

Alibaba Cloud Apsara Stack Enterprise

Operations Guide

Version: 1901

Issue: 20190528

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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Note: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other contents.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK .
Courier font	It is used for commands.	Run the <code>cd /d C:/windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is a optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	It indicates that it is a required value, and only one item can be selected.	<code>switch {stand slave}</code>

Contents

Legal disclaimer	I
Generic conventions	I
1 Basic platform operations	1
1.1 Apsara Stack Operations (ASO).....	1
1.1.1 Apsara Stack Operations overview.....	1
1.1.2 Log on to Apsara Stack Operations.....	3
1.1.3 Web page introduction.....	5
1.1.4 Operation and maintenance dashboard.....	5
1.1.5 Alarm management.....	5
1.1.5.1 Overview.....	5
1.1.5.2 Alarm events.....	6
1.1.5.3 Alarm history.....	7
1.1.5.4 Alarm configuration.....	8
1.1.5.5 Alarm overview.....	9
1.1.5.6 Alarm subscription and push.....	9
1.1.6 Resource management.....	11
1.1.6.1 Overview.....	11
1.1.6.2 Physical servers.....	11
1.1.7 Inventory management.....	12
1.1.7.1 Overview.....	12
1.1.7.2 View the ECS inventory.....	12
1.1.7.3 View the SLB inventory.....	20
1.1.7.4 View the RDS inventory.....	20
1.1.7.5 View the OSS inventory.....	21
1.1.7.6 View the Table Store inventory.....	21
1.1.7.7 View the Log Service inventory.....	21
1.1.8 Products.....	22
1.1.9 ITIL management.....	22
1.1.9.1 Overview.....	22
1.1.9.2 Dashboard.....	23
1.1.9.3 Services.....	24
1.1.9.3.1 Basic functions.....	24
1.1.9.3.1.1 Overview.....	24
1.1.9.3.1.2 Manage requests.....	24
1.1.9.3.1.3 Manage tasks.....	25
1.1.9.3.2 Manage incidents.....	26
1.1.9.3.2.1 Create an incident request.....	26
1.1.9.3.2.2 Manage incident requests.....	27
1.1.9.3.2.3 Manage incident tasks.....	29
1.1.9.3.3 Manage problems.....	30
1.1.9.3.3.1 Create a problem request.....	30

1.1.9.3.3.2 Manage problem requests.....	32
1.1.9.3.3.3 Manage problem tasks.....	33
1.1.9.4 Version control.....	35
1.1.9.5 Process template configuration.....	35
1.1.9.6 Configure CAB or ECAB.....	38
1.1.10 API management.....	39
1.1.10.1 Overview.....	39
1.1.10.2 Category.....	39
1.1.11 Configurations.....	39
1.1.11.1 Overview.....	39
1.1.11.2 Modify a configuration item of a product.....	40
1.1.11.3 Restore the configuration value of a modified configuration item.....	40
1.1.11.4 Kernel configurations list.....	41
1.1.11.5 Kernel configurations actions.....	41
1.1.12 Offline backup.....	42
1.1.13 NOC.....	45
1.1.13.1 Overview.....	45
1.1.13.2 Dashboard.....	45
1.1.13.3 Network topology.....	46
1.1.13.4 Physical network integration.....	47
1.1.13.5 Password management.....	48
1.1.13.6 Configuration comparison.....	48
1.1.14 Full stack monitor.....	49
1.1.14.1 Overview.....	49
1.1.14.2 SLA.....	49
1.1.14.2.1 Overview.....	49
1.1.14.2.2 View the current state of a cloud product.....	49
1.1.14.2.3 View the history data of a cloud product.....	50
1.1.14.2.4 View the availability of an instance.....	50
1.1.14.3 ECS operations full link logs.....	51
1.1.14.4 Correlation diagnosis and alarm.....	51
1.1.14.4.1 Full stack correlation alert.....	51
1.1.14.4.2 Server.....	52
1.1.14.4.3 Network equipment.....	53
1.1.14.4.4 ECS.....	53
1.1.14.4.5 RDS.....	54
1.1.14.4.6 SLB.....	55
1.1.15 Pangu monitoring.....	55
1.1.15.1 Overview.....	55
1.1.15.2 Pangu grail.....	56
1.1.15.3 Cluster information.....	56
1.1.15.4 Node information.....	57
1.1.16 System management.....	58
1.1.16.1 Overview.....	58

1.1.16.2 Department management.....	58
1.1.16.3 Role management.....	59
1.1.16.4 Logon policy management.....	60
1.1.16.5 User management.....	61
1.1.16.6 Two-factor authentication.....	63
1.1.16.7 Application whitelist.....	66
1.1.16.8 Server password management.....	67
1.1.16.9 Operation logs.....	68
1.2 Apsara Stack Doctor (ASD).....	69
1.2.1 Apsara Stack Doctor introduction.....	69
1.2.2 Log on to Apsara Stack Doctor.....	70
1.2.3 Product dependency.....	72
1.2.4 Apsara Stack Inspection System.....	72
1.2.4.1 Apsara Stack Inspection System introduction.....	72
1.2.4.2 Access Apsara Stack Inspection System.....	73
1.2.4.3 Apsara Stack Inspection System overview.....	74
1.2.4.4 Platform Inspection.....	75
1.2.4.4.1 Platform Inspection introduction.....	75
1.2.4.4.2 Basic Inspection.....	75
1.2.4.4.3 Apsara System Inspection.....	76
1.2.4.4.4 Inspection in Other Systems.....	77
1.2.4.4.5 Inventory Inspection.....	78
1.2.4.4.6 Cloud Product Inspection.....	79
1.2.4.4.7 Middleware Inspection.....	80
1.2.4.4.8 Big Data Inspection.....	80
1.2.4.4.9 Inspection reports.....	80
1.2.4.5 Inspection history.....	81
1.2.4.6 Work Reports.....	81
1.2.4.7 Products.....	82
1.2.4.8 End-to-end ECS links.....	82
1.2.5 ASA.....	83
1.2.5.1 Apsara Stack Assistant (ASA) introduction.....	83
1.2.5.2 RPM check.....	83
1.2.5.3 Virtual IP check.....	84
1.2.5.4 Volume check.....	85
1.2.5.5 NTP check.....	86
1.2.5.6 IP conflict check.....	86
1.2.5.7 DNS check.....	87
1.2.5.8 IP details.....	87
1.2.5.9 Quota check.....	88
1.2.5.10 Error diagnostics.....	88
1.2.5.11 Versions.....	89
1.2.6 Support tools.....	89
1.2.6.1 OS tool.....	89

1.2.6.2 Apsara Distributed File System Diagnostics.....	90
1.2.7 Update Monitoring Dashboard.....	91
1.3 Operation Access Manager (OAM).....	92
1.3.1 OAM introduction.....	92
1.3.2 Instructions.....	93
1.3.3 Quick start.....	94
1.3.3.1 Log on to OAM.....	94
1.3.3.2 Create a group.....	95
1.3.3.3 Add group members.....	96
1.3.3.4 Add group roles.....	96
1.3.3.5 Create a role.....	97
1.3.3.6 Add inherited roles to a role.....	97
1.3.3.7 Add resources to a role.....	98
1.3.3.8 Add authorized users to a role.....	99
1.3.4 Manage groups.....	100
1.3.4.1 Modify the group information.....	100
1.3.4.2 View group role details.....	100
1.3.4.3 Delete a group.....	101
1.3.4.4 View assigned groups.....	101
1.3.5 Manage roles.....	101
1.3.5.1 Search for roles.....	101
1.3.5.2 Modify the role information.....	102
1.3.5.3 View the role inheritance tree.....	102
1.3.5.4 Transfer roles.....	102
1.3.5.5 Delete a role.....	103
1.3.5.6 View assigned roles.....	103
1.3.5.7 View all roles.....	104
1.3.6 Search for resources.....	104
1.3.7 View the personal information.....	104
1.3.8 Appendix.....	105
1.3.8.1 Default roles and their functions.....	105
1.3.8.1.1 OAM default role.....	105
1.3.8.1.2 Apsara Infrastructure Management Framework default roles.....	105
1.3.8.1.3 Webapp-rule default roles.....	107
1.3.8.1.4 Workflow (grandcanal) console default roles.....	107
1.3.8.1.5 Tianjimom default role.....	108
1.3.8.2 Operation permissions of operations platforms.....	108
1.3.8.2.1 Apsara Infrastructure Management Framework permission list.....	108
1.3.8.2.2 Webapp-rule permission list.....	117
1.3.8.2.3 Workflow (grandcanal) console permission list.....	118
1.3.8.2.4 Tianjimom permission list.....	118
1.4 Apsara Infrastructure Management Framework.....	118
1.4.1 What is Apsara Infrastructure Management Framework?.....	118
1.4.1.1 Overview.....	118

1.4.1.2 Basic concepts.....	119
1.4.2 Log on to Apsara Infrastructure Management Framework.....	121
1.4.3 Web page introduction.....	122
1.4.3.1 Introduction on the Home page.....	122
1.4.3.2 Introduction on the left-side navigation pane.....	124
1.4.4 Cluster operations.....	126
1.4.4.1 View cluster configurations.....	126
1.4.4.2 View the cluster dashboard.....	128
1.4.4.3 View the cluster operation and maintenance center.....	131
1.4.4.4 View the service final status.....	134
1.4.4.5 View operation logs.....	136
1.4.5 Service operations.....	136
1.4.5.1 View the service list.....	136
1.4.5.2 View the service instance dashboard.....	137
1.4.5.3 View the server role dashboard.....	139
1.4.6 Server operations.....	142
1.4.6.1 View the server dashboard.....	142
1.4.7 Monitoring center.....	144
1.4.7.1 Modify an alarm rule.....	144
1.4.7.2 View the status of a monitoring instance.....	144
1.4.7.3 View the alarm status.....	145
1.4.7.4 View the alarm history.....	146
1.4.7.5 View alarm rules.....	146
1.4.8 Tasks and deployment summary.....	147
1.4.8.1 View rolling tasks.....	147
1.4.8.2 View running tasks.....	148
1.4.8.3 View history tasks.....	149
1.4.8.4 View the deployment summary.....	149
1.4.9 Reports.....	151
1.4.9.1 View reports.....	151
1.4.9.2 Add a report to favorites.....	152
1.4.10 Appendix.....	153
1.4.10.1 IP list.....	153
1.4.10.2 Info of project component report.....	154
1.4.10.3 Auto healing - install approval pending report.....	154
1.4.10.4 Machine info report.....	154
1.4.10.5 Rolling info report.....	156
1.4.10.6 Machine RMA approval pending list.....	158
1.4.10.7 Registration vars of service.....	159
1.4.10.8 Virtual machines map.....	160
1.4.10.9 Service inspector report.....	160
1.4.10.10 Machine power on or off state of cluster.....	160
1.4.10.11 Resource apply report.....	162
1.4.10.12 State of project component.....	164

1.4.10.13 Relationship of service dependency.....	165
1.4.10.14 Check report of network topology.....	166
1.4.10.15 State of machine clone.....	166
1.4.10.16 Action of machine SR.....	167
1.5 Network operations.....	168
1.5.1 Apsara Network Intelligence.....	168
1.5.1.1 What is Apsara Network Intelligence.....	168
1.5.1.2 Log on to the Apsara Network Intelligence console.....	169
1.5.1.3 Query information.....	170
1.5.1.4 Manage cloud service instances.....	170
1.5.1.5 Tunnel VIP.....	171
1.5.1.5.1 Apply for layer-4 listener VIPs.....	171
1.5.1.5.2 Query the tunnel VIP of a cloud service.....	172
1.5.1.6 Apply for Direct Any Tunnel VIPs.....	172
1.5.1.7 Leased line connection.....	172
1.5.1.7.1 Overview.....	172
1.5.1.7.2 Manage access points.....	173
1.5.1.7.3 Manage access devices.....	174
1.5.1.7.4 Establish leased lines.....	174
1.5.1.7.5 Create VBRs.....	175
1.5.1.7.6 Create router interfaces.....	176
1.5.1.7.7 Create routing tables.....	176
1.5.1.8 Manage Business Foundation System flows in a VPC.....	177
1.5.1.9 Configure reverse access to cloud services.....	178
2 Cloud product operations.....	179
2.1 Elastic Compute Service (ECS).....	179
2.1.1 ECS overview.....	179
2.1.2 Log on to Apsara Stack Operations.....	180
2.1.3 ECS Operations and Maintenance System.....	181
2.1.3.1 Overview.....	181
2.1.3.2 VMs.....	182
2.1.3.2.1 Overview.....	182
2.1.3.2.2 Search for VMs.....	182
2.1.3.2.3 Start a VM.....	182
2.1.3.2.4 Stop a VM.....	183
2.1.3.2.5 Restart a VM.....	184
2.1.3.2.6 Downtime migration.....	184
2.1.3.2.7 Hot migration.....	185
2.1.3.2.8 Reset a disk.....	186
2.1.3.3 Disks.....	186
2.1.3.3.1 Overview.....	186
2.1.3.3.2 Search for disks.....	186
2.1.3.3.3 View snapshots.....	187
2.1.3.3.4 Mount a disk.....	187

2.1.3.3.5 Unmount a disk.....	188
2.1.3.3.6 Create a snapshot.....	188
2.1.3.4 Snapshots.....	188
2.1.3.4.1 Overview.....	188
2.1.3.4.2 Search for snapshots.....	189
2.1.3.4.3 Delete a snapshot.....	189
2.1.3.4.4 Create an image.....	190
2.1.3.5 Images.....	190
2.1.3.5.1 Overview.....	190
2.1.3.5.2 Search for images.....	190
2.1.3.5.3 Delete images.....	191
2.1.3.6 Security groups.....	191
2.1.3.6.1 Overview.....	191
2.1.3.6.2 Search for security groups.....	191
2.1.3.6.3 Add security group rules.....	192
2.1.4 VM hot migration.....	193
2.1.4.1 Overview.....	193
2.1.4.2 Hot migration usage restrictions.....	194
2.1.4.3 Complete hot migration on AG.....	194
2.1.4.4 Modify the position of the NC for a VM.....	196
2.1.4.5 FAQs.....	197
2.1.5 Hot migration of cloud disks.....	199
2.1.5.1 Overview.....	199
2.1.5.2 Usage restrictions.....	199
2.1.5.3 O&M after hot migration.....	200
2.1.6 Upgrade solution.....	200
2.1.6.1 Overview.....	200
2.1.6.2 GPU cluster restrictions.....	200
2.1.6.3 FPGA cluster restrictions.....	201
2.1.7 Disk of instance maintenance solution.....	201
2.1.7.1 Overview.....	201
2.1.7.2 Maintenance procedure.....	202
2.1.7.3 Additional instructions.....	211
2.1.8 Handle routine alarms.....	212
2.1.8.1 Overview.....	212
2.1.8.2 API proxy.....	213
2.1.8.3 API server.....	213
2.1.8.4 RegionMaster.....	214
2.1.8.5 RMS.....	215
2.1.8.6 PYNC.....	216
2.1.8.7 Zookeeper.....	216
2.1.8.8 AG.....	217
2.1.8.9 Server groups.....	218
2.1.9 Inspection.....	218

2.1.9.1 Overview.....	218
2.1.9.2 Cluster basic health inspection.....	218
2.1.9.2.1 Overview.....	218
2.1.9.2.2 Monitoring inspection.....	218
2.1.9.2.3 Basic software package version inspection.....	218
2.1.9.2.4 Basic public resources inspection.....	218
2.1.9.3 Cluster resource inspection.....	219
2.1.9.3.1 Overview.....	219
2.1.9.3.2 Cluster inventory inspection.....	219
2.1.9.3.3 VM inspection.....	221
2.2 Auto Scaling (ESS).....	222
2.2.1 Log on to Apsara Stack Operations.....	222
2.2.2 Product resources and services.....	223
2.2.2.1 Application deployment.....	223
2.2.2.2 Troubleshooting.....	224
2.2.3 Inspection.....	225
2.2.3.1 Overview.....	225
2.2.3.2 Monitoring inspection.....	226
2.2.3.3 Basic software package version inspection.....	226
2.3 Object Storage Service (OSS).....	226
2.3.1 Usage in Apsara Stack Operations.....	226
2.3.1.1 Log on to Apsara Stack Operations.....	226
2.3.1.2 Business data.....	227
2.3.1.2.1 User data.....	227
2.3.1.2.1.1 User data overview.....	227
2.3.1.2.1.2 Data monitoring.....	227
2.3.1.2.1.3 Basic bucket information.....	229
2.3.1.2.2 Cluster data.....	230
2.3.1.2.2.1 Inventory monitoring.....	230
2.3.1.2.2.2 Check disk space.....	231
2.3.1.2.2.3 Data monitoring.....	232
2.3.1.2.2.4 Resource usage rankings.....	233
2.3.1.2.2.5 Bucket statistics.....	235
2.3.1.2.2.6 Object statistics.....	236
2.3.2 Use of tools.....	237
2.3.2.1 Typical commands for tsar.....	237
2.4 Table Store.....	237
2.4.1 Storage Operations and Maintenance System.....	237
2.4.1.1 Overview.....	237
2.4.1.2 User data.....	237
2.4.1.2.1 Manage instances.....	237
2.4.1.3 Cluster management.....	240
2.4.1.3.1 Cluster information.....	240
2.4.1.4 Inspection center.....	243

2.4.1.4.1 Abnormal usage.....	243
2.4.1.5 Monitoring center.....	244
2.4.1.5.1 Cluster monitoring.....	244
2.4.1.5.2 Application monitoring.....	244
2.4.1.5.3 Top requests.....	245
2.4.1.5.4 Request log search.....	246
2.4.1.6 System management.....	246
2.4.1.6.1 Manage tasks.....	246
2.4.1.6.2 View tasks.....	248
2.4.1.7 Platform audit.....	248
2.4.1.7.1 Operation logs.....	248
2.4.2 Cluster environment description.....	249
2.4.3 System role description.....	249
2.4.4 Pre-partition a table.....	250
2.4.4.1 Pre-partitioning.....	250
2.4.4.2 View partitions.....	252
2.5 Distributed File System (DFS).....	253
2.5.1 Introduction.....	253
2.5.1.1 Apsara Distributed File System introduction.....	253
2.5.1.2 Overview of Apsara Infrastructure Management Framework.....	254
2.5.1.2.1 Basic concepts of Apsara Infrastructure Management Framework.....	255
2.5.2 Configuration update.....	256
2.5.2.1 Overview.....	256
2.5.2.2 Operation method.....	256
2.5.2.3 Configuration file structure.....	257
2.5.2.4 Validation of configuration changes.....	258
2.5.2.5 Overwriting relationship between configurations.....	259
2.5.2.6 Configuration validity checks.....	259
2.5.3 Apsara Distributed File System cluster O&M.....	259
2.5.3.1 Global flag settings of Apsara Distributed File System.....	259
2.5.3.2 Operations on Apsara Distributed File System files.....	260
2.5.3.3 Common commands of puadmin.....	260
2.5.3.4 GC functions of Apsara Distributed File System.....	261
2.5.3.5 Cluster rebalance.....	263
2.5.3.6 Directory quota operations.....	264
2.5.3.7 Directory pin operations.....	266
2.5.4 Master O&M.....	267
2.5.4.1 Overview.....	267
2.5.4.2 Primary master switchover.....	268
2.5.4.3 Status check for multiple masters.....	268
2.5.4.3.1 View the election status.....	269
2.5.4.3.2 View the status of log synchronization among multiple masters.....	269
2.5.4.4 Rules for generating .cpt and .log files.....	270
2.5.4.5 Master replacement.....	270

2.5.4.5.1 Procedure.....	270
2.5.4.5.2 Procedure for replacing a master.....	272
2.5.4.5.3 Manual replacement procedure.....	273
2.5.4.6 Manually synchronize logs between the primary master and a secondary master.....	274
2.5.4.7 Rename a chunkserver online.....	274
2.5.4.8 Multi-master tools.....	275
2.5.5 Chunkserver O&M.....	276
2.5.5.1 Set the chunkserver status.....	276
2.5.5.2 Set the disk status.....	276
2.5.5.3 Chunkserver scale-out and scale-in.....	277
2.5.5.3.1 Procedure.....	277
2.5.5.3.2 Capacity expansion procedure.....	278
2.5.5.3.3 Capacity reduction procedure.....	278
2.5.5.3.4 Disable manual chunkserver-based capacity reduction using puadmin....	280
2.5.6 Cluster status tracking.....	280
2.5.6.1 View the cluster status.....	280
2.5.6.2 Submission history.....	280
2.5.7 FAQs.....	281
2.5.7.1 Identify the machine where pangu_supervisor runs and the log location....	281
2.5.7.2 Adjust the flag of pangu_supervisor.....	281
2.5.7.3 Supervisor approval problems.....	282
2.5.7.3.1 Supervisor approval prerequisites.....	282
2.5.7.3.2 Failure to approve master replacement.....	283
2.5.7.3.3 Unable to approve chunkserver disconnection.....	283
2.5.7.4 Accelerate chunkserver disconnection.....	284
2.5.7.5 Rolling failure during a hot upgrade of Apsara Distributed File System.....	284
2.5.7.6 Manually replace binary files in emergency.....	284
2.5.7.6.1 Manual overwrite.....	284
2.5.7.6.2 Use the overwrite tool of Apsara Infrastructure Management Framework..	285
2.6 ApsaraDB for RDS.....	286
2.6.1 Service architecture.....	286
2.6.1.1 System architecture.....	286
2.6.1.1.1 Backup System.....	286
2.6.1.1.2 Monitoring system.....	287
2.6.1.1.3 Control system.....	288
2.6.1.1.4 Task scheduling system.....	289
2.6.2 RDS O&M overview.....	289
2.6.3 Log on to Apsara Stack Operations.....	289
2.6.4 Instance management.....	291
2.6.5 Host management.....	294
2.6.6 Security maintenance.....	294
2.6.6.1 Network security maintenance.....	294
2.6.6.2 Account password maintenance.....	295

2.7 KVStore for Redis.....	295
2.7.1 O&M tools.....	295
2.7.2 Service architecture.....	295
2.7.2.1 System architecture.....	295
2.7.2.1.1 Backup system.....	295
2.7.2.1.2 Data migration system.....	295
2.7.2.1.3 Monitoring system.....	296
2.7.2.1.4 Control system.....	296
2.7.2.1.5 Task scheduling system.....	296
2.7.3 Log on to Apsara Stack Operations.....	296
2.7.4 Instance management.....	298
2.7.5 Host management.....	298
2.7.6 Security maintenance.....	299
2.7.6.1 Network security maintenance.....	299
2.7.6.2 Account password maintenance.....	300
2.8 ApsaraDB for MongoDB.....	300
2.8.1 Service architecture.....	300
2.8.1.1 System architecture.....	300
2.8.1.1.1 Backup system.....	300
2.8.1.1.2 Data migration system.....	301
2.8.1.1.3 Monitoring system.....	301
2.8.1.1.4 Control system.....	302
2.8.1.1.5 Task scheduling system.....	302
2.8.2 ApsaraDB for MongoDB O&M overview.....	302
2.8.3 Log on to Apsara Stack Operations.....	302
2.8.4 Instance management.....	303
2.8.5 Host management.....	304
2.8.6 Security maintenance.....	305
2.8.6.1 Network security maintenance.....	305
2.8.6.2 Account password maintenance.....	305
2.9 Apsara Stack Security.....	305
2.9.1 Log on to the Apsara Infrastructure Management Framework console.....	305
2.9.2 Routine operations and maintenance of Server Guard.....	306
2.9.2.1 Check the service status.....	306
2.9.2.1.1 Check the client status.....	306
2.9.2.1.2 Check the status of Aegiserver.....	307
2.9.2.1.3 Check the Server Guard Update Service status.....	308
2.9.2.1.4 Check the Defender module status.....	308
2.9.2.2 Restart Server Guard.....	309
2.9.3 Routine operations and maintenance of Network Traffic Monitoring System...311	
2.9.3.1 Check the service status.....	311
2.9.3.1.1 Basic inspection.....	311
2.9.3.1.2 Advanced inspection.....	311
2.9.3.2 Common operations and maintenance.....	313

2.9.3.2.1 Restart the Network Traffic Monitoring System process.....	313
2.9.3.2.2 Uninstall Network Traffic Monitoring System.....	313
2.9.3.2.3 Disable TCP blocking.....	313
2.9.3.2.4 Enable TCPDump.....	314
2.9.4 Routine operations and maintenance of Anti-DDoS Service.....	314
2.9.4.1 Check the service status.....	314
2.9.4.1.1 Basic inspection.....	314
2.9.4.1.2 Advanced inspection.....	315
2.9.4.2 Common operations and maintenance.....	317
2.9.4.2.1 Restart Anti-DDoS Service.....	317
2.9.4.2.2 Troubleshoot common faults.....	318
2.9.5 Routine operations and maintenance of the vulnerability analysis service.....	321
2.9.5.1 Check service status.....	321
2.9.5.1.1 Basic inspection.....	321
2.9.5.1.2 Advanced inspection: Checks the status of the Cactus-batch service.....	322
2.9.5.1.3 Advanced inspection: Check the status of the Cactus-keeper service.....	323
2.9.5.2 Restart the vulnerability analysis service.....	324
2.9.5.2.1 Restart Cactus-batch.....	324
2.9.5.2.2 Restart the Beaver_server service.....	325
2.9.5.2.3 Restart Cactus-keeper.....	325
2.9.6 Routine operations and maintenance of Threat Detection Service.....	326
2.9.6.1 Check the service status.....	326
2.9.6.1.1 Basic inspection.....	326
2.9.6.1.2 Advanced inspection.....	326
2.9.6.2 Restart TDS.....	327
2.9.7 Routine operations and maintenance of WAF.....	328
2.9.7.1 Check service status.....	328
2.9.7.1.1 Basic inspection.....	328
2.9.7.1.2 Advanced inspection: Check the Tengine service status.....	328
2.9.7.1.3 Advanced inspection: Check the status of the tmd_server service.....	329
2.9.7.1.4 Advanced inspection: Check the status of the gf_server service.....	330
2.9.7.1.5 Check the etcd service status.....	331
2.9.7.1.6 Advanced inspection: Checks the status of the logcenter service.....	332
2.9.7.2 Restart WAF.....	332
2.9.8 Routine operations and maintenance of Security Audit.....	334
2.9.8.1 Check service status.....	334
2.9.8.1.1 Basic inspection.....	334
2.9.8.1.2 Advanced inspection: check the status of the security-auditlog-app service.....	335
2.9.8.1.3 Advanced inspection: Check the security-auditlog-syslog service status..	336
2.9.8.2 Restart Security Audit.....	337
2.9.9 Routine operations and maintenance of Apsara Stack Security Center.....	338
2.9.9.1 Check service status.....	338
2.9.9.1.1 Basic inspection.....	338

2.9.9.1.2 Advanced inspection.....	339
2.9.9.2 Restart the secure-console service.....	340
2.9.10 Routine operations and maintenance of secure-service.....	340
2.9.10.1 Check the service status.....	340
2.9.10.1.1 Basic inspection.....	340
2.9.10.1.2 Advanced inspection: Check the secure-service status.....	341
2.9.10.1.3 Check the Dolphin service status.....	342
2.9.10.1.4 Check the data-sync service status.....	343
2.9.10.2 Restart secure-service.....	343
2.10 Key Management Service (KMS).....	345
2.10.1 Operations and maintenance of KMS components.....	345
2.10.1.1 Overview.....	345
2.10.1.2 KMS_HOST.....	345
2.10.1.3 HSA.....	349
2.10.1.4 Etcd.....	352
2.10.1.5 Rotator.....	354
2.10.1.5.1 Primary IDC.....	354
2.10.1.5.2 Secondary IDC.....	355
2.10.2 Log analysis.....	356
2.10.2.1 Overview.....	356
2.10.2.2 Request IDs.....	357
2.10.2.3 Common KMS errors.....	358
2.10.2.3.1 Overview.....	358
2.10.2.3.2 Error codes 4xx.....	358
2.10.2.3.3 Error code 500.....	358
2.10.2.3.4 Error code 503.....	358
2.10.2.3.5 Dependent service degradation.....	359
2.10.3 View and process internal data.....	359
2.11 Log Service.....	362
2.11.1 Components operations and maintenance.....	362
2.11.2 O&M and troubleshooting.....	363
2.11.2.1 Nginx.....	363
2.11.2.2 Console.....	363
2.11.2.3 Service.....	364
2.12 Domain Name System (DNS).....	364
2.12.1 Introduction to Apsara Stack DNS.....	364
2.12.2 Maintenance.....	364
2.12.2.1 View running logs.....	364
2.12.2.2 Enable and disable a service.....	365
2.12.2.3 Data backup.....	365
2.12.3 DNS API.....	365
2.12.3.1 Manage the API system.....	365
2.12.3.2 Troubleshooting.....	366
2.12.4 DNS system.....	367

2.12.4.1 Check whether the service role is correct.....	367
2.12.4.2 Troubleshooting.....	367
2.12.4.3 Errors and exceptions.....	368
2.12.5 Log analysis.....	368
2.12.6 View and process data.....	368
2.13 API Gateway.....	369
2.13.1 API Gateway introduction.....	369
2.13.2 Routine maintenance.....	369
2.13.2.1 View operation logs.....	369
2.13.2.2 Enable and disable API services.....	369
2.13.3 API Gateway console O&M.....	370
2.13.3.1 System O&M.....	370
2.13.3.2 Troubleshooting.....	371
2.13.4 API Gateway O&M.....	371
2.13.4.1 System O&M.....	371
2.13.4.2 Troubleshooting.....	372
2.13.5 Log analysis.....	372

1 Basic platform operations

1.1 Apsara Stack Operations (ASO)

1.1.1 Apsara Stack Operations overview

Apsara Stack Operations (ASO) is an operations management system developed for the Apsara Stack operations management personnel, such as field operations engineers, operations engineers on the user side, and operations management engineers, operations security personnel, and audit personnel of the cloud platform. ASO allows the operations engineers to master the operating conditions of the system in time and perform Operation & Maintenance (O&M) operations.

ASO also supports the multi-region function and provides centralized operations management for multiple regions.

ASO has the following main functions:

- [*Operation and maintenance dashboard*](#)

The **Operation and Maintenance** dashboard displays the local version list, inventory overview, alarms breakdown, and inventory curve of the cloud platform, which allows you to know the current usage of resources.

- [*Alarm management*](#)

Alarm Management allows operations engineers to quickly know the information of alarms generated by the system, locate the problems based on the alarm information, track the problem processing, and configure the alarms.

- [*Resource management*](#)

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

- [*Inventory management*](#)

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

- [*Products*](#)

Products allows you to access the operations and maintenance services of other products on the cloud platform. You are redirected to the corresponding operations and maintenance page of a product by using Single Sign-On (SSO) and redirection.

- [ITIL management](#)

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

- [API management](#)

API Management provides a unified encapsulation for the operations APIs of all cloud products on the cloud platform, which facilitates the secondary development of the operations platform for third-parties, allows the fine-grained access control and security audit of the operations APIs, and provides the centralized management in terms of Apsara Stack versions and APIs.

- [Configurations](#)

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

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

- [Offline backup](#)

Offline Backup is used to back up the key metadata of Apsara Stack. Metadata backup is used for the fast recovery of Apsara Stack faults.

- [NOC](#)

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

- [Full stack monitor](#)

Full Stack Monitor allows you to perform an aggregate query on the system alarm events, query and retrieve all the alarm data in the link based on the host IP address, instance ID, and time range, and view the end-to-end topology.

- [Pangu monitoring](#)

Pangu Monitoring displays the **Pangu Grail**, **Cluster Information**, and **Node Information**.

- [System management](#)

System Management consists of the user management, two-factor authentication, role management, department management, logon policy management, application whitelist, server password management, and operation logs. As the module for centralized management of accounts, roles, and permissions, system management supports the SSO function of ASO. After logging on to ASO, you can perform O&M operations on all components of the cloud platform or be redirected to the operations and maintenance page without providing the username or password.

1.1.2 Log on to Apsara Stack Operations

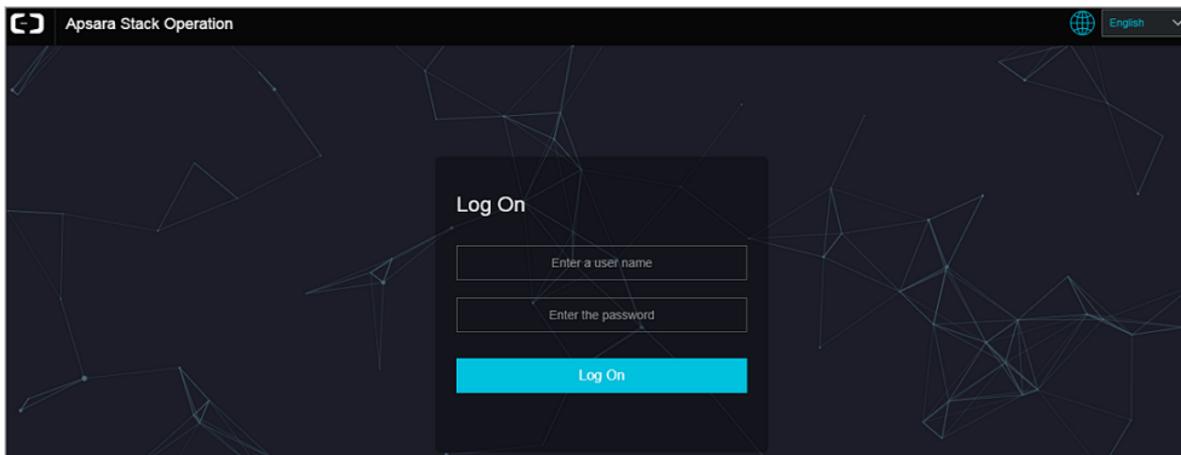
This topic describes how to log on to Apsara Stack Operations (ASO) as users, such as operations engineers.

Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 1-1: Log on to ASO**Note:**

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.

- The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
- You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.

4. Click **Log On** to go to the **Apsara Stack Operations** page.

1.1.3 Web page introduction

After you log on to Apsara Stack Operations (ASO), the home page appears. This topic allows you to get a general understanding of the basic operations and functions of the ASO page.

- : Select a region from the drop-down list to switch between different regions, which provides the centralized operations management for multiple regions.
- : In the Help Center, you can view the ASO instructions and other topics related to operations.
- : Select the language from the drop-down list to change the language of ASO.
- : Click  and select **Personal Information** to view the information of the current user or change the password.
- : Click this to expand the left-side navigation pane.

1.1.4 Operation and maintenance dashboard

Apsara Stack Operations (ASO) displays the current usage and monitoring information of system resources in graphs, which allows you to know the current operating conditions of the system.

[Log on to ASO](#). In the left-side navigation pane, click **Operation and Maintenance**. The operation and maintenance dashboard displays the current product version, inventory statistics, and alarm statistics of the cloud platform. By viewing the dashboard, the operations engineers can know the overall operating conditions of Apsara Stack products in time.

1.1.5 Alarm management

1.1.5.1 Overview

Alarm management allows operations engineers to quickly know the information of alarms generated by the system, locate the problems based on the alarm information, track the problem processing, and configure the alarms.

[Log on to ASO](#). Click **Alarm Management** in the left-side navigation pane to view the alarm overview.

1.1.5.2 Alarm events

The **Alarm Events** page displays the information of all alarms generated by the system. The alarm information is aggregated by alarm item and alarm source. You can also search for alarms based on filter conditions, such as product, severity, status, and time range when the alarm is triggered, and then perform Operation & Maintenance (O&M) operations on the alarms.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Alarm Management > Alarm Events**. The **Alarm Events** page appears.



Note:

- On this page, alarms are sorted by severity, time, and status, which makes sure that the most urgent alarms to be processed are displayed at the top.
- Alarms displayed on this page are aggregated by alarm source and alarm item. You can view the aggregated alarms by clicking the number in the **Alarms** column.

3. You can perform the following operations on this page:

- Search for alarms

On top of the page, you can search for alarms by **Monitoring Item Type, Region, Product, Service, Severity, Status, Start Date - End Date**, and/or search content.

- View alarm sources

Click an alarm source name in blue in the **Alarm Source** column to view the details of the alarm source.

- View alarm details

Click an alarm name in blue in the **Alarm Details** column. On the displayed **Alarm Details** page, you can view the alarm information, such as the summary, reference, scope, and resolution.

- View the number of alarms

Click a number in blue in the **Alarms** column. On the displayed **Alarms** page, you can view the alarm time and the alarm level of all aggregated alarms.

- Process an alarm

- If an alarm is being processed by operations engineers, click **Actions > Process** to set the alarm status to **In process**.
- If the processing of an alarm is finished, click **Actions > Processed** to set the alarm status to **Processed**.
- To view the whole processing flow of an alarm, click **Actions > Alarm Tracing**.
- If an alarm is considered as an incident when being processed, click **Actions > Report to ITIL**. Then, an incident request is created in ITIL to track the issue. For more information, see [Manage incidents](#).

To manage alarms in batches, select multiple alarms and then click **Process, Complete, Report, Shield, or Remove Shield**.

- Export a report

Click **Export Report** in the upper-right corner to generate a list of alarms.

1.1.5.3 Alarm history

The **Alarm History** page displays all the alarms generated by the system and the corresponding information in chronological order.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Alarm Management > Alarm History**.
3. On the **Alarm History** page, you can perform the following operations:

- Search for alarms

On top of the page, you can search for alarms by **Monitoring Item Type, Region, Product, Service, Severity, Status, Start Date - End Date**, and/or search content.

- Export a list of alarms

Click **Export Report** in the upper-right corner to generate a list of alarms.

- View alarm sources

Click an alarm source name in blue in the **Alarm Source** column to view the details of the alarm source.

- View alarm details

Click an alarm name in blue in the **Alarm Details** column. On the displayed **Alarm Details** page, you can view the alarm information, such as the summary, reference, scope, and resolution.

- View the original alarm information

Click a row of alarm to display the original alarm information. Click to view the original information of the alarm.

1.1.5.4 Alarm configuration

On the **Alarm Configuration** page, you can search for, add, modify, or delete an alarm contact or alarm contact group.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Alarm Management > Alarm Configuration**.
3. On the **Alarm Configuration** page, you can perform the following operations:

- **Contacts:** Click the **Contacts** tab.

- Search for alarm contacts

On top of the page, configure the corresponding product name, contact name, and/or phone number and then click **Search**. The alarm contacts that meet the search condition are displayed in the list.

- Add an alarm contact

Click **Add Contact** in the upper-left corner. In the displayed **Add Contact** dialog box, complete the configurations and then click **OK**.

- Modify an alarm contact

Find the alarm contact to be modified and then click **Modify** in the **Actions** column. In the displayed **Modify Contact** dialog box, modify the information and then click **OK**.

- Delete an alarm contact

Find the alarm contact to be deleted and then click **Delete** in the **Actions** column. Click **OK** in the displayed dialog box.

