alert rules	◦ Delete a single system event-triggered alert rule a. Find the system event-triggered alert rule that you want to delete and click **Delete** in the **Actions** column. b. In the message that appears, click **OK**. ◦ Delete multiple system event-triggered alert rules at a time a. Select the check boxes of system event-triggered alert rules that you want to delete, and then click **Delete** in the lower-left corner of the alert rule list. b. In the message that appears, click **OK**.

View the records of system event-triggered alerts	After an event-triggered alert rule is created or enabled, CloudMonitor records the information about alerts that are generated based on the rule. i. Click the **Event-triggered Alert History** tab. ii. Optional. Enter a rule ID or rule name, specify a time range, and click the icon to search for alerts that are generated based on the rule. An alert record includes the following information: **product**, **instance name**, **point in time when the alert occurred**, **event level**, and **event name**.

10.8. Application group

Application groups allow you to manage Elastic Compute Service (ECS) instances by group based on your business requirements. You can use an application group to manage business-specific ECS instances in a centralized manner. You can manage the alert rules and view monitoring data of different application groups. This greatly simplifies management complexity and improves efficiency.

10.8.1. Create an application group

The application group feature allows you to manage resources from different cloud services and regions by group. You can use application groups to manage business-specific alert rules and view business-specific monitoring data to improve O&M efficiency.

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Application Group**.
4. On the **Application Groups** page, click **Create Application Group**.
5. In the **Create Application Group** panel, configure the parameters. The following table describes the parameters.

Parameter	Description
Creation Method	◦ **Manually Create**: If you want to manage existing instances based on your business requirements, you can manually create a standard application group. ◦ **Dynamic creation**: If you have configured a resource set, you can create a resource group for the resource set, specify the Elastic Compute Service (ECS) instances that you do not want to manage, and then create a standard application group.
Resource Group	Select a resource set that belongs to the current organization to create a resource group.
Application Group Name	Specify a name for the application group that you want to create.

Add Resource	**Note** This parameter is valid and required when you set the Creation Method parameter to Manually Create. Select one or more ECS instances in the list of ECS instances that belong to the current resource set. The application group contains all ECS instances that you select.
Exclude resource items	**Note** This parameter is valid and required when you set the Creation Method parameter to Dynamic creation. Select one or more ECS instances in the list of ECS instances that belong to the current resource set. The application group does not contain the ECS instances that you select.

6. Click **OK**.

10.8.2. Manage application groups

You can modify or delete application groups on the Application Groups page in CloudMonitor.

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Application Group**. On the **Application Groups** page, manage the application groups. The following table describes supported operations.

Operation	Procedure

View the information about an application group	i. Optional. In the top navigation bar, select the organization in which you want to manage application groups. On the **Application Groups** page, enter an application group name or ID, and then click **Search** to search for the application group whose information you want to view. ii. Click the name of the application group whose information you want to view. On the details page of the application group, click tabs in the left-side navigation pane to view the content. ▪ **Group Overview**: You can view the basic information about the application group and the associated cloud services. ▪ **Group Resources**: You can view the details of the resources in the application group. ▪ **Availability Monitoring**: You can view the availability monitoring data of all Elastic Compute Service (ECS) instances in the current application group.
Modify an application group	i. Optional. In the top navigation bar, select the organization in which you want to manage application groups. On the **Application Groups** page, enter an application group name or ID, and then click **Search** to search for the application group that you want to modify. ii. Click the ••• icon of the application group that you want to modify, and then click **Modify Group**. iii. In the **Modify application Group** panel, modify the **Creation Method** and **Application Group Name** parameters, and resources in the application group. iv. Click **OK**.
Modify the basic information about an application group	i. Click the name of the application group that you want to manage. In the left-side navigation pane, click the **Group Overview** tab. ii. In the **Basic Information** section, modify the information. ▪ Modify the **Group Name** parameter. a. Click the icon next to **Group Name** to modify the name of the application group. b. After the modification is complete, click **OK**. ▪ Modify the **Alert Contact Group** parameter. a. Click the icon next to **Alert Contact Group** and select one or more alert contact groups from the drop-down list. b. After the modification is complete, click **OK**.