- **Contact Groups:** Click the **Contact Groups** tab.

- Search for an alarm contact group

On top of the page, enter the group name and then click **Search**. The alarm contact group that meets the search condition is displayed in the list.

— Add an alarm contact group

Click **Add Contact Group** in the upper-left corner. In the displayed **Add Contact Group** dialog box, enter the group name and select the contacts to add to the contact group. Then, click **OK**.

— Modify an alarm contact group

Find the alarm contact group to be modified and then click **Modify** in the **Actions** column. In the displayed **Modify Contact Group** dialog box, modify the information and then click **OK**.

— Delete an alarm contact group

Find the alarm contact group to be deleted and then click **Delete** in the **Actions** column. Click **OK** in the displayed dialog box.

1.1.5.5 Alarm overview

By viewing the alarm overview, you can know the distribution of different levels of alarms for Apsara Stack products.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Alarm Management > Alarm Overview**. The **Alarm Overview** page appears.
 - The list on the left displays the numbers of remind alarms, minor alarms, major alarms, critical alarms, cleared alarms, and system alarms for various products.
 - The pie chart in the upper-right corner displays the distribution proportion of all alarms at different levels.
 - The column chart in the lower-right corner displays the statistics of alarms newly added per day in the past seven days.

1.1.5.6 Alarm subscription and push

The alarm subscription and push function allows you to configure the alarm notification channel and then push the alarm to operations engineers in certain ways.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).

2. In the left-side navigation pane, choose **Alarm Management > Subscribe/Push**.
3. Click **Add Channel**.
4. On the **Add Subscription** page, complete the following configurations.

Configuration item	Description
Channel Name	The name of the subscription channel.
Subscribed Language	Select Chinese or English .
Subscription Region	Select the region where the subscription is located.
Product Name	Select the name of the subscribed product.
Protocol	Currently, only HTTP is supported.
Push Interface Address	The IP address of the push interface.
Port Number	The port number of the push interface.
URI	The URI of the push interface.
HTTP Method	Currently, only POST is supported.
Push Cycle (Minutes)	The push cycle, which is calculated by minute.
Pushed Alerts	The number of alarms pushed each time.
Push Mode	Select one of the following methods: <ul style="list-style-type: none"> • ALL: All of the alarms are pushed in each push cycle. • TOP: Only alarms with high priority are pushed in each push cycle.
Push Template	Select one of the following templates: <ul style="list-style-type: none"> • ASO: The default template. • ANS: Select this template to push alarms by DingTalk, SMS, or email. Currently, you can only configure one channel of this type. <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: A preset ANS template exists if the system already connects with the ANS product. To restore the initial </div>

Configuration item	Description
	configurations of the template with one click, click Reset ANS Channel .
Custom JSON Fields	The person who receives the push can use this field to configure the identifier in a custom way. The format must be JSON.
Push Switch	Select whether to push the alarms.

- After completing the configurations, click **OK**.

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

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



Note:

For ANS push channels, you must enter the mobile phone number, email address, and/or DingTalk to which alarms are pushed after clicking **Test** in the **Actions** column.

What's next

After configuring the push channel and turning on the push switch, you can click the **Push** tab to view the push records.

1.1.6 Resource management

1.1.6.1 Overview

Resource management monitors and manages hardware devices in the data center, including the physical servers, physical switches, and network security devices. The major monitoring information includes the overall status information, monitoring metrics, alarm delivery status, and port traffic of devices.

1.1.6.2 Physical servers

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

Procedure

- [Log on to Apsara Stack Operations \(ASO\)](#).
- In the left-side navigation pane, choose **Resource Management > Physical Servers**.



Note:

This page displays physical servers in two dimensions: **Product** and **Server**. You can click the corresponding tab as required to view the details of a physical server.

- Expand the left-side navigation tree level by level based on regions, data centers, and cabinets until all products under a cabinet are displayed. Select a product, such as RDS, to view a list of physical servers where services in RDS are located on the right.
- Find a product and click **Details** in the **Operation** column. On the displayed **Physical Server Details** page, view the basic information, monitoring details, and alarm information of the physical server.

You can switch tabs to view the monitoring details and the alarm information, or select different time ranges to observe the monitoring values in different time periods. The monitoring metrics are CPU usage, memory usage, system load, network throughput, disk usage, and disk I/O.

1.1.7 Inventory management

1.1.7.1 Overview

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

1.1.7.2 View the ECS inventory

By viewing the Elastic Compute Service (ECS) inventory, you can know the current usage and surplus of ECS product resources and perform Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

- [Log on to Apsara Stack Operations \(ASO\)](#).
- In the left-side navigation pane, choose **Inventory Management > ECS Instances**.



Note:

Click  in the upper-right corner to configure the inventory threshold.

- CPU Inventory Details (Core)** and **Memory Inventory Details (G)** display the used and available CPU (core) and memory (GiB) of all ECS instance type families in the last five days.
- ECS Instances Inventory Details** allows you to perform a paging query on the inventory of a certain type of ECS instances at a certain date by **Zone**, **Instance Type**, and **Date**.

For more information about the mapping between instance type families and CPU/memory configurations of instances, see [Instance type](#).

Table 1-1: Instance type

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
N4	ecs.n4.small	None	1	2.0	1
	ecs.n4.large	None	2	4.0	1
	ecs.n4.xlarge	None	4	8.0	2
	ecs.n4.2xlarge	None	8	16.0	2
	ecs.n4.4xlarge	None	16	32.0	2
	ecs.n4.8xlarge	None	32	64.0	2
MN4	ecs.mn4.small	None	1	4.0	1
	ecs.mn4.large	None	2	8.0	1
	ecs.mn4.xlarge	None	4	16.0	2
	ecs.mn4.2xlarge	None	8	32.0	3
	ecs.mn4.4xlarge	None	16	64.0	8
	ecs.mn4.8xlarge	None	32	128.0	8
E4	ecs.e4.small	None	1	8.0	1
	ecs.e4.large	None	2	16.0	1
	ecs.e4.xlarge	None	4	32.0	2
	ecs.e4.2xlarge	None	8	64.0	3

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.e4.4xlarge	None	16	128.0	8
XN4	ecs.xn4.small	None	1	1.0	1
gn5	ecs.gn5-c4g1.xlarge	440	4	30.0	2
	ecs.gn5-c8g1.2xlarge	440	8	60.0	3
	ecs.gn5-c4g1.2xlarge	880	8	60.0	3
	ecs.gn5-c8g1.4xlarge	880	16	120.0	8
	ecs.gn5-c28g1.7xlarge	440	28	112.0	8
	ecs.gn5-c8g1.8xlarge	1760	32	240.0	8
	ecs.gn5-c28g1.14xlarge	880	56	224.0	8
	ecs.gn5-c8g1.14xlarge	3520	56	480.0	8
d1	ecs.d1.2xlarge	4 * 5500	8	32.0	3
	ecs.d1.4xlarge	8 * 5500	16	64.0	8
	ecs.d1.6xlarge	12 * 5500	24	96.0	8
	ecs.d1-c8d3.8xlarge	12 * 5500	32	128.0	8
	ecs.d1.8xlarge	16 * 5500	32	128.0	8

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.d1-c14d3.14xlarge	12 * 5500	56	160.0	8
	ecs.d1.14xlarge	28 * 5500	56	224.0	8
gn4	ecs.gn4-c4g1.xlarge	None	4	30.0	2
	ecs.gn4-c8g1.2xlarge	None	8	60.0	3
	ecs.gn4.8xlarge	None	32	48.0	8
	ecs.gn4-c4g1.2xlarge	None	8	60.0	3
	ecs.gn4-c8g1.4xlarge	None	16	60.0	8
	ecs.gn4.14xlarge	None	56	96.0	8
ga1	ecs.ga1.xlarge	1*87	4	10.0	2
	ecs.ga1.2xlarge	1*175	8	20.0	3
	ecs.ga1.4xlarge	1*350	16	40.0	8
	ecs.ga1.8xlarge	1*700	32	80.0	8
	ecs.ga1.14xlarge	1*1400	56	160.0	8
se1ne	ecs.se1ne.large	None	2	16.0	1
	ecs.se1ne.xlarge	None	4	32.0	2

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.se1ne.2xlarge	None	8	64.0	3
	ecs.se1ne.4xlarge	None	16	128.0	8
	ecs.se1ne.8xlarge	None	32	256.0	8
	ecs.se1ne.14xlarge	None	56	480.0	8
sn2ne	ecs.sn2ne.large	None	2	8.0	1
	ecs.sn2ne.xlarge	None	4	16.0	2
	ecs.sn2ne.2xlarge	None	8	32.0	3
	ecs.sn2ne.4xlarge	None	16	64.0	8
	ecs.sn2ne.8xlarge	None	32	128.0	8
	ecs.sn2ne.14xlarge	None	56	224.0	8
sn1ne	ecs.sn1ne.large	None	2	4.0	1
	ecs.sn1ne.xlarge	None	4	8.0	2
	ecs.sn1ne.2xlarge	None	8	16.0	3
	ecs.sn1ne.4xlarge	None	16	32.0	8
	ecs.sn1ne.8xlarge	None	32	64.0	8

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
gn5i	ecs.gn5i-c2g1.large	None	2	8.0	1
	ecs.gn5i-c4g1.xlarge	None	4	16.0	2
	ecs.gn5i-c8g1.2xlarge	None	8	32.0	2
	ecs.gn5i-c16g1.4xlarge	None	16	64.0	2
	ecs.gn5i-c28g1.14xlarge	None	56	224.0	2
g5	ecs.g5.large	None	2	8.0	2
	ecs.g5.xlarge	None	4	16.0	3
	ecs.g5.2xlarge	None	8	32.0	4
	ecs.g5.4xlarge	None	16	64.0	8
	ecs.g5.6xlarge	None	24	96.0	8
	ecs.g5.8xlarge	None	32	128.0	8
	ecs.g5.16xlarge	None	64	256.0	8
	ecs.g5.22xlarge	None	88	352.0	15
c5	ecs.c5.large	None	2	4.0	2
	ecs.c5.xlarge	None	4	8.0	3
	ecs.c5.2xlarge	None	8	16.0	4

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.c5.4xlarge	None	16	32.0	8
	ecs.c5.6xlarge	None	24	48.0	8
	ecs.c5.8xlarge	None	32	64.0	8
	ecs.c5.16xlarge	None	64	128.0	8
r5	ecs.r5.large	None	2	16.0	2
	ecs.r5.xlarge	None	4	32.0	3
	ecs.r5.2xlarge	None	8	64.0	4
	ecs.r5.4xlarge	None	16	128.0	8
	ecs.r5.6xlarge	None	24	192.0	8
	ecs.r5.8xlarge	None	32	256.0	8
	ecs.r5.16xlarge	None	64	512.0	8
	ecs.r5.22xlarge	None	88	704.0	15
se1	ecs.se1.large	None	2	16.0	2
	ecs.se1.xlarge	None	4	32.0	3
	ecs.se1.2xlarge	None	8	64.0	4
	ecs.se1.4xlarge	None	16	128.0	8
	ecs.se1.8xlarge	None	32	256.0	8
	ecs.se1.14xlarge	None	56	480.0	8

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
d1ne	ecs.d1ne.2xlarge	4 * 5500	8	32.0	4
	ecs.d1ne.4xlarge	8 * 5500	16	64.0	8
	ecs.d1ne.6xlarge	12 * 5500	24	96.0	8
	ecs.d1ne.8xlarge	16 * 5500	32	128.0	8
	ecs.d1ne.14xlarge	28 * 5500	56	224.0	8
f3	ecs.f3-c16f1.4xlarge	None	16	64.0	8
	ecs.f3-c16f1.8xlarge	None	32	128.0	8
	ecs.f3-c16f1.16xlarge	None	64	256.0	16
ebmg5	ecs.ebmg5.24xlarge	None	96	384.0	32
i2	ecs.i2.xlarge	1 * 894	4	32.0	3
	ecs.i2.2xlarge	1 * 1788	8	64.0	4
	ecs.i2.4xlarge	2 * 1788	16	128.0	8
	ecs.i2.8xlarge	4 * 1788	32	256.0	8
	ecs.i2.16xlarge	8 * 1788	64	512.0	8
re5	ecs.re5.15xlarge	None	60	990.0	8
	ecs.re5.30xlarge	None	120	1980.0	15
	ecs.re5.45xlarge	None	180	2970.0	15

1.1.7.3 View the SLB inventory

By viewing the Server Load Balancer (SLB) inventory, you can know the current usage and surplus of SLB product resources and perform Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Inventory Management > SLB Instances**.



Note:

Click  in the upper-right corner to configure the inventory threshold.

- The section in the upper-left corner displays the frozen, assigned, protected, and released internal VIP history inventory and public VIP history inventory in the last five days.
- The section in the upper-right corner displays the current proportions of used internal VIP inventory, available internal VIP inventory, used public VIP inventory, and available public VIP inventory.
- The section at the bottom displays the SLB inventory details, which allows you to perform a paging query on the SLB inventory by **Type** and **Date**.

1.1.7.4 View the RDS inventory

By viewing the Relational Database Service (RDS) inventory, you can know the current usage and surplus of RDS product resources and perform Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Inventory Management > RDS Instances**.



Note:

Click  in the upper-right corner to configure the inventory threshold.

- **RDS Inventory** displays the inventories of different types of RDS instances in the last five days. Different colors represent different types of RDS instances.
- **RDS Inventory Details** allows you to perform a paging query on the RDS inventory by **Engine** and **Date**.

1.1.7.5 View the OSS inventory

By viewing the Object Storage Service (OSS) inventory, you can know the current usage and surplus of OSS product resources and perform Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Inventory Management > OSS Buckets**.



Note:

Click  in the upper-right corner to configure the inventory threshold.

- **Inventory Availability History (G)** displays the available OSS inventory in the last five days.
- **Current Inventory Usage (G)** displays the percentage of used OSS inventory.
- **OSS Bucket Inventory Details** allows you to perform a paging query on the OSS inventory by **Date**.

1.1.7.6 View the Table Store inventory

By viewing the Table Store inventory, you can know the current usage and surplus of Table Store product resources and perform Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Inventory Management > Table Store**.
 - **Table Store Instance Usage Details (G)** displays the capacity consumed by each Table Store instance.
 - **Table Store Inventory Details** allows you to perform a paging query on the Table Store inventory by **Date**.

1.1.7.7 View the Log Service inventory

By viewing the Log Service inventory, you can know the current usage and surplus of Log Service product resources and perform Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Inventory Management > Log Service**.

**Note:**

Click  in the upper-right corner to configure the inventory threshold.

- **History Inventory Records (G)** displays the available and total Log Service inventory in the last five days by using the line graph.
- **Current Quota Details (G)** displays the capacity consumed by each Log Service instance.
- **Log Service Inventory Details** allows you to perform a paging query on the Log Service inventory by **Date**.

1.1.8 Products

Products allows you to access the operations and maintenance services of other products on the cloud platform. You are redirected to the corresponding operations and maintenance page of a product by using Single Sign-On (SSO) and redirection.

[Log on to Apsara Stack Operations \(ASO\)](#). In the left-side navigation pane, click **Products**.

On the **Product List** page, you can view the operations and maintenance icons of different products based on your permissions. For example, a Table Store operations engineer can only view the **OTS Storage Operations and Maintenance System** icon. Click **OTS Storage Operations and Maintenance System** to go to the Table Store operations and maintenance console. An operations system administrator can view all the operations and maintenance components of the cloud platform. The read and write permissions for product operations and maintenance are separated. Therefore, different operation permissions can be dynamically assigned based on different roles.

1.1.9 ITIL management

1.1.9.1 Overview

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

ITIL has the following functions:

- **Dashboard**

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

- **Services**

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

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

- Incident management: used to recover from exceptions and guarantee the normal production by a series of recovery operations, including diagnosis, processing, resolution, and confirmation. Incident management provides a unified mode and standardizes the process for incident processing, and supports automatically collecting or manually recording the incident information.

- Problem management: Temporarily resolved incidents or incidents whose root cause is not clear can be transformed to problems for further analysis and thorough troubleshooting. Incidents aim to resume the production, whereas problems aim to be completely solved to make sure the problems do not recur. Problem management allows you to find the root cause of incidents, thoroughly troubleshoot the incidents, and reduce repeated incidents.

- **Version Control**

Version Control displays the version information of Apsara Stack products.

- **Process Template Configuration**

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

- **CAB/ECAB Configuration**

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

1.1.9.2 Dashboard

Dashboard allows you to view the summary of incident requests, problem requests, and change requests, namely the total numbers of incident requests, problem requests, and change requests, the numbers of new and closed incident requests, problem requests, and change requests, and

their change trend. You can also view the distribution of request fulfillment and the information of version management.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Dashboard**.

1.1.9.3 Services

1.1.9.3.1 Basic functions

1.1.9.3.1.1 Overview

This topic focuses on the basic functions of requests and tasks.

Services is composed of requests and tasks.

- **Requests**

A request is the complete process of an incident request or problem request. For example, the process of an incident request is a complete request that may consist of **Diagnose**, **Resolve**, and **Confirm** phases.

- **Tasks**

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

1.1.9.3.1.2 Manage requests

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **Request** tab.
3. You can perform the following operations on the **Request** tab:

- **Create a request**

Click **New** and then select a request type. Configure the parameters and then click **Confirm** to create a request. This topic takes incident requests and problem requests as examples. For more information, see [Create an incident request](#) and [Create a problem request](#).

- **Filter requests**

Click  at the right of the first drop-down list and then select a request type to display the corresponding requests in the list.

- **Search for requests**

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

- **View request details**

Find the request that you want to view the details, and then click **Detail**. The request details page is composed of the following sections:

- **Function:** the function buttons for the request processing. For more information, see [Manage incident requests](#) and [Manage problem requests](#).
- **Request Flow:** the current processing flow of this request.
- **Basic Information:** the basic information of this request, which is generally the information configured when you create the request.
- **Track:** each phase of the request processing and their corresponding time point.
- **Detail Tabs:** the task list and comments related to this request.

1.1.9.3.1.3 Manage tasks

After a request is created, the system automatically goes to the **Diagnose** phase. In the **Diagnose** phase, the system automatically generates a task. Each task corresponds to a specific processing phase.

Context

Tasks are currently divided into the following three types:

- **My Task:** tasks waiting to be processed by you.
- **Task Pool:** a collection of tasks that are not assigned to related person in charge. You can check out the tasks in the **Task Pool** to make the tasks exclusive to you. Others cannot process the tasks that you have checked out. You can view the checked out tasks in **My Task**.
- **Processed by me:** the history tasks that have been processed by you. After you process the tasks in **My Task**, they are displayed in **Processed by me**.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **My Task** tab.

3. You can perform the following operations on the **My Task** tab:

- **Search for tasks**

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

- **View task details**

Find the task that you want to view the details, and then click **Detail**. On the task details page, you can view the request details related to the task. For more information, see the "View request details" section of the [Manage requests](#) topic.

1.1.9.3.2 Manage incidents

1.1.9.3.2.1 Create an incident request

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

Context

Currently, ITIL management supports creating incident requests in the following two ways:

- **Automatically created**

The incident information comes from the alarm information in Apsara Stack Operations (ASO). The alarm module transfers the alarm information to the ITIL module to generate the incident request based on the actual conditions, such as the alarm level and the alarm filtering.

- **Manually created**

You can manually create incident requests, which is supplementary to the automatic way. For example, you can manually create an incident request if the incident is not automatically recognized. This topic describes how to manually create an incident request.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **Request** tab.
3. Click **New** and then select **Incident**. Configure the parameters on the displayed page.

For more information about the parameters, see [Parameter descriptions](#).

Table 1-2: Parameter descriptions

Parameter	Description
Report Object	The person who is required to process the request.
Callback Email	The email address of the person who records the request.
Callback Telephone	The telephone number of the person who records the request.
Product	The product to which the request belongs. Select a specific product from the drop-down list.
Service Name	The service related to the selected product. Select a specific service from the drop-down list.
Happen Date	The time when the request happens.
Priority	The priority of processing this request. The priority indicates the urgency of the request. The higher the urgency is, the higher priority the request must have. The priority has the following levels, from high to low, based on the urgency: <ul style="list-style-type: none"> • Critical • Major • Minor • Remind • Cleared • System
Alarm Code	The alarm ID.
Summary	The summary of this request.
Description	The detailed description about the request.
Suggestion	Optional. The suggestion about the request processing.

4. After configuring the preceding parameters, click **Confirm**.

1.1.9.3.2.2 Manage incident requests

After creating an incident request, you can change the priority of, comment, suspend, and resume the created incident request.

Prerequisites

An incident request is created. For more information about how to create an incident request, see [Create an incident request](#).

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **Request** tab.
3. Click  at the right of the first drop-down list and then select **Incident** to display the incident requests in the list.
4. Find the incident request that you want to manage, and then click **Detail**.
5. You can perform the following operations on the request details page.

- **Change Priority**

Click **Change Priority**. In the displayed dialog box, select the new priority. Perform this operation for temporary adjustment or correcting the error in priority.



Note:

You can only change the priority of incident requests in the **Diagnose** phase.

- **Comment**

Click **Comment**. In the displayed dialog box, enter the comment for this incident request. Perform this operation for collaborative scenarios. For example, users can comment the incident request to share the information with each other and guide each other when they process the same incident.

- **Suspend**

Click **Suspend**. In the displayed dialog box, enter the remarks. Perform this operation for incident requests that currently do not require to be processed.

- **Resume**

Click **Resume**. In the displayed dialog box, enter the remarks. Perform this operation for suspended incident requests that require to be processed.

- **Recycle**

Perform this operation for incident requests in the **In Processing** list. Click **Recycle** to cancel or logically delete the incident request. The incident request is in the **Recycle Bin** list after being recycled.

- **Restore**

Perform this operation for incident requests in the **Recycle Bin** list. Click **Restore** to restore the recycled incident request. After being restored, the incident request is in the **In Processing** list and restored to the status before the request is recycled.

- **Delete**

Perform this operation for incident requests in the **Recycle Bin** list. Click **Delete** to delete the incident request. After being deleted, the incident request is physically deleted and cannot be restored.

1.1.9.3.2.3 Manage incident tasks

After being created, an incident request is divided into different tasks based on the incident processing flow. Different tasks are to be processed by different people in charge.

Context

The processing of an incident task is divided into the following three steps:

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **My Task** tab.
3. Click the  (**My Task**) button.



Note:

To check out the tasks in the **Task Pool** to the current username, click the  (**Task Pool**) button and then click **Detail** at the right of the task. Click **Check Out**. In the displayed dialog box, enter the description and then click **OK**.

4. Find the task that you want to manage and then click **Detail**.
5. On the task details page, click **Diagnose**. In the displayed **Diagnose** dialog box, complete the configurations and then click **OK**.
 - **Diagnose Step:** analyzes the task steps.

- **Solution Type:** Select **Permanent Solution** or **Temporary Solution**. If you select **Temporary Solution**, you may have to create a problem request in the **Confirm** phase for further troubleshooting and locating the root cause of the problem.
 - **Is Complete:** Select **Yes** or **No** to indicate whether the task processing is complete. Sometimes the incident has been processed after being reported because of the time difference. In this case, you can directly select **Yes** and configure the resolved date. Then, the **Resolve** phase is skipped and the system goes to the **Confirm** phase directly.
 - **Remarks:** Enter the information about the task.
6. The system goes to the **Resolve** phase after the **Diagnose** phase. The **Resolve** phase consists of the incident troubleshooting and solving. ITIL only tracks this step in a standardized way and processes the log records. After processing the incident offline, click **Resolve** on the page. In the displayed **Resolve** dialog box, configure the resolved date and the handling steps. Then, click **OK**.
 7. The system goes to the **Confirm** phase after the **Resolve** phase. This phase reviews the processing result of the incident. Then, click **Confirm**. In the displayed **Confirm** dialog box, select the review result from the **Is Pass** drop-down list. Then, click **OK**.

The review results have the following three statuses:

- **Solved:** The incident is completely solved.
- **Unsolved, re-analysis:** The incident cannot be solved effectively because of an error in the reason analysis. The task is sent back to the **Diagnose** phase to restart the processing until the incident is solved.
- **Unsolved, reprocessing:** The reason of the incident is clear. The incident cannot be solved effectively because the incident is not effectively processed. The task is sent back to the **Resolve** phase to restart the processing until the incident is solved.

1.1.9.3.3 Manage problems

1.1.9.3.3.1 Create a problem request

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

Context

Temporarily resolved incidents or incidents whose root cause is not clear can be transformed to problems for further analysis and thorough troubleshooting. Problem management allows you

to find the root cause of incidents, thoroughly troubleshoot the incidents, and reduce repeated incidents.

Compared with the incident processing, problems have lower timeliness. The occurrence rate of repeated incidents is used to determine whether the problem management is good. The lower the occurrence rate is, the more effective the problem processing is.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management** > **Services**. Click the **Request** tab.
3. Click **New** and then select **Problem**. Configure the parameters on the displayed page.

For more information about the parameters, see [Parameter descriptions](#).

Table 1-3: Parameter descriptions

Parameter	Description
Report Object	The person who is required to process the request.
Callback Email	The email address of the person who records the request.
Callback Telephone	The telephone number of the person who records the request.
Product	The product to which the request belongs. Select a specific product from the drop-down list.
Service Name	The service related to the selected product. Select a specific service from the drop-down list.
Happen Date	The time when the request happens.
Priority	The priority of processing this request. The priority indicates the urgency of the request. The higher the urgency is, the higher priority the request must have. The priority has the following levels, from high to low, based on the urgency: <ul style="list-style-type: none"> • Critical • Major • Minor • Remind • Cleared • System
Alarm Code	The alarm ID.
Summary	The summary of this request.

Parameter	Description
Description	The detailed description about the request.
Suggestion	Optional. The suggestion about the request processing.

4. After configuring the preceding parameters, click **Confirm**.

1.1.9.3.3.2 Manage problem requests

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

Prerequisites

A problem request is created. For more information about how to create a problem request, see [Create a problem request](#).

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **Request** tab.
3. Click  at the right of the first drop-down list and then select **Problem** to display the problem requests in the list.
4. Find the problem request that you want to manage, and then click **Detail**.
5. You can perform the following operations on the request details page.

- **Change Priority**

Click **Change Priority**. In the displayed dialog box, select the new priority. Perform this operation for temporary adjustment or correcting the error in priority.



Note:

You can only change the priority of problem requests in the **Diagnose** phase.

- **Comment**

Click **Comment**. In the displayed dialog box, enter the comment for this problem request. Perform this operation for collaborative scenarios. For example, users can comment the problem request to share the information with each other and guide each other when they process the same problem.

- **Suspend**

Click **Suspend**. In the displayed dialog box, enter the remarks. Perform this operation for problem requests that currently do not require to be processed.

- **Resume**

Click **Resume**. In the displayed dialog box, enter the remarks. Perform this operation for suspended problem requests that require to be processed.

- **Recycle**

Perform this operation for problem requests in the **In Processing** list. Click **Recycle** to cancel or logically delete the problem request. The problem request is in the **Recycle Bin** list after being recycled.

- **Restore**

Perform this operation for problem requests in the **Recycle Bin** list. Click **Restore** to restore the recycled problem request. After being restored, the problem request is in the **In Processing** list and restored to the status before the request is recycled.

- **Delete**

Perform this operation for problem requests in the **Recycle Bin** list. Click **Delete** to delete the problem request. After being deleted, the problem request is physically deleted and cannot be restored.

1.1.9.3.3.3 Manage problem tasks

After being created, a problem request is divided into different tasks based on the problem processing flow.

Context

The processing of a problem task is divided into the following three steps:

- **Diagnose**: analyzes the reason of the problem.
- **Resolve**: The system goes to the **Resolve** phase after the **Diagnose** phase. The problem is repaired in this phase.
- **Confirm**: The system goes to the **Confirm** phase after the **Resolve** phase and reviews if the problem processing is reasonable.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **My Task** tab.

3. Click the  (**My Task**) button.



Note:

To check out the tasks in the **Task Pool** to the current username, click the **Task Pool**  button and then click **Detail** at the right of the task. Click **Check Out**. In the displayed dialog box, enter the description and then click **OK**.

4. Find the task that you want to manage and then click **Detail**.
5. On the task details page, click **Diagnose**. In the displayed **Diagnose** dialog box, complete the configurations and then click **OK**.
- **Diagnose Step:** analyzes the task steps.
 - **Solution Type:** Select **Permanent Solution** or **Temporary Solution**. If you select **Temporary Solution**, you may have to create a problem request in the **Confirm** phase for further troubleshooting and locating the root cause of the problem.
 - **Is Complete:** Select **Yes** or **No** to indicate whether the task processing is completed. Sometimes the problem has been processed after being reported because of the time difference. In this case, you can directly select **Yes** and configure the resolved date. Then, the **Resolve** phase is skipped and the system goes to the **Confirm** phase directly.
 - **Remarks:** Enter the information about the task.
6. The system goes to the **Resolve** phase after the **Diagnose** phase. The **Resolve** phase includes the problem troubleshooting and solving. ITIL only tracks this step in a standardized way and processes the log records. After processing the problem offline, click **Resolve** on the page. In the displayed **Resolve** dialog box, configure the resolved date and handling steps. Then, click **OK**.
7. The system goes to the **Confirm** phase after the **Resolve** phase. This phase reviews the processing result of the problem. Then, click **Confirm**. In the displayed **Confirm** dialog box, select the review result from the **Is Pass** drop-down list. Then, click **OK**.

The review results have the following three statuses:

- **Solved:** The problem is completely solved.
- **Unsolved, re-analysis:** The problem cannot be solved effectively because of an error in the reason analysis. The task is sent back to the **Diagnose** phase to restart the processing until the problem is solved.

- **Unsolved, reprocessing:** The reason of the problem is clear. The problem cannot be solved effectively because the problem is not effectively processed. The task is sent back to the **Resolve** phase to restart the processing until the problem is solved.

1.1.9.4 Version control

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > Version Control**.

Select the product in the tree or enter the product name in the search box. The version and cluster information are displayed on the right.



Note:

Before the search, click  to synchronize the information to ASO.

1.1.9.5 Process template configuration

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

After [logging on to Apsara Stack Operations \(ASO\)](#), choose **ITIL Management > Process Template Configuration** in the left-side navigation pane. On this page, you can view the following three sections: **Process**, **Process Template**, and **Regulation**.

Process

Currently, the following processes are supported:

- Incident
- Problem
- Change Role
- Create Identity
- Reset Password
- Logout Identity
- Change

- Version Upgrade
- Hotfix Upgrade
- Configuration Upgrade

Process Template

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

-  is the start node of the process. A process usually starts with the request creation.
-  indicates the gateway. The gateway defines the process trend in different branches. In the BPMN specification, gateways are classified into different types, such as inclusive gateway, exclusive gateway, parallel gateway, and hybrid gateway. Here it is the exclusive gateway, indicating that multiple routes have only one valid path.
-  is the end node of the process. A process usually ends with archiving.
-  indicates the phase. A phase is usually composed of roles with specific functions.
-  is the route, indicating the process trend. A phase contains one or more egress routes and ingress routes.

The templates can be classified into the following three types:

- **Incidents and problems**

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

- **Request fulfillment**

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

- **Change management**

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

Regulation

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

Currently, the system supports four regulations:

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

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

- **Assign By Role**

Select **Assign By Role** and then select roles from the drop-down list.

- If no role is selected, all the users can view the current task in the **Task Pool** by default.
- If the selected role has only one user, only that user can view the current task in **My Task**.
- If the selected role has more than one user, all the users under the selected role can view the current task in the **Task Pool**.



Note:

If no regulation is configured in this phase, all the users can view the current task in the **Task Pool** by default.

- **Assign By User**

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

- If no user is selected, all the users can view the current task in the **Task Pool** by default.
- If only one user is selected, only that user can view the current task in **My Task**.
- If more than one user is selected, all the selected users can view the current task in the **Task Pool**.



Note:

If no regulation is configured in this phase, all the users can view the current task in the **Task Pool** by default.

- **Assign By Owner**

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

**Note:**

If no regulation is configured in this phase, all the users can view the current task in the **Task Pool** by default.

- **CAB/ECAB Configuration**

CAB/ECAB Configuration only appears if you click the **CAB Audit** or **ECAB Audit** phase in a change management process.

Click **CAB/ECAB Configuration** to go to the **CAB Configuration** or **ECAB Configuration** page. For more information, see [Configure CAB or ECAB](#).

1.1.9.6 Configure CAB or ECAB

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

Context

CAB and ECAB are terminologies of ITIL specifications. CAB is abbreviated from Change Advisory Board and ECAB is abbreviated from Emergency Change Advisory Board.

In all the process templates, the CAB configuration of the **CAB Audit** phase is similar to the ECAB configuration of the **ECAB Audit** phase. In this section, use the CAB configuration as an example.

If no regulation is configured, all the users can generate the current task in **My Task** by default.

With one or more users configured, each configured user can generate the current task in **My Task**, and the task can go to the next phase only after all the users configured in this phase finish the current task.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **ITIL Management > CAB/ECAB Configuration**.
3. Select the check boxes on the left or right and then click >> or << to configure the CAB. Users listed on the right are the current user configuration.

1.1.10 API management

1.1.10.1 Overview

API management provides a unified encapsulation for the operations APIs of all cloud products on the cloud platform, which facilitates the secondary development of the operations platform for third-parties, allows the fine-grained access control and security audit of the operations APIs, and provides the centralized management in terms of Apsara Stack versions and APIs.

1.1.10.2 Category

Category allows you to view all the APIs published by products.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **API Management** > **Category**.
3. On the **Category** page, you can perform the following operations:

- **Search for APIs**

Select a product from the **Select a product** drop-down list and then enter the API name in the search box to search for APIs. Fuzzy search is supported.

- **Edit an API**

To edit an API, click **Edit** in the **Actions** column. In the displayed dialog box, edit the **Basic Information** and **Parameters** of the API. Then, click **Save** to submit the changes.

- **Delete an API**

To delete an API, click **Delete** in the **Actions** column, and then click **Confirm** in the displayed dialog box.

1.1.11 Configurations

1.1.11.1 Overview

Configurations allows you to modify the related configuration items of each product as required.

To modify a configuration item of a product, you can modify the configuration value in Apsara Stack Operations (ASO) and then apply the modifications. To restore the configuration value of a modified configuration item, you can roll back the configuration value with one click.

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

1.1.11.2 Modify a configuration item of a product

You can modify a configuration item of a product as required.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Configurations > Configuration Items**.
3. On the **Configuration Items** page, enter the name of the product or configuration item in the **Product** or **Configuration Name** field. Click **Search** to check if the configuration item already exists in the list.
 - **The configuration item already exists in the list**
 - a. Click **Get** in the **Actions** column to load the actual data from the product to your local host.
 - b. Click **Modify** in the **Actions** column. In the displayed **Modify Configurations** dialog box, modify the values and then click **OK** to modify the configuration item locally.
 - **The configuration item does not exist in the list**

You must add a configuration item. Click **Add** in the upper-right corner. In the displayed **Add Configuration** dialog box, configure the information, such as **Product**, **Configuration Name**, **Default Value**, and **Data Source Type**, for the configuration item. Click **OK** and then this configuration item is displayed in the list. You can search for or modify this configuration item.
4. After the configuration item is modified, click **Apply** in the **Actions** column to make the modifications take effect.
5. Optional: To import or export configuration items as a file, click **Import** or **Export**.

1.1.11.3 Restore the configuration value of a modified configuration item

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Configurations > Restore**.

3. On the **Restore** page, enter the name of the configuration item whose configuration value you want to roll back in the **Configuration Name** field and then click **Search**. All modification records of the configuration item appear in the list.
4. Find the record to be rolled back, and then click **Restore** in the **Actions** column. Click **OK** in the displayed dialog box to restore the configuration value of the configuration item.

1.1.11.4 Kernel configurations list

You can add or modify the kernel configurations.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Configurations > Kernel Configurations**.
3. On the **Kernel Configurations** page, you can perform the following operations:

- **Add a kernel configuration**

Click **Add** at the top of the page. In the displayed dialog box, enter the **Configuration Name**, **Read Command**, and **Modify Command**. Then, click **Submit**.

- **Modify a kernel configuration**

Find the kernel configuration that you want to modify. Click **Modify** in the **Actions** column. Modify the **Kernel Configuration**, **Read Command**, and **Modify Command**. Then, click **Save**.

- **Delete a kernel configuration**

Find the kernel configuration that you want to delete. Click **Delete** in the **Actions** column. In the displayed dialog box, click **OK**.

1.1.11.5 Kernel configurations actions

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

Prerequisites

The kernel configurations to be scanned are added in the [Kernel configurations list](#) before the scanning.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Configurations > Kernel Configurations Actions**.

3. On the **Kernel Configurations Actions** page, enter the hostname or IP address in the search box and then click **Scan Configuration**.

The scan results are displayed in the list.

4. Optional: To modify the scanned configuration value, click **Modify** to modify the **Configuration Value**. Click **Save** to modify the local value of the kernel configuration.

After the modification, click **Apply** to apply the local value of the kernel configuration to the corresponding host. To read the value of the kernel configuration on the host again, click **Obtain**.

1.1.12 Offline backup

You can view the backup information by using offline backup.

Context

Offline backup is used to back up the key metadata of Apsara Stack. Metadata backup is used for the fast recovery of Apsara Stack faults. Offline backup services include:

- Backup service: provides backup configuration, backup details, and service status.
- Service configuration: provides backup service configuration and product management.
- Service status: queries the current status of backup services, including backup products, completed backup items, timeout backup items, and failed backup items, and displays the status of the current backup server in graphs.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, click **Offline Backup**.
3. On the **Offline Backup** page, you can perform the following operations:
 - **Backup Service**

Table 1-4: Description of backup service

Function menu	Description
Backup Configuration	The left part of the Backup Configuration page displays backup configurations in a hierarchical tree structure. The root node is a product list and displays backup products provided by the current backup system.

Function menu	Description
	<p>Currently, only pangu metadata backup is provided.</p> <p>The backup item is the minimum unit of backup. You can back up the metadata of different pangus, such as ecs pangu, rds pangu, and ots pangu based on Apsara Stack. The preceding configurations are added in Service Configuration > Product Management.</p> <p>Click a product under pangu to view the product details, namely Product, Backup Items, Backup Script, Product Cluster Location, Backup File Folder, Script Execution Folder, Script Parameters, Backup Schedule, Backup Schedule Unit, and Time-out.</p> <p>In the upper-right corner, click Modify to modify the configurations.</p>
Backup Details	<p>Displays the current backup status. The backup details include Product, Backup Items, File Name (files that require to be backed up), Start Time, and State (not started, in process, timeout, and error).</p> <p>You can configure the search conditions and then click Search to obtain the backup details.</p>
Service Status	<p>Displays the status of the current backup server and provides usage graphs for the memory and disk.</p>

- **Service Configuration**

Table 1-5: Description of service configuration

Function menu	Description
Backup Service Configuration	<p>Provides backup server configurations.</p> <p>— Backup Server IP Address: Configure the IP address of the backup server. The server must be an independent physical server managed by Apsara Infrastructure Management Framework,</p>

Function menu	Description
	<p>with the network connected with other servers in Apsara Stack. Pangu cannot be deployed on the server, at least cannot be deployed on its disk that stores the backup metadata.</p> <ul style="list-style-type: none"> — Backup Server Monitoring Path: The backup server detects new backup files by monitoring the specified folder on the backup server and determines whether the backup is successful by comparing the MD5 values of the backup file and the original file. The monitoring path is the file storage path on the backup server. — Backup Retention (day): the file storage time on the backup server. The backup file that exceeds the time will be deleted. <p>Click Modify in the Actions column to modify the configurations.</p>
Product Management	<p>Provides the basic management of backup products as follows:</p> <ol style="list-style-type: none"> a. Click Add in the upper-right corner. In the displayed Add Product dialog box, enter the Product, Backup Items, and Backup Script, select the Retry Times, and then click OK. The added product is displayed in Backup Configuration of Backup Service. b. The status of current backup products is displayed in the table. You can modify or delete a product by clicking Modify or Delete in the Actions column. Modifying a product is similar to adding a product. When modifying a product, you can click Delete to delete a backup item.

- **Service Status**

The current backup status. The status at the top of the table includes **In Process**, **Complete**, **Time-out**, and **Failed**. The following table lists the status of the latest backup items. The single record indicates the current product, the numbers of completed backup items and failed backup items, and the status of the latest backup items. The backup status includes success, not started, in process, timeout, and failure.

The **Backup Server Status** on the right displays the status of memory and disk of the backup server in graphs.

1.1.13 NOC

1.1.13.1 Overview

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

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

1.1.13.2 Dashboard

The dashboard is mainly used to monitor the current devices, network, and traffic.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **NOC > Dashboard**.

Item		Description
Device Management	Device Overview	The model distribution of used network devices.
	Ports Usage	<ul style="list-style-type: none"> • Ports Utilization: the proportion of ports in use to the total ports in the network devices. • Error Packets by Port (Top 5): the total number of error packets generated by device ports within a certain time range,

Item		Description
		of which the top 5 are displayed.
	Configuration Management	<ul style="list-style-type: none"> • Automatic Backup: the backup of startup configurations for all network devices. • Configuration Sync: the synchronization of running configurations and startup configurations for all network devices.
Network Monitoring	Alarms	The total number of alarms generated by network devices .
	Alarming Devices	The number of network devices that generate alarms and the total number of network devices.
	Alarm Details	The details of the alarm.
Traffic Dashboard	SLB Overview	The bandwidth utilization of SLB clusters.
	XGW Overview	The bandwidth utilization of XGW clusters.

1.1.13.3 Network topology

The **Network Topology** page allows you to view the topology of physical networks.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **NOC > Network Topology**. The **Network Topology** page displays the physical network topology of a physical data center.



Note:

The colors of the connections between network devices represent the connectivity between the network devices.

- Green: The connection works normally.

- Red: The connection has an error.
- Grey: The connection is not enabled.

3. Click **Detail** in the upper-right corner to view the **Device Properties** and **Port Status**.
4. Click a physical network device in the network topology. The **Device Properties** and **Port Status** of the device are displayed on the right.

1.1.13.4 Physical network integration

The **Physical Network Integration** allows network operations engineers to perform automated integration of physical networks on the user interface by entering the integration parameters. Network Operation Center (NOC) automatically generates and issues the configurations to specific devices and the network integration test is started automatically.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **NOC > Network Reconfiguration > Physical Network Integration**.

3. Create a project file. Enter the project name and click **Create**.

The operations engineer creates a project file for this change to store the parameters related to the change and this project file can be imported for later usage.

4. Click **Next**.
5. Find the device required by this change on the left, and then click **Add to Target**.

To change the device, click **Manage > Delete** on the right. You can also click **Manage > Set Username and Password** to modify the logon username and password of the device.

6. Click **Next**.
7. Configure the interface parameters. Click **Edit**. Complete the parameter configurations and then click **Add**.
8. Click **Next**.
9. Configure the route parameters. Click **Edit**. Complete the parameter configurations and then click **Add**.
10. Click **Next**.
11. Configure the route policies. Click **Edit**. Complete the parameter configurations and then click **Add**.
12. Click **Next**.

13. Generate the combination configurations. Click **Generate**.

Operations engineers can automatically generate the configurations of each device based on the configured parameters, view, and export the generated configurations.

1.1.13.5 Password management

You can manage the account password of a device.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
 2. In the left-side navigation pane, choose **NOC > Resource Management > Password Management**.
 3. Enter the name of the device whose password you want to modify in the search box and then click **Search**.
 4. Select the device and then click **Add to Target**.
- Then, the device is displayed under **Target Devices**.
5. The system must verify the password before you modify it. Enter the **Username** and **Old Password** in the lower-right corner and then click **Verify**.
 6. After the verification is passed, click **Modify** to modify the password.

Select one or more devices to modify the password as required. You can verify and modify the password of all physical devices under **Target Devices**.

1.1.13.6 Configuration comparison

For a device, you can compare its current configuration with its configuration at startup and check if they are consistent.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **NOC > Configuration Management > Config Comparison**.
3. Find and select the device whose configurations you want to compare and then click **Compare Configuration**.

After the comparison, click **Export Results** to export the differences.

1.1.14 Full stack monitor

1.1.14.1 Overview

Full stack monitor allows you to perform an aggregate query on the system alarm events, query and retrieve all the alarm data in the link based on the host IP address, instance ID, and time range, and view the end-to-end topology.

1.1.14.2 SLA

1.1.14.2.1 Overview

SLA allows you to view the current state, history data, and instance availability of each cloud product.

You can view the current and history fault state of products to obtain the SLA values and unavailable events of product instances within a certain time period.

1.1.14.2.2 View the current state of a cloud product

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > SLA**.
3. Click the **Current State** tab.

The current state and the state in the last 24 hours of each cloud product are displayed on this page. Different colors represent different states:

- Green: normal. The service is running normally.
 - Yellow: warning. The service has some latency, but can still work normally.
 - Red: hitch. The service is temporarily interrupted and cannot work normally.
4. Find the product whose running state you want to view. Click **Check** in the **Operation** column.
 - **Overall Availability** displays the availability of a product. You can view the product availability by hour, day, or minute.
 - **Related Events** displays the current exception events. Click **Show Details** to view the event details.

1.1.14.2.3 View the history data of a cloud product

The **History Data** tab allows you to view the history state of a cloud product and the details of exception events.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > SLA**.
3. Click the **History Data** tab.

The product availability of each cloud product in the last two weeks are displayed on this page. Different colors represent different statuses:

- Green: normal. The service is running normally.
 - Yellow: warning. The service has some latency, but can still work normally.
 - Red: hitch. The service is temporarily interrupted and cannot work normally.
4. Find the product whose history state you want to view. Click **Check** in the **Operation** column.
 - **Overall Availability** displays the history availability of a product. You can view the history availability by hour, day, or minute.
 - **Related Events** displays the history exception events. Click **Show Details** to view the event details.

1.1.14.2.4 View the availability of an instance

The **Availability of Instance** tab allows you to view the availability of an instance and know the instance damages.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > SLA**.
3. Click the **Availability of Instance** tab.
4. Enter the **Instance ID** and **Belonged to User**, and/or select the **Time Range**. Then, click **Search**.
5. Click the instance ID to view the following information of the instance.
 - **Basic Information**: the instance ID and the user to whom the instance belongs.
 - **Availability**: the availability ratio of the instance.
 - **Damage Event**: the exception event list.

1.1.14.3 ECS operations full link logs

The **ECS Operations Full Link Logs** allows you to search for logs of ECS-related applications.

Context

Currently, you can search for logs of multiple product components, such as POP, OpenAPI, PYNC, and OPSAPI. You can enter any string in the **Query** field as the search condition, such as the instance ID, request ID, or the keyword "error".

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > ECS Operations Full Link Logs**.
3. Enter the keyword in the **Query** field, such as the instance ID. Select the time range in the **Time** field and then click **Search**.
4. Select **Abnormal logs only** to only display the abnormal logs.

If code != 200, success=false, or error exists in a log, the log is an abnormal log.
5. Enter the keyword in the search box to search for the related information in the search results.

1.1.14.4 Correlation diagnosis and alarm

1.1.14.4.1 Full stack correlation alert

Full Stack Correlation Alert allows you to view the alarm event list after the aggregation and the corresponding details.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **Full Stack Correlation Alert** tab.

The **Full Stack Correlation Alert** tab displays the alarm events aggregated from the abnormal events in the current system by using the correlation diagnosis.

4. Enter the instance ID, such as a physical machine name, instance name of a cloud product, and network device name, in the search box, select the time range, and then click **Search**.
5. Optional: In the displayed alarm list, click  at the right of **Alert Type** and **Alert Level** to filter the alarm results.

6. Click **Details** to view the details of the abnormal event.

The event details page displays the **Alert Basic Information**, **Associated Event Information**, **Impacted Instances in ECS**, and **Impacted Instances in RDS**.

1.1.14.4.2 Server

You can use the server IP address or server name to query the end-to-end topology, basic information, and real-time diagnosis information of a server, the alarm information of the network where a server is located, and the full stack correlation alarm information.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **Server** tab.
4. Enter the host IP address or ECS instance ID in the search box, select the time range, and then click **Search**.

Click **+** at the right of the search box and then another search box is displayed. You can query the network topology from a server to another target server as required.

5. You can view the following information on this page.
 - **Topology** displays the uplink network topology of the host, which directly shows the alarms of network devices (blue indicates the normal status and red indicates the abnormal status).
 - **Title Message** displays the basic operating data for the operating system of the host.
 - **NC Diagnostics Info** displays the real-time diagnosis and alarm data of the host.
 -  indicates the diagnosis is passed.
 -  indicates the detection does not obtain results.
 -  indicates abnormal alarm level exists.
 -  indicates fatal exception exists.
 -  indicates the running diagnosis items.
 - **Network Warning Information** displays the alarm data of the network devices that are included in the uplink network topology of the host.

- **Full Stack Alert** displays the list of alarm events after the aggregation and the corresponding details.

1.1.14.4.3 Network equipment

You can use the network equipment IP address or network equipment name to query the basic information, real-time diagnosis information, and full stack correlation alarm information of a network equipment.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **Network Equipment** tab.
4. Enter the network equipment ID in the search box, select the time range, and then click **Search**.
5. You can view the following information on this page.
 - **Essential Information** displays the basic operating data for the operating system of the network equipment.
 - **Diagnostic Information** displays the diagnosis information of the network equipment.
 - **Full Stack Alert** displays the full stack correlation alarm information.

1.1.14.4.4 ECS

You can use the ECS instance ID to query the basic information, bandwidth charts of physical network devices and virtual network devices, and full stack correlation alarm information of an ECS instance.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **ECS** tab.
4. Enter the ECS instance ID in the search box, select the time range, and then click **Search**.
5. You can view the following information on this page.
 - **Topology** displays the uplink network topology of the host, which directly shows the alarms of network devices (blue indicates the normal status and red indicates the abnormal status).

- **ECS Basic Info** displays the basic data of the ECS instance and the host.
- **HostNC Basic Info** displays the basic information of the physical machine of the host.
- **ECS Diagnosis Info** displays the diagnosis and alarm data of the ECS instance and the host to which the ECS instance belongs.
- **HostNC Diagnosis Info** displays the health diagnosis information of all physical devices deployed by Apsara Stack, including the discovered problems and the corresponding fixes.
- The operating water level of the ECS instance.
- **netdev** displays the traffic and packet information of the virtual NIC netdev on the host to which the ECS instance belongs. You can display the traffic or packet information by switching between the two tabs.
- **vport** displays the traffic, number of connections, and packet information of the virtual switch port vport on the host to which the ECS instance belongs. You can display the traffic, number of connections, or packet information by switching among the tabs.
- **Network Warning Information** displays the alarm data of the network devices that are included in the uplink network topology of the host to which the ECS instance belongs.
- **Full Stack Alert** displays the aggregated alarm events on the ECS instance and the uplink devices of the ECS instance.

1.1.14.4.5 RDS

You can use the RDS instance ID to query the full stack information, availability diagnosis results, and full stack correlation alarm information of an RDS instance.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **RDS** tab.
4. Enter the RDS instance ID in the search box, select the time range, and then click **Search**.
5. You can view the following information on this page.
 - **Topology** displays the uplink network topology of the host where the RDS instance is located, which directly shows the alarms of network devices (blue indicates the normal status and red indicates the abnormal status).
 - **Basic Info** displays the basic information of the RDS instance, including the primary database IP address, secondary database IP address, SLB ID, and Proxy IP address.

- **Diagnosis Info** displays the availability detection results of the RDS instance in the selected time range.
- **Network Alert Info** displays the alarm data of the network devices that are included in the uplink network topology of physical machines in the primary database.
- **Full Stack Alert** displays the aggregated alarm events on the RDS instance and the uplink devices of the RDS instance.

1.1.14.4.6 SLB

You can use the SLB instance ID to query the deployment information, traffic diagnosis results, and bandwidth chart of an SLB instance.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **SLB** tab.
4. Enter the SLB instance ID in the search box, select the time range, and then click **Search**.
5. You can view the following information on this page.
 - **SLB Clusters** displays the deployment information of the SLB cluster, including the service name and host ID.
 - **Diagnostics** displays the availability detection results of the SLB instance in the selected time range.
 - **SLB Bandwidth Chart** displays the bandwidth chart of the SLB instance in the selected time range.

1.1.15 Pangu monitoring

1.1.15.1 Overview

Pangu Monitoring displays the **Pangu Grail**, **Cluster Information**, and **Node Information**.

- **Pangu Grail** allows you to view the overview, heatmap of health, and data of the top 5 healthiest clusters of a product.
- **Cluster Information** allows you to view the overview and run chart of a cluster.
- **Node Information** allows you to view the master information and chunk server information in a cluster.

1.1.15.2 Pangu grail

Pangu Grail allows you to view the overview, heatmap of health, and data of the top 5 healthiest clusters of a product.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Pangu Monitoring > Pangu Grail**.
3. Select the product that you want to view from the **Service** drop-down list.

Pangu grail displays the current overview, heatmap of health, and top 5 data of each accessed cloud product. By default, data of the first product in the returned product data columns of the current environment is displayed. You can select products as required.

- **Overview**

Overview displays the storage space, server information, and health information of the selected product. Values of abnormal disks, abnormal masters, abnormal chunk servers, and abnormal water levels in the **Health** section are displayed in red if they are larger than zero.

- **Heatmap of Health**

Heatmap of Health displays the health information of all the clusters in the selected product. Clusters in different health statuses are displayed in different colors. Green indicates the normal status, yellow indicates a warning, red indicates the abnormal status, dark red indicates a major mistake, and grey indicates the closed status. Click the name of a cluster that is not in the closed status to go to the corresponding cluster information page.

- **Data of Top 5 Services**

Data of Top 5 Services displays the data of the top 5 healthiest clusters in the time range between zero o'clock and the current time in the current day for the selected product. Click the cluster name to go to the corresponding cluster information page.

1.1.15.3 Cluster information

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Pangu Monitoring > Cluster Information**.

By default, data of the first cluster in the returned cluster names is displayed.

3. Select the cluster that you want to view from the **Cluster Name** drop-down list.

You can select all the accessed clusters that are not in the closed status in the current environment from the **Cluster Name** drop-down list.

- **Overview** displays the storage space, server information, and health information of the selected cluster. Values of abnormal water levels, abnormal masters, abnormal chunk servers, and abnormal disks in the **Health** section are displayed in red if they are larger than zero.
- **Run Chart of Clusters** displays the charts of historical water levels, predicted water levels, number of files, number of chunk servers, number of disks, number of servers in rack, and storage for the selected cluster.
 - **Predicted Water Levels** predicts the run chart of the next seven days.



Note:

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

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

1.1.15.4 Node information

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **Pangu Monitoring > Node Information**.

By default, data of the first cluster in the returned cluster names is displayed, including the master information and chunk server information.

3. Select the cluster that you want to view from the **Cluster Name** drop-down list.

You can select all the accessed clusters that are not in the closed status in the current environment from the **Cluster Name** drop-down list.

- **Master Info** displays the master information in the selected cluster. Partial refresh is supported. You can click **Refresh** to refresh the master information in the selected cluster.

- **Chunk Server Info** displays the chunk server information in the selected cluster. Partial refresh is supported. You can click **Refresh** to refresh the chunk server information in the selected cluster. Click **+** to display the disk overview and SSDCache overview in the current chunk server. Fuzzy search is supported.

1.1.16 System management

1.1.16.1 Overview

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

1.1.16.2 Department management

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

Context

After Apsara Stack Operations (ASO) is deployed, a root department is generated by default. You can create other departments under the root department. The departments are displayed in hierarchy and you can create subdepartments under each department level.

Procedure

1. [Log on to ASO](#).
2. In the left-side navigation pane, choose **System Management > Departments**.

On the **Department Management** page, you can view the tree structure of all created departments, and the user information under each department.

3. You can perform the following operations on this page:

- **Add a department**

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

- **Modify a department**

Select the department to be modified in the tree catalog and click **Modify Department** at the top of the page. In the displayed **Modify Department** dialog box, enter the **Department Name** and click **OK**.

- **Delete a department**

Select the department to be deleted in the tree catalog and click **Delete Department** at the top of the page. Click **OK** in the displayed dialog box.

1.1.16.3 Role management

You can add custom roles on Apsara Stack Operations (ASO) to better allocate permissions to users.

Context

A role is a collection of access permissions. When creating users, you must assign roles to users to meet their access control requirements on the system. Roles are classified into basic roles and user-created roles. The basic roles, also known as atomic roles, are preset by the Operation Access Management (OAM) system and cannot be modified or deleted by users. The user-created roles can be modified, updated, and deleted.

Procedure

1. [Log on to ASO](#).
2. In the left-side navigation pane, choose **System Management > Roles**.
3. On the **Role Management** page, you can perform the following operations:

- **Search for roles**

**Note:**

Both the ASO security officer and the system administrator can search for roles.

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

- **Add a role**

**Note:**

In ASO, only the ASO security officer can create roles.

Click **Add Role** at the top of the page. In the displayed **Add Role** dialog box, enter the **Role Name**, **Role Description**, and **Base Role**, and then click **OK**.

- **Modify a role**

**Note:**

In ASO, only the ASO security officer can modify roles.

Find the role to be modified, and then click **Modify** in the **Actions** column. In the displayed **Modify Role** dialog box, modify the information and then click **OK**.

- **Delete a role**

Find the role to be deleted, and then click **Delete** in the **Actions** column. Click **OK** in the displayed dialog box.

1.1.16.4 Logon policy management

The administrator can configure the logon policies to control the read/write permissions of user logon.

Context

During the system initialization, the system has a default policy for the read/write permissions of users. After configuring the logon policies, you can better guarantee the read/write permissions of users, which improves the system security.

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **System Management > Logon Policies**.
3. On the **Logon Policy Management** page, you can perform the following operations:

- **Search for policies**

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

- **Add a policy**

Click **Add Policy**. In the displayed dialog box, configure the **Policy Name**, **Start Time**, **End Time**, and allowed logon address. Then, click **OK**.

- **Modify a policy**

Find the policy to be modified, and then click **Modify** in the **Actions** column. In the displayed **Update Policy** dialog box, modify the information and then click **OK**.

- **Delete a policy**

Find the policy to be deleted, and then click **Delete** in the **Actions** column. Click **OK** in the displayed dialog box.

1.1.16.5 User management

The administrator can create users and assign roles to users to meet their access control requirements on the system.

Prerequisites

Before you create a user, make sure that:

- A department is created. For more information, see [Department management](#).
- A custom role is created, if required. For more information, see [Role management](#).

Context

Role management provides different operation permissions for different users. During the system initialization, the system creates three default users: asosysadmin, asosecurity, and asoauditor. The default users are respectively bound to the following default roles: system administrator, security officer, and security auditor. The three roles have the same default password: AliOS%1688. The permissions of these three roles are as follows:

- The system administrator can view, modify, delete, and add the operation and maintenance dashboard, alarm management, resource management, inventory management, backup service, configurations, help center, and application whitelist, and can view users, roles, departments, logon policies, and server passwords in system management.
- The security officer can view, modify, delete, and add the users, roles, departments, logon policies, and server passwords in system management.
- The security auditor can read and write Apsara Stack Operations (ASO) system logs.

Procedure

1. [Log on to ASO](#).
2. In the left-side navigation pane, choose **System Management > Users**. Click the **Users** tab.

The **Users** tab allows you to view a list of all created users. In the list, you can search for, add, modify, and delete users and bind logon policies to users.

- **Search for users**



Note:

Both the ASO security officer and the system administrator can search for users.

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

- **Add a user**



Note:

In ASO, only the security officer can add users.

At the top of the page, click **Add**. In the displayed **Add User** dialog box, configure the information, such as **User Name** and **Password**, and then click **OK** to add the user.

- **Modify a user**



Note:

In ASO, only the security officer can modify the user information.

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

- **Delete a user**

Find the user to be deleted, and then click **Delete** in the **Actions** column. Click **OK** in the displayed dialog box.



Note:

Deleted users are in the recycle bin. To restore a deleted user, click the **Recycled** tab.

Find the user to be restored, click **Cleared** in the **Actions** column, and then click **OK** in the displayed dialog box.

- **Bind a logon policy**

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

- **View personal information of the current user**

In the upper-right corner, click  next to the logon username and then select **Personal Information** from the drop-down list. The **Personal Information** dialog box appears, indicating the personal information of the current user.

1.1.16.6 Two-factor authentication

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

Context

Currently, Apsara Stack Operations (ASO) supports three authentication methods. Select one method to configure the authentication:

- **Google Two-Factor Authentication**

Use the password and mobile phone to provide double protection for accounts if you select this method. You can obtain the logon key after configuring users in ASO, and then enter the key in the Google authenticator app of your mobile phone. The app dynamically generates a verification code based on the time and key for logon.

- **USB Key Authentication**

Install the drive and browser controls (currently, only Windows + IE 11 environment is supported) according to the third-party manufacturer instructions if you select this method. The third-party manufacturer provides the hardware USB key and the service that the backend authenticates and verifies the certificates. The hardware USB key includes the serial number and certificate information. Before the authentication, bind the serial number with a user account, configure the authentication server provided by the third-party manufacturer, and enable the USB key authentication for the user when you configure the authentication method in ASO.

Upon logon, if the account enables the USB key authentication, the ASO frontend calls the browser controls, reads the certificate in USB key, obtains the random code from the backend, encrypts the information, and sends the information to the backend. The backend calls the authentication server to parse the encrypted strings, verifies the certificate and serial number, and then completes the other logon processes if the verification is passed.

- **PKI Authentication**

Enable the ASO HTTPS mutual authentication and change the certificate provided by the user if you select this method. The third-party manufacturer makes the certificate and provides the service that the backend verifies the certificate. After the mutual HTTPS authentication is enabled, the request carries the client certificate upon logon to send the certificate to the backend, and the backend calls the parsing and verification service of the third-party manufacturer to verify the certificate. The certificate includes the name and ID number of a

user. Therefore, bind the name and ID number with a user account when you configure the authentication method in ASO.

Authentication server

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

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

Procedure

1. [Log on to ASO](#).
2. In the left-side navigation pane, choose **System Management > Two Factor Authentication**.
3. On the **Two Factor Authentication** page, you can perform the following operations:
 - **Google Two-Factor Authentication**
 - a. Select **Google Two-Factor Authentication** as the **Current Authentication Method**.
 - b. Click **Add User** in the upper-right corner. The added user is displayed in the user list.
 - c. Find the user that you want to enable the Google two-factor authentication, and then click **Create Key** in the **Actions** column. After the key is created, **No Key** is changed to **Show Key**. Click **Show Key**, the created key is displayed in plain text.
 - d. Enter the key in the Google authenticator app of your mobile phone. The app dynamically generates a verification code based on the time and key for logon. With the two-factor authentication enabled, you are required to enter the verification code on your app when logging on to the system.



Note:

Both Google two-factor authentication app and server generate the verification code based on the public algorithms of time and keys, and can work offline without connecting to the Internet or Google server. Therefore, keep your key confidential.

- e. To disable the two-factor authentication, click **Delete Key** in the **Actions** column. After the successful deletion, **Show Key** is changed to **No Key**.
- **USB Key Authentication**

- a. Select **USB Key Authentication** as the **Current Authentication Method**.
- b. In **Authentication Server Configuration**, click **Add Server**. In the displayed dialog box, enter the IP address and port of the server, and then click **OK**. The added server is displayed in **Authentication Server Configuration**. Click **Test** to test the connectivity of the authentication server.
- c. In **User List**, click **Add User**. The added user is displayed in the user list.
- d. Find the user that you want to enable the USB key authentication, and then click **Bind Serial Number** in the **Actions** column. In the displayed dialog box, enter the serial number to bind the user account with this serial number.

**Note:**

When adding an authentication in ASO, ASO calls the browser control to automatically enter the serial number. If the serial number fails to be entered, you must enter it manually. The serial number of USB key authentication is contained in the USB key hardware. Therefore, you must insert the USB key, install the drive and browser control, and then read the serial number by calling the browser control.

- e. Then, click **Enable Authentication** in the **Actions** column.
- **PKI Authentication**
 - a. Select **PKI Authentication** as the **Current Authentication Method**.
 - b. In **Authentication Server Configuration**, click **Add Server**. In the displayed dialog box, enter the IP address and port of the server, and then click **OK**. The added server is displayed in **Authentication Server Configuration**. Click **Test** to test the connectivity of the authentication server.
 - c. In **User List**, click **Add User**. Enter the **Username**, **Full Name**, and **ID Card Number**, and then click **OK**. The added user is displayed in the user list.
 - d. Find the user that you want to enable the PKI authentication, and then click **Bind** in the **Actions** column. Enter the full name and ID card number of the user to bind the user account with the name and ID number.
 - e. Then, click **Enable Authentication** in the **Actions** column.

- **No Authentication**

Select **No Authentication** as the **Current Authentication Method**. Then, the two-factor authentication is disabled. All the two-factor authentication methods become invalid.

1.1.16.7 Application whitelist

You can perform operations on the application whitelist.

Context

All the Apsara Stack Operations (ASO) services are accessed based on Operation Access Management (OAM) permission management. Therefore, if your account does not have the corresponding role, your access requests are rejected. The application whitelist function allows you to access ASO in scenarios where no permissions are assigned. With the whitelist function enabled, the application can be accessed by all users who have successfully logged on. The application whitelist permissions include read-only and read/write. The configured value is the logon user permission.

The application whitelist is managed by a super administrator or system administrator. You can access this page after logging on as a super administrator.

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

Procedure

1. [Log on to ASO](#).
2. In the left-side navigation pane, choose **System Management > Application Whitelist**.
3. On the **Application Whitelist** page, you can perform the following operations:

- **Add a whitelist**

In the upper-right corner, click **Add to Whitelist**. In the displayed **Add to Whitelist** dialog box, complete the configurations and then click **OK**.

- **Modify the permission**

Select the product permission from the **Permission** drop-down list.

- **Delete a whitelist**

Find the whitelist to be deleted, and then click **Delete** in the **Actions** column. Click **OK** in the displayed dialog box.

1.1.16.8 Server password management

Server Password allows you to configure and manage server passwords and search for history passwords.

Context

Server password management allows you to manage passwords of all the servers in the Apsara Stack environment.

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **System Management > Server Password**.
3. You can perform the following operations:
 - **Password Management**
 - a. Click the **Password Management** tab. This tab displays passwords of all the servers in the Apsara Stack environment.
 - b. After clicking **Show** in the **Password** column, the system displays the host password in plain text, which turns into cipher text after 10 seconds. Alternatively, directly click **Hide** to display cipher text.
 - c. Click **Update Password** in the **Actions** column. In the displayed **Update Password** dialog box, enter the **Password** and **Confirm Password**, and then click **OK** to update the server password.
 - d. Select one or more servers that you want to update the password, and then click **Batch Update**. Enter the **Password** and **Confirm Password**, and then click **OK** to update the passwords of the selected servers.
 - e. Click **Configuration**. In the displayed **Configuration Item** dialog box, enter the **Password Expiration Period** and select the **Unit**, and then click **OK**. Server passwords are updated immediately and will be updated again after an expiration period.

- **History Password**

The **History Password** tab shows the history of password updates for each server. You can search for the history passwords of servers by product, hostname, and/or IP address.

- **Configuration**

The **Configuration** tab displays the metadata, namely the initial password, password length, and retry times, of server password management.

- The initial password is the one when server password management is deployed in the Apsara Stack environment. This parameter is important, which is used to update the password of a server in the Apsara Stack environment.
- The password length is the length of passwords updated automatically in the system.
- Retry times is the number of retries when the password fails to be updated.

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

1.1.16.9 Operation logs

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

Context

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

Procedure

1. [Log on to Apsara Stack Operations \(ASO\)](#).
2. In the left-side navigation pane, choose **System Management > Operation Logs**.
3. On the **Log Management** page, you can perform the following operations:

- **Search for logs**

In the upper-left corner, configure the **User Name** and **Time Period**, and then click **Search** to view the log information in the list.

- **Delete logs**

Select one or more logs to be deleted. Click **Delete Logs** and then click **OK** in the displayed dialog box.

- **Export logs**

Select one or more logs to be exported, and then click **Export**.

1.2 Apsara Stack Doctor (ASD)

1.2.1 Apsara Stack Doctor introduction

Apsara Stack Doctor (ASD) checks the health of services for Apsara Stack Management Console and troubleshoots faulty services. Data in Apsara Stack Doctor comes from Apsara Infrastructure Management Framework SDK. The data includes the raw data of deployed Apsara Stack products, network topology metadata, and monitoring data.

Basic features

- Provides data filtering, analysis, and processing for O&M data consumers.
- Provides encapsulation, orchestration, and rights management of O&M operations.
- Provides O&M experience accumulation and archiving capabilities.
- Provides troubleshooting, pre-diagnosis, health check, and early warning capabilities.
- Records O&M experience, prescriptions, monitoring data, and log data to support intelligent O&M.

Benefits

- Provides unified management of Apsara Stack O&M data.
- Complements on-site O&M tools.
- Provides a unified tool for automated inspection of Apsara Stack.
- Allows you to perform O&M through Web interfaces, eliminating highly risky black screen operations.
- Allows you to have a periodic offline backup of Apsara Stack metadata, providing out-of-band support for metadata recovery.

Terms

Apsara Stack has five levels of release granularity.

- **system**

The greatest granularity at which Apsara Stack is available to external users. It is a collection of one or more Apsara Stack products.

- **product**

A category of product visible to users in Apsara Stack. It provides users with a kind of relatively independent features. For example, both ECS and SLB are products. Each product provides one or more features. Each product feature may be provided by one or more types of clusters.

- **service**

A type of software that provides independent features. It represents a product module or component. Each service can be managed separately or combined with other services into a product. If a service provides a complete set of features, it can also serve as a separate product alone.

- **server role (sr)**

A service component. A service can contain multiple server roles, each of which serves as a submodule of the service and provides a separate feature. Server role is also the smallest granularity monitored during Apsara Infrastructure Management Framework deployment and O&M. Some examples of server roles include PanguMaster and PanguChunkserver. Server roles are mapped to servers. Applications can be deployed to servers by their server role. A server role can contain multiple applications. Multiple applications belonging to a server role are packaged together for deployment. Different applications in a single server role can only be deployed to the same server. Multiple server roles are combined into a server role group (srg) for software deployment purposes. Only one server role group can be deployed to a server.

- **application (app)**

An independent process. Applications are one component of a server role, the other two being docker and file. All applications are built from source code.

- docker: a Docker image that is built from source code.
- file: a file that is placed on a server.
- application: a piece of software that is built from source code files and can be started directly from a start executable.

1.2.2 Log on to Apsara Stack Doctor

This topic describes how to log on to Apsara Stack Doctor.

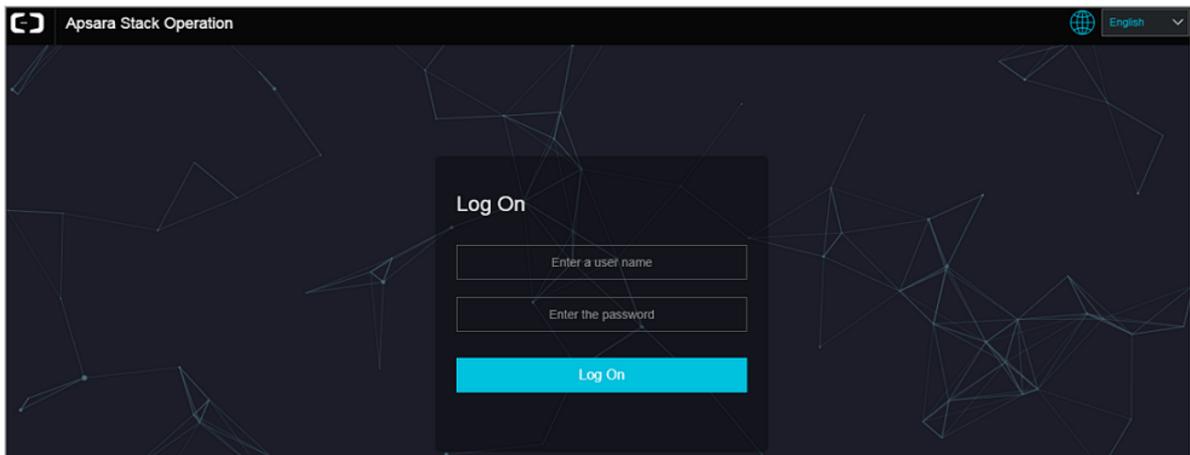
Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 1-2: Log on to ASO



Note:

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.
5. In the left-side navigation pane, click **Products**. On the Product List page that appears, click **Apsara Stack Doctor**.

1.2.3 Product dependency

You can use the product dependency function in Apsara Stack Doctor to view the dependencies between products, services, and server roles.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **Dependencies > Product Dependency** to view the dependencies between different products.

Click the name of a product (such as RDS) in the dependency graph. Information about products that have dependency relationships with the selected product is displayed. The **Basic Product Information** for the selected product is displayed on the right side.

3. In the left-side navigation pane, choose **Dependencies > Service Dependency** to view the dependency relationships between a service component and its surrounding components.
4. In the left-side navigation pane, choose **Dependencies > Server Role Dependency** to view dependencies between server roles.

After selecting a cloud service, you can view the server roles contained in the cloud service. All server role names end with a number sign (#). You can click the name of a server role to view its basic information.

1.2.4 Apsara Stack Inspection System

1.2.4.1 Apsara Stack Inspection System introduction

Apsara Stack Inspection System is an automated inspection tool available in Apsara Stack V3. It allows you to perform inspections with a single click to locate issues quickly and improve O&M efficiency. The following inspection functions are available:

- Basic Inspection: inspects OPS servers, hardware devices, and MiniRDS clusters and instances.
- Apsara System Inspection: inspects Apsara Distributed File System and Apsara Name Service and Distributed Lock Synchronization System.
- Inspection in Other Systems: inspects Apsara Stack Assistant (ASA).
- Inventory Inspection: provides the statistics of available resources of Apsara Stack products.
- Cloud Product Inspection: inspects server roles and ECS and RDS instances.
- Middleware Inspection: inspects Butler-related services.

After the inspections have been completed, you can generate an inspection report summarizing the inspection results.

**Note:**

You must submit tickets for any alarms reported or issues found during the inspection process in a timely manner. You are not allowed to handle alarms or issues without customer approval and solution verification.

1.2.4.2 Access Apsara Stack Inspection System

This topic describes how to access Apsara Stack Inspection System.

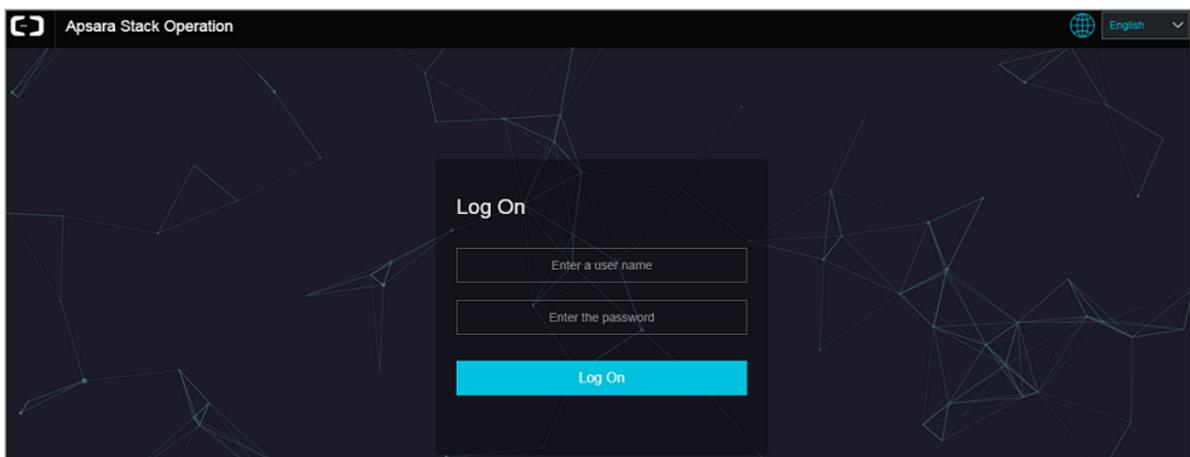
Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 1-3: Log on to ASO

**Note:**

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.

- The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.
 5. In the left-side navigation pane, click **Products**. On the Product List page that appears, click **Apsara Stack Doctor**.
 6. In the left-side navigation pane, click **Apsara Stack Inspection System**.

1.2.4.3 Apsara Stack Inspection System overview

On the Overview page of Apsara Stack Inspection System, you can view the current inspection status, resource usage, and issue resolution status in Apsara Stack Inspection System.

The Overview page consists of the following parts:

Quantity statistics

Statistics on the physical devices in Apsara Stack:

- Physical Machines: the total number of physical machines in Apsara Stack.
- Products: the number of cloud products in Apsara Stack.
- Hosts: the number of Docker hosts in Apsara Stack.
- Networking Devices: the number of networking devices in Apsara Stack.
- Online Containers: the number of Docker containers currently online.
- Total Containers: the total number of Docker containers in Apsara Stack.

Physical machine statistics

The number of physical machines used for Apsara Stack products.

Issues detected

- **Issues Detected per Inspection:** the total number of issues detected per inspection over a recent period of time.
- **Issues Detected Last Inspection:** the number of issues detected during the last inspection and their distribution.

Issues resolved

- **Issues Resolved per Day:** the total number of issues resolved per day over a recent period of time.
- **Issues Resolved Today:** the number of different types of issues resolved today.

Usage

You can view the CPU, memory, disk, Internet, and intranet metrics for ECS, RDS, and SLB instances only.