Delete one or more resources from the application group	i. Click the name of the application group that you want to manage. In the left-side navigation pane, click the **Group Resources** tab. ii. In the list of ECS instances, delete one or more instances. ▪ Delete a single instance a. Find the instance that you want to delete, and click **Delete** in the **Actions** column. b. In the message that appears, click **OK**. ▪ Delete multiple application groups at a time a. Select the check boxes of the instances that you want to delete and click **Delete** in the lower-left corner of the instance list. b. In the message that appears, click **OK**.
Delete one or more application groups	**Note** If you delete an application group, all alert rules and event-triggered alerts of the application group are deleted. The instances in the application group are automatically migrated back to the pool that contains all instances. i. Optional. In the top navigation bar, select an organization in which you want to manage application groups. On the **Application Groups** page, enter an application group name or ID, and then click **Search** to search for the application group that you want to manage. ii. On the Application Groups page, delete one or more application groups. ▪ Delete a single application group a. Click the ••• icon of the application group that you want to delete and click **Delete Group**. b. In the **Delete Group** message, click **OK**. ▪ Delete multiple application groups at a time a. Select the check boxes of the application groups that you want to delete and click **Delete** in the lower-left corner of the application group list. b. In the **Delete Group** message, click **OK**.

10.8.3. Availability monitoring

10.8.3.1. Create an availability monitoring task

CloudMonitor provides the availability monitoring feature to help you check whether a local service or cloud server responds as expected. You can specify the path and port for the local service or cloud server. If the local service or cloud server times out or returns an error code, CloudMonitor sends an alert notification for you to troubleshoot the issue at the earliest opportunity.

Procedure

> **Note**
>
> You can create availability monitoring tasks for only Elastic Compute Service (ECS) instances.

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Application Group**.
4. On the **Application Groups** page, click the name of the application group that you want to manage. On the details page of the application group, click the **Availability Monitoring** tab in the left-side navigation pane.
5. Click **Create Availability Monitoring Task**. In the **Create/Modify Availability Monitoring Task** panel, complete the configuration steps as prompted by the wizard. The following table describes the parameters.

Step and parameter		Description
	Task Name	The name of the availability monitoring task. The name must be 4 to 50 characters in length, and can contain letters, digits, underscores (_), and hyphens (-).
	Monitoring Nodes	Select one or more instances that you want to monitor in the application group. **Note** If you create over 500 tasks for an instance, some of the tasks may fail.
	Detection Target	Default value: **URL or IP Address**.
	Detection Type	◦ **HTTP(S):** Enter the URL of the object that you want to monitor. ◦ **TELNET:** Enter the IP address of the object that you want to monitor. ◦ **PING:** Enter the IP address of the object that you want to monitor.

Create Task	Request Method	The HTTP request method. Valid values: **HEAD**, **GET**, and **POST**. **Note** This parameter is required if you set the **Detection Target** parameter to **URL or IP Address** and the **Detection Type** parameter to HTTP(S).
	Monitoring Frequency	The interval at which detection requests are sent. Default value: 1 Minutes. **Note** This parameter is required if you set the **Monitored Target** parameter to **URL or IP Address** and the **Detection Type** parameter to HTTP(S).
	Headers	The headers of the HTTP or HTTPS request. Format: **parameter1:value1** . Separate multiple headers with commas (,). **Note** This parameter is required if you set the **Detection Target** parameter to **URL or IP Address** and the **Detection Type** parameter to **HTTP(S)**.
	Match Response Content	The response content that you want to match and the method to match the response content. If you specify response content, the monitoring task checks whether the first 64 KB of the HTTP response body contains the response content that you specify. Valid values: ◦ **Generate Alerts If Response Contains** ◦ **Generate Alerts If Response Does Not Contain** **Note** This parameter is available in the **Advanced Configuration** section if you set the **Detection Target** parameter to **URL or IP Address** and the **Detection Type** parameter to **HTTP(S)**.

Configure Alert Rule	Status Code	The range of the status codes that you want to match. CloudMonitor generates an alert if the status code that is returned meets the condition. For more information about status codes, see Status codes. If the value of the **Status Code** parameter or the **Response Time** parameter reaches the threshold, an alert is triggered. CloudMonitor sends alert notifications to the alert group of a specified application group.
	Response Time	The periods of the response time that you want to match. CloudMonitor generates an alert if the response time meets the condition. Unit: milliseconds. If the value of the **Status Code** parameter or the **Response Time** parameter meets the threshold, an alert is triggered. CloudMonitor sends alert notifications to the alert group of a specified application group.
	Notification Methods	Select the alert level and the methods that are used to send alert notifications. Valid alert levels: ◦ **Critical** ◦ **Warning** ◦ **Info**
	Alert Callback	The URL that can be accessed over the Internet. CloudMonitor sends HTTP POST requests to push alert notifications to the specified URL. **Note** You can enter only an HTTP URL.
	Mute For	The interval at which you want CloudMonitor to send alert notifications again if the alert is not cleared. Valid values: **5 Minutes**, **10 Minutes**, **15 Minutes**, **30 Minutes**, **60 Minutes**, **3 Hours**, **6 Hours**, **12 Hours**, and **24 Hours**. An alert is triggered if the conditions of an alert rule are met. CloudMonitor does not resend an alert notification when the alert is triggered within the mute period. If the alert is not cleared after the mute period ends, CloudMonitor resends alert notifications. **Note** You can click **Advanced Configuration** to configure this parameter for the alert rule.