Apsara Distributed File System

Statistics on Apsara Distributed File System utilization of Apsara Stack products.

1.2.4.4 Platform Inspection

1.2.4.4.1 Platform Inspection introduction

The Platform Inspection module provides several major functions such as Basic Inspection, Apsara System Inspection, Inspection in Other Systems, Inventory Inspection, Cloud Product Inspection, Middleware Inspection, and .

You can select different functions to perform corresponding inspections. You can also click **Start All Inspections** in the upper-right corner of the Platform Inspection page to complete all inspections.

1.2.4.4.2 Basic Inspection

Basic Inspection is designed to check the existence of hardware faults in clusters and the health statuses of some basic services.

Context

Basic Inspection consists of four parts:

- **OPS inspection:** checks the operating statuses of primary and secondary servers, including the health statuses of memory modules, Docker containers, CPUs, and disks.

- Hardware inspection: checks the health statuses of hardware devices such as disks, CPUs, and fans.
- DNS inspection: checks the health statuses of DNS servers in Apsara Stack.
 - status: This metric is used to check whether a DNS server is available.
 - master_slave: This metric is used to check whether the DNS data of the primary and secondary DNS servers is the same.
- MiniRDS inspection: checks the health statuses of MiniRDS clusters and instances.
 - minirids: This metric is used to check information about the disks in MiniRDS clusters and instances.
 - master_slave: This metric is used to check the statuses of primary and secondary MiniRDS clusters and instances.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. In the upper part of the Platform Inspection page, click **Basic Inspection**. Then, click **Basic Inspection** in the center of the page. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message "Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.
4. When you click an item with issues, a message appears, indicating the specific information. Locate and hover over the **detail** column to display the specific alarm information.

1.2.4.4.3 Apsara System Inspection

Apsara System Inspection is designed to check the statuses of Apsara Distributed File System and Apsara Name Service and Distributed Lock Synchronization System.

Context

Apsara System Inspection consists of two parts:

- Apsara Distributed File System inspection: checks several Apsara Distributed File System metrics of Apsara Stack products, such as the read/write performance, data backup status, and version status.

- Apsara Name Service and Distributed Lock Synchronization System inspection: checks several Apsara Name Service and Distributed Lock Synchronization System metrics of Apsara Stack products, such as the disk usage and queuing status.

**Note:**

- Apsara Distributed File System and Apsara Name Service and Distributed Lock Synchronization System are important parts of the Apsara system. You must pay special attention to any issues found during the inspection. You must strictly abide by O&M rules and regulations when performing online operations on Apsara Distributed File System.
- Apsara Stack Inspection System cannot display Apsara system inspection information in detail. If an alarm is reported, you need to access the admin gateway of the corresponding product based on the IP address, and use Apsara-related commands to view the alarm.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. In the upper part of the Platform Inspection page, click **Apsara System Inspection**. Then, click **Apsara System Inspection** in the center of the page. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message "Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.
4. When you click an item with issues, a message appears, indicating the specific information. Locate and hover over the **detail** column to display the specific alarm information.

1.2.4.4.4 Inspection in Other Systems

Inspection in Other Systems is used to inspect ASA.

Context

Inspection in Other Systems provides the following metrics:

- ntp: This metric is used to check whether the system time of all machines (including Docker VMs and NCs) is synchronized with the NTP time. If no, the time offset (measured in ms) is reported.
- ip_conflict: This metric is used to check for IP address conflicts in the current environment.
- rpm: This metric is used to check whether the RPM service is available on all machines, including Docker VMs and NCs.

- `dns_bind`: This metric is used to check whether the IP address bound to a domain name is the same as the requested IP address.
- `mem_quota`: This metric is used to check whether the memory capacity used by Docker hosts exceeds the threshold.
If memory usage exceeds the threshold, contact O&M engineers to confirm whether memory capacity needs to be scaled up.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. In the upper part of the Platform Inspection page, click **Inspection in Other Systems**.
Then, click **Inspection in Other Systems**. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message "Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.
4. When you click an item with issues, a message appears, indicating the specific information. Locate and hover over the **detail** column to display the specific alarm information.

1.2.4.4.5 Inventory Inspection

Inventory Inspection inspects the inventory status of ECS, OSS, RDS, and SLB, Apsara Distributed File System utilization of Apsara Stack products, and remaining resources of each Apsara Stack product.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. (Optional) Click **Set Thresholds** in the upper-right corner to set the CPU and memory thresholds of several products.
If the inventory of a product exceeds the threshold, an alarm is reported. The thresholds are measured in percentage. Both the CPU and memory thresholds are 80% by default. You can set the thresholds as required or use the default values.
4. In the upper part of the Platform Inspection page, click **Inventory Inspection**. Then, click **Inventory Inspection**. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message

"Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.

5. When you click an item with issues, a message appears, indicating the specific information. Locate and hover over the **detail** column to display the specific alarm information.

1.2.4.4.6 Cloud Product Inspection

Cloud Product Inspection inspects the computing, memory, and disk resource utilizations of Apsara Stack products, and uses these metrics to determine whether cluster resources are insufficient.

Context

Cloud Product Inspection consists of three parts:

- Component inspection: checks for any server roles which are not running properly. For example, if a server role is in upgrading state, it does not reach the final status.
- ECS inspection: checks for any ECS instances which are not running properly.
- RDS inspection:
 - exception: This metric is used to check for the alarms reported for ApsaraDB Operations and Maintenance System.
 - task: This metric is used to check for interrupted tasks in ApsaraDB Operations and Maintenance System.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. In the upper part of the Platform Inspection page, click **Cloud Product Inspection**. Then, click **Cloud Product Inspection**. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message "Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.
4. When you click an item with issues, a message appears, indicating the specific information. Locate and hover over the **detail** column to display the specific alarm information.

1.2.4.4.7 Middleware Inspection

Middleware Inspection inspects Butler-related services, such as middleware-dncc, edas-customer, middleWare, drds, dauthProduct, mq, edas, butler, and tlog.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. In the upper part of the Platform Inspection page, click **Middleware Inspection**. Then, click **Middleware Inspection** in the center of the page. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message "Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.
4. When you click an item with issues, a message appears, indicating the specific information. Locate and hover over the **message** column to display the specific alarm information.

1.2.4.4.8 Big Data Inspection

Big Data Inspection inspects big data products such as ads, dataworks, dataphin, iplus, asap, quickbi, odps, pai, datahub, biggraph, streamcompute, bcc, and es.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.
3. In the upper part of the Platform Inspection page, click **Big Data Inspection**. Then, click **Big Data Inspection** in the center of the page. On the inspection page, two messages appear successively: Wait a moment and Inspecting. If issues have been detected during the inspection, the message "Issues Detected" is displayed after the inspection is completed. Items with and without issues detected are displayed on the page.
4. When you click an item with issues, a message appears, indicating the specific information.

1.2.4.4.9 Inspection reports

Inspection reports summarize inspection results by inspection type. You can view all detected issues on the Start All Inspections Reports page.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Platform Inspection**.

3. In the upper part of the Platform Inspection page, click **Start All Inspections Reports**.

On the Start All Inspections Reports page, you can click different types of inspections as required, and view detected issues and causes. In RDS Inspection Report, you can click **View Report** to view RDS inspection details. Click **Download JSON File** to download the JSON file to your local PC.

1.2.4.5 Inspection history

On the Inspection History page, you can specify a time frame and view the number of inspections per day, number of issues detected per inspection, and issues detected.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Inspection History**.

You can view the following information:

- **Inspections Per Day:** You can view the number of inspections per day over a specified time frame.
- **Issues Detected per Inspection:** You can view the number of issues detected per inspection.
- **Logons:** You can view user logon statistics.
- **Detected Issues:** You can search for and view the issues detected per inspection over a time frame.

1.2.4.6 Work Reports

On the Work Reports page, the project information, current utilization information, and work reports are displayed.

Context

- **Project Information:** You can enter project information to search for relevant work reports. Project information includes the project name, version, number of physical devices, number of services, and TAM.
- **Current Utilization:** You can view the resource and Apsara Distributed File System utilization data. When the utilization level exceeds 75% or the absolute value of resource growth rate is greater than 5%, the data is displayed in red.
- **Reports:** You can switch between tabs to view issues submitted today, unresolved issues, and historical data. Historical data includes daily, weekly, and monthly reports.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **Work Reports**.
3. On the page that appears, you can perform the following operations:
 - Click **Edit**, set parameters, and click **Submit** to edit the project information.
 - View current resource and Apsara Distributed File System utilization information.
 - View and edit work reports. Click **Add**, enter information about an issue, and click **Save** to submit the issue.
 - Click **Download EXCEL File** in the upper-right corner of the page to download the work report data as an Excel file.



Note:

You must click **Start All Inspections** to use this function.

1.2.4.7 Products

The Products module allows you to view the current resource usage of Apsara Stack products.

Context

Information can be displayed using either of the following methods:

- The resource information of each Apsara Stack product is displayed by project.
- The resource information is displayed by product.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, choose **Products > Product Overview**.

On the Product Overview page, the resource information of all Apsara Stack products is displayed by project. You can click **Export Excel File** in the upper-right corner of the page to export the information as a spreadsheet.

3. In the left-side navigation pane, choose **Products > Database Services > RDS**.

Relevant information of ApsaraDB Operations and Maintenance System is displayed.

1.2.4.8 End-to-end ECS links

The E2E ECS Links module searches logs for issues in ECS instances.

Context

With the E2E ECS Links module, you can check for errors reported in ECS instance and request logs, and view error details. You can locate ECS issues based on the error details.

Procedure

1. [Access Apsara Stack Inspection System](#).
2. In the left-side navigation pane, click **E2E ECS Links**.
3. In the search bar, select the product name and service name. Then, enter the service instance ID or request ID, and click **Run**.

A row of data is added to the task list. In the running status column, different colors represent different statuses:

- Green: indicates that the task has been completed.
 - Red: indicates that the task fails to be completed.
 - Blue: indicates that the task is in progress.
4. When the running status becomes Inspected, click **Download**.

The log file in JSON format is downloaded to your local PC.



Note:

The log file can only be downloaded after the running status has changed to Inspected.

5. Click **Learn More**.

On the **Link Logs** page, you can click each log item to view log details. Logs with errors are marked with red exclamation marks (!) to make troubleshooting easier.

1.2.5 ASA

1.2.5.1 Apsara Stack Assistant (ASA) introduction

ASA is a tool provided to help you test, perform O&M on, and release Apsara Stack products while ensuring the stability of version qualities. ASA has also retained the inspection, scanning, and version tracking capabilities of Apsara Stack V2.

1.2.5.2 RPM check

The RPM Check module allows you to check whether the RPM service is available on all machines, including Docker VMs and NCs.

Procedure

1. [Log on to Apsara Stack Doctor](#).

2. In the left-side navigation pane, choose **ASA > RPM Check**.

Table 1-6: Description of parameters on the RPM check page

Parameter	Description
Host	A hostname.
Status	The status of a machine. Valid values: <ul style="list-style-type: none"> • normal: indicates that the machine is operating normally. • unavailable: indicates that the machine is not operating normally and unavailable.

1.2.5.3 Virtual IP check

Context

The Virtual IP Check module allows you to check whether a virtual IP addresses is bound to a backend IP address properly.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > Virtual IP Check**.

Table 1-7: Description of parameters on the virtual IP check page

Parameter	Description
Virtual IP Address	A virtual IP address.
Virtual Port	The port corresponding to the virtual IP address.
Port	The port corresponding to the backend IP address.
Backend IP Address	The IP address of the backend server.
Cluster	The cluster to which the backend IP address is assigned.
Service	The service to which the backend IP address is assigned.
Server Role	The server role to which the backend IP address is assigned.

Parameter	Description
Status	<p>The health status indicating whether the virtual IP address is bound to the backend IP address properly. Valid values:</p> <ul style="list-style-type: none"> normal: indicates that the virtual IP address is bound to the backend IP address properly. abnormal: indicates that the virtual IP address is not bound to the backend IP address properly.

1.2.5.4 Volume check

Context

The Volume Check module allows you to view detailed information about the disk capacity of Docker hosts.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > Volume Check**.

Table 1-8: Description of parameters on the volume check page

Parameter	Description
Container ID	The unique ID of a Docker container.
Container Name	The name of the Docker container.
Host IP Address	The IP address of the Docker host. Typically , Docker hosts include physical hosts and Docker VMs.
Path	The disk partition mount point of a Docker volume.
Disk Quota	The quota of a disk.
Total Partition Space	The total available space of a mount point calculated by using the df command.
Partition Space Used	The space used by a mount point directory.
Directory Space Used	The total used space of a mount point calculated by using the du command.

1.2.5.5 NTP check

Context

The NTP Check module allows you to check whether the system time of all machines (including Docker VMs and NCs) is synchronized with the NTP time. If no, the time offset (measured in ms) is reported.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > NTP Check**.

Table 1-9: Description of parameters on the NTP check page

Parameter	Description
Host	A hostname.
Time Offset	The time offset, in ms.

1.2.5.6 IP conflict check

Context

The IP Conflict Check module allows you to check for IP address conflicts in the current environment.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > IP Conflict Check**.

Table 1-10: Description of parameters on the IP conflict check page

Parameter	Description
IP	A conflicting IP address.
Physical Host	The name of the physical host with the conflicting IP address.
Server Role	The server role that requests the resource.
Type	The IP address type. Valid values: docker, vm, and physical.
Virtual Host	The hostname of the Docker VM.

1.2.5.7 DNS check

Context

The DNS Check module allows you to check whether the IP address bound to a domain name is the same as the requested IP address.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > DNS Check**.

Table 1-11: Description of parameters on the DNS check page

Parameter	Description
Domain	The domain name requested by Apsara Infrastructure Management Framework.
Virtual IP Address	The IP address that is bound to the domain name requested by Apsara Infrastructure Management Framework.
Owner	The application that requests the DNS resource.
IP	The physical IP address bound to the domain name.

1.2.5.8 IP details

Context

The IP Details module allows you to check detailed information for all IP addresses in the current environment (including the IP addresses of physical hosts, Docker containers, and VMs).

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > IP Details**.

Table 1-12: Description of parameters on the IP details page

Parameter	Description
IP	The IP address of a resource.
Virtual Host	The name of the VM.

Parameter	Description
Type	The resource type. Valid values: <ul style="list-style-type: none"> • physical • docker • vm
Physical Host	The name of the physical host.
Server Role	The server role that requests the resource.

3. Hover over **Server Role Information** in the **Server Role** column. The server role information is displayed.

1.2.5.9 Quota check

The Quota Check module allows you to check the memory, CPU, and disk quotas of containers.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > Quota Check**.
3. On the quota check page, you can view memory, CPU, and disk quota information.
 - Memory quota check

Click the **Memory** tab to view the memory allocation of specified machines.
 - CPU quota check

Click the **CPU** tab to view the CPU allocation of specified machines.
 - Disk quota check

Click the **Disk** tab to view the disk allocation of specified machines.

1.2.5.10 Error diagnostics

Context

The error diagnostics page consists of the following tabs:

- Error with Source: displays resource issues.
- Error with Self: displays ASA issues.
- Error with Dependency: displays dependency issues.
- Normal: displays resources with no issues.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > Error Diagnostics**.
3. Switch between tabs to view the corresponding information.

1.2.5.11 Versions

The Versions module allows you to obtain version information and upgrade information of all products in the current environment.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **ASA > Versions**.
3. You can perform the following operations:
 - Click the **Product Versions** tab. On the Product Versions tab page, view information related to product versions, such as the IDC, product, and version.
 - Click the **Server Role Versions** tab. On the Server Role Versions tab page, view information related to server role versions, such as the IDC, product, version, server role, and type.
 - Click the **Version Tree** tab. On the Version Tree tab page, view information related to version trees.

1.2.6 Support tools

1.2.6.1 OS tool

Context

The OS tool enables system-level diagnostics of physical machines in Apsara Stack. This tool allows you to query the physical machine list, deploy and run OS diagnostics tools on physical machines, and collect, display, and download diagnostic results.

Metrics that can be diagnosed by the OS tool include: disk file metadata usage, memory usage, process statuses, time synchronization, kernel faults, risky operations, system load, fstab files, read-only file systems, kdump services, kdump configurations, conman configurations, domain name resolution, disk I/O load, file deletion exceptions, system errors, RPM databases, fgc, tair, route_curing, default routes, abnormal network packets, TCP connection status exceptions, TCP queuing exceptions, network packet loss, bonding exceptions, NIC exceptions, SN retrieval exceptions, OOB IP address retrieval exceptions, sensor exceptions, sensor record exceptions, SEL record exceptions, Docker status exceptions, and RAID exceptions.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **Support Tools > OS Tool**.
3. Click **Get Physical Machine List**. All physical machines in the system are displayed in the list below.
4. Optional: Enter the name of a physical machine in the search box, and click **Search**. The physical machine is displayed in the list.
5. Select the physical machine to be diagnosed, and click **Run Script** in the upper-right corner of the page.

When **Script Execution Status** changes from **Not Executed** to **Decompressed**, you can view the health score of the physical machine in the **Health Score** column.

6. After the diagnostics are completed, you can click **Report** in the Actions column to view diagnostic results. For more information, click **URL Link** or **Download**.

1.2.6.2 Apsara Distributed File System Diagnostics

Context

Apsara Distributed File System Diagnostics is used to collect and analyze the health statuses of Apsara Distributed File System services, and information about the services and environments on which Apsara Distributed File System depends. This tool allows you to perform diagnostics and generate diagnostic reports when Apsara Distributed File System is abnormal.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, choose **Support Tools > Apsara Distributed File System Diagnostics**.
3. Click **Get Admin Gateway List**. The admin gateways obtained are displayed in the list below.
4. Optional: Select a product from the drop-down list, and click **Search**. The admin gateway for the product is displayed in the list.
5. Select the admin gateway to be diagnosed, and click **Diagnose** in the upper-right corner of the page.

When **Script Execution Status** changes from **Waiting for Diagnostic** to **Diagnosed**, you can view the health status of the product in the **Health Status** column.

6. After the diagnostics are completed, you can click **Report** in the Actions column to view diagnostic results.

1.2.7 Update Monitoring Dashboard

Update Monitoring Dashboard allows you to view information about alarms and monitoring items for different products.

Procedure

1. [Log on to Apsara Stack Doctor](#).
2. In the left-side navigation pane, select **Update Monitoring Dashboard**.

On the Monitoring Dashboard page, alarm information is displayed by product. You can perform the following operations:

- **Search**

Select a product from the Product drop-down list, and click **Search**. The alarms reported for the selected product are displayed in the alarm list.

- **Refresh**

Within a few minutes after Apsara Stack Doctor restarts, data is acquired based on the enabled scheduled tasks in the backend. If a request failure warning message appears in the upper-right corner, click **Refresh**. Wait a few minutes for the backend to reacquire the product list and alarm information.

The backend refreshes and caches all alarm information every three minutes by default.

The frontend acquires the cached information once every minute. To refresh alarm information manually, click **Refresh**.

3. Click a product name in the product list. The product alarm summary page appears.

Select a service or cluster from the **Service** or **Cluster** drop-down list, and click **Search**. Click **Refresh** in the upper-right corner of the page to refresh the product alarm information.

4. Click the blue instance name in the **Monitoring Instance** column. The instance alarm metric page appears.

1.3 Operation Access Manager (OAM)

1.3.1 OAM introduction

Overview

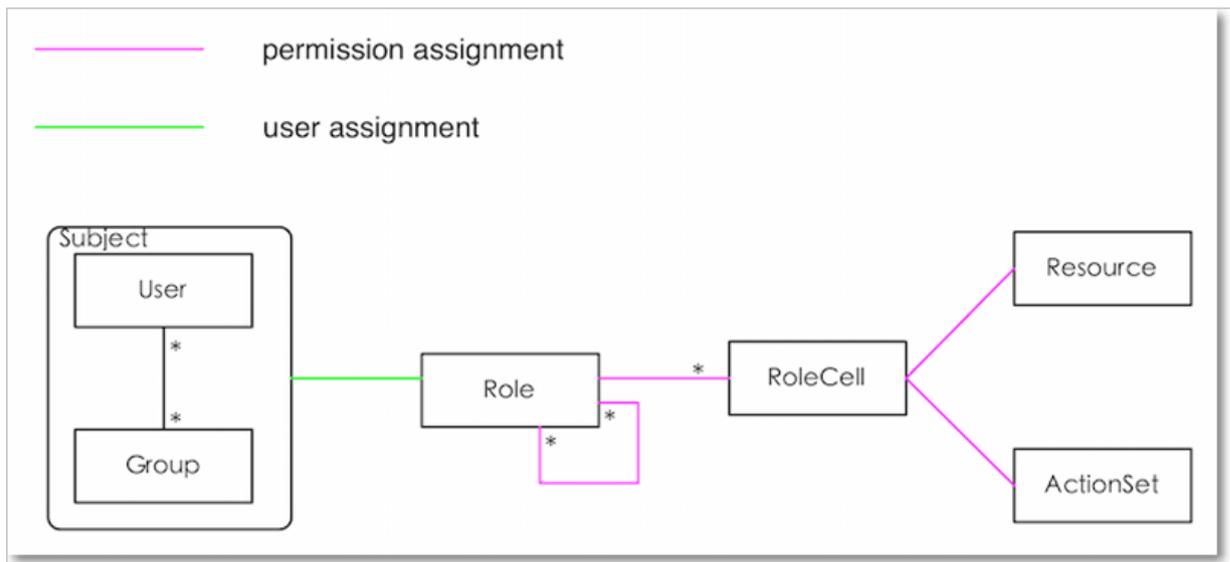
Operation Access Manager (OAM) is a centralized permission management platform of Apsara Stack Operations (ASO). OAM uses a simplified role-based access control (RBAC) model. Administrators can use OAM to assign roles to operations personnel, granting them corresponding operation permissions to operations systems.

OAM permission model

In RBAC, administrators do not directly grant system operation permissions to users. Instead, they create a collection of roles between a collection of users and a collection of permissions. Each role corresponds to a group of permissions. If a role is assigned to a user, the user is granted all the operation permissions of that role. Therefore, when creating a user, administrators are only required to assign a role to the user, saving the trouble to grant specific permissions to the user. In addition, the frequency of role permission changes is less than that of user permission changes, simplifying the user permission management and reducing the system overhead.

See the [OAM permission model](#) as follows.

Figure 1-4: Permission model



1.3.2 Instructions

Before using Operation Access Manager (OAM), you must know the following basic concepts about permission management.

subject

Operators of the access control system. OAM subjects include users and groups.

user

Administrators and operators of operations systems.

group

A collection of users.

role

The core of the role-based access control (RBAC) system.

Generally, a role can be regarded as a collection of permissions. A role can contain multiple RoleCells and/or roles.

RoleHierarchy

In the OAM system, a role can contain other roles to form RoleHierarchy.

RoleCell

The specific description of a permission. A RoleCell consists of resources, ActionSets, and WithGrantOptions.

resource

The description of an authorized object. For more information about resources of operations platforms, see [Operation permissions of operations platforms](#).

ActionSet

The description of authorized actions. An ActionSet can contain multiple actions. For more information about actions of operations platforms, see [Operation permissions of operations platforms](#).

WithGrantOption

The maximum number of authorizations in the cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.

For example, if administrator A sets `WithGrantOption` to 5 when granting a permission to administrator B, the permission can be granted for another five times at most. When administrator B grants the permission to administrator C, the value of `WithGrantOption` cannot be greater than 4. If `WithGrantOption` is set to 0 when administrator B grants the permission to operator D, operator D can only use the permission but cannot grant it to others.



Note:

Currently, OAM does not support the cascaded revocation for cascaded authorization. Therefore, administrator C and operator D still have the permission even if the permission is revoked for administrator B.

1.3.3 Quick start

1.3.3.1 Log on to OAM

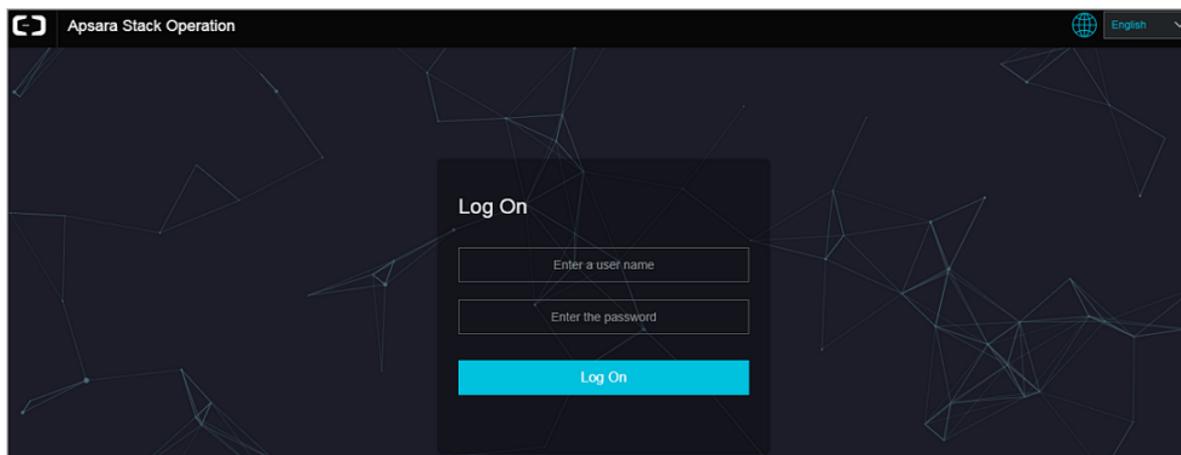
This topic describes how to log on to Operation Access Manager (OAM).

Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 1-5: Log on to ASO**Note:**

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.
5. In the left-side navigation pane, select **Products** and then select **Operation Access Management**.

1.3.3.2 Create a group

Create a user group for centralized management.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. In the upper-right corner, click **Create Group**. In the displayed dialog box, enter the **Group Name** and **Memo**.
4. Then, click **Confirm**.

You can view the created group on the **Owned Groups** page.

1.3.3.3 Add group members

Add members to an existing group to grant permissions to the group members in a centralized way.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group and then click **Manage** in the **Operation** column.
4. Click **Add Member** in the **Group Members** section.
5. Select the search mode, enter the corresponding information, and then click **Detail**. The user details are displayed.

Three search modes are available:

- **RamAliasName**: Search for the user in the format of *RAM username@primary account ID*. Use this mode for users who have activated Resource Access Management (RAM).
 - **AliyunPk**: Search for the user by using the unique ID of the user's cloud account.
 - **AliyunId**: Search for the user by using the logon name of the user's cloud account.
6. Click **Add**.
 7. You can repeat the preceding steps to add more group members.

To remove a member from the group, click **Remove** at the right of the member in the **Group Members** section.

1.3.3.4 Add group roles

You can add roles to an existing group, that is, assign roles to the group.

Prerequisites

- The role to be added is created. For more information about how to create a role, see [Create a role](#).

- You are the owner of the group and the role.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group and then click **Manage** in the **Operation** column.
4. Click **Add Role** in the **Role List** section.
5. Search for roles by **Role Name**. Select one or more roles and then configure the expiration time.
6. Then, click **Confirm**.

To delete a role, click **Remove** at the right of the role in the **Role List** section.

1.3.3.5 Create a role

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. In the upper-right corner of the **Owned Roles** page, click **Create Role**.
4. In the displayed dialog box, enter the **Role Name** and **Memo**, and then select the **Role Type**.
5. Optional: Configure the role tags, which can be used to filter roles.
 - a) Click **Edit Tag**.
 - b) In the displayed **Edit Tags** dialog box, click **Create**.
 - c) Enter the **Key** and the corresponding **Value** of the tag and then click **Confirm**.
 - d) Repeat the preceding step to create more tags.

The created tags are displayed in the dotted box.

- e) Click **Confirm** to create the tags.
6. Click **Confirm** to create the role.

1.3.3.6 Add inherited roles to a role

Add inherited roles to a role to grant the permissions of the former to the latter.

Prerequisites

You are the owner of the current role and the inherited role to be added.

For more information about how to search for your roles, see [Search for roles](#).

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role and then click **Manage** in the **Operation** column.
4. Click the **Inherited Role** tab and then click **Add Role**.
5. Search for roles by **Role Name** and then select one or more roles.
6. Click **Confirm**.

1.3.3.7 Add resources to a role

You must add resources to a created role.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role and then click **Manage** in the **Operation** column.
4. Click the **Resource List** tab.
5. Click **Add Resource**.
6. Complete the configurations. For more information, see [Configurations](#).

Table 1-13: Configurations

Configuration item	Description
BID	The deployment region ID.
Product	<p>The cloud product to be added, for example, rds.</p> <div style="background-color: #f0f0f0; padding: 5px;">  Note: The cloud product name must be lowercase. For example, enter rds, instead of RDS. </div>
Resource Path	For more information about resources of cloud products and operations platforms, see Operation permissions of operations platforms .
Actions	An ActionSet, which can contain multiple actions.

Configuration item	Description
	For more information about actions of operations platforms, see Operation permissions of operations platforms .
Grant Option	The maximum number of authorizations in the cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.
Memo	The description of the resource.

7. Click **Add**.

1.3.3.8 Add authorized users to a role

You can assign an existing role to users or user groups.

Prerequisites

The corresponding users or user groups are created. Users are created in the Apsara Stack console. For more information about how to create user groups, see [Create a group](#).

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role and then click **Manage** in the **Operation** column.
4. Click the **Operator List** tab.
5. Click **Add Operator**.
6. Select the search mode and enter the corresponding information.

Four search modes are available:

- **RamAliasName**: search in the format of *RAM username@primary account ID*. Use this mode for users who have activated Resource Access Management (RAM).
- **AliyunPk**: search by using the unique ID of the user's cloud account.
- **AliyunId**: search by using the logon name of the user's cloud account.
- **Group Name**: search by group name.



Note:

You can search for a single user or user group. For more information about how to create a user group, see [Create a group](#).

7. Configure the expiration time.

After the expiration time is reached, the user does not have the permissions of the role. To authorize the user again, the role creator must click **Renew** at the right of the authorized user on the **Operator List** tab, and then configure the new expiration time.

8. Click **Add** to assign the role to the user.

To cancel the authorization, click **Remove** at the right of the authorized user on the **Operator List** tab.

1.3.4 Manage groups

1.3.4.1 Modify the group information

After creating a group, you can modify the group name and memo on the **Group Information** page.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group and then click **Manage** in the **Operation** column.
4. Click **Edit** in the upper-right corner.
5. In the displayed **Edit Group** dialog box, modify the **Group Name** and **Memo**.
6. Click **Confirm**.

1.3.4.2 View group role details

You can view the information about the inherited role, resource list, and inheritance tree of a group role.

Prerequisites

A role is added to the group.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group and then click **Manage** in the **Operation** column.
4. In the **Role List** section, click **Detail** at the right of the role.

5. On the **Role Detail** page, you can perform the following operations:

- Click the **Inherited Role** tab to view the information about the inherited roles.

To view the detailed information about an inherited role, click **Detail** at the right of the inherited role.

- Click the **Resource List** tab to view the resource information of the role.

To add other resources to this role, see [Add resources to a role](#).

- Click the **Inheritance Tree** tab to view the basic information and resource information of the role and its inherited roles by using the inheritance tree on the left.

1.3.4.3 Delete a group

You can delete a group that is no longer in use as required.

Prerequisites

The group to be deleted does not contain members.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > Owned Groups**.
3. Find the group to be deleted and then click **Delete** in the **Operation** column.

1.3.4.4 View assigned groups

You can view the groups to which you are assigned on the **My Groups** page.

Context

You can only view the groups to which you belong, but cannot view groups of other users.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Group Management > My Groups**.
3. On the **My Groups** page, view the name, owner, memo, and modified time of the group to which you belong.

1.3.5 Manage roles

1.3.5.1 Search for roles

You can view your owned roles on the **Owned Roles** page.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Optional: Enter the **Role Name**.
4. Click **Search** to search for roles that meet the search condition.

**Note:**

If the role you want to search for has a tag, you can click **Tag** and select the tag key to search for the role based on the tag.

1.3.5.2 Modify the role information

After creating a role, you can modify the role information.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role and then click **Manage** in the **Operation** column.
4. Click **Edit** in the upper-right corner.
5. In the displayed **Edit Role** dialog box, modify the **Role Name**, **Memo**, **Role Type**, and **Tag**.
6. Then, click **Confirm**.

1.3.5.3 View the role inheritance tree

You can view the role inheritance tree to know the basic information and resource information of a role and its inherited roles.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role and then click **Manage** in the **Operation** column.
4. Click the **Inheritance Tree** tab.

View the basic information and resource information of this role and its inherited roles by using the inheritance tree on the left.

1.3.5.4 Transfer roles

You can transfer roles to other groups or users according to business requirements.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Configure the search condition and search for the roles to be transferred.
4. Select one or more roles in the search results and click **Transfer**.
5. In the displayed **Transfer** dialog box, select the search mode, enter the corresponding information, and then click **Detail**. The user details or group details are displayed.

Four search modes are available:

- **RamAliasName**: search in the format of *RAM username@primary account ID*. Use this mode for users who have activated Resource Access Management (RAM).
- **AliyunPk**: search by using the unique ID of the user's cloud account.
- **AliyunId**: search by using the logon name of the user's cloud account.
- **Group Name**: search by group name.

6. Click **Transfer** to transfer the roles to the user or group.

1.3.5.5 Delete a role

You can delete a role that is no longer in use according to business requirements.

Prerequisites

The role to be deleted does not contain inherited roles, resources, or authorized users.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role to be deleted and then click **Delete** in the **Operation** column.

1.3.5.6 View assigned roles

You can view the roles assigned to you and permissions granted to the roles.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, choose **Role Management > My Roles**.
3. On the **My Roles** page, you can view the name, owner, memo, modified time, and expiration time of the role assigned to you.

Click **Detail** at the right of the role to view the inherited roles, resources, and inheritance tree information of this role.

1.3.5.7 View all roles

You can view all the roles in Operation Access Manager (OAM) on the **All Roles** page.

Procedure

1. [Log on to OAM](#).
2. In the left-side navigation pane, choose **Role Management > All Roles**.
3. On the **All Roles** page, view all the roles in the system.

You can search for roles by **Role Name** on this page.

4. At the right of the role, click **Detail** to view the inherited roles, resources, and inheritance tree information of this role.

1.3.6 Search for resources

You can search for resources to view the roles to which the resources are assigned.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, select **Search Resource**.
3. Enter the **Resource** and **Action** in the search boxes, and then click **Search** to search for roles that meet the conditions.
4. At the right of the search result, click **Detail** to view the inherited roles, resources, and inheritance tree information of the role.

1.3.7 View the personal information

You can view your personal information and perform permission tests on the **Personal Information** page.

Procedure

1. [Log on to Operation Access Manager \(OAM\)](#).
2. In the left-side navigation pane, select **Personal Information**.
3. In the **Basic Information** section, you can view your user name, user type, created time, AccessKey ID, and AccessKey Secret.



Note:

Click **Display** or **Hide** to display or hide the AccessKey Secret.

4. In the **Test Access** section, test if you have a certain permission.

- a) Enter the resource information in the **Resource** field.
- b) Enter the permissions in the **Action** field, such as create, read, and write. Separate multiple permissions with commas (,).

1.3.8 Appendix

1.3.8.1 Default roles and their functions

1.3.8.1.1 OAM default role

Role name	Role description	Resource	Actions	GrantOption
Super administrator	Root permission administrator	*.*	*	10

1.3.8.1.2 Apsara Infrastructure Management Framework default roles

Role name	Role description	Resource	Actions	GrantOption
Tianji_Project read-only	Has the read-only permission to Apsara Infrastructure Management Framework projects, which allows you to view the configurations and statuses of all projects and clusters	*:tianji:projects	["read"]	0
Tianji_Project administrator	Has all the permissions to Apsara Infrastructure Management Framework projects, which allows you to view and modify the configurations and statuses of all projects and clusters	*:tianji:projects	["*"]	0

Role name	Role description	Resource	Actions	GrantOption
Tianji_Service read-only	Has the read-only permission to Apsara Infrastructure Management Framework services, which allows you to view the configurations and templates of all services	*:tianji:services	["read"]	0
Tianji_Service administrator	Has all the permissions to Apsara Infrastructure Management Framework services, which allows you to view and modify the configurations and templates of all services	*:tianji:services	["*"]	0
Tianji_IDC administrator	Has all the permissions to Apsara Infrastructure Management Framework data centers, which allows you to view and modify the data center information	*:tianji:idcs	["*"]	0
Tianji_administrator	Has all the permissions to Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastruc	*:tianji	["*"]	0

Role name	Role description	Resource	Actions	GrantOption
	ture Management Framework configurations			

1.3.8.1.3 Webapp-rule default roles

Role name	Role description	Resource	Actions	GrantOption
Webapp-rule operations administrator	Has all the permissions to Webapp-rule projects, which allows you to view, modify, add, and delete all the configurations and statuses	26842:webapp-rule:*	["read", "write"]	0
Webapp-rule read-only	Has the read-only permission to Webapp-rule projects, which allows you to view all the configurations and statuses	26842:webapp-rule:*	["read"]	0

1.3.8.1.4 Workflow (grandcanal) console default roles

Role name	Role description	Resource	Actions	GrantOption
grandcanal.ADMIN	The workflow console administrator, who can query the workflow and activity details, and retry, roll back, stop, and restart a workflow	26842:grandcanal	["write", "read"]	0
grandcanal.Reader	Has the read-only permission to the workflow console and can only	26842:grandcanal	["read"]	0

Role name	Role description	Resource	Actions	GrantOption
	perform the read operation			

1.3.8.1.5 Tianjimon default role

Role name	Role description	Resource	Actions	GrantOption
Tianjimon operations	Has all Tianjimon permissions, which allows you to perform basic monitoring and operations	26842:tianjimon:*	["*"]	0

1.3.8.2 Operation permissions of operations platforms

1.3.8.2.1 Apsara Infrastructure Management Framework permission list

Resource	Action	Description
*:tianji:services:[sname]:tjmontemplates:[tplname]	delete	DeleteServiceTjmonTmpl
*:tianji:services:[sname]:tjmontemplates:[tplname]	write	PutServiceTjmonTmpl
*:tianji:services:[sname]:templates:[tplname]	write	PutServiceConfTmpl
*:tianji:services:[sname]:templates:[tplname]	delete	DeleteServiceConfTmpl
*:tianji:services:[sname]:serviceinstances:[sname]:tjmontemplate	read	GetServiceInstanceTjmonTmpl
*:tianji:services:[sname]:serviceinstances:[sname]:tssessions	terminal	CreateTsSessionByService
*:tianji:services:[sname]:serviceinstances:[sname]:template	write	SetServiceInstanceTmpl

Resource	Action	Description
*:tianji:services:[sname]: serviceinstances:[siname]: template	delete	DeleteServiceInstanceTpl
*:tianji:services:[sname]: serviceinstances:[siname]: template	read	GetServiceInstanceTpl
*:tianji:services:[sname]: serviceinstances:[siname]:tags :[tag]	delete	DeleteServiceInstanc eProductTagInService
*:tianji:services:[sname]: serviceinstances:[siname]:tags :[tag]	write	AddServiceInstancePr oductTagInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: resources	read	GetServerroleResourc eInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	write	OperateSRMachineInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	read	GetMachineSRInfoInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	delete	DeleteSRMachineActio nInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	read	GetMachinesSRInfoInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	delete	DeleteSRMachinesActi onInService
*:tianji:services:[sname]: serviceinstances:[siname	write	OperateSRMachinesInService

Resource	Action	Description
]:serverroles:[serverrole]: machines		
*:tianji:services:[sname]: serviceinstances:[sname]: serverroles:[serverrole]:apps:[app]:resources	read	GetAppResourceInService
*:tianji:services:[sname]: serviceinstances:[sname]: serverroles:[serverrole]:apps :[app]:machines:[machine]: tianjilogs	read	TianjiLogsInService
*:tianji:services:[sname]: serviceinstances:[sname]: serverroles	read	GetServiceInstanceSe rverrolesInService
*:tianji:services:[sname]: serviceinstances:[sname]: schema	write	SetServiceInstanceSchema
*:tianji:services:[sname]: serviceinstances:[sname]: schema	delete	DeleteServiceInstanceSchema
*:tianji:services:[sname]: serviceinstances:[sname]: rollings:[version]	write	OperateRollingJobInService
*:tianji:services:[sname]: serviceinstances:[sname]: rollings	read	ListRollingJobInService
*:tianji:services:[sname]: serviceinstances:[sname]: resources	read	GetInstanceResourceInService
*:tianji:services:[sname]: serviceinstances:[sname]: machines:[machine]	read	GetMachineAllSRInfoInService
*:tianji:services:[sname]: serviceinstances:[sname]	write	DeployServiceInstanc eInService
*:tianji:services:[sname]: serviceinstances:[sname]	read	GetServiceInstanceConf

Resource	Action	Description
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps:[app] :files:name	read	GetMachineAppFileLis tInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps:[app] :files:download	read	GetMachineAppFileDow nloadInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps:[app] :files:content	read	GetMachineAppFileCon tentInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps:[app] :filelist	read	GetMachineFileListInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps:[app] :dockerlogs	read	DockerLogsInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps:[app] :debuglog	read	GetMachineDebugLogIn Service
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps	read	GetMachineAppListInService
*:tianji:services:[sname]: serverroles:[serverrole]:apps:[app]:dockerinspect	read	DockerInspect
*:tianji:services:[sname]: schemas:[schemaname]	write	PutServiceSchema
*:tianji:services:[sname]: schemas:[schemaname]	delete	DeleteServiceSchema
*:tianji:services:[sname]: resources	read	GetResourceInService
*:tianji:services:[sname]	delete	DeleteService

Resource	Action	Description
*:tianji:services:[sname]	write	CreateService
*:tianji:projects:[pname]: machinebuckets:[bname]: machines:[machine]	read	GetMachineBucketMachineInfo
*:tianji:projects:[pname]: machinebuckets:[bname]: machines	read	GetMachineBucketMachines
*:tianji:projects:[pname]: machinebuckets:[bname]	write	CreateMachineBucket
*:tianji:projects:[pname]: machinebuckets:[bname]	write	OperateMachineBucketMachines
*:tianji:projects:[pname]: machinebuckets:[bname]	delete	DeleteMachineBucket
*:tianji:projects:[pname]: machinebuckets:[bname]	read	GetMachineBucketMachinesLegacy
*:tianji:projects:[pname]: machinebuckets	read	GetMachineBucketList
*:tianji:projects:[pname]: projects:[pname]:clusters:[cname]:tssessions:[tssessionn ame]:tsses	terminal	UpdateTsSessionTssByCluster
*:tianji:projects:[pname]: projects:[pname]:clusters:[cname]:tssessions	terminal	CreateTsSessionByCluster
*:tianji:projects:[pname]: clusters:[cname]:serviceins tances:[sname]:tjmontemplate	read	GetServiceInstanceTjmonTplInCluster
*:tianji:projects:[pname]: clusters:[cname]:serviceins tances:[sname]:template	delete	DeleteServiceInstanceTplInCluster
*:tianji:projects:[pname]: clusters:[cname]:serviceins tances:[sname]:template	write	SetServiceInstanceTplInCluster
*:tianji:projects:[pname]: clusters:[cname]:serviceins tances:[sname]:template	read	GetServiceInstanceTplInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	write	AddServiceInstanceProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanceProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	GetServerroleResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:name	read	GetMachineAppFileList
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:download	read	GetMachineAppFileDownload
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:content	read	GetMachineAppFileContent
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:filelist	read	GetMachineFileList
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:dockerlogs	read	DockerLogsInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[read	GetMachineDebugLog

Resource	Action	Description
serverrole]:machines:[machine]:apps:[app]:debuglog		
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]:apps	read	GetMachineAppList
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	read	GetMachineSRInfoInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	write	OperateSRMachineInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	delete	DeleteSRMachineActionInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	write	OperateSRMachinesInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	delete	DeleteSRMachinesActionInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	read	GetAllMachineSRInfoInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverrole	read	GetAppResourceInCluster

Resource	Action	Description
s:[serverrole]:apps:[app]:resources		
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	TianjiLogsInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:dockerinspect	read	DockerInspectInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:serverroles	read	GetServiceInstanceServerrolesInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	delete	DeleteServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	write	SetServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:resources	read	GetInstanceResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	delete	DeleteServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	write	CreateServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	read	GetServiceInstanceConfInCluster
*:tianji:projects:[pname]:clusters:[cname]:rollings:[version]	write	OperateRollingJob

Resource	Action	Description
*:tianji:projects:[pname]: clusters:[cname]:rollings	read	ListRollingJob
*:tianji:projects:[pname]: clusters:[cname]:resources	read	GetResourceInCluster
*:tianji:projects:[pname]: clusters:[cname]:quota	write	SetClusterQuotas
*:tianji:projects:[pname]: clusters:[cname]:machinesinfo	read	GetClusterMachineInfo
*:tianji:projects:[pname]: clusters:[cname]:machines:[machine]	read	GetMachineAllSRInfo
*:tianji:projects:[pname]: clusters:[cname]:machines:[machine]	write	SetMachineAction
*:tianji:projects:[pname]: clusters:[cname]:machines:[machine]	delete	DeleteMachineAction
*:tianji:projects:[pname]: clusters:[cname]:machines	write	OperateClusterMachines
*:tianji:projects:[pname]: clusters:[cname]:difflist	read	GetVersionDiffList
*:tianji:projects:[pname]: clusters:[cname]:diff	read	GetVersionDiff
*:tianji:projects:[pname]: clusters:[cname]:deploylogs:[version]	read	GetDeployLogInCluster
*:tianji:projects:[pname]: clusters:[cname]:deploylogs	read	GetDeployLogListInCluster
*:tianji:projects:[pname]: clusters:[cname]:builds:[version]	read	GetBuildJob
*:tianji:projects:[pname]: clusters:[cname]:builds	read	ListBuildJob
*:tianji:projects:[pname]: clusters:[cname]	write	OperateCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]	delete	DeleteCluster
*:tianji:projects:[pname]:clusters:[cname]	read	GetClusterConf
*:tianji:projects:[pname]:clusters:[cname]	write	DeployCluster
*:tianji:projects:[pname]	write	CreateProject
*:tianji:projects:[pname]	delete	DeleteProject
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	CreateRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	SetRackunitAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	delete	DeleteRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	SetRackAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	CreateRack
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	delete	DeleteRack
*:tianji:idcs:[idc]:rooms:[room]	write	CreateRoom
*:tianji:idcs:[idc]:rooms:[room]	delete	DeleteRoom
*:tianji:idcs:[idc]:rooms:[room]	write	SetRoomAttr
*:tianji:idcs:[idc]	delete	Deleteldc
*:tianji:idcs:[idc]	write	SetldcAttr
*:tianji:idcs:[idc]	write	Createldc

1.3.8.2.2 Webapp-rule permission list

Resource	Action	Description
26842:webapp-rule:*	write	Adds, deletes, and updates configuration resources
26842:webapp-rule:*	read	Queries configuration resources

1.3.8.2.3 Workflow (grandcanal) console permission list

Resource	Action	Description
26842:grandcanal	read	Queries the workflow activity details and summary
26842:grandcanal	write	Restarts, retries, rolls back, and stops a workflow

1.3.8.2.4 Tianjimon permission list

Resource	Action	Description
26842:tianjimon:monitor-manage	manage	Monitoring and operations

1.4 Apsara Infrastructure Management Framework

1.4.1 What is Apsara Infrastructure Management Framework?

1.4.1.1 Overview

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

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

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

Core functions

- Network initialization in data centers
- Server installation and maintenance process management

- Deployment, expansion, and upgrade of cloud products
- Configuration management of cloud products
- Automatic application for cloud product resources
- Automatic repair of software and hardware faults
- Basic monitoring and business monitoring of software and hardware

1.4.1.2 Basic concepts

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

- project

A collection of clusters that provide service capabilities for external entities.

- cluster

A collection of physical machines that logically provide services and are used to deploy product software.

— A cluster can only belong to one product.

— Multiple services can be deployed on a cluster.

- service

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

- service instance

A service that is deployed on a cluster.

- server role

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

- server role instance

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

- application

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

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

- rolling

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

- service configuration template

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

- associated service template

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

- final status

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

- dependency

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

- upgrade

A way of aligning the current status with the final status of a service. After a user submits the version change, Apsara Infrastructure Management Framework upgrades the service version

to the target version. With the server role as the processing unit, upgrade aims to update the versions of all servers to the target version.

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

1.4.2 Log on to Apsara Infrastructure Management Framework

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

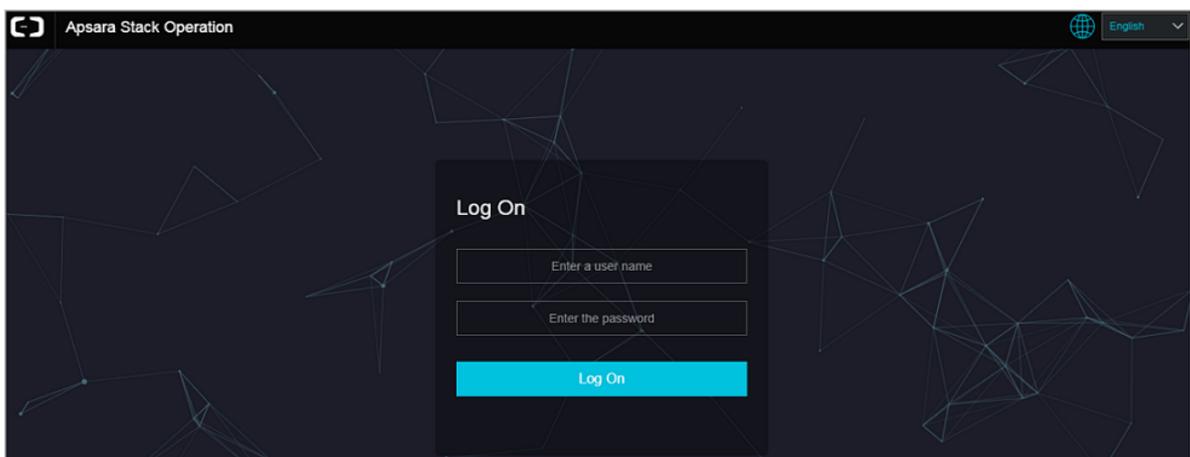
Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 1-6: Log on to ASO



Note:

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.

- The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.
 5. In the left-side navigation pane, select **Products**. In the **Product List**, select **Apsara Infrastructure Management Framework**.

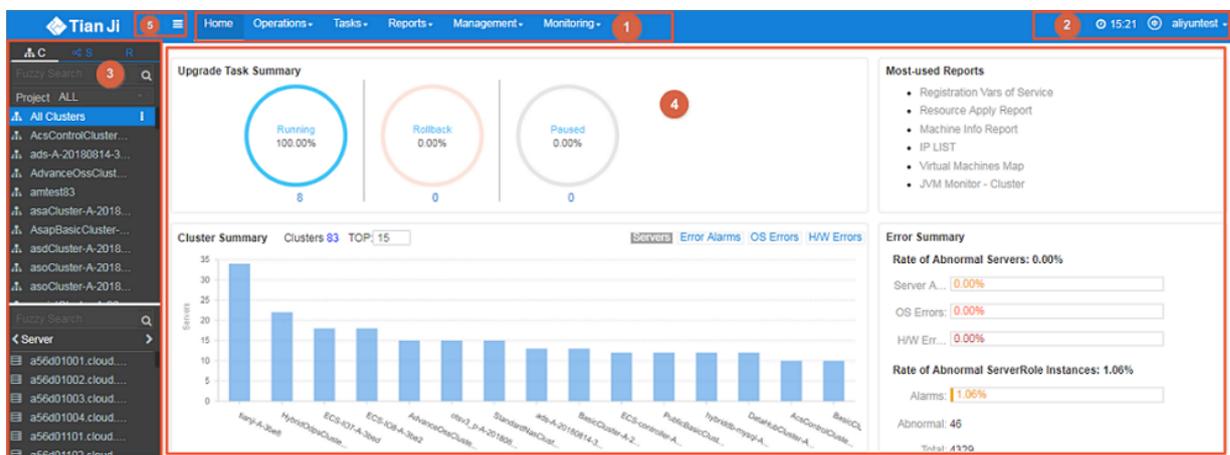
1.4.3 Web page introduction

1.4.3.1 Introduction on the Home page

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

[Log on to Apsara Infrastructure Management Framework](#), the **Home** page appears, as shown in [Home page of Apsara Infrastructure Management Framework](#).

Figure 1-7: Home page of Apsara Infrastructure Management Framework



For more information about the descriptions of functional areas on the **Home** page, see [Descriptions of functional areas](#).

Table 1-14: Descriptions of functional areas

Functional area		Description
1	Navigation bar	<ul style="list-style-type: none"> • Operations: the quick entrance of Operation & Maintenance (O&M) operations, which allows operations engineers to quickly find the corresponding operations and operation objects. This menu includes: <ul style="list-style-type: none"> — Cluster Operations: performs O&M operations on and manages clusters with the project permissions, such as viewing the cluster status. — Service Operations: manages services with the service permissions, such as viewing service list information. — Server Operations: maintains and manages all servers in Apsara Infrastructure Management Framework, such as viewing the server status. • Tasks: A rolling task is generated after you modify the configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects. • Reports: displays the monitoring data in lists and provides the function of searching for different reports. • Monitoring: effectively monitors metrics in the process of system operation and sends alarm notifications for abnormal conditions. This menu includes the functions of displaying alarm status, modifying alarm rules, and querying alarm history.
2	Function buttons in the upper-right corner	<ul style="list-style-type: none"> • : <ul style="list-style-type: none"> — TJDB synchronization time: the generated time of the data that is displayed on the current page. — Final-status calculation time: the calculation time of the final-status data that is displayed on the current page. <p>After data is generated, the system processes the data at the maximum speed. As an asynchronous system, Apsara Infrastructure Management Framework has some latency. The latency helps explain the data results and determine whether the current system has a problem.</p>

Functional area		Description
		<ul style="list-style-type: none"> : In the English environment, click this button to switch to the Chinese environment. : The logon account information. Click  and select Logout to log off from Apsara Infrastructure Management Framework.
3	Left-side navigation pane	<p>In the left-side navigation pane, you can directly view the logical structure of the Apsara Infrastructure Management Framework model.</p> <p>You can view the corresponding detailed data analysis and operations by selecting different levels of nodes in the left-side navigation pane. For more information, see Introduction on the left-side navigation pane.</p>
4	Home page	<p>Displays the summary of related tasks or information as follows:</p> <ul style="list-style-type: none"> Upgrade Task Summary: the number and proportion of running, rolled back, and paused upgrade tasks. Cluster Summary: the numbers of servers, error alarms, operating system errors, and hardware errors for different clusters. Error Summary: the metrics for the rate of abnormal servers and the rate of abnormal server role instances. Most-used Reports: links of the most commonly used statistics reports, which facilitate you to view the report information.
5	Button used to collapse /expand the left-side navigation pane	<p>If you are not required to use the left-side navigation pane when performing O&M operations, click  to collapse the left-side navigation pane and increase the space of the content area.</p>

1.4.3.2 Introduction on the left-side navigation pane

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

Cluster

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

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

- Enter the cluster name in the search box to search for the target cluster quickly. Fuzzy search is supported.
- Select a project from the **Project** drop-down list to filter all the clusters in the project.
- At the right of the cluster, move the pointer over  to perform operations on the cluster as instructed.
- Click a cluster and all the servers and services in this cluster are displayed in the lower-left corner. At the right of the server or service, move the pointer over  to perform operations on the server or service as instructed.
- Click the **Server** tab. Double-click a server to view all the server roles in the server. Double-click a server role to view the applications and then double-click an application to view the log files.
- Click the **Service** tab. Double-click a service to view all the server roles in the service. Double-click a server role to view the servers, double-click a server to view the applications, and double-click an application to view the log files.
- Double-click the log file or click **View** at the right of the log file to view the log details based on time. To download the log file, place your cursor on  and select **Download** after double-clicking the log file. On the **Log Viewer** page, enter the keyword to search for logs.

Service

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

.

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

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

Report

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

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

- Enter the report name in the search box to search for the target report quickly. Fuzzy search is supported.
- Click **All Reports** or **Favorites** to display groups of different categories in the lower-left corner. Double-click a group to view all the reports in this group. Double-click a report to view the report details on the right pane.

1.4.4 Cluster operations

1.4.4.1 View cluster configurations

By viewing the cluster configurations, you can view the basic information, deployment plan, and configurations of a cluster.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Operations > Cluster Operations**.
 - **Cluster**: the cluster name. Click the cluster name to go to the [Cluster Dashboard](#) page.
 - **Scale-Out/Scale-In**: the number of servers or server roles that are scaled out or in. Click the link to go to the [Cluster Operation and Maintenance Center](#) page.
 - **Abnormal Server Count**: the number of servers whose status is not **Good** in the cluster. Click the link to go to the [Cluster Operation and Maintenance Center](#) page.
 - **Final Status of Normal Servers**: whether or not the cluster reaches the final status. Select **Clusters Not Final** to display clusters that do not reach the final status. Click the link to go to the [Service Final Status Query](#) page.
 - **Rolling**: whether the cluster has a running rolling task. Select **Rolling Tasks** to display clusters that have rolling tasks. Click the link to go to the [Rolling Task](#) page.
3. Select a project from the **Project** drop-down list and/or enter the cluster name in the **Cluster** field to search for the corresponding cluster.
4. Find the cluster whose configurations you want to view and then click **Cluster Configuration** in the **Actions** column.
5. On the **Cluster Configuration** page, complete the following configurations as described in [Cluster configurations](#).

Table 1-15: Cluster configurations

Category	Item	Description
Basic Information	Cluster	The cluster name.
	Project	The project to which the cluster belongs.
	Clone Switch	<ul style="list-style-type: none"> • Pseudo Clone: The system is not cloned when a server is added to the cluster. • Real Clone: The system is cloned when a server is added to the cluster.
	Servers	The number of servers in the cluster. Click View Clustering Servers to view the server list.
	Security Verification	The access control among processes. Generally, the non-production environment uses the default configurations and does not perform the verification. In other cases, customize the configurations based on actual requirements to enable or disable the verification.
	Cluster Type	<ul style="list-style-type: none"> • RDS • NETFRAME • T4: a special type that is required by the mixed deployment of e-commerce. • Default: other conditions.
Deployment Plan	Service	The service deployed in the cluster.
	Dependency Service	The service that the current service depends on.
Service configurations	Service Info	Select a service from the Service Info drop-down list and then the configurations of this service are displayed.
	Service Template	The template associated with the service.
	Monitoring Template	The monitoring template used by the service.
	Server Mappings	The servers contained by different server roles in the service.
	Software Version	The software version of the server role in the service.

Category	Item	Description
	Availability Configuration	The availability configuration percentage of the server role in the service.
	Deployment Plan	The deployment plan of the server role in the service.
	Configuration Information	The configuration file used in the service.

6. Click **Operation Logs** in the upper-right corner to view the release changes. For more information, see [View operation logs](#).

1.4.4.2 View the cluster dashboard

On the **Cluster Dashboard** page, you can view the basic information and related statistics of a cluster.

Procedure

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

Table 1-16: Cluster dashboard descriptions

Item	Description
Cluster Basic Information	<p>Displays the basic information of the cluster as follows:</p> <ul style="list-style-type: none"> • Project Name: the project name. • Cluster Name: the cluster name. • IDC: the data center to which the cluster belongs. • Final Status Version: the latest version of the cluster.

Item	Description
	<ul style="list-style-type: none"> • Cluster in Final Status: whether the cluster reaches the final status. • Servers Not In Final Status: the number of servers that do not reach the final status in the cluster when the cluster does not reach the final status. • Real/Pseudo Clone: whether to clone the system when a server is added to the cluster. • Expected Servers: the number of expected servers in the cluster. • Actual Servers: the number of servers in the current environment. • Servers Not Good: the number of servers whose status is not Good in the cluster. • Actual Services: the number of services that are actually deployed in the cluster. • Actual Service Roles: the number of server roles that are actually deployed in the cluster. • Cluster Status: whether the cluster is starting or shutting down servers.
Server Status Overview	The statistical chart of the server status in the cluster.
Servers in Final Status	The numbers of servers that reach the final status and those that do not reach the final status in each service of the cluster.
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
Disk_usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP state-system	The TCP request status chart.
TCP retrans-System	The chart of TCP retransmission amount.
Disk_IO-System	The statistical table of the disk input and output.
Service Instances List	<p>Displays the service instances deployed in the cluster and the related final status information.</p> <ul style="list-style-type: none"> • Service Instance: the service instance deployed in the cluster. • Final Status: whether or not the service instance reaches the final status.

Item	Description
	<ul style="list-style-type: none"> • Expected Server Roles: the number of server roles that the service instance expects to deploy. • Server Roles In Final Status: the number of server roles that reach the final status in the service instance. • Server Roles Going Offline: the number of server roles that are going offline in the service instance. • Actions: Click Details to go to the Service Instance Dashboard page.
Upgrade Tasks	<p>Displays the upgrade tasks related to the cluster.</p> <ul style="list-style-type: none"> • Cluster Name: the name of the upgrade cluster. • Type: the type of the upgrade task. The options include app (version upgrade) and config (configuration change). • Git Version: the change version to which the upgrade task belongs. • Description: the description about the change. • Rolling Result: the result of the upgrade task. • Submitted By: the person who submits the change. • Submitted At: the time when the change is submitted. • Start Time: the time to start the rolling. • End Time: the time to finish the upgrade. • Time Used: the time used for the upgrade. • Actions: Click Details to go to the Rolling Task page.
Cluster Resource Request Status	<ul style="list-style-type: none"> • Version: the resource request version. • Msg: the exception information. • BeginTime: the time to start analyzing the resource request. • EndTime: the time to finish analyzing the resource request. • Build Status: the build status of resources. • Resource Process Status: the resource request status in the version.
Cluster Resource	<ul style="list-style-type: none"> • Service: the service name. • Service Role: the server role name. • App: the application of the server role. • Name: the resource name. • Type: the resource type. • Status: the resource request status. • Error Msg: the exception message. • Parameters: the resource parameters.

Item	Description
	<ul style="list-style-type: none"> • Result: the resource request result. • Res: the resource ID. • Reprocess Status: the status of interaction with Business Foundation System during the VIP resource request. • Reprocess Msg: the exception message of interaction with Business Foundation System during the VIP resource request. • Reprocess Result: the result of interaction with Business Foundation System during the VIP resource request. • Refer Version List: the version that uses the resource.
VM Mappings	<p>Displays the information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.</p> <ul style="list-style-type: none"> • VM: the hostname of the virtual machine. • Currently Deployed On: the hostname of the physical machine where the virtual machine is deployed. • Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.
Service Dependencies	<p>Displays the dependencies of service instances and server roles in the cluster, and the final status information of the dependent service or server role.</p> <ul style="list-style-type: none"> • Service: the service name. • Service Role: the server role name. • Dependent Service: the service on which the server role depends. • Dependent Service Role: the server role on which the server role depends. • Dependent Cluster: the cluster to which the dependent server role belongs. • Dependency in Final Status: whether the dependent server role reaches the final status.

1.4.4.3 View the cluster operation and maintenance center

On the **Cluster Operation and Maintenance Center** page, you can view the status or statistics of services or servers in the cluster.

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)

2. Currently, you have three ways to go to the **Cluster Operation and Maintenance Center** page:

- In the left-side navigation pane, click the **C** tab. At the right of the cluster, move the pointer over  and then select **Cluster Operation and Maintenance Center**.
- In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click **Monitoring > Cluster Operation and Maintenance Center** in the **Actions** column at the right of the cluster.
- On the **Cluster Dashboard** page, click **Operations Menu > Cluster Operation and Maintenance Center**.

3. For more information about the **Cluster Operation and Maintenance Center** page, see [Cluster operation and maintenance center descriptions](#).

Table 1-17: Cluster operation and maintenance center descriptions

Item	Description
SR not in Final Status	Displays all server roles that do not reach the final status in the cluster. Click the number to expand a server role list, and click a server role in the list to display the information of servers contained by the server role.
Running Tasks	Displays whether the cluster has running rolling tasks. Click Rolling to go to the Rolling Task page.
Head Version Commit Time	Displays the committed time of the head version. Click the time to view the submission details.
Head Version Analysis	The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to updated details. The head version analysis has the following statuses: <ul style="list-style-type: none"> • Preparing: No new version is available now. • Waiting: The latest version is found. The analysis module has not started up yet. • Doing: The module is analyzing the application that requires update. • Done: The head version analysis is successfully completed. • Failed: An error occurred while analyzing the final status. The updated contents cannot be parsed.

Item	Description
	<p>If the status is not done, Apsara Infrastructure Management Framework cannot detect the updated contents of server roles in the latest version.</p> <p>Click the status to view the relevant information.</p>
Service	<p>Select a service deployed in the cluster from the Service drop-down list.</p>
Service Role	<p>Select a server role of a service in the cluster from the Service Role drop-down list.</p> <div data-bbox="564 689 1442 853" style="background-color: #f0f0f0; padding: 5px;">  Note: After you select the service and server role, the information of servers related to the service or server role is displayed in the list. </div>
Total Servers	<p>The total number of servers in the cluster, or the total number of servers contained by a specific server role of a specific service.</p>
Scale-In/Scale-Out	<p>The number of servers or server roles that are scaled in or out.</p>
Abnormal Server	<p>The number of abnormal servers that encounter each type of the following faults.</p> <ul style="list-style-type: none"> • Ping Failed: A ping_monitor error is reported, and TianjiMaster cannot successfully ping the server. • No Heartbeat: TianjiClient on the server does not regularly report data to indicate the status of this server, which may be caused by the TianjiClient problem or network problem. • Status Error: The server has an error reported by the monitor or a fault of the critical or fatal level. Check the alarm information and accordingly solve the issue.
Abnormal Service	<p>The number of servers with abnormal services. To determine if a service reaches the final status, see the following rules:</p> <ul style="list-style-type: none"> • The server role on the server is in the GOOD status. • Each application of the server role on the server must keep the actual version the same as the head version. • Before the Image Builder builds an application of the head version, Apsara Infrastructure Management Framework cannot determine the value of the head version and the service final status is unknown. This process is called the change preparation process. The service final status cannot be determined during the preparation process or upon a preparation failure.

Item	Description
Server List	<p>Displays all the servers in the cluster or the servers contained by a specific server role of a specific service.</p> <ul style="list-style-type: none"> • Server Search: Click the search box to enter the server in the displayed dialog box. Fuzzy search is supported and multiple servers can be queried at a time. • Click the server name to view the physical information of the server in the displayed Server Information dialog box. Click DashBoard to view the server details. • Click Details in the Final Status column to view the status and exception information of services on the server. <ul style="list-style-type: none"> — Normal. — Server scale-in: The server is being removed from the cluster for scale-in purpose. — Server scale-out: The server is being added to the cluster for scale-out purpose. — SR scale-in: A server role is being removed from the server for scale-in purpose. • Click Details in the Running Status column to view the running status or exception information of the server. • Click Error, Warning, or Good in the Monitoring Statistics column to view the server monitoring items and server role monitoring items. • Click Actions in the Actions column to restart, out-of-band restart, reclone, or open a ticket for the server.

1.4.4.4 View the service final status

By querying the service final status, you can view if a service in a cluster reaches the final status and the final status information.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. Currently, you have two ways to go to the **Service Final Status Query** page:
 - In the left-side navigation pane, click the **C** tab. At the right of the cluster, move the pointer over  and then choose **Monitoring > Service Final Status Query**.
 - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click **Monitoring > Service Final Status Query** in the **Actions** column at the right of the cluster.

3. For more information about the **Service Final Status Query** page, see [Service final status query descriptions](#).

Table 1-18: Service final status query descriptions

Item	Description
Project Name	The name of the project to which the cluster belongs.
Cluster Name	The cluster name.
Head Version Commit Time	The committed time of the head version.
Head Version Analysis	<p>The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to updated details. The head version analysis has the following statuses:</p> <ul style="list-style-type: none"> • Preparing: No new version is available now. • Waiting: The latest version is found. The analysis module has not started up yet. • Doing: The module is analyzing the application that requires update. • Done: The head version analysis is successfully completed. • Failed: An error occurred while analyzing the final status. The updated contents cannot be parsed. <p>If the status is not done, Apsara Infrastructure Management Framework cannot detect the updated contents of server roles in the latest version.</p>
Cluster Rolling Status	Displays the information of the current rolling task in the cluster, if any. The rolling task may not be of the head version.
Cluster Server Final Status Statistics	The status of all servers in the cluster. Click View Details to go to the Cluster Operation and Maintenance Center page and view the detailed information of all servers.
Cluster SR Version Final Status	<p>The final status of cluster service version.</p> <div style="background-color: #f0f0f0; padding: 10px;">  Note: Take statistics of services that do not reach the final status, which is caused by version inconsistency or status exceptions. If services do not reach the final status because of server problems, go to Cluster Server Final Status Statistics to view the statistics. </div>

Item	Description
SR Version Final Status	The number of servers that do not reach the final status when a server role has tasks.

1.4.4.5 View operation logs

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. Currently, you have two ways to go to the **Cluster Operation Log** page:
 - In the left-side navigation pane, click the **C** tab. At the right of the cluster, move the pointer over  and then choose **Monitoring > Operation Logs**.
 - In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page, click **Monitoring > Operation Logs** in the **Actions** column at the right of the cluster.
3. Find the log and then click **View Release Changes** in the **Actions** column.
4. On the **Version Difference** page, perform the following operations:
 - **Select Base Version:** Select a base version.
 - Select a **Configuration Type**:
 - **Extend Configuration:** displays the configuration differences after the configuration on the cluster is combined with the configuration in the template.
 - **Cluster Configuration:** displays the configuration differences on the cluster.
5. Click **Obtain Difference**.

The differences are displayed in the differential file list.

1.4.5 Service operations

1.4.5.1 View the service list

By querying the service list, you can view the list of all services and the related information.

Procedure

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

- For more information about the **Service Operations** page, see [Service operations descriptions](#).

Table 1-19: Service operations descriptions

Item	Description
Service	The service name.
Service Instances	The number of service instances in the service.
Service Configuration Templates	The number of templates configured by the service.
Monitoring Template Count	The number of monitoring templates.
Service Schemas	The number of service configuration validation templates.
Actions	Click Management to view the service instances, service templates, monitoring templates, monitoring instances, service schemas, and detection scripts.

1.4.5.2 View the service instance dashboard

On the **Service Instance Dashboard** page, you can view the basic information and statistics of a service instance.

Procedure

- [Log on to Apsara Infrastructure Management Framework](#).
- In the left-side navigation pane, click the **S** tab.
- Enter the service name in the search box. Services that meet the criterion are displayed.
- Click a service name and then service instances in the service are displayed in the lower-left corner.
- At the right of the service instance, move the pointer over  and then select **Dashboard**.
- For more information about the **Service Instance Dashboard** page, see [Service instance dashboard descriptions](#).

Table 1-20: Service instance dashboard descriptions

Item	Description
Service Instance Summary	<p>Displays the basic information of the service instance as follows:</p> <ul style="list-style-type: none"> • Cluster Name: the name of the cluster to which the service instance belongs. • Service Name: the name of the service to which the service instance belongs. • Actual Servers: the number of servers in the current environment. • Expected Servers: the number of servers that the service instance expects. • Template Name: the name of the service template used by the service instance. • Template Version: the version of the service template used by the service instance. • schema: the name of the service schema used by the service instance. • Apsara Infrastructure Monitor Template: the name of the Apsara Infrastructure Management Framework Monitor template used by the service instance.
Service Role Statuses	The statistical chart of the current status of server roles in the service instance.
Server Statuses for Service Roles	The status statistics of servers where server roles are located.
Service Monitoring List	<ul style="list-style-type: none"> • Monitor: the monitoring item name. • Level: the monitoring item level. • Description: the description about the monitoring contents. • Updated At: the time when the data is updated.
Service Alarm Status	<ul style="list-style-type: none"> • Alarm Name: the alarm name. • Instance Information: the instance information. • Alarm Start: the start time of the alarm. • Alarm End: the end time of the alarm. • Alarm Duration: the alarm duration. • Severity Level: the severity level of the alarm. • Occurrences: the number of times the alarm is triggered.
Service Role List	<ul style="list-style-type: none"> • Service Role: the server role name.

Item	Description
	<ul style="list-style-type: none"> • Current Status: the current status. • Expected Servers: the number of expected servers. • Servers In Final Status: the number of servers that reach the final status. • Servers Going Offline: the number of servers that are going offline. • Rolling Task State: the status of the rolling task. • Actions: Click Details to go to the Service Role Dashboard page.
Service Alarm History	<ul style="list-style-type: none"> • Alarm Name • Alarm Time • Instance Information • Severity Level • Contact Group
Service Dependencies	<p>Displays the dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role.</p> <ul style="list-style-type: none"> • Service Role: the server role name. • Dependent Service: the service on which the server role depends. • Dependent Service Role: the server role on which the server role depends. • Dependent Cluster: the cluster to which the dependent server role belongs. • Dependency in Final Status: whether the dependent server role reaches the final status.

1.4.5.3 View the server role dashboard

On the **Service Role Dashboard** page, you can view the statistics of a server role.

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the left-side navigation pane, click the **S** tab.
3. Enter the service name in the search box. Services that meet the criterion are displayed.
4. Click a service name and then service instances in the service are displayed in the lower-left corner.

5. At the right of the service instance, move the pointer over  and then select **Dashboard**.
6. On the **Service Instance Dashboard** page, click **Details** in the **Service Role List** section.
7. For more information about the **Service Role Dashboard** page, see [Server role dashboard descriptions](#).

Table 1-21: Server role dashboard descriptions

Item	Description
Service Role Summary	Displays the basic information of the server role as follows: <ul style="list-style-type: none"> • Project Name: the name of the project to which the server role belongs. • Cluster Name: the name of the cluster to which the server role belongs. • Service Instance: the name of the service instance to which the server role belongs. • Service Role: the server role name. • In Final Status: whether the server role reaches the final status. • Expected Servers: the number of expected servers. • Actual Servers: the number of actual servers. • Servers Not Good: the number of servers whose status is not Good. • Servers with Role Status Not Good: the number of server roles whose status is not Good. • Servers Going Offline: the number of servers that are going offline. • rolling: whether a running rolling task exists. • Rolling Task State: the current status of the rolling task. • Time Used: the time used for running the rolling task.
Server Final Status Overview	The statistical chart of the current status of the server role.
Service Role Monitor Information	<ul style="list-style-type: none"> • Updated At: the time when the data is updated. • Monitor: the monitoring item name. • Level: the monitoring item level. • Description: the description of the monitoring item.
Server Information	<ul style="list-style-type: none"> • Machine Name: the hostname of the server. • IP: the IP address of the server.

Item	Description
	<ul style="list-style-type: none"> • Server Status: the server status. • Server Action: the action that the server is performing. • Service Role Status: the status of the server role. • Service Role Action: the action that the server role is performing. • Current Version: the current version of the server role on the server. • Target Version: the expected version of the server role on the server. • Error Message: the exception message. • Actions: Click Details to go to the Server Details page.
Service Role for Servers	<ul style="list-style-type: none"> • Updated At: the time when the data is updated. • Machine Name: the hostname of the server. • Monitor: the monitoring item name. • Level: the monitoring item level. • Description: the description of the monitoring item.
VM Mappings	<p>The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.</p> <ul style="list-style-type: none"> • VM: the hostname of the virtual machine. • Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed. • Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.
Service Dependencies	<p>The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role.</p> <ul style="list-style-type: none"> • Dependent Service: the service on which the server role depends. • Dependent Service Role: the server role on which the server role depends. • Dependent Cluster: the cluster to which the dependent server role belongs. • Dependency in Final Status: whether the dependent server role reaches the final status.

1.4.6 Server operations

1.4.6.1 View the server dashboard

By viewing the server dashboard, you can view the statistics of a server.

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the left-side navigation pane, click the **C** tab.
3. On the **Server** tab in the lower-left corner, enter the server name in the search box. Servers that meet the criterion are displayed.
4. At the right of the server, move the pointer over  and then select **Dashboard**.
5. On the **Server Details** page, view all the information of this server. For more information, see [Server dashboard descriptions](#).

Table 1-22: Server dashboard descriptions

Item	Description
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
DISK Usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-System	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
DISK IO-System	The statistical table of the disk input and output.
Server Summary	<ul style="list-style-type: none"> • Project Name: the name of the project to which the server belongs. • Cluster Name: the name of the cluster to which the server belongs. • Machine Name: the server name. • sn: the serial number of the server. • IP: the IP address of the server. • IDC: the data center of the server. • Room: the room in the data center where the server is located. • Rack: the rack where the server is located.

Item	Description
	<ul style="list-style-type: none"> • Unit in Rack: the location of the rack. • Warranty: the warranty of the server. • Purchase Date: the date when the server is purchased. • Server Status: the running status of the server. • Status: the hardware status of the server. • CPUs: the number of CPUs for the server. • Disks: the disk size. • Memory: the memory size. • Manufacturer: the server manufacturer. • os: the operating system of the server. • part: the disk partition.
Service Role Status by Server	The distribution of the current status of all server roles on the server.
Server Monitoring Information	<ul style="list-style-type: none"> • Monitor: the monitoring item name. • Level: the monitoring item level. • Description: the description of the monitoring contents. • Updated At: the time when the monitoring information is updated.
Server Service Role Status	<ul style="list-style-type: none"> • Service Instance: the service instance name. • Service Role: the server role name. • Service Role Status: the server role status. • Service Role Action: the server role action. • Error Message: the error message. • Target Version: the expected version. • Current Version: the current version. • Actual Version Update Time: the updated time of the current version. • Actions: Click Details to go to the Service Role Dashboard page.
APP Status with Service Role	<ul style="list-style-type: none"> • Application Name: the application name. • Process Number: the process number. • Status: the application status. • Current Build ID: the ID of the current package version. • Target Build ID: the ID of the expected package version. • Git Version • Start Time

Item	Description
	<ul style="list-style-type: none"> • End Time • Interval: the interval between the time when Apsara Infrastructure Management Framework detects that the process exits and the time when Apsara Infrastructure Management Framework repairs the process. • Info Message: the normal output logs. • Error Message: the abnormal logs.