	Effective Period	The period during which the alert rule is effective. CloudMonitor sends alert notifications within the effective period and only records alerts beyond the effective period. **Note** You can click **Advanced Configuration** to configure this parameter for the alert rule.

10.8.3.2. Manage availability monitoring tasks

You can modify, disable, enable, and delete availability monitoring tasks on the Availability Monitoring tab of an instance in CloudMonitor.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, move your pointer over Products and choose **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Application Group**.
4. On the **Application Groups** page, click the name of the application group that you want to manage. On the page that appears, click the **Availability Monitoring** tab in the left-side navigation pane. On the **Availability Monitoring** tab, perform required operations. The following table describes the supported operations.

Operation	Procedure
Modify an availability monitoring task	i. Find the availability monitoring task that you want to modify and click **Edit** in the **Actions** column. ii. In the **Create/Modify Availability Monitoring** panel, modify the parameters as prompted by the wizard. For more information about the parameters, see Create an availability monitoring task. iii. After the modification is complete, click **OK**.
Disable or enable one or more availability monitoring tasks	**Note** You can disable one or more availability monitoring tasks in the **Enabled** state, or enable one or more availability monitoring tasks in the **Disabled** state. ◦ Disable or enable a single availability monitoring task a. Click **Disable** or **Enable** in the **Actions** column of the task. b. In the **Disable** or **Enable** message, click **OK**. ◦ Disable or enable multiple availability monitoring tasks at a time a. Select the check boxes of the availability monitoring tasks that you want to disable or enable, and click **Batch Disable** or **Batch Enable** in the lower-left corner of the task list. b. In the **Disable** or **Enable** message, click **OK**.

Delete one or more availability monitoring tasks	◦ Delete a single availability monitoring task a. Click the ••• icon in the **Actions** column of the availability monitoring task and click **Delete**. b. In the **Delete** message that appears, click **OK**. ◦ Delete multiple availability monitoring tasks at a time a. Select the check boxes of the availability monitoring tasks that you want to delete and click **Batch Delete** in the lower-left corner of the task list. b. In the **Delete** message that appears, click **OK**.

10.8.3.3. Status codes

If an error occurs when you use the availability monitoring feature of CloudMonitor to check the availability of a specified path or port, the corresponding status code is returned.

The following table describes status codes.

Protocol type	Status code	Description
HTTP	610	The status code returned because the connection times out. If CloudMonitor does not receive a response within 5 seconds after CloudMonitor sends an HTTP request, CloudMonitor determines that the connection times out.
	611	The status code returned because the monitoring fails.
Telnet	630	The status code returned because the connection times out. If CloudMonitor does not receive a response within 5 seconds after CloudMonitor sends a Telnet request, CloudMonitor determines that the connection times out.
	631	The status code returned because the monitoring fails.

10.9. Download a report

You can export reports on the alert history, alert rules, and alert templates, and then download the reports to your on-premises machine.

Background information

- The downloaded reports are in the **.xlsx** format. You can find the downloaded reports in the default download directory of your browser.
- You can export and download the following types of reports:
  - Alert history report: The report contains the product namespace, instance that triggers the alert, alert rule name, alert rule ID, point in time when the alert occurs, report export status, alert level, alert contact group, metric name, alert expression, and monitoring data of each alert.

- Alert rule report: The report contains the product namespace, rule name, rule ID, dimension (instance-specific or not), monitored instance, group (resource set), rule status (enabled or disabled), report export status, alert contact group, metric, and alert expression of each alert rule.
- Alert template report: The report contains the template name, template ID, description, resource set, modification time, and template type (system template or not) of each alert template.

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, choose **Products** > **Monitoring and O&M** > **Cloud Monitor**.
3. In the left-side navigation pane, click **Export Report**.
4. On the **Report download** page, perform required operations. The following table describes supported operations.

Operation	Procedure
Download a report	i. Click the **Alarm history**, **Alarm rules**, or **Alarm Template** tab. ii. Optional. Select a state from the **Report status** drop-down list to filter reports that you want to download. iii. Click **Download report** in the **Actions** column of the report that you want to download. iv. In the **Download report task** message, click **OK**.
Delete a report	i. Click the **Alarm history**, **Alarm rules**, or **Alarm Template** tab. ii. Optional. Select a state from the **Report status** drop-down list to filter reports that you want to delete. iii. Click **Delete** in the **Actions** column of the report that you want to delete. iv. In the message that appears, click **OK**.