1.4.7 Monitoring center

1.4.7.1 Modify an alarm rule

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the top navigation bar, choose **Operations > Service Operations**.
3. Enter the service name in the search box.
4. At the right of the service, click **Management** in the **Actions** column.
5. Click the **Monitoring Template** tab.
6. At the right of the template, click **Edit** in the **Actions** column.
7. On the **Alarm Rules** tab, find the alarm rule and then click **Edit**.
8. Configure the monitoring parameters based on actual conditions.
9. Click **Preview Change** to view the changes.
10. Click **Save Change**.

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

1.4.7.2 View the status of a monitoring instance

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the top navigation bar, choose **Operations > Service Operations**.
3. Enter the service name in the search box.

4. At the right of the service, click **Management** in the **Actions** column.
5. Click the **Monitoring Instance** tab.

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

1.4.7.3 View the alarm status

By querying the alarm status, you can view the alarms generated in different services and the corresponding alarm details.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Monitoring > Alarm Status**.
3. Search for the alarms based on the service name, cluster name, alarm name, and/or the time range when the alarm is triggered.
4. For more information about the alarm status, see [Alarm status descriptions](#).

Table 1-23: Alarm status descriptions

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Alarm Instance	The name of the service instance being monitored. Click the alarm instance to view the alarm history of this instance.
Alarm Status	The alarm has two statuses: Restored and Alarming .
Alarm Level	Alarms have four levels according to the effect on services. The smaller the number, the bigger the effect. <ul style="list-style-type: none"> • P1 • P2 • P3 • P4
Alarm Name	The name of the generated alarm. Click the alarm name to view the alarm details.
Happen Time	The time when the alarm is triggered and how long the alarm has lasted.
Show	Click Show in the Actions column to show the data before and after the alarm is triggered.

1.4.7.4 View the alarm history

By querying the alarm history, you can view all the history alarms generated in different services and the corresponding alarm details.

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the top navigation bar, choose **Monitoring > Alarm History**.
3. Search for the alarms based on the service name, cluster name, period, and/or time range.
4. For more information about the alarm history descriptions, see [Alarm history descriptions](#).

Table 1-24: Alarm history descriptions

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Alarm Instance	The name of the resource where an alarm is triggered.
Status	The alarm has two statuses: Restored and Alarming .
Alarm Level	Alarms have four levels according to the effect on services. The smaller the number, the bigger the effect. <ul style="list-style-type: none"> • P1 • P2 • P3 • P4
Alarm Name	The name of the generated alarm. Click the alarm name to view the alarm details.
Happen Time	The time when the alarm is triggered.
Notification	The groups and members that are notified when an alarm is triggered.
Show	Click Show in the Actions column to show the data before and after the alarm is triggered.

1.4.7.5 View alarm rules

You can view the configured alarm rules.

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the top navigation bar, choose **Monitoring > Alarm Rules**.

3. Search for the alarm rules based on the service name, cluster name, and/or alarm name.
4. For more information about the alarm rule descriptions, see [Alarm rule descriptions](#).

Table 1-25: Alarm rule descriptions

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Alarm Name	The name of the generated alarm.
Alarm Conditions	The conditions met when the alarm is triggered.
Period	The frequency (in seconds) with which an alarm rule is run.
Receiver	The groups and members that are notified when an alarm is triggered.
Status	The current status of the alarm rule. <ul style="list-style-type: none"> • Running • Stopped

1.4.8 Tasks and deployment summary

1.4.8.1 View rolling tasks

You can view running rolling tasks and the corresponding status.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Operations > Cluster Operations**.
3. Select **Rolling Tasks**.
4. Click **rolling** in the **Rolling** column. The **Rolling Task** page appears.
5. For more information about the **Rolling Task** page, see [Rolling task descriptions](#).

Table 1-26: Rolling task descriptions

Item	Description
Change Version	The version that triggers the change of the rolling task.
Description	The description about the change.
Head Version Analysis	The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the

Item	Description
	<p>version to updated details. The head version analysis has the following statuses:</p> <ul style="list-style-type: none"> • Preparing: No new version is available now. • Waiting: The latest version is found. The analysis module has not started up yet. • Doing: The module is analyzing the application that requires update. • Done: The head version analysis is successfully completed. • Failed: An error occurred while analyzing the final status. The updated contents cannot be parsed. <p>If the status is not done, Apsara Infrastructure Management Framework cannot detect the updated contents of server roles in the latest version.</p>
Blocked Service Role	Server roles blocked in the rolling task. Generally, server roles are blocked because of dependencies.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.
Actions	Click View Difference to go to the Version Difference page. For more information, see View operation logs .
Service Name	The name of the service where a change occurs.
Status	<p>The current status of the service.</p> <ul style="list-style-type: none"> • succeeded: The task is successfully completed. • Downloading: The task is being downloaded. • Rolling: The rolling task is running. • RollingBack: The rolling task failed and is rolling back.
Server Role Status	The server role status. Click  at the left of the service name to show the rolling task status of each server role in the service.

1.4.8.2 View running tasks

By searching for running tasks, you can view all the running tasks or filter the current tasks to view the running status of each task.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Tasks > Running Tasks**.
3. You can search for the running tasks based on the cluster name, role name, status, submitter, Git version, and/or time range.

4. Find the task and click **View Tasks** in the **Rolling Task Status** column. The *Rolling Task* page appears.

1.4.8.3 View history tasks

You can view the historical running conditions of completed tasks.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Tasks > History Tasks**.
3. You can search for the history tasks based on the cluster name, Git version, submitter, and/or time range.
4. Find the task and click **Details** in the **Actions** column. The *Rolling Task* page appears.

1.4.8.4 View the deployment summary

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Tasks > Deployment Summary**.
 - View the deployment status and the duration of a certain status for each project.
 - Gray: wait to be deployed. It indicates that some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed.
 - Blue: being deployed. It indicates that the project has not reached the final status for one time yet.
 - Green: has reached the final status. It indicates that all clusters in the project have reached the final status.
 - Orange: not reaches the final status. It indicates that a server role does not reach the final status for some reason after the project reaches the final status for the first time.
 - Configure the global clone switch.
 - normal: Clone is allowed.
 - block: Clone is prohibited.
 - Configure the global dependency switch.

- normal: All configured dependencies are checked.
- ignore: The dependency is not checked.
- ignore_service: None of the service-level dependencies, including the server role dependencies across services, are checked, and only the server role-level dependencies are checked.

3. Click the **Deployment Details** tab.

For more information, see [Deployment details descriptions](#).

Table 1-27: Deployment details descriptions

Item	Description
Status Statistics	<p>The general statistics of deployment conditions, including the total number of projects that are currently available. Click each status to display the projects in the corresponding status in the list. The projects have five deployment statuses:</p> <ul style="list-style-type: none"> • Final: All the clusters in the project have reached the final status. • Deploying: The project has not reached the final status for one time yet. • Waiting: Some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed. • Non-final: A server role does not reach the final status for some reason after the project reaches the final status for the first time. • Inspector Warning: An error is detected on service instances in the project during the inspection.
Start Time	The time when Apsara Infrastructure Management Framework starts the deployment.
Progress	The proportion of server roles that reach the final status to all the server roles in all projects.
Deployment Status	<p>The time indicates the deployment duration for the following statuses: Final, Deploying, Waiting, and Inspector Warning.</p> <p>The time indicates the duration before the final status is reached for the Non-final status.</p> <p>Click the time to view the detailed information.</p>
Deployment Progress	The proportion of clusters, services, and server roles that reach the final status to the total clusters, services, and server roles in the project.

Item	Description
	Click Details to view the deployment statuses of clusters, services, and server roles. The deployment statuses are indicated by icons, which are the same as those used for status statistics.
Resource Application Progress	Total indicates the total number of resources related to the project. <ul style="list-style-type: none"> • Done: the number of resources that have been successfully applied for. • Doing: the number of resources that are being applied for. The number of retries (if any) is displayed next to the number of resources. • Block: the number of resources whose applications are blocked by other resources. • Failed: the number of resources whose applications failed.
Inspector Error	The number of inspection alarms for the current project.
Monitoring Information	The number of alarms generated for the machine monitor and the machine server role monitor in the current project.
Dependency	Click the icon to view the project services that depend on other services , and the current deployment status of the services that are depended on.

1.4.9 Reports

1.4.9.1 View reports

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

Context

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

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. You can go to the report list in the following three ways:
 - In the top navigation bar, choose **Reports > System Reports**.
 - In the top navigation bar, choose **Reports > All Reports**.

- In the left-side navigation pane, click the **R** tab. Move the pointer over  and then select **View**.

For more information about the report descriptions, see [Report descriptions](#).

Table 1-28: Report descriptions

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

3. Enter the report name in the search box.
4. Click the report name to go to the corresponding report details page.

For more information about the reports, see [Appendix](#).

1.4.9.2 Add a report to favorites

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. You can go to the report list in the following three ways:
 - In the top navigation bar, choose **Reports > System Reports**.
 - In the top navigation bar, choose **Reports > All Reports**.

- In the left-side navigation pane, click the **R** tab. Move the pointer over  and then select

View.

3. Enter the report name in the search box.
4. At the right of the report, click **Add to Favorites** in the **Actions** column.
5. In the displayed **Add to Favorites** dialog box, enter tags for the report.
6. Click **Add to Favorites**.

1.4.10 Appendix

1.4.10.1 IP list

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

IP List of Physical Machine

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The hostname of the server.
IP	The IP address of the server.

IP List of Docker App

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Service Role	The server role name.
Machine Name	The hostname of the server.
Docker host	The Docker hostname.
Docker IP	The Docker IP address.

1.4.10.2 Info of project component report

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

Item	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Service Role	The name of a server role in the service.
Service Role Status	The running status of the server role on the server.
Service Role Action	The action that the server role performs on the server. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the server.
IP	The IP address of the server.
Server Status	The running status of the server.
Server Action	The action that Apsara Infrastructure Management Framework asks the server to perform, such as the clone action.

1.4.10.3 Auto healing - install approval pending report

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

1.4.10.4 Machine info report

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

Machine Status

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

Item	Description
Machine Name	The server name.

Item	Description
IP	The IP address of the server.
Server Status	The server status.
Server Action	The action currently performed by the server.
Server Action Status	The action status.
State Desc	The description about the server status.

Expected SR List of Machine

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

Item	Description
Machine Name	The server name.
Service Role	The name of the server role that is expected on the server.

Abnormal Monitoring Status

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

Item	Description
Machine Name	The server name.
Monitor	The name of the monitoring item.
Level	The level of the monitoring item.
Description	The description about the monitoring item contents.
Updated At	The updated time of the monitoring item.

SR's Version and Status of Machine

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

Item	Description
Machine Name	The server name.
Service Role	The server role name.
Service Role Status	The status of the server role.
Target Version	The expected version of the server role on the server.

Item	Description
Current Version	The current version of the server role on the server.
State Desc	The description about the status.
Error Message	The exception information of the server role.

Monitor of Machine

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

Item	Description
Machine Name	The server name.
Service Role	The server role name.
Monitor	The name of the monitoring item.
Level	The level of the monitoring item.
Description	The description about the monitoring item contents.
Updated At	The updated time of the monitoring item.

1.4.10.5 Rolling info report

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

Choose a Job

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

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

SR in Job

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

Item	Description
Service Role	The server role name.
Service Role Status	The rolling status of the server role.
Error Message	The exception information of the rolling task.
Git Version	The version of change to which the rolling task belongs.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Approve Rate	The proportion of servers that have the rolling task approved by the decider.
Failure Rate	The proportion of servers that have the rolling task failed.
Success Rate	The proportion of servers that have the rolling task succeeded.

SR Rolling Build Info

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

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

SR State in Cluster

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

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

1.4.10.6 Machine RMA approval pending list

Some Apsara Infrastructure Management Framework actions (such as restart) on servers and server roles can be triggered by users, but this type of actions must be reviewed and approved.

This report is used to process the actions that must be reviewed and approved.

- **Machine:** the basic information of the server.

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

- **Machine Serverrole:** the information of server roles on servers.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the server.

Item	Description
IP	The IP address of the server.
Serverrole	The server role name.
State	The running status of the server role.
Action Name	The action on the server role.
Action Status	The status of the action on the server role.
Actions	The approval button.

- **Machine Component:** the hard disk information of servers.

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

1.4.10.7 Registration vars of service

This report displays values of all service registration variables.

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

1.4.10.8 Virtual machines map

Use the global filter to display the virtual machines of a specific cluster in the **VM Mappings** section.

VM Mappings displays the information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.

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

1.4.10.9 Service inspector report

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

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

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

1.4.10.10 Machine power on or off state of cluster

After the cluster power-on/power-off operation is triggered, you can read the related information in this report.

- **Cluster Running State:** If a cluster is performing the power-on/power-off operation, the corresponding data is available in this list. No data indicates that the power-off operation is not performed by any cluster.

Item	Description
Project	The project name.

Item	Description
Cluster	The cluster name.
Machine Live State	The power-on/power-off operation that is being performed by the cluster.

- **SR Power On or Off State:** the power-on/power-off status of the server roles in the cluster selected in the **Cluster Running State** section.

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

Item	Description
Cluster	The cluster name.
Service Role	The server role name.
Machine Live State	The power-on/power-off status of the server role.

- **State on Machine:** displays the running status of the selected server role on the server.

Select a row in the **SR Power On or Off State** section to display the information of the corresponding server role in the list.

Item	Description
Cluster	The cluster name.
Service Role	The server role name.
Machine Name	The server name.
Service Role Status	The running status of the server role.
Service Role Action	The action currently performed by the server role.
Service Role Action Status	The action status.
Error Message	The exception information.

- **Machine State:** displays the running status of servers in the selected cluster.

Select a row in the **State on Machine** section to display the information of the corresponding server in the list.

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

1.4.10.11 Resource apply report

In the global filter on the top of the page, select the project, cluster, and server from the **project**, **cluster**, and **machine** drop-down lists and then click **Filter** on the right to display the corresponding resource application data.

Commit List

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

Commit Resource Map

Item	Description
Res	The resource ID.
Type	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.

Item	Description
Parameters	The resource parameters.
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource State

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

1.4.10.12 State of project component

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

Error State Component Table

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

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

Machine SR Monitor Info

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

Item	Description
Cluster	The cluster name.
Service	The service name.
Service Role	The server role name.
Machine Name	The server name.
Monitor	The monitoring name of the server role.
Level	The alarm level.
Description	The description about the alarm contents.
Updated At	The updated time of the alarm information.

Machine Monitor Info

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

Item	Description
Cluster	The cluster name.
Machine Name	The server name.
Monitor	The monitoring name of the server role.
Level	The alarm level.
Description	The description about the alarm contents.
Updated At	The updated time of the alarm information.

Service Inspector Info

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

Item	Description
Cluster	The cluster name.
Service	The service name.
Service Role	The server role name.
Monitor	The monitoring name of the server role.
Level	The alarm level.
Description	The description about the alarm contents.
Updated At	The updated time of the alarm information.

1.4.10.13 Relationship of service dependency

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

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Service Role	The server role name.

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

1.4.10.14 Check report of network topology

This report checks if wirecheck alarms are generated for the network devices and servers.

- **Check Report of Network Topology:** checks if wirecheck alarms are generated for network devices.

Item	Description
Cluster	The cluster name.
Network instance	The network device name.
Level	The alarm level.
Description	The description about the alarm information.

- **Check Report of Server Topology:** checks if wirecheck alarms are generated for servers (machines).

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

1.4.10.15 State of machine clone

This report displays the server clone status.

The Progress of Machine Clone

Item	Description
Project	The project name.

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

The State of Machine Clone

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The server name.
Server Action	The action performed by the server, such as the clone action.
Server Action Status	The status of the action performed by the server.
Server Status	The running status of the server.
Level	Whether the clone action performed by the server is normal.
State of Machine Clone	The current status of the clone action performed by the server.

1.4.10.16 Action of machine SR

Apsara Infrastructure Management Framework manages information of all servers that are performing the Apsara Infrastructure Management Framework actions, such as the clone action. If the server is a host, you can view the virtual machine status on the server and the server role status on the virtual machine.

Action of Machine SR

Only displays the servers with actions.

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The server name.
IP	The IP address of the server.

Item	Description
Server Status	The running status of the server.
Server Action	The action performed by the server, such as the clone action.
Service Role	The server role name.
Service Role Status	The running status of the server role.
Service Role Action	The action performed by the server role on the server, such as the rolling, restart, and offline actions.

VM SR Action on Host

Select a row in the **Action of Machine SR** section to display the corresponding information in the list.

Item	Description
VM	The virtual machine name.
IP	The IP address of the virtual machine.
Server Status	The running status of the virtual machine.
Server Action	The action performed by the virtual machine, such as the clone action.
Service Role	The server role name.
Service Role Status	The running status of the server role.
Service Role Action	The action performed by the server role on the server, such as the rolling, restart, and offline actions.

1.5 Network operations

1.5.1 Apsara Network Intelligence

1.5.1.1 What is Apsara Network Intelligence

Apsara Network Intelligence is a system designed for network traffic analysis. It provides data to facilitate resource planning, diagnostic functions, monitoring, system management, and user behavior analysis.

Apsara Network Intelligence allows you to:

- Manage cloud service types.
- Query SLB and VPC instance details with a single click.

- Implement reverse access to cloud services.
- Configure leased lines through Web interfaces and set up active and standby routers
- Query tunnel VIPs of cloud services
- Create layer-4 listeners

1.5.1.2 Log on to the Apsara Network Intelligence console

This topic describes how to log on to the Apsara Network Intelligence console.

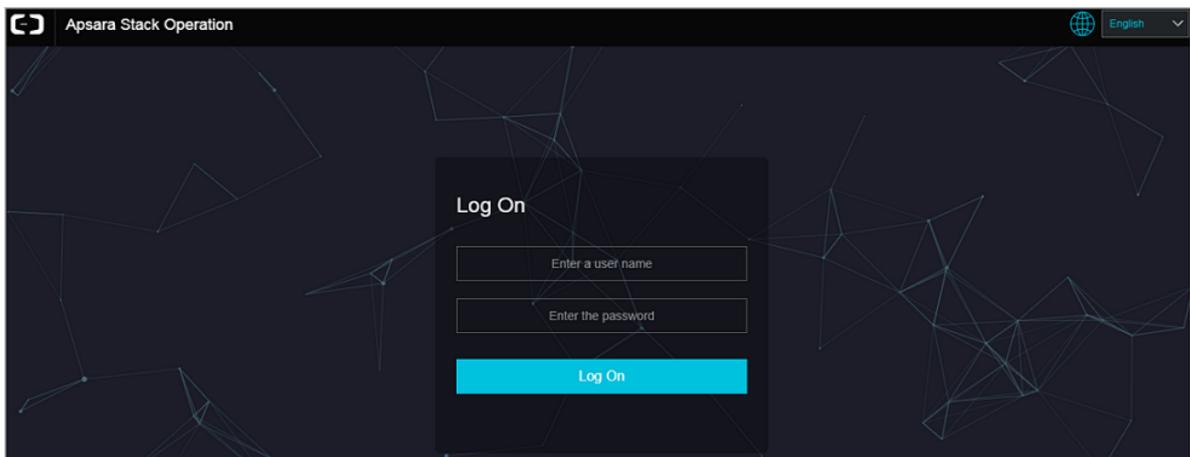
Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 1-8: Log on to ASO



Note:

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.

- System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.
 5. In the left-side navigation pane, click **Products**. On the right side of the Web page, click **Apsara Network Intelligence**.

1.5.1.3 Query information

You can enter an instance ID to query VPC, VRouter, and VSwitch details.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. Enter a VPC or SLB instance ID to query instance details.
 - Enter a VPC instance ID to query VPC, VRouter, and VSwitch details.
 - Instance details
 - Information about VRouters, routing tables, router interfaces, and VSwitches
 - Enter an SLB instance ID to query instance details.
 - Information about instance configurations, VIPs, specifications, and users
 - Listener information

Click **Show** in the **Backend Server/Health Check** column to view backend server details.

1.5.1.4 Manage cloud service instances

You can create a cloud service in a region or query the instance information of a region.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Virtual Private Cloud > VPC Instance Type Management**.
3. Select a region to which a target cloud service instance belongs from the **Select Region** drop-down list. All cloud service instances in the specified region are displayed.

4. Click **Add** to add a cloud service type.

1.5.1.5 Tunnel VIP

1.5.1.5.1 Apply for layer-4 listener VIPs

You can apply for layer-4 listener VIPs for cloud services in your VPC to allow traffic forwarding.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Server Load Balancer > VIP Management**.
3. Click **Create VIP**.
4. On the **Create VPC Cloud Instance** tab page, select Cloud Service, CIDR Type, and Tunnel Type.

The values of Tunnel Type are listed as follows:

- **singleTunnel**: indicates a single tunnel VIP that allows ECS instances in a single VPC to access external cloud services.
 - **anyTunnel**: indicates a tunnel VIP that allows ECS instances in all VPCs to access a specified cloud service.
5. Click **Create**. On the **Create LB Instance** tab page, select a Primary Data Center or use the default data center.
 6. Click **Create**. On the **Add Backend Server to LB Instance** tab page, configure the following parameters as needed:
 - **VPC ID**: indicates the ID of the VPC to which target ECS instances belong. This parameter must be configured when the network type of the ECS instances is VPC.
 - **Backend Servers**: indicates the backend servers that you want to add. You can only enter the information of one backend server on each line. A backend server information entry contains server IP address and weight. You can separate IP addresses and weight values with either a space or a comma (,). If no weight value is specified, the default value is 100.
 7. Click **Create**. On the **Create LB Instance** tab page, select a Primary Data Center or use the default data center.
 8. Click **OK**. On the **Create Listener** tab page, click **Add** to configure a UDP or TCP listener. Then, click **Submit**.
 9. On the **Publish Online** tab page, click **Yes** and click **OK**.

Result

The cloud services for which you have applied for VIPs can forward traffic through the created layer-4 listener.

1.5.1.5.2 Query the tunnel VIP of a cloud service

You can query information about cloud services that have Server Load Balancer (SLB) VIPs, such as creation time, connectivity, and VIP.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, select **Server Load Balancer > VIP Management**.
3. On the **Tunnel VIP Management** page, configure Region ID, Cloud Service, and Status, and click **Search**.

1.5.1.6 Apply for Direct Any Tunnel VIPs

You can apply for Direct Any Tunnel VIPs for cloud services in your VPC to allow traffic forwarding through XGW.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Server Load Balancer > Direct Any Tunnel VIP Management**.
3. On the **Direct Any Tunnel VIP Management** page, click **Create Direct Any Tunnel VIP**.
4. On the **Create Direct Any Tunnel VIP** page, configure the parameters for the Direct Any Tunnel VIP.
5. Click **Create**. Cloud service instances that have Direct Any Tunnel VIPs can forward traffic through XGW.

1.5.1.7 Leased line connection

1.5.1.7.1 Overview

You can connect a VPC to an IDC through a leased line.

Before you connect to a VPC through a leased line, you must confirm initial CSW configurations meet the following conditions:

- You have downloaded licenses required for VLAN functions onto the CSWs.
- You have set management IP addresses of the CSWs to loopback IP addresses on the loopback 100 interface.

- You have configured CSW uplink interfaces to ensure compatibility with the layer-3 interfaces used by VPC APIs.
- You have deleted the default configuration of bridge-domain.
- You have enabled NETCONF and STelnet for CSWs. The configuration details are described in the CSW initial configuration template.
- You have configured the service type of CSW interfaces to tunnel.

You also need to obtain the following account information:

- BID: indicates the ID of an account group. The BID for Mainland China users is 26842, and the BID for international users is 26888.
- UID: indicates the ID of the account to which a target VPC belongs.

1.5.1.7.2 Manage access points

Access points are Alibaba Cloud data centers in different regions. There are one or more access points in each region. This topic describes how to query and modify access point information of a region.

Query access point information

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Express Connect > Daily Operation and Maintenance Management**.
3. Enter Region and Access Point ID of an access point that you want to query.
4. Click **Search**.

Modify access point information

Perform the following operations to modify access point information:

1. Click **Modify** in the **Actions** column.
2. Modify access point information.
3. Click **Modify**.

The parameters are described as follows:

- **Access point location:** indicates the physical location of an access point. You can customize this parameter.
- **Access point machine operator:** indicates the operator name.

1.5.1.7.3 Manage access devices

This topic describes how to query and modify access device information of a region.

Query access device information

Perform the following operations to query access device information:

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Express Connect > Daily Operation and Maintenance Management**.
3. Click **Access Device Management**.
4. Enter the region and device ID of an access device that you want to query.

**Note:**

If Device ID is not entered, information about all devices in a region is queried.

5. Click **Search**.
6. Click **Show Details** in the **Actions** column to view access device details.

Modify access device information

Perform the following operations to modify access device information:

1. Click **Modify** in the **Actions** column.
2. Follow the on-screen prompts to modify device information.
3. Click **Modify**.

1.5.1.7.4 Establish leased lines

A leased line can be obtained from a telecom operator to establish a physical connection between your local IDC and an Alibaba Cloud access point. This topic describes how to establish a leased line and query leased line information of a region.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Express Connect > Network Environment Management**.
3. Select **Network Environment Management > Physical Connection Management**.
4. Follow the on-screen prompts to configure the leased line information and click **Create**.

The parameters are described as follows:

- **Device Name**: optional. If specified, the device name must be the CSW host name.

- **Device Port:** optional. If specified, the device port must be the CSW port number.
- **UID:** the ID of the account to which a target VPC belongs.
- **Access Point ID:** the ID of the region where your IDC is located.
- **Redundant Leased Lines:** a leased line, previously applied for, to act as a redundancy for the connection you are creating.

If the leased line status is **Confirmed**, the line is created.

5. On the **Physical Connection Management** page, locate the created leased line and select **Actions > Enable**.

If the allocation process for a leased line persists for several minutes after you click Enable, choose **Network Controller > Business Foundation System Flow** from the Products menu. On the displayed page, set Instance ID to the leased line ID, set **Step Status** to **All**, and click Search. A flow in red indicates that a corresponding step has failed. Click **Resend** to try again, and then requery the flow status.

If the flow fails, run the `vpcregiondb -e "select * from xnet_publish_task order by id desc limit 5"` command on the ECS AG. If an error is reported, you can check the xnet service logs to troubleshoot the issue based on the reported errors.

1.5.1.7.5 Create VBRs

A VBR is a router between customer-premises equipment (CPE) and a VPC, and functions as a data forwarding bridge from a VPC to a local IDC. This topic describes how to create a VBR in a region and query VBR information of the region.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Express Connect > Network Environment Management**.
3. Select **Network Environment Management > VBR Management**.
4. Click **Create VBR**.
5. Follow the on-screen prompts to configure the VBR parameters.

The parameters are described as follows:

- **Leased Line ID:** indicates the ID of the leased line that a VBR connects.
- **VLAN ID:** indicates the VLAN ID of a VBR. Value range: 0–2999.

When creating router interfaces, you can use VLAN IDs to identify subsidiaries or departments that use the leased line, thus implementing layer-2 network isolation between them.

- **Local Gateway IP:** indicates the local IP address of the router interface for the leased line.
- **Peer Gateway IP:** indicates the peer IP address of the router interface for the leased line.
- **Subnet Mask:** indicates the subnet mask for the leased line between the local IP address and peer IP address.

Only two IP addresses are required. Therefore, you can enter a longer subnet mask.

6. Click **Create**.

The created VBR status is **Active**.

You can click **Release**, **Modify**, **Terminate**, or **Show Details** in the **Actions** column to manage a VBR.

1.5.1.7.6 Create router interfaces

After you create a VBR, you must create a pair of router interfaces to connect the VBR and VPC. The connection initiator must be a VBR.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Express Connect > Network Environment Management**.
3. Select **Network Environment Management > Route Interface Management**.
4. Click **Create Router Interface**.
5. Configure router interface parameters and click **Submit**.

Set **Create Router Interface** to **Double**. Configure the local router interface based on the created VBR information, and configure the peer router interface based on the target VPC information.

When the router interface status is **Active**, it has been successfully created.

1.5.1.7.7 Create routing tables

A routing table is a list of route entries on a VRouter. This topic describes how to create routing tables in a region and query the routing table information of a region.

Procedure

1. Perform the following operations to add routes on a virtual border router (VBR) destined for a VPC and an IDC :
 - a) [Log on to the Apsara Network Intelligence console](#).
 - b) From the **Products** menu, choose **Express Connect > Network Environment Management**.
 - c) Select **Network Environment Management > Routing Table Management**.
 - d) Set Region, BID, UID, Router Type, Routing Table ID, or Router ID, and click **Search** to query routing tables.
 - e) Click **Add Route Entry** in the **Actions** column.
 - f) Add a route entry destined for the CIDR block of a target VPC, and click **Create**.

The parameters are described as follows:

- **Destination CIDR Block:** the destination CIDR block.
 - **Next Hop Type:** the next hop type.
 - **Next Hop Instance ID:** the ID of a next hop instance to receive traffic based on the next hop type.
- g) Repeat the preceding steps to add a route destined for a target IDC.



Note:

You can navigate to the VBR Management page and locate the **VLAN Interface ID** area to obtain next hop router interface information.

2. Add a route destined for the router interface of a VBR in the VPC.
3. Configure a route on the local IDC gateway destined for the target VPC.

1.5.1.8 Manage Business Foundation System flows in a VPC

You can view operation execution status of a VPC and restart tasks.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. From the **Products** menu, choose **Network Controller > Business Foundation System Flow**.
3. Query the flow status of the task you want to view.

Enter a leased line ID in **Instance ID** and set **Step Status** to **All** to check flow status. A flow in red indicates that a corresponding step has failed. Click **Resend** to try again, and then requery the flow status.

1.5.1.9 Configure reverse access to cloud services

Cloud services cannot be accessed directly through external networks. Therefore, you must configure reverse access to allow external networks to access cloud services through ECS instances.

Prerequisites

Log on to Apsara Stack Management Console. Navigate to the **Personal Information** page and obtain **AccessKey ID** and **AccessKey Secret**.

Procedure

1. [Log on to the Apsara Network Intelligence console](#).
2. In the Apsara Network Intelligence console, enter **AccessKey ID** and **AccessKey Secret**, and click **OK**. The **Cloud Service Reverse Access Management** page is displayed.
3. Click **Create Cloud Service Reverse Access**.
4. On the **Allocate App ID** tab page, set Region, Name, and Description.
5. Click **Continue**. The following information is automatically created and displayed on the **Create Address Pool** tab page: the app IDs of cloud services that allow reverse access and the address pools that are used for reverse access to the cloud services.
6. Click **Continue**. On the **Add Server Address** tab page, configure an ECS instance to be used for reverse access.
 - **VPC ID**: indicates the ID of a VPC, ECS instance, or cloud service instance whose Tunnel Type is single tunnel.
 - **Server IP**: indicates the IP address of the ECS instance that is used for reverse access.
7. Click **Continue**. On the **Create Mapping IP** tab page, configure VSwitch ID to Mapping IP of the ECS instance of the target VPC.
8. Click **Continue**. On the **Complete Authorization** tab page, configure VPC ID, ECS Instance IP, and Instance Port for reverse access.

An Instance Port is defined as an integer value. You can separate multiple values using commas (,). For example, 10,20,30. You can configure up to 10 instance ports.

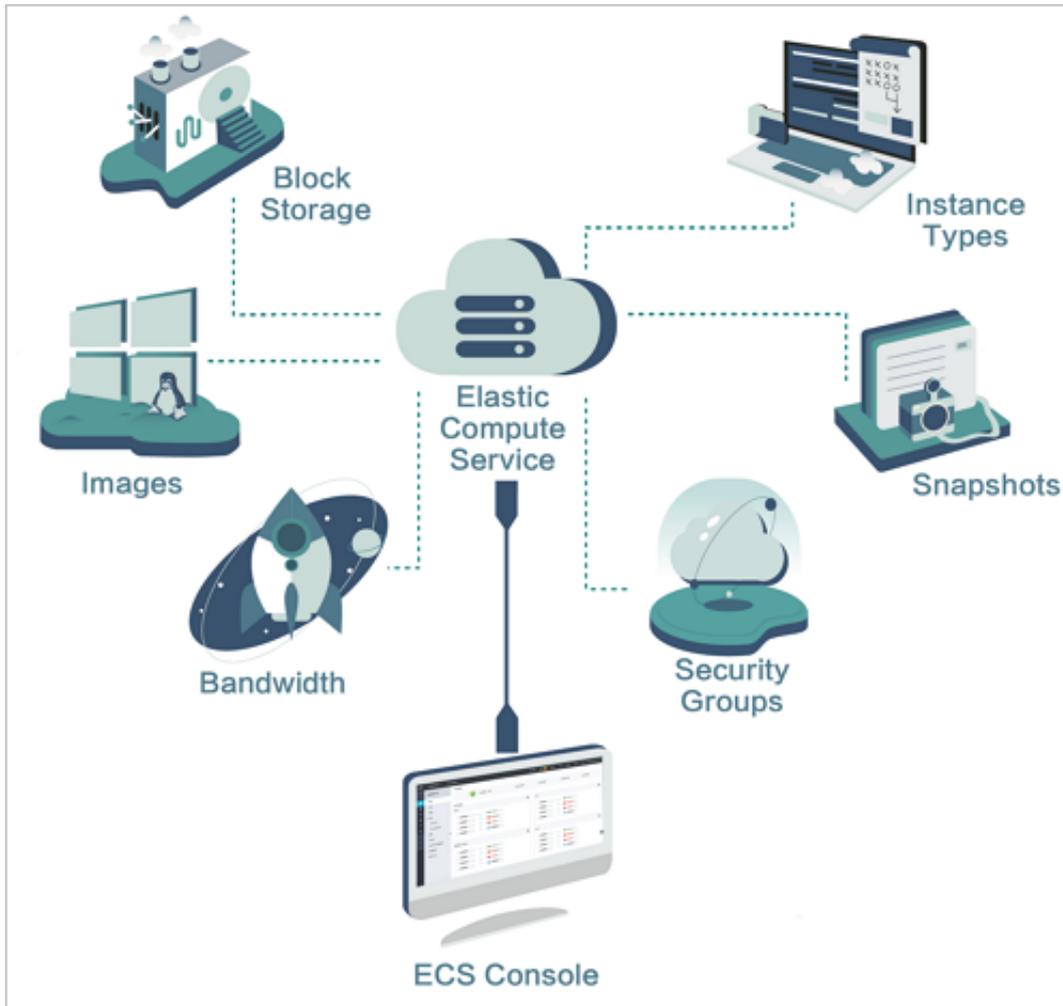
2 Cloud product operations

2.1 Elastic Compute Service (ECS)

2.1.1 ECS overview

Elastic Compute Service (ECS) is a type of computing service that features elastic processing capabilities. Compared with physical servers, ECS is more user-friendly and can be managed more efficiently. You can create instances, resize disks, and add or release any number of ECS instances at any time based on your business demands.

An ECS instance is a virtual computing environment made up of the basic components such as the CPU, memory, and storage. Users perform operations on ECS instances. It is the core concept of ECS and you can perform operations on ECS instances through the ECS console. Other resources such as block storage, images, and snapshots can be used only after they are integrated with ECS instances. For more information, see [Figure 2-1: Concept of an ECS instance](#).

Figure 2-1: Concept of an ECS instance

2.1.2 Log on to Apsara Stack Operations

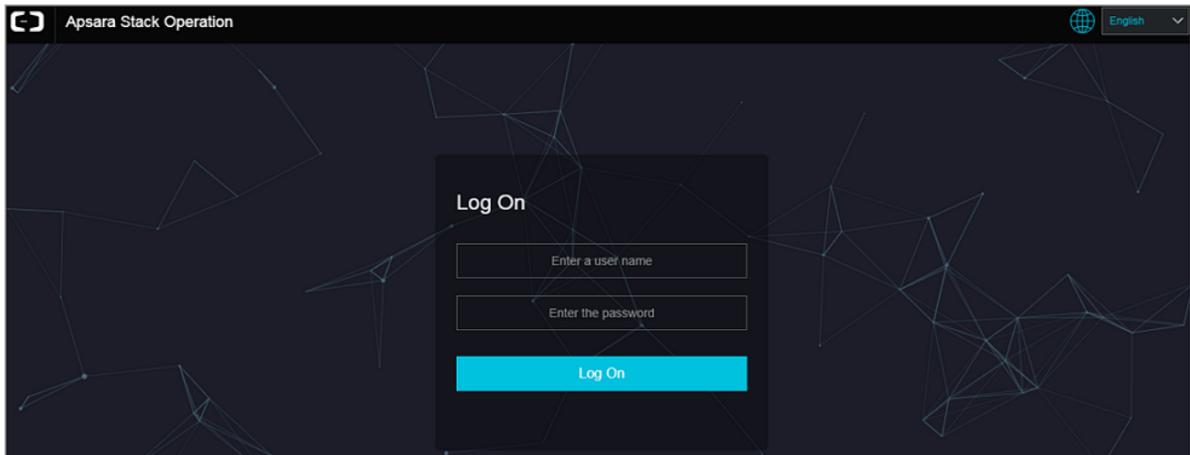
This section describes how to log on to Apsara Stack Operations (ASO).

Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 2-2: Log on to ASO**Note:**

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.

2.1.3 ECS Operations and Maintenance System

2.1.3.1 Overview

The ECS Operations and Maintenance System is an operations and maintenance platform designed for Apsara Stack. Operations and maintenance engineers can use this system to

operate and monitor ECS instances, help users solve problems, and ensure smooth operation and usage of ECS instances.

2.1.3.2 VMs

2.1.3.2.1 Overview

On the ECS Operations and Maintenance System, the existing ECS VM information and executable operations and maintenance functions are displayed. You can search for, start, and migrate a VM as needed.

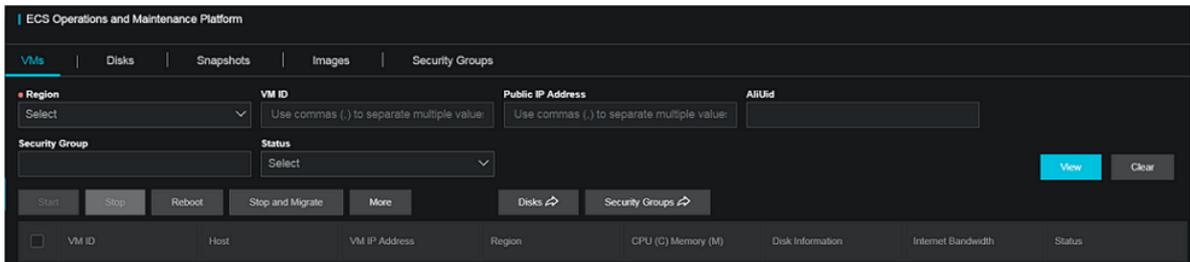
2.1.3.2.2 Search for VMs

On the Apsara Stack Operations Console, you can view the list of existing VMs and their related information.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose **Products > ECS**.

You have logged on to the **ECS Operations and Maintenance System**.



3. On the **VMs** page, enter the filter criteria, click **Search** to filter the VMs that need to be managed.

The **Region** item is a required filter condition.

4. Click VM ID and the **VM Details** dialog box appears on the right side of the page.

2.1.3.2.3 Start a VM

On the Apsara Stack Operations Console, you can start a VM as if it were a real server.

Prerequisites

The VM is currently in **stopped** state.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose **Products > ECS**.

3. On the **ECS Operations and Maintenance System** page, enter the filter criteria, and click **Search** to filter the VMs that need to be managed.
4. On the VM list, select those you want to start and click **Start** above the list.
5. On the displayed page, set **start Mode** to **Normal** or **Repair**.

**Note:**

If you need to reset the network settings for the VM, select the **repair mode** when you start the VM. Otherwise, select the **normal mode**.

6. Enter the **operation reason** and click **OK**.

2.1.3.2.4 Stop a VM

On the Apsara Stack Operations Console, you can stop a VM as if it were a real server.

Prerequisites

- The VM is currently in the **running** state.
- This operation may result in interruption of the program running on the VM. Perform this operation at off-peak periods where the impact on services is the lowest.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, enter the filter criteria, and click **Search** to filter the VMs that need to be managed.
4. On the VM list, select those you want to stop and click **Stop** above the list.
5. On the displayed page, set the **shutdown Policy** to **Normal Shutdown** or **Forced Shutdown**.

**Note:**

When **Forced Shutdown** is selected, the VM is shut down regardless of whether its processes are stopped. We recommend that you do not select **Forced Shutdown** unless **Normal Shutdown** does not work.

6. Enter the **operation reason** and click **OK**.

2.1.3.2.5 Restart a VM

You can restart a VM on the Apsara Stack Operations Console. The operations are similar to the operations on a real server.

Prerequisites

- The VM is currently in the `running` state.
- This operation may result in interruption of the program running on the VM. Perform this operation at off-peak periods where the impact to service is the lowest.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, enter the filtering criteria and click **Search** to filter the VMs that need to be managed.
4. On the VM list, select the VMs you want to restart and click **Restart** above the list.
5. In the **VM Restart** dialog box that appears, specify **Start Mode** and **Shutdown Policy** as needed.



Note:

- For the **Start Mode**, you can select **Normal Mode** or **Repair Mode**.
- For the **Shutdown Policy**, you can select **Normal Shutdown** or **Forced Shutdown**.

6. Set `operation reason` and click **OK**.

2.1.3.2.6 Downtime migration

On the Apsara Stack Operations Console, you can perform downtime migration on a VM.

Prerequisites

- Downtime migration is cold migration. Check that the VM is in `stopped` state before starting downtime migration.
- Downtime migration has to be performed within the same region. Cross-region migration is not allowed.

Context

In case of a faulty VM or NC, you need to shut down the VM and migrate it from one NC to another. This is an example of downtime migration.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, enter the filter criteria and click **Search** to filter the VMs that need to be managed.
4. In the list of VMs, select the VMs that need to be migrated and click **Downtime Migration**.
5. In the dialog box that appears, select **NC Switching** and then select **Switching Policy**, **Start Mode**, and **Recovery Mode** respectively.
6. Enter the **operation reason** and click **OK**.

2.1.3.2.7 Hot migration

On the Apsara Stack Operations Console, you can perform hot migration of a VM.

Context

- Due to the high load on the current NC or other business considerations, you may need to migrate the VM in the **running** state from one NC to another. This is called hot migration. If a failure occurs, you need to perform [downtime migration](#).
- Hot migration is a high-risk operation. Exercise caution when you perform hot migration.
- Services will not be interrupted during hot migration.

Prerequisites

- Hot migration has to be performed within the same available zone, and cross-zone migration is not available.
- The target VM of hot migration must be in the **running** state.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, enter the filtering criteria and click **Search** to filter the VMs that need to be managed.
4. From the VM list, select the target VM for hot migration, and then choose **More Operations > Hot Migration**.
5. In the dialog box that appears, select **NC Switching**, and then set the traffic limit.

**Note:**

The traffic range is 1–1,000 Mbit/s, and the default value is 20 Mbit/s.

6. Set `operation reason`, and click **OK**.

2.1.3.2.8 Reset a disk

You can reset a disk to restore the disk to the initial state as required.

Prerequisites

- If any applications are installed after a VM is created, these applications will be lost. Before this operation, back up your data properly.
- The VM whose disk is to be reset must be in the `stopped` state.

Context

Resetting a disk does not result in disk formatting. It just restores the disk to the initial state. If an image is used during the creation of a disk, the image still exists after the disk is reset.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. The **ECS Operations and Maintenance System** page, enter the filter criteria and click **Search** to filter the VMs that need to be managed.
4. On the VM list, select the VM whose disk is to be reset.
5. Choose **More Operations > Reset Disk**.
6. Select the disk to be reset. Set `operation reason` and click **OK**.

2.1.3.3 Disks

2.1.3.3.1 Overview

Cloud disks can be considered as physical disks in an ECS instance. You can mount and unmount a disk and create snapshots for it.

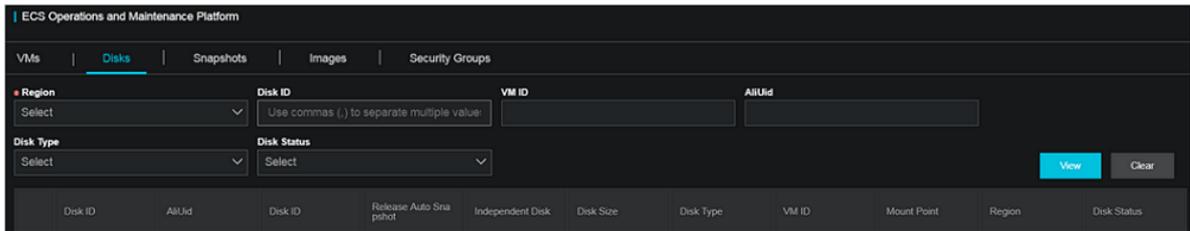
2.1.3.3.2 Search for disks

You can view the list of existing disks and their related information on the Apsara Stack Operations Console.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Disks** tab.
4. On the **Disks** page, enter the filter criteria and click **Search** to filter the specified disks.

Region is a required filter condition.



2.1.3.3.3 View snapshots

You can view the list of existing snapshots and their information on the Apsara Stack Operations Console.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Disks** tab.
4. Enter the filtering criteria and click **Search** to filter the specified disks.
5. Click the  icon on the left side of a disk and select **View Snapshots**.

The information of all snapshots on the disk is displayed on the page.

2.1.3.3.4 Mount a disk

After a disk is created, you need to mount the disk. You can only mount independent cloud disks to ECS instances.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Disks** tab.
4. Enter the filtering criteria and click **Search** to filter the specified disks.
5. On the left side of a disk, click the  icon and select **Mount** from the drop-down list.
6. In the dialog box that appears, enter the **VM ID** and **operation reason**, and then click **Confirm**.

2.1.3.3.5 Unmount a disk

You can only unmount data disks on the Apsara Stack Operations Console. You cannot unmount system disks or local disks.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Disks** tab.
4. Enter the filtering criteria and click **Search** to filter the specified disks.
5. Click the  icon on the left side of a disk and select **Unmount**.
6. In the dialog box that appears, enter the **operation reason** and click **OK**.

2.1.3.3.6 Create a snapshot

You can manually create a snapshot of a disk as needed on the Apsara Stack Operations Console.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Disks** tab.
4. Enter the filtering criteria and click **Search** to filter the specified disks.
5. Click  on the left side of a disk and select **Create Snapshot**.
6. In the dialog box that appears, enter the **snapshot name**, **snapshot description**, and **operation reason**. Click **OK**.

2.1.3.4 Snapshots

2.1.3.4.1 Overview

Snapshots can save the state of disk data at a certain point in time for data backup or custom image creation.

Note the following points when using disks:

- Use the data on a disk as the basic data of another disk when writing or saving data to that disk .
- Although the disk provides a secure storage mode, make sure that the stored data is complete . However, if the data stored on the disk is incorrect (for example, due to an application error, or

a vulnerability in the application was exploited for malicious uses), a mechanism is required to ensure that your data can be recovered to the desired state when it encounters a problem.

Alibaba Cloud allows you to create a snapshot to retain a copy of the data on a disk at a point in time. You can create disk snapshots on a scheduled basis to guarantee continuous operation of your business.

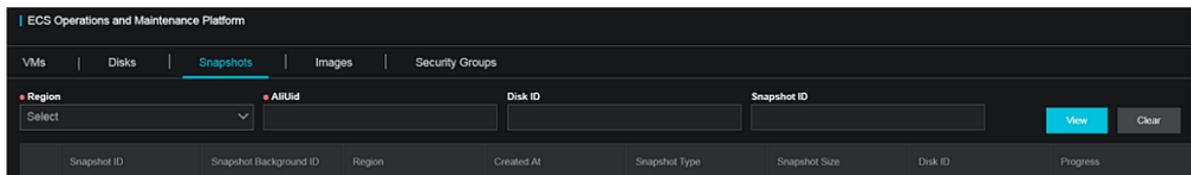
2.1.3.4.2 Search for snapshots

You can view the list of existing snapshots and their related information on the Apsara Stack Operations Console.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products** > **ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Snapshots** tab.
4. On the **Snapshots** page, enter the filter criteria and click **Search** to filter the specified workflows.

Region and **AliUid** are required filter conditions.



2.1.3.4.3 Delete a snapshot

On the Apsara Stack Operations Console, you can delete a snapshot that is no longer used.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products** > **ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Snapshots** tab.
4. Enter the filtering criteria and click **Search** to filter the specified snapshot.
5. Click the  icon on the left side of a snapshot and select **Delete**.
6. In the dialog box that appears, enter the **operation reason** and click **OK**.

2.1.3.4.4 Create an image

You can create a custom image with a snapshot, which contains the operating system and data environment information of the snapshot in the image.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Snapshot** tab.
4. Enter the filtering criteria and click **Search** to filter the specified snapshots.
5. On the left side of a snapshot, click  and select **Create Image** from the drop-down list.
6. In the dialog box that appears, specify **Image Name**, **Image Version**, **Image Description** and other information, and then click **OK**.

2.1.3.5 Images

2.1.3.5.1 Overview

An ECS image is a template that contains the software configurations such as the operating system, application server, and application programs of the ECS instance. To create an instance, you need to specify an ECS image. The operating system and software provided by the image will be installed on the instance that you create.

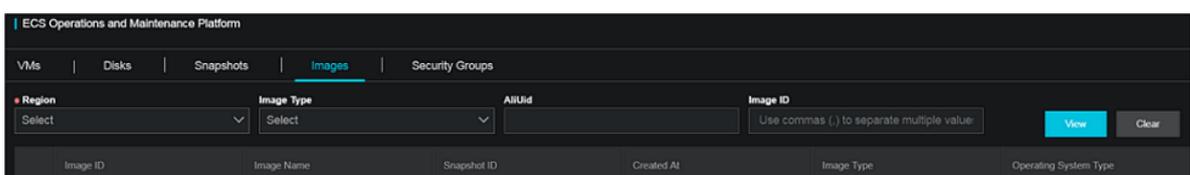
2.1.3.5.2 Search for images

You can view the list of existing mirrors and their information on the Apsara Stack Operations Console.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Images** tab.
4. On the **Images** page, enter the filtering criteria and click **Search**.

The **Region** item is a required filtering condition.



2.1.3.5.3 Delete images

You can delete images that are no longer used on the Apsara Stack Operations Console.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products** > **ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Images** tab.
4. On the **Images** page, enter the filter criteria and click **Search** to filter the specified mirrors.
5. Click the  icon on the left side of an image and select **Delete**.
6. In the dialog box that appears, enter the **operation reason** and click **OK**.

2.1.3.6 Security groups

2.1.3.6.1 Overview

A security group is a virtual firewall that provides Stateful Packet Inspection (SPI). The security group provides virtual firewall-like functionality and is used for network access control for one or more cloud servers. It is an important means of network security isolation, and is used to divide security domains on the cloud.

Security group rules can allow or deny inbound or outbound Internet and intranet traffic for ECS instances.

You can authorize or cancel security group rules at any time. Changes in security group rules are automatically applied to ECS instances associated with the security group.

Make sure the rules are concise when you configure security group rules. If you associate an instance with multiple security groups, hundreds of rules may apply to the instance, which may cause connection errors when you access the instance.

2.1.3.6.2 Search for security groups

On the Apsara Stack Operations Console, you can view the list of current security groups and their information.

Context

You can modify security group rules to allow or deny inbound or outbound traffic of the Internet and intranet for ECS instances associated with security groups. You can authorize or cancel security group rules at any time. Changes in security group rules are automatically applied to ECS instances associated with security groups.

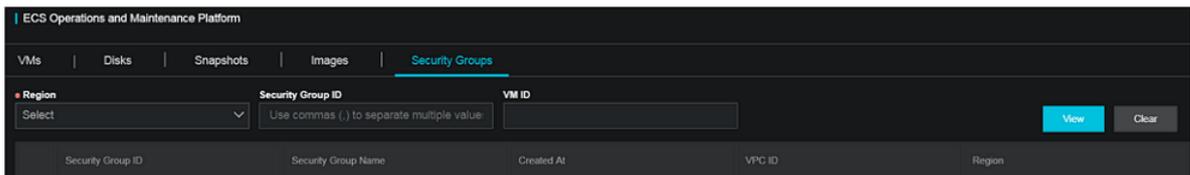
**Note:**

- If two security groups have identical rules but different access rules, access deny rules are valid and access allow rules are invalid.
- No security group rule allows outbound access while denying inbound access to ECS instances.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Security Groups** tab.
4. On the **Security Groups** page, enter the filtering criteria and click **Search**.

The **Region** item is a required filtering condition.



5. Click **Authorize** to authorize a rule for this security group.

2.1.3.6.3 Add security group rules

You need to add appropriate security group rules to security groups as needed.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance System** page, click the **Security Groups** tab.
4. On the **Security Groups** page, set the filtering criteria and click **Search**.
5. On the left side of a security group, click the  icon and select **Add Rules** from the drop-down list.
6. In the dialog box that appears, configure the relevant parameters.

[Table 2-1: Security group rule parameters](#) describes the parameter configurations.

Table 2-1: Security group rule parameters

Parameter	Description
Protocol	ALL TCP UDP ICMP GRE.
Rule priority (1–100)	The smaller the number, the higher the priority.
Network type	Internet Intranet.
Authorization policies	Accept access Abandon the packet on access Deny the packet on access.
Port range	1 to 65535, for example, 1/200, 80/80, and -1/-1.
Access direction	In Out.
IP address segment	If the authorization type is address segment access, enter the IP address or CIDR segment in the authorized object field, such as 10.0.0.0, 0.0.0.0/0, or 192.168.0.0/24. Only IPv4 addresses are supported.
Associate security group IDs	Enter the ID of the associated security group.
Operation reason	Optional. You can enter related operation reasons.

7. When the parameter configuration is completed, click **OK**.

2.1.4 VM hot migration

2.1.4.1 Overview

During hot migration, a VM in the running state is migrated from one host to another. During migration, the VM runs normally and the services within the VM do not sense the migration or can detect a very short service interruption (100–1000 ms).

Scenarios

During the operations and maintenance of the system, hot migration is typically used in three types of scenarios:

- **Active O&M:** The host encounters a fault and needs maintenance, but the fault does not affect the operation of the system. You can use hot migration to migrate the VM to another host and perform repairs on the faulty host in offline mode.
- **Load balancing:** When the load on a host is relatively high, you can migrate some VMs to other idle hosts through hot migration. This reduces resource contention on the source host.
- Other scenarios where VM migration is required without affecting the internal business operations of the VM.

2.1.4.2 Hot migration usage restrictions

Before you trigger hot migration, you need to understand the following usage restrictions.

Apsara Stack's hot migration is subject to these usage restrictions:

- You can only run the `go2hyapi` command to implement hot migration in the KVM virtualization environment. The ECS Operations and Maintenance System does not support hot migration interfaces.
- It supports only hot migration of ECS standard images. ECS provides a list of images that can be migrated. If you migrate a VM that is not included in the list of images that can be migrated, Alibaba Cloud will not be responsible for troubleshooting.
- If a VM is used as an RS to provide services to SLB or as a client to access SLB, the previous session will be closed after hot migration. New sessions created after migration are not affected.
- Migration can only be performed between hosts of the same type. Furthermore, the hosts must have the same software versions.
- Hot migration is not supported in DPDK avs scenarios.
- VMs using local storage solutions do not support hot migration. This is because after the VM is migrated to another host, it will no longer have access to the storage.
- VMs that use GPU, FPGA or other (passthrough or SRIOV) devices do not support hot migration.



Note:

VMs created in versions earlier than V3.3 do not support hot migration. Hot migration becomes available when you restart the VMs.

2.1.4.3 Complete hot migration on AG

On the Apsara Stack Operations Console, you can trigger or cancel hot migration as needed through the command line interface.

Trigger hot migration

After hot migration is successfully triggered, you can run the `go2which` command or use the ECS Operations and Maintenance System to check whether the VM state is `migrating`. When hot migration is completed, the VM is in the `running` state.

The `go2which` command output is as follows:

```
go2hyapi live_migrate_vm == Functions usage: == |- live_migrate_vm <
vm_name> [nc_id] [rate] [no_check_image] [no_check_load] [downtime]==
Usage: == houyi_api.sh <function_name> [--help|-h] [name=value]
```

Table 2-2: Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
nc_id	Designates an NC as the target of the migration.	The NC may not support the specifications of the VM, which results in a migration failure.	N/A
rate	The bandwidth consumed for migration.	The migration will use the bandwidth resources of hosts.	<ul style="list-style-type: none"> 10-Gigabit network : 80 MB 1-Gigabit network: 40 MB
downtime	The maximum allowable downtime for migration. The default value is 300 ms.	The migration service downtime is affected.	200 ms to 2000 ms
no_check_image	Forcibly migrates the images not in the list.	SLA cannot be ensured during migration.	false
no_check_load	Forcibly migrates the images even when the threshold requirements are not met.	The downtime is not controllable.	false

Cancel hot migration

Run the following command to cancel hot migration:

```
go2hyapi cancel_live_migrate_vm == Usage: == houyi_api.sh <
function_name> [--help|-h] [name=value] == Functions usage: == |-
cancel_live_migrate_vm <region_id> <vm_name>
```

Table 2-3: Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated	N/A	N/A
region_id	The ID of the region in which the VM is located	N/A	N/A

2.1.4.4 Modify the position of the NC for a VM

When an exception occurs during hot migration and the migration cannot be rolled back in the ECS Operations and Maintenance System, you can modify the VM state to trigger the rollback.

Trigger rollback

If an exception occurs during hot migration, run the following command to trigger rollback:

```
go2hyapi call_api manually_change_migration_status == Functions
usage: == |- call_api manually_change_migration_status <vm_name> <
region_id> <where>
```

Table 2-4: Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated	N/A	N/A
region_id	The ID of the region where the VM is located	N/A	N/A
where	The ID of the NC in which the VM is located	N/A	N/A

2.1.4.5 FAQs

For any problems encountered during hot migration, refer to this topic first.

- **Which parameters are required when I call the Server Controller API to perform hot migration?**

- Vm_name: VM name
- nc_id: (You do not need to designate the destination NC in the new version)

- **What preparations are required for hot migration?**

- Confirm that the VM is in the running state.
- Confirm the destination of the VM migration.

- **Can hot migration be canceled? How can I cancel hot migration?**

Yes. If the API has been executed successfully, run the `go2hyapi cancel_live_migrate_vm vm_name=[vm_name] region_id=[region_id]` command to cancel hot migration. If the VM has been migrated to the destination NC, you cannot cancel hot migration regardless of whether the VM has been started.

You can get the region_id by running the `go2which [vm_name]` command to view region_info.

- **The VM remains in the migrating state after hot migration, and the cancel_live_migrate_vm command is not working. What should I do?**

You can run the `virsh query-migrate [domid]` command on the source NC of the VM to check whether the VM is still being migrated. If the VM is still being migrated, a piece of JSON information will be returned. Then, run the following command on the AG to modify the state of the VM:

```
go2hyapi manually_change_migration_status vm_name=[vm_name] where=[nc_id for the VM] region_id=[region_id]
```

domid is the name of the VM instance. You can run the `virsh list|grep vm_name` command to view it.

- **How can I confirm whether the VM migration is successful?**

On the destination NC of the VM, run the `sudo virsh list|grep [vm_name]` command. If the VM instance exists and is not in the running state, the migration is successful.

- **What logs should I refer to when an exception occurs during hot migration?**

- View the Libvirt bottom plane migration log (NC).

Run the `/var/log/libvirt/libvirt.log` command to view information about the entire migration process, such as vport offline, detach, delete, and relay route.

- Run the following command to view the API management log of Server Controller (AG):

```
/var/log/houyi/pync/houyipync.log
```

- View the Qemu log.

- Run the following command to view the regionmaster log (VM):

```
regionmaster/logs/regionmaster/error.log
```

- **After hot migration, the VM fails to start. Is it still in the pending state?**

If **error vport update nc conf by vpc master fails dest_nc_id:xxx** is prompted, it usually indicates a problem with the VPC and that the underlying task is interrupted.

- **During hot migration, the API prompts the following error message: distributed lock fail. What are the possible causes of this issue?**

APIs are often called frequently. Try again after several minutes.

- **What are the commonly seen scenarios of migration failures? How can I fix them?**

Table 2-5: Hot migration issues

Scenario	Cause	Solution
The load is too high and the VM fails to pass pressure predetermination.	Long service interruption.	You can run <code>no_check_load=true</code> to skip this inspection.
The VM fails to pass image determination.	It is not the image designated by Alibaba Cloud.	You can run <code>no_check_image=true</code> to skip this inspection. Be aware of the risks involved.

2.1.5 Hot migration of cloud disks

2.1.5.1 Overview

Hot migration aims to facilitate operations and maintenance of online clusters, provide online migration capability for virtual disks, and improve service operability. Furthermore, this function also enhances service flexibility and provides fast online copy capabilities for virtual disks.

2.1.5.2 Usage restrictions

Understand the following restrictions before you perform hot migration of a cloud disk.

Restrictions

- Only river disks are supported.
- The source cluster and destination cluster for hot migration must have the same OSS domain.
- Disk sharing is not supported
- GPT disks are not supported.
- Format and capacity changes are not supported
- Only intra-region migration is supported.
- Due to the internal implementation, only the storage clusters with their clustername length less than 15 bytes can be migrated.



Note:

- After migration, the data on the original disk will be retained. You can use the pu tool to delete the retained data. At present, job recycling is unavailable.
- During the migration, there will be an I/O latency of less than 1 second. This is a normal situation.
- Rollback is not supported now.
- Migration will consume network bandwidth. Therefore, monitor and control concurrent traffic.

Migration operation

For more information about the disk hot migration API, see "**Disk hot migration**" in *ECS Developer Guide*.

2.1.5.3 O&M after hot migration

After successful hot migration and data copy of the cloud disk is completed, the data still exists on the source cloud disk. To release disk space, delete the data from the source cloud disk. After the data is deleted from the source cloud disk, the space will be released at a later time.

Procedure

The procedure for deleting a cloud disk is as follows:

1. On the computing cluster AG, run the `go2houyiregiondbrnd -e 'select task_id from device_migrate_log where status="complete"'` command to obtain `task: allTaskIds`.
2. Run the `go2riverdbrnd -e 'select task_id,src_pangu_path,dst_pangu_path from migration_log where task_id in ($allTaskIds) and status=2 and src_recycled=0 and DATE(gmt_finish) < DATE_ADD(CURDATE(), INTERVAL -1 DAY)'` command on the computing cluster AG.
3. Perform the following operations for each set of `<task_id,src_pangu_path,dst_pangu_path>`:
 - a) Run the `/apsara/deploy/bsutil rlm --dir=$dst_pangu_path|grep 'not-loaded'|wc -l` command on the bstools role host in the storage cluster. If the returned information is not 0, proceed to the next step.
 - b) Run the `/apsara/deploy/bsutil delete-image --dir=$src_pangu_path` command on the storage cluster bstools role host.
 - c) Run the `/apsara/river/river_admin migrate recycle $task_id` command on the river role machine in the storage cluster.

2.1.6 Upgrade solution

2.1.6.1 Overview

For both hot and cold migration of GPU clusters and FPGA clusters, you need to understand the relevant upgrade restrictions before upgrading a cluster.

2.1.6.2 GPU cluster restrictions

Understand the following restrictions before you upgrade the GPU cluster.

The upgrade of Apsara Stack GPU clusters is subject to the following restrictions:

- GPU clusters are only supported in Apsara Stack versions 3.3 or later.
- The GPU cluster does not support Apsara Stack's hot upgrade feature. The VM has to be shut down to upgrade the GPU cluster.

- VMs that use GPU, FPGA or other (passthrough or SRIOV) devices do not support hot migration.
- GPU clusters without specifications of local disk instances (GN5i, GN5e, and GN4) support cold migration.
- When you perform forced cold migration on GPU clusters with local disk instance specifications (GN5 and GA1), the local disk will be reformatted, which results in data loss. You need to back up the data before migration.

2.1.6.3 FPGA cluster restrictions

Understand the following restrictions before you upgrade the FPGA cluster.

FPGA cluster upgrades on Apsara Stack are subject to the following restrictions:

- FPGA clusters are only supported in Apsara Stack versions 3.5 or later.
- The FGPA cluster does not support Apsara Stack's hot upgrade feature. The VM has to be shut down to upgrade the FPGA cluster.
- Due to the strong dependence of the FPGA service on Redis, if the Redis service is interrupted during the hot upgrade of Apsara Stack, the FPGA service will be interrupted. The FPGA service is restored automatically when the Redis service is restored. However, if you fail to create a Redis instance, you need to restart the FPGA service after the Redis service is restored.

2.1.7 Disk of instance maintenance solution

2.1.7.1 Overview

This topic describes the limits, procedure, and related information of disk of instance maintenance.

Application scope

- Applicable to D1 disks only.
- Applicable to disks whose mount point is /apsarapangu/disk* only.
- Before and after maintenance, the mount point of the physical disk remains unchanged on the NC.
- Applicable to Apsara Stack versions 3.1 to 3.6.
- Currently only applicable to the N41S1-6T model.

Scenarios

A disk is damaged, but you want to restore the physical disk without migrating data and restore the data disk.

Impact: You need to shut down the VM associated with the damaged disk to restore the physical disk without migrating data.

Potential risks

- The original data on the replaced physical disk is all lost.
- Impact on fstab: If you write the uuid of the disk of the instance to the fstab in the VM, a problem will occur during the next startup. This problem can occur in any scenario where the disk-mounting relationship is changed.
- Follow the operation procedure precisely.

Environment inspection

Use a tool to inspect the entire cluster environment.

2.1.7.2 Maintenance procedure

This topic describes the specific maintenance procedure for the disk of instance.

Procedure

1. Log on to the AG with the admin account to search for NC-related information.

Run the following command to search for NC IDs based on NC IP addresses:

```
go2ncinfo {nc_ip}
```

{nc_ip} is the IP address of the physical machine with the damaged disk.

Example:

- Physical machine IP: 10.10.3.5
- Host name: c43b07003.cloud.b07.amtest1221
- File name and mount point of the physical machine with the damaged disk: /dev/sdb1/apsarapangu/disk1
- AG: vm010010016025
- Run the go2ncinfo 10.10.3.5 command to search for the NC ID.
- nc_id: 21765-26

```
[root@ecs-io11-a-5505-cn-neimeng-env10-d01:io11:vpc:21765]
$ go2ncinfo 10.10.3.5
    nc_id: 21765-26
      ip: 10.10.3.5
hostname: c43b07003.cloud.b07.amtest1221
biz_status: free
  priority: 8
    health: 5
      cpu: 56
```

2. Query the VMs on the AG that are affected by this physical disk (via the Server Controller).

- We recommend that you use the following APIs to check the affected VMs:

```
$ go2hyapi query_vm_list format=json region_id={region_id} nc_id={nc_id} nc_storage_device_id={mount_point}
```

{region_id} is the region where the host is located. You can run the `go2which {vm_id}` command on the AG to check the region. {nc_id} is the ID of the host, and {mount_point} is the mount point of the disk on the host.

- You can also use the script in `/etc/houyi/script/local_disk_ops.py` or the following API (the API may not be supported on the AG):

```
$/home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/tmp.log nc_id={nc_id} storage_device_id={mount_point}
```

{nc_id} is the ID of the host, and {mount_point} is the mount point of the disk on the host.

Example:

```
go2hyapi query_vm_list format=json region_id=cn-neimeng-env10-d01 nc_id=21765-26 nc_storage_device_id=/apsarapangu/disk1
```

```
[admin@ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ go2hyapi query_vm_list format=json region_id=cn-neimeng-env10-d01 nc_id=21765-26 nc_storage_device_id=/apsarapangu/disk1
[ERROR] [2018-05-10 16:41:36] The function 'query_vm_list' doesn't exist!
Usage:
houyi_api.sh <function_name> [name=value]
Available functions:
```

An error is reported if you use the API, so you need to use the following script instead (the `local_disk_ops.py` script for this environment is in `/home/admin` directory):

```
/home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1
```

```
[admin@ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ ls | grep local
local_disk_ops.py

[admin@ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$/home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1
[{'vm_name': 'i-5wf05ykw7mic5aq65dv2', 'status': 'running'}]
```

You can see that only the instance `i-5wf05ykw7mic5aq65dv2` runs on this disk and is in the running state.

3. Shut down the VM on the AG (via the Server Controller).

- a) If the VM is in the running state, you need to shut it down first.

Run the following command:

```
go2hyapi stop_vm vm_name={vm_name}
```

{vm_name} is the ID of the VM in the running state found earlier.

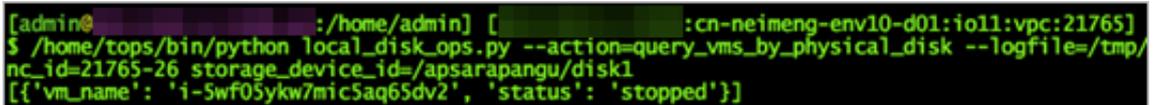
Example:

```
go2hyapi stop_vm vm_name=i-5wf05ykw7mic5aq65dv2
```



```
[admin@ :/home/admin] [ :cn-neimeng-e
$ go2hyapi stop_vm vm_name=i-5wf05ykw7mic5aq65dv2
<?xml version="1.0" encoding="utf-8"?>
<rsp>
  <code>200</code>
  <msg>successful</msg>
  <data>
    <vm_name>i-5wf05ykw7mic5aq65dv2</vm_name>
    <vm_status>Shutting</vm_status>
  </data>
</rsp>
```

Wait till the VM state becomes stopped.



```
[admin@ :/home/admin] [ :cn-neimeng-env10-d01:io11:vpc:21765]
$ /home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/
nc_id=21765-26 storage_device_id=apsarapangu/disk1
[{'vm_name': 'i-5wf05ykw7mic5aq65dv2', 'status': 'stopped'}]
```

- b) If the VM is in the pending or stopped state, no shutdown is required.
- c) If it is in other states, you need to wait till its state becomes pending or stopped, or carry out an inspection.
4. Check the local data disk associated with the physical disk (via the Server Controller).

Run the following command on the AG:

```
$/home/tops/bin/python local_disk_ops.py --action=query_loca
l_disks_by_physical_disk --logfile=/tmp/tmp.log nc_id={nc_id}
storage_device_id={mount_point}
```

{nc_id} is the ID of the host found earlier, and {mount_point} is the mount point of the disk on the host. The disk_id and the name of the VM to which the disk is mounted are obtained with this command.

Example:

```
/home/tops/bin/python local_disk_ops.py --action=query_local_disks_by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1
```

```
[admin@ ~ :/home/admin] [root@ ~ :cn-neimeng-env10-d01:io11:vpc:21765]
$ /home/tops/bin/python local_disk_ops.py --action=query_local_disks_by_physical_disk --logfile=
[{'vm_name': 'i-5wf05ykw7mic5aq65dv2', 'disk_id': '1000-3388'}]
```

Only the local data disk with disk_id = 1000-3388 is associated.

5. Replace the damaged disk on the NC.**a) Check the device file name of the damaged physical disk on the NC.**

Run the following command on the NC:

```
df -h
```

Example:

The device file name corresponding to /apsarapangu/disk1 is /dev/sdb1.

b) Check the SNs of the NC and the hard disk.

A. On the Apsara Infrastructure Management Framework, check the SN of the NC in the corresponding cluster O&M center. The SN of the NC is used to locate the machine if the disk is replaced on site.

Example: CVXKB7CD00J

B. Check the SN of the hard disk.

Run the following command:

```
smartctl -a {device_file_name} | grep 'Serial Number'
```

{device_file_name} is the device file name.

Example:

```
smartctl -a /dev/sdb1 | grep 'Serial Number'
```

```
[root@ ~ :/proc/scsi]
#smartctl -a /dev/sdb1 | grep 'Serial Number'
Serial Number:      K1K3EPKD
```

The SN of /dev/sdb1: K1K3EPKD

c) Remove the original disk.

The on-site engineer should locate the physical disk of the preceding NC based on the preceding information and the actual server model.



Note:

(The physical slot may vary with manufacturers and specific configuration.) Existing disk of instance model: N41S1-6T and V53. N41S1-6T supports HDD hot-swapping. V53 is an SSD model, and you need to shut down the whole machine before replacing the disk.

The following operations are only applicable to N41S1-6T.

Example:

C4-3.NT12	B07	06	CVXKB7CD001	N41S1-6T.22
-----------	-----	----	-------------	-------------

N41S1-6T supports hot-swapping. It uses the M.2 card as its system disk. The 12 hard disks can be seen on the front panel.

The disk order is as follows:

- /dev/sdb : 1 /dev/sde : 4 ...
- /dev/sdc : 2 /dev/sdf : 5 ...
- /dev/sdd : 3 /dev/sdg : 6 ...



You should remove /dev/sdb1, that is, the hard disk in slot 1. Check that the SN of the hard disk is consistent with the SN found earlier.



- Insert a new disk.
- Partition and mount the disk and modify the label and fstab file (the new disk must be mounted to the original mount point).

A. Check whether the hard disk is installed correctly.

Run the `fdisk -l` command to view the hard disk ID.

Example:

```
Disk /dev/sdb: 6001.2 GB, 6001175126016 bytes, 11721045168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disk label type: dos
Disk identifier: 0x00000000

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1             1  4294967295  2147483647+  ee    GPT
Partition 1 does not start on physical sector boundary.
```

You can see that the new hard disk is identified as SDB.

B. Partition the hard disk.

Run the `fdisk` command if the hard disk capacity is not greater than 2 TB.

```
fdisk /dev/sdb
```

Run the `parted` command if the hard disk capacity is greater than 2 TB:

```
parted /dev/sdb
```

(The `part` command is used to partition the 5.5 TB hard disk)

```
mklabel gpt
```

(Use the `gpt` command to make a 5.5 TB partition)

```
(parted) mklabel gpt
Warning: The existing disk label on /dev/sdb will be destroyed and all data on this disk will be
lost. Do you want to continue?
Yes/No? Yes
```

`mkpart primary 1049k -1` (Configure a primary partition of 5.5 TB, starting at 1,049 KB and ending at the end of the hard disk).

`print` (Display the space of the configured partition), `quit` (exit the parted program).

```
[root@ ~]# lsblk | grep sdb
sdb      8:16    0    5.5T  0 disk
└─sdb1   8:17    0    5.5T  0 part
```

C. Format the partition.

```
mkfs -t {filesystem_type} {device_name}
```

{filesystem_type} is the type of the file system to be formatted, and {device_name} is the name of the partition to be formatted.

Example:

```
mkfs -t ext4 /dev/sdb1
```

```
[root@ ~]# mkfs.ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
183144448 inodes, 1465130240 blocks
73256512 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=3613392896
44713 block groups
32768 blocks per group, 32768 fragments per group
4096 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848, 512000000, 550731776, 644972544

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

```
[root@ ~]# lsblk -T
NAME        FSTYPE LABEL        UUID                                 MOUNTPOINT
sda
├─sda1 ext4 /boot        1fd12aa3-8f54-4bb0-a1d3-a29595f391b8 /boot
├─sda2 ext4 /            3ac491f4-c2a4-4372-a4c3-3b3605b8a6da /
├─sda3 swap SWAP        57955bd2-1038-4f7e-8e85-f3b16d95794d
├─sda4
├─sda5 ext4 /apsarapangu 0e287a91-1e95-47a9-a815-c9a6b80d821e /apsarapangu
├─sda6 ext4 /ansara     67aec0b4-9bd0-4601-96ea-973d006c0979 /ansara
└─sdb
   └─sdb1 ext4 disk1       fd10be5a-efac-4cc7-8ecd-f1dd8df7824d
sdc
├─sdc1 ext4 disk2       a3b778c6-dc3e-40fe-89fe-6593d48db54e /apsarapangu/disk2
sdd
├─sdd1 ext4 disk3       b7c2c0c3-379d-41f2-9a09-dbc6add14093 /apsarapangu/disk3
sde
├─sde1 ext4 disk4       369a120f-4cd0-4249-b6cb-17c995a662cc /apsarapangu/disk4
sdf
```

D. Mount the hard disk to its original directory.

The server supports hot-swapping. If you remove and insert the same hard disk, it will be automatically mounted into the original directory. If a new disk is inserted, you need to manually mount it. Here you need to mount it manually:

```
mount {device_name} {mount_point}
```

{ergonomic_name} is the name of the device to be mounted, and {mount_point} is the target mount point.

Example:

```
mount /dev/sdb1 /apsarapangu/disk1
```

E. Modify the label.

The device files in the `/etc/fstab` directory are identified by their labels, so you need to modify the label of a new disk.

```
e2label {device_name} {label_name}
```

{ergonomic_name} is the device file name, and {label_name} is the label name.

Example:

The label of the old disk is `disk1`, so you need to change the label of the new disk to `disk1`.

```
[root@ ~]# cat /etc/fstab | grep 'disk1'
LABEL=disk1 /apsarapangu/disk1 ext4 noatime,nodiratime,nobarrier 0 0
```

```
e2label /dev/sdb1 disk1
```

```
[root@ ~]# blkid
/dev/sdb1: LABEL="disk1" UUID="65ce9f79-ab6f-48ea-8e28-84a2bb3ff420" TYPE="ext4" PARTLABEL="primary"
PARTUUID="6a0c5246-1002-4b5b-be24-c8d2ae20eff3"
```

F. Perform the mounting defined in fstab.

The label and mount point are consistent with those of the old disk, so you do not need to modify `/etc/fstab`. Run the following command to mount the new disk:

```
sudo mount -a
```

G. Run the `df -h` command to check disk information. It includes information such as mount information and disk capacity.

```
[root@...:/apsarapangu/disk1]
#ls
3388 1ost+found
```

6. Reset the data disk found earlier (via the Server Controller).

```
$/home/tops/bin/python local_disk_ops.py --action=reset_local_disk_after_change_physical_disk --logfile=/tmp/tmp.log disk_id={disk_id}
```



Note:

Exercise caution when performing the operation. The parameter {disk_id} must be the data disk found earlier based on the damaged disk.

Example:

```
/home/tops/bin/python local_disk_ops.py --action=reset_local_disk_after_change_physical_disk --logfile=/tmp/tmp.log disk_id=1000-3388
```

```
[admin@...:/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ /home/tops/bin/python local_disk_ops.py --action=reset_local_disk_after_change_physical_disk --logfile=/tmp/tmp.log disk_id=1000-3388
OK
```

OK indicates the resetting is successful.

7. Start the VM (via the Server Controller).

The Server Controller sends a command to rebuild the disks. Run the following command on the VM that needs to be started:

```
go2hyapi start_vm vm_name={vm_name}
```

{vm_name} is the ID of the VM that you want to start.

Example:

```
go2hyapi start_vm vm_name=i-5wf05ykw7mic5aq65dv2
```

```
[admin@...:/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ go2hyapi start_vm vm_name=i-5wf05ykw7mic5aq65dv2
<?xml version="1.0" encoding="utf-8"?>
<rsp>
  <code>200</code>
  <msg>successful</msg>
  <data>
    <vm_name>i-5wf05ykw7mic5aq65dv2</vm_name>
    <vm_status>Starting</vm_status>
  </data>
</rsp>
```

Result

You can log on to the VM through SSH, format the device corresponding to the new disk, and mount it to the mount point. Check the disk capacity, and whether data read/write operations are successful.

2.1.7.3 Additional instructions

Related scripts are described here for specific solutions of local disk maintenance.

local_disk_ops script usage instructions

- For more information on how to use the scripts, see:

```
/home/tops/bin/python local_disk_ops.py -h
```

- Log description:**

When a script is executed, a detailed log is recorded in the log file. If an error occurs, the error log is also output to the current shell. You can specify a log file. Otherwise, the default log file is used. The default log file is in the same directory as the script. It is a log file with the same name as the script (a different extension).

For example: The logs of `/home/tops/bin/python local_disk_ops.py --action=xxx arg1=value1` are recorded in the `local_disk_ops.log` file.

- Error description:**

If an error occurs when you execute a script, the error log is output to the current shell. Perform inspections based on specific error information. Error message format:

Error time Error (script error line)-error message.

Example 1: `$/home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk nc_id=xxx`

```
2018-03-13 21:12:37,864 ERROR (local_disk_ops.py:98) - storage_device_id can not be empty.
```

The preceding error indicates that `storage_device_id` is not specified.

Example 2

```
$/home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk nc_id=1-1 storage_device_id=/apsarapangu/disk20
```

```
2018-03-13 21:23:42, 764 error (FIG: 174)-check NC record error, should have one record. resource_info: {'nc_id': '1-1'}
```

The preceding error indicates that an error occurred during the NC resource check because an inbound `nc_id` is incorrect.

- For more information about this error, see [Maintenance procedure](#).

2.1.8 Handle routine alarms

2.1.8.1 Overview

This topic describes the definition of each key metric and how to handle alarms.

The metrics of ECS can be categorized into three types:

- **Basic metrics:** These metrics are used to monitor the CPU, memory, and correlated service processes of physical machines.
- **Connectivity metrics:** These metrics are used to monitor the connectivity between different components and the connectivity between different networks.
- **Service metrics:** These metrics are used for service monitoring. For example, the state of various types of API requests.

Table 2-6: Description of metric types

Metric type	Function	Solution
Basic metric/ service availability metric	Monitors the basic performance of the host and the availability of the services on the host. For example, CPU, memory, and handle count.	CPU utilization too high: check for processes which consume large amounts of CPU resources. If it is a key process, evaluate whether it can be restarted.
		Memory utilization too high (for key services): dump the data on the memory, request the backend R&D team to analyze this issue, and restart the application.
Connectivity metric	Checks the connectivity between a module and the modules related to it.	<ul style="list-style-type: none"> • First, check the health of the corresponding modules. For example, whether the host works normally and whether service, ports, and domain names are normal. • If both modules are healthy, troubleshoot the network connectivity.
Service metric	Monitors key request calls. For example, the latency, total number	<ul style="list-style-type: none"> • In case of an API request failure, you need to view the corresponding log to identify the cause of failure.

Metric type	Function	Solution
	, and failures of API requests and database SQL exceptions.	<ul style="list-style-type: none"> In case of a database SQL failure, check whether it is caused by a database exception (system breakdown or high connection count) or a problem with the application. If it is an application problem, forward the error information to the backend R&D team for troubleshooting.

2.1.8.2 API proxy

This topic describes the metrics of API proxy.

Table 2-7: Metric description

Metric	Alarm	Description
check_apiproxy_dns	Database HA switchover or not	Checks whether Server Controller database switchover occurs. If so, nginx will be reloaded automatically.
check_apiproxy_conn_new	check_apiproxy_conn_new	Checks the connectivity to the Server Controller database.
		Checks the connectivity to the API server: <ul style="list-style-type: none"> Checks whether the API server is down. Checks the network connectivity.
check_apiproxy_proc_new	check_apiproxy_proc_new	Checks the memory and CPU utilization for nginx and memcache processes.

2.1.8.3 API server

The topic describes the metric items about the API server.

Table 2-8: Metric description

Metric	Alarm	Solution
check_API_Server_proc_new	The process does not exist or is abnormal.	Check the state of the Java process: whether it exists and the utilization of the CPU and memory.

Metric	Alarm	Solution
check_API Server_conn_new	Check the connectivity between the API Server and Server Controller database.	Check whether the corresponding component is down. If the corresponding component is down, fix the issue in line with the relevant O&M approach. If the database is down, contact DBA to fix the issue. Check whether the VIP is connected to the corresponding component. If not, contact the network engineer to fix it.
	Check the connectivity between the API Server and TAIR.	
	Check the connectivity between the API Server and RegionMaster.	
	Check the connectivity between the API Server and RMS.	
check_API Server_perf	Check the status of API requests, such as the latency, total number of API requests, and number of failed API requests.	It is primarily used to identify faults.
check_API Server_errorlog	Mainly check database exceptions and failed instance creation cases.	<ul style="list-style-type: none"> If the database is abnormal, contact DBA to check whether the database is normal. If the creation of an instance fails, locate the cause of failure.

2.1.8.4 RegionMaster

This topic describes the metric items of RegionMaster.

Table 2-9: Metric description

Metric	Alarm	Description
check_regionmaster_proc	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists and the utilization of the CPU and memory.
check_regionmaster_work	rms_connectivity	Checks the connectivity to RMS.
	regiondb_connectivity	Checks the connectivity to hoiuyiregiondb.

Metric	Alarm	Description
	houyi_connectivity	Checks the connectivity to the Server Controller database.
	tair_connectivity	Checks the connectivity of TAIR.
check_zookeeper_work	status	Checks the operating state of the Zookeeper on the Server Controller.
check_regionmaster_errorlog	errorlog_for_db	Checks whether the SQL statement is properly executed.
	check_regionmaster_errorlog	
check_workflow_master	Checks the operating state of the master in the workflow.	
check_workflow_worker	Checks the operating state of the worker in the workflow.	

2.1.8.5 RMS

This topic describes the metrics of RMS.

Table 2-10: Metric description

Metric	Alarm	Description
check_rms_proc	Checks the process state , and CPU and memory utilization of RMS.	
check_rabbitmq_proc	Checks the process state , and CPU and memory utilization of the rabbitmq cluster.	
check_rabbitmq_status	Checks the number of queues, exchanges, and bindings of the rabbitmq cluster.	Follow the maintenance guide for the rabbitmq cluster.
check_rabbitmq_queues	Whether messages are accumulated.	If messages are accumulated, check the root cause of message accumulation.
	Check whether there are consumers.	If there are no consumers, check whether Regionmaster and APIServer are operating normally. If they are operating

Metric	Alarm	Description
		normally, check whether there is a problem with the rabbitmq cluster.

2.1.8.6 PYNC

This topic describes the metrics of PYNC.

Table 2-11: Metric description

Metric	Alarm	Description
check_vm_start_failed	Checks the reasons for failed VM startup.	You do not have to handle it immediately. It is typically caused by customized images.
check_pync	Checks the utilization of CPU and memory for the PYNC.	
	PYNC has opened too many file handles.	
	PYNC process count.	PYNC must have four processes.
	pyncVmMonitor.LOG has not been updated for a long time. Last update occurred at \${pync_monitor_log_last_updated}.	<p>Check why the log has not been updated for a long time.</p> <ul style="list-style-type: none"> Whether the PYNC process has a problem. Whether the NC has a key process (Uninterruptible Sleep).

2.1.8.7 Zookeeper

This topic describes the metrics of Zookeeper.

Table 2-12: Metric description

Metric	Alarm	Description
check_zookeeper_proc	proc	The process does not exist.
		Memory and CPU utilization is too high

2.1.8.8 AG

This topic describes the metrics of AGs.

Table 2-13: Metric description

Metric	Alarm	Description
disk_usage	apsara_90	Utilization of the <i>/apsaradisk</i> .
	homeadmin_90	Utilization of <i>/home/admin</i> .
check_system_ag	mem_85	Memory utilization.
	cpu_98	CPU utilization.
	df_98	Utilization of the root directory disk.
check_ag_disk_usage	check_ag_disk_usage	Checks the disk utilization.
check_nc_down_new	check_recover_failed	Checks the cause of failed VM migration. Possible causes include: <ul style="list-style-type: none"> No resources are available in the cluster. There exists a VM which does not belong to any cluster.
	check_repeat_recovered	Continuous VM migration.
	check_continuous_nc_down	Checks continuous NC downtime.
	check_nc_down_with_vm	The state of the NC in the database is <i>nc_down</i> , but there are still VMs operating on the NC. Checks the NC for hardware faults: <ul style="list-style-type: none"> You are required to perform operations and maintenance if a hardware fault is detected. If no hardware fault is detected, restore the NC, and change its state to <i>locked</i>.
check_ag_fhtd_new	Checks whether the FHT downtime migration and handling tool works normally.	It is mainly used by local disks. If the tool does not exist, download the downtime migration tool.

2.1.8.9 Server groups

This topic describes the metrics of server groups.

Table 2-14: Metric description

Metric	Alarm	Description
check_pync	pync_mem	Monitors the memory utilization of PYNC.
	pync_cpu	Monitors the CPU utilization of PYNC.
	pync_nofile	Monitors the count of PYNC handles.
	pync_nproc	Monitors the count of PYNC processes.
	pync_monitor_log_not_updated	Monitors the state of the scheduled tasks of PYNC.

2.1.9 Inspection

2.1.9.1 Overview

ECS inspection can be divided into cluster basic health inspection and cluster resources inspection.

2.1.9.2 Cluster basic health inspection

2.1.9.2.1 Overview

Cluster basic health inspection mainly includes basic software package version inspection, and basic public resources inspection.

2.1.9.2.2 Monitoring inspection

This topic covers inspection on basic monitoring and connectivity monitoring.

2.1.9.2.3 Basic software package version inspection

This topic describes the version inspection of Server Controller components, the Apsara system, virtualization packages, and basic service packages.

2.1.9.2.4 Basic public resources inspection

This topic includes ISO inspection and basic image inspection.

ISO inspection

Currently, the ECS Operations and Maintenance System provides two basic ISOs for each region:

- linux-virt-release-xxxx.iso

- windows-virt-release-xxxx.iso

You can run the following command to search the database for relevant information.

```
$ houyiregiondb
mysql>select name,os_type,version,path,oss_info from iso_resource
where os_type! ='\G
```

Parameters in the command are as follows:

- *name*: name of the ISO. For example, xxxx.iso.
- *os_type*: image type of ISO service.
- *path*: Apsara Distributed File System path on the cloud disk in which the ISO is stored. You can run the `/apsara/deploy/pu meta $path` command to check whether the ISO exists in the files of the Apsara Distributed File System.
- *oss_info*: OSS path in which the ISO is stored on the local disk cluster. To search for this path, you need to provide relevant information to OSS operations and maintenance engineer for inspection.

Basic image inspection

- Run the following command to check the state of a basic image in the database:

```
houyiregiondb
mysql>select image_no,status,visibility,platform,
region_no from image;
```

- Check whether the basic image is usable. You can call the `create_instance` API to use the relevant image to create a VM and manually check whether it can operate normally.

2.1.9.3 Cluster resource inspection

2.1.9.3.1 Overview

Cluster resource inspection mainly includes cluster inventory inspection and VM inspection.

2.1.9.3.2 Cluster inventory inspection

Cluster inventory resources mainly refer to the remaining resources in the cluster, which can be used to create VMs with different specifications. You can use the database to search for cluster inventory resources.

To check the inventory for 16-core 64 GB VMs:

```
$ houyiregiondb
```

```
mysql> select sum( least ( floor(available_cpu/16),floor(available_
memory/64/1024))) from nc_resource,nc where nc.cluster_id=$id and nc.
biz_status='free' and nc.id=nc_resource.id;
```

If there is a large VM in the current cluster, make sure that the cluster has enough free resources, and a host with sufficient resources is available for backup. This host will be the migration destination of the large VM in case the current host goes down. Otherwise, the large VM will not be migrated when its host goes down. You will have to use hot migration to transfer resources, or release redundant VMs in the cluster.

NC state inspection

NC state inspection mainly checks whether the state of a host is normal in the database and Apsara Infrastructure Management Framework.

- A host can be in one of the following states in Apsara Infrastructure Management Framework:
 - Good: indicates that the host is in a normal working state.
 - Error: indicates that the host has a monitoring alarm.
 - Probation: indicates that the host is in the probationary period and is likely to fail.
 - OS_error: indicates that the host has failed and is being cloned.
 - Hw_error: indicates that the host has a hardware failure and is being repaired.
 - OS_probation: indicates the host is recovering from a fault or hardware failure and is in the probationary period. If it passes the probationary period, the state will be changed to probation. If it fails to pass the probationary period (a monitor reports an error during the period), the state will be changed to OS_error.



Note:

Good is referred to as the stable state, while others are collectively called the unstable state.

- Cluster definitions for Apsara Infrastructure Management Framework:
 - Default cluster: the cluster where NCs are placed when they go offline.
 - Non-default cluster: the cluster for online NCs.

A NC that is operating normally is placed in a non-default cluster, and is in the GOOD state.

The mapping between the host states in the ECS database and in Apsara Infrastructure Management Framework is as shown in [Table 2-15: Mapping between host states in ECS database and in Apsara Infrastructure Management Framework](#).

Table 2-15: Mapping between host states in ECS database and in Apsara Infrastructure Management Framework

State in ECS database	Cluster	Host state	Scenario
Mlock	Non-default cluster	Unstable	A host that has just come online is proactively locked.
locked	Non-default cluster	Unstable	Unlock the NC.
free	Non-default cluster	Stable	Operating normally.
nc_down	Non-default cluster	Unstable	Operating normally or at downtime.
offline	Default cluster	Unstable	Offline from business attributes.

2.1.9.3.3 VM inspection

This topic describes how to pend VM inspection, VM state inspection, and VM resource inspection.

Pending VM inspection

This type of inspection focuses on the VMs that have been in the pending state in the cluster for a long time. If a VM has been in the pending state for a long time, it is believed to be a redundant resource. Contact the user to handle it.

VM state inspection

This type of inspection focuses on the VM state consistency. For example, the VM is displayed as stopped in the database, but is displayed as running in NC. The inspection compares the state of the VM recorded in the database and the state on the physical host, and handles the VMs which have inconsistent states.

- Obtain the VM state in a database

```
houyiregiondb -Ne "select status from vm where name='$name'"
```

- Obtain the VM state of a host

```
sudo virsh list | grep $name
```

VM resource inspection

After the configuration of the VM is changed, the system checks whether the configuration of the VM recorded in the database is consistent with that used on the host.

- Obtain the VM state in a database

```
houyiregiondb -Ne "select vcpu, memory from vm where name='$name'"
```

- Obtain the VM state of a host

```
sudo virsh list | grep $name
```

View the corresponding field to acquire information about CPU and memory.

2.2 Auto Scaling (ESS)

2.2.1 Log on to Apsara Stack Operations

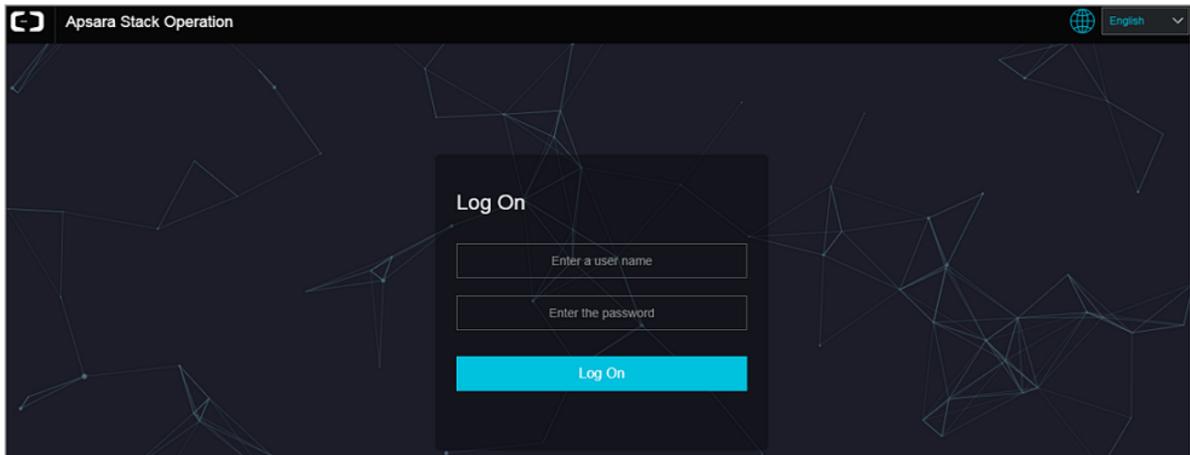
This section describes how to log on to Apsara Stack Operations (ASO).

Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 2-3: Log on to ASO**Note:**

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.

2.2.2 Product resources and services

2.2.2.1 Application deployment

All the applications in the ESS Business Foundation System are stateless. You must restart the applications by running the docker restart command.

- **ess-init**

It first initializes the database service, and then pushes all API configuration files of ESS to the pop configuration center to initialize OpenAPI Gateway.

- **Trigger (dependent on ess-init)**

- Trigger executes tasks such as checking health status, checking the maximum and minimum instance numbers, and deleting scaling groups.

- It triggers scheduled tasks and alarms.

- **Coordinator**

Coordinator is the OpenAPI layer that provides public-facing services. It maintains persistent requests and issues tasks.

- **Worker**

- Worker executes all scaling-related tasks, such as creating ECS instances, adding instances to SLB backend server groups and RDS whitelists, and synchronizing CloudMonitor or group information.

- It retries failed tasks and provides the rollback mechanism.

- **service_test**

It is used for regression tests on the overall application running status. It contains over 60 regression test cases to test the integrity of functions.

2.2.2.2 Troubleshooting

This topic describes how to troubleshoot issues of product resources and services.

Prerequisites

When issues related to Business Foundation System occur, you can submit tickets on the [Alibaba Cloud Business Support Platform](#) and check related service status in the Apsara Infrastructure Management Framework console.

Procedure

1. Submit a ticket.
2. Check the status of services that depend on Business Foundation System in the Apsara Infrastructure Management Framework console.

If a service cannot be executed, it affects the running of ESS Business Foundation System.

[Table 2-16: Failed services and their impacts](#) describes the details.

Table 2-16: Failed services and their impacts

Service	Impact
middleWare.dubbo	Deployment is affected. The service is unavailable.
middleWare.tair	Deployment is affected. The service is unavailable.
middleWare.metaq (message middleware)	Deployment is affected.
middleWare.zookeeper	Deployment is affected. The service is unavailable.
middleWare.jmenvDiamondVips	Deployment is affected, the Diamond configuration item cannot be obtained.
ram.ramService (RAM users)	The RAM-user service is unavailable.
webapp.pop (API gateway)	The OpenAPI service is unavailable.
ecs.yaochi (ECS Business Foundation System)	All ECS creation requests become invalid.
slb.yaochi (SLB Business Foundation System)	All SLB association requests become invalid.
rds.yaochi (RDS Business Foundation System)	All RDS association requests become invalid.
tianjimom (Monitoring System of Apsara Infrastructure Management Framework)	Some services are unavailable.

2.2.3 Inspection

2.2.3.1 Overview

ESS inspection monitors the basic health conditions of the clusters.

The basic health conditions inspected include the following aspects:

- [Monitoring inspection](#)
- [Basic software package version inspection](#)

2.2.3.2 Monitoring inspection

This topic describes basic monitoring and connectivity monitoring inspection.

2.2.3.3 Basic software package version inspection

Version inspection for trigger, coordinator, worker, and base services.

2.3 Object Storage Service (OSS)

2.3.1 Usage in Apsara Stack Operations

2.3.1.1 Log on to Apsara Stack Operations

This section describes how to log on to Apsara Stack Operations (ASO).

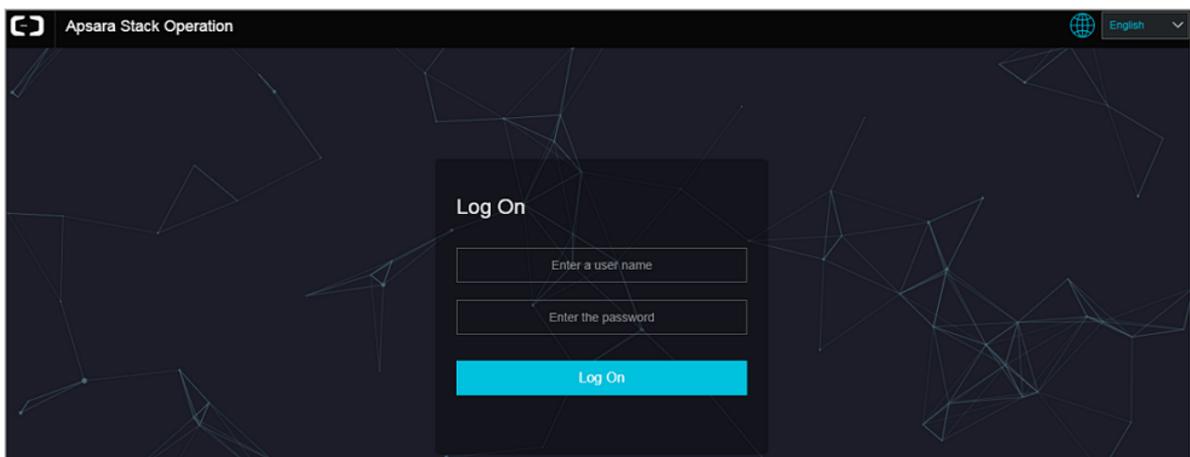
Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 2-4: Log on to ASO



Note:

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.

- The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.

2.3.1.2 Business data

2.3.1.2.1 User data

2.3.1.2.1.1 User data overview

You can query data statistics and trends including resource usage and basic attributes of resources by UID, account, or bucket.

Context

The User Data Overview page is displayed only when you search by UID or account. On the User Data Overview page, you can specify a date or day to view total usage of various resources in all buckets owned by the user account.

Resource statistics can be collected by storage capacity, inbound traffic (intranet and Internet), total outbound traffic (intranet and Internet), total number of requests, and SLA statistics.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data > User Data** in the left-side navigation pane. Select **UID** or **Alibaba Cloud Account**.
3. Set **Date**, and click **View**.

2.3.1.2.1.2 Data monitoring

Context

You can query resource running status and usage such as the storage capacity, traffic, SLA, HTTP status, latency, QPS, and image processing capacity by UID or bucket. You can also query the resource usage and trends based on the time range.

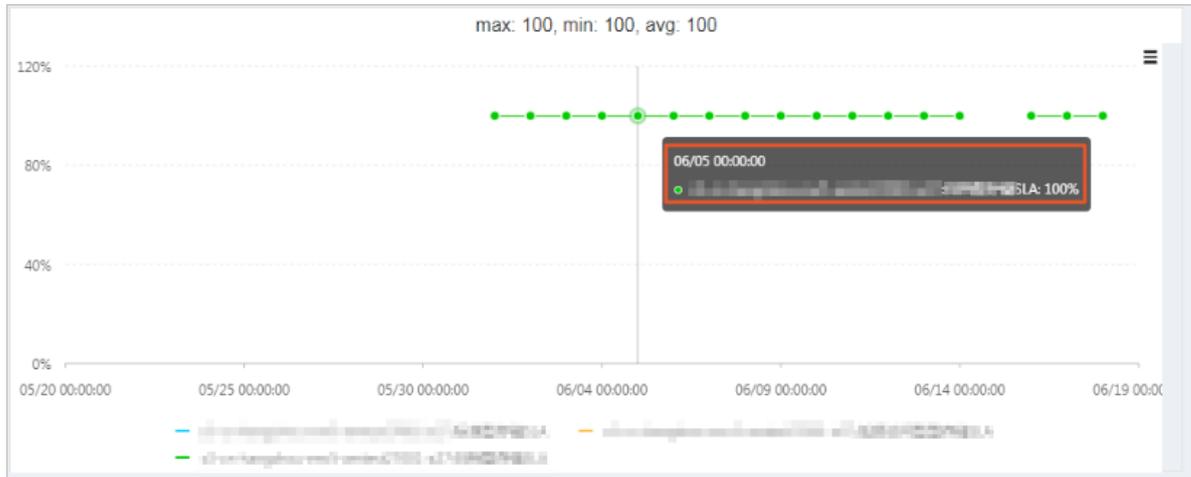
Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data > User Data** in the left-side navigation pane. Click the **Data Monitoring** tab.
3. Set Bucket, Time Range, and Sampling Time. Click **View**.
4. If you query data by user, you can click the name of a bucket on the trend chart to show or hide the curve of the bucket.

Figure 2-5: Data monitoring 1



5. Hover over a point on the trend chart to display data at a specific point in time.

Figure 2-6: Data monitoring 2

Metric descriptions:

- SLA: indicates the service level availability metric for OSS. Formula: $SLA = \frac{\text{Non-5XX request count per 10s or hour}}{\text{Total valid request count}}$
- HTTP status: collects statistics on the counts of 5XX, 403, 404, 499, 4XX_others, 2XX, and 3XX status codes returned as well as the request ratios.
- Latency: collects statistics on the latency of APIs such as put_object, get_object, and upload_part as well as the maximum latency.
- Storage capacity: collects statistics on the storage capacity of standard, infrequent, and archive storage and their increments.
- Image processing capacity: collects statistics on the number of processed images.
- Traffic: collects statistics on the inbound and outbound Internet and intranet, CDN, and the synchronized inbound and outbound traffic.
- QPS: collects statistics on the request counts of billing, copy, get object, put object, upload part, post object, append object, head object, and get object info.

2.3.1.2.1.3 Basic bucket information

Context

You can query basic bucket information such as the cluster deployment location, configuration information, current capacity, and object count of a bucket. You can also view this information in a table.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data** > **User Data** in the left-side navigation pane. Click the **Basic Bucket Information** tab.
3. Select a bucket you want to view.

Figure 2-7: Basic bucket information

The screenshot displays the 'Basic Bucket Information' page in the Apsara Stack Operations console. The left navigation pane shows 'Business Data' expanded, with 'User Data' selected. The main content area has three tabs: 'Bucket Basic Information' (active), 'User Data Overview', and 'Data Monitoring'. Under 'View Bucket', the bucket 'bajesjar' is selected, and a 'View' button is visible. The bucket details are as follows:

Bucket Name:	bajesjar
User Account:	aliyuntest(999999999)
Enterprise/Individual Name:	
Application:	file
BID:	26842
Current Capacity:	Standard: 34462B, IA: 0B, AR: 0B
Storage Type:	standard
Objects:	Standard: 21, IA: 0, AR: 0
Log Service:	Inactivated

2.3.1.2.2 Cluster data

2.3.1.2.2.1 Inventory monitoring

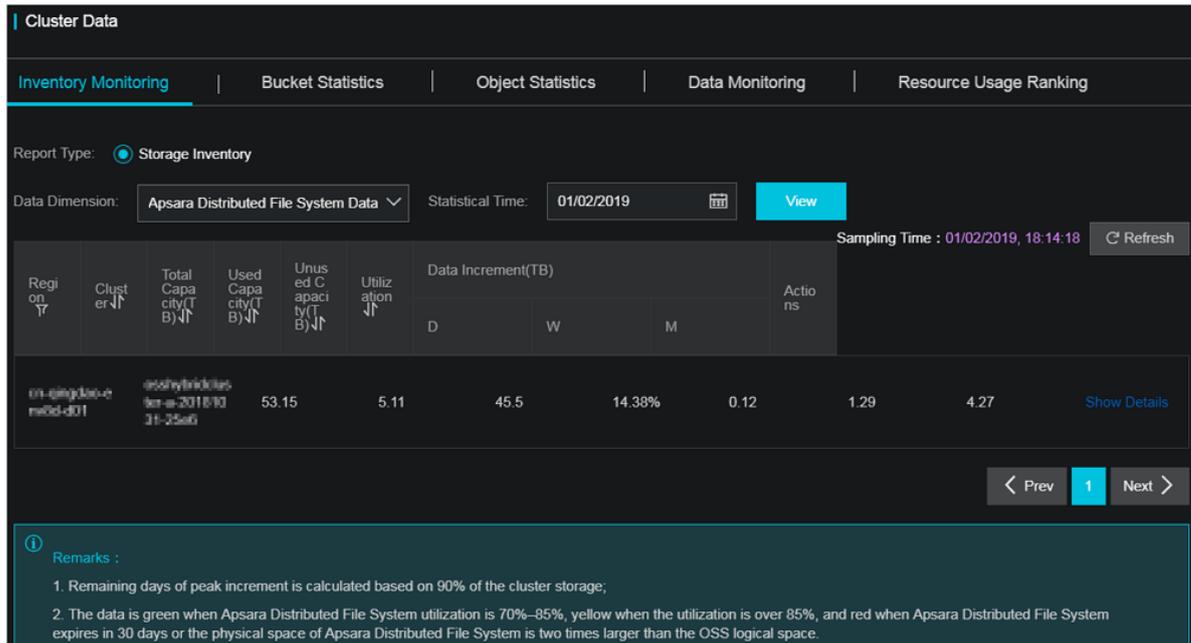
Context

Metrics of inventory monitoring include the total capacity, available capacity, used capacity, and storage backup ratio.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data** > **Cluster Data** in the left-side navigation pane.
3. Click **Inventory Monitoring**.

Figure 2-8: Inventory monitoring



Aside from basic cluster information such as the cluster name and region, you can also view the following metrics:

- **Apsara Distributed File System Data:** includes the actual total capacity for storage (including the total capacity for multiple data backups), used capacity, remaining capacity (available), usage, data increment (by day, week, or month), days available based on the average increment (reference value), and days available based on the maximum increment (reference value).
- **Metric Data:** includes the statistical capacity of ECS and non-ECS.
- **KV Data:** includes the logic KV data, KV data in the recycle bin, and data increment (by day, week, or month).

2.3.1.2.2.2 Check disk space

Context

Inventory Monitoring allows you to check the disk space.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. Choose **Business Data > Cluster Data** in the left-side navigation pane.
3. Click **Inventory Monitoring**.

4. On the inventory information page that appears, click **View Details**.
5. On the details page, check **Uncompleted GC in OSS**, **Uncompleted Merge in KV**, and **Completed Merge and Uncompleted GC in KV** to check the disk space.

2.3.1.2.2.3 Data monitoring

Context

The cluster operation metric is identical to the user data metric except that the object of cluster operation metric is the data collected by cluster.

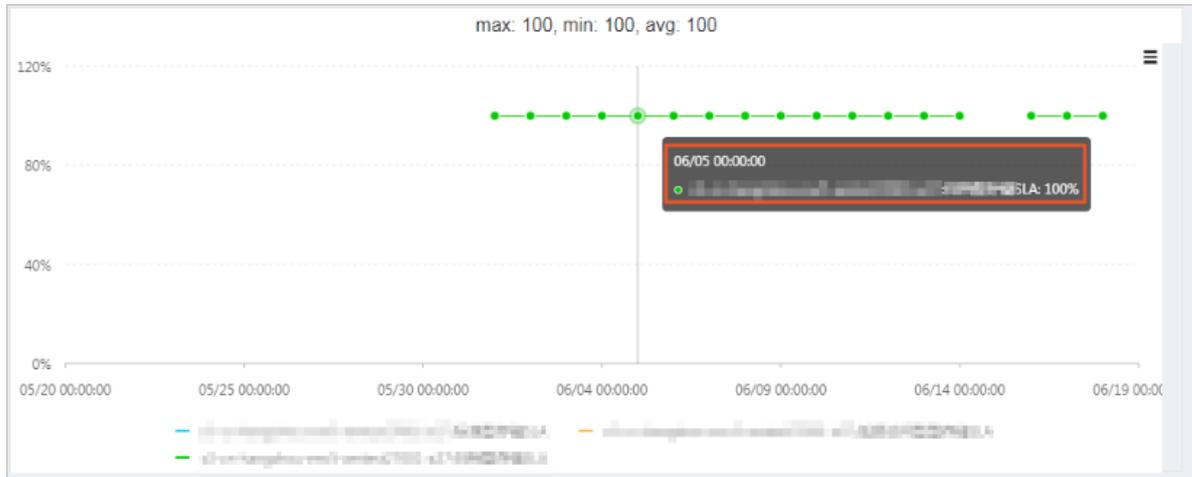
Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data** > **Cluster Data** in the left-side navigation pane.
3. Click **Data Monitoring**.
4. Set Bucket, Time Range, and Sampling Time. Click **View**.
5. If you query data by user, you can click the name of a bucket on the trend chart to show or hide the curve of the bucket.

Figure 2-9: Data monitoring 1



6. Hover over a point on the trend chart to display data at a specific point in time.

Figure 2-10: Data monitoring 2

Metric descriptions:

- SLA: indicates the service level availability metric for OSS. Formula: $SLA = \frac{\text{Non-5XX request count per 10s or hour}}{\text{Total valid request count}}$
- HTTP status: collects statistics on the counts of 5XX, 403, 404, 499, 4XX_others, 2XX, and 3XX status codes returned as well as the request ratios.
- Latency: collects statistics on the latency of APIs such as put_object, get_object, and upload_part as well as the maximum latency.
- Storage capacity: collects statistics on the storage capacity of standard, infrequent, and archive storage and their increments.
- Image processing capacity: collects statistics on the number of processed images.
- Traffic: collects statistics on the inbound and outbound Internet and intranet, CDN, and the synchronized inbound and outbound traffic.
- QPS: collects statistics on the request counts of billing, copy, get object, put object, upload part, post object, append object, head object, and get object info.

2.3.1.2.2.4 Resource usage rankings

Context

You can view the top 10 to 100 users and buckets in each cluster based on resource usage. These statistics help administrators monitor the users who consume the most resources.

Data resource ranking items include:

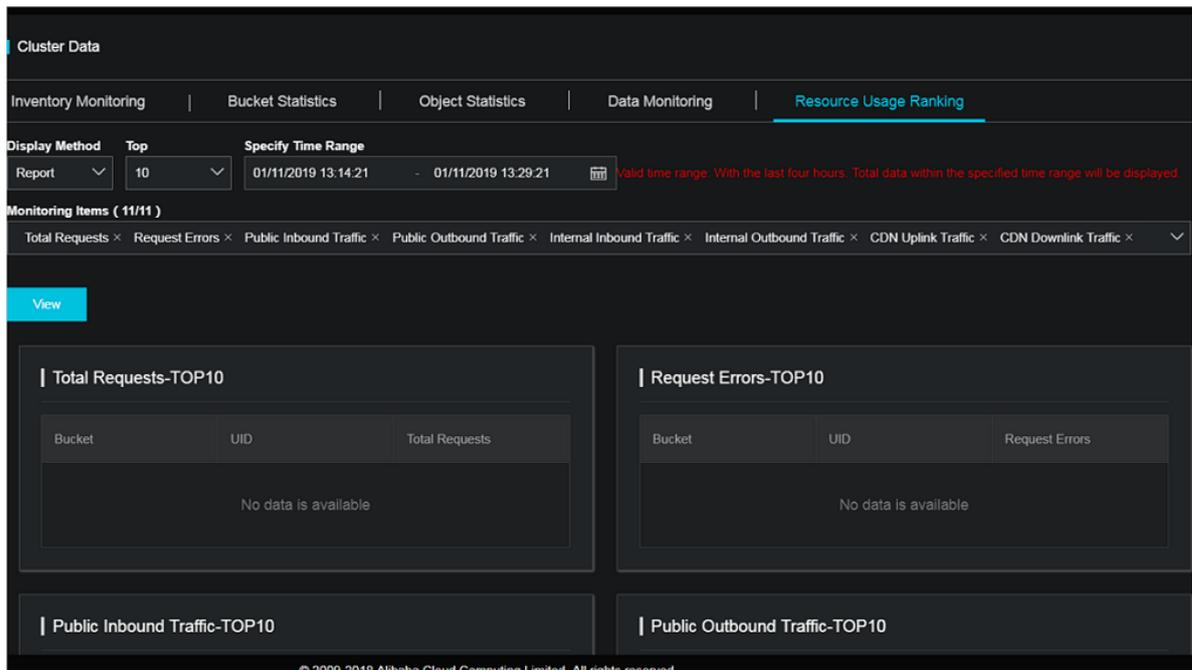
- Total Requests

- Request Errors
- Public Inbound Traffic
- Public Outbound Traffic
- Internal Inbound Traffic
- Internal Outbound Traffic
- CDN Uplink Traffic
- CDN Downlink Traffic
- Storage Capacity

Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data > Cluster Data** in the left-side navigation pane.
3. Click **Resource Usage Ranking**.
4. You can select **Report** or **Trend** to view rankings.
5. Set the time range, metrics, and top number of projects you want to view.
6. You can click a bucket and choose **User Data > Data Monitoring** to view bucket reports and trend charts.
7. You can click UID and choose **User Data > User Data Overview** to view user data reports and trend charts.

Figure 2-11: Resource usage rankings



2.3.1.2.2.5 Bucket statistics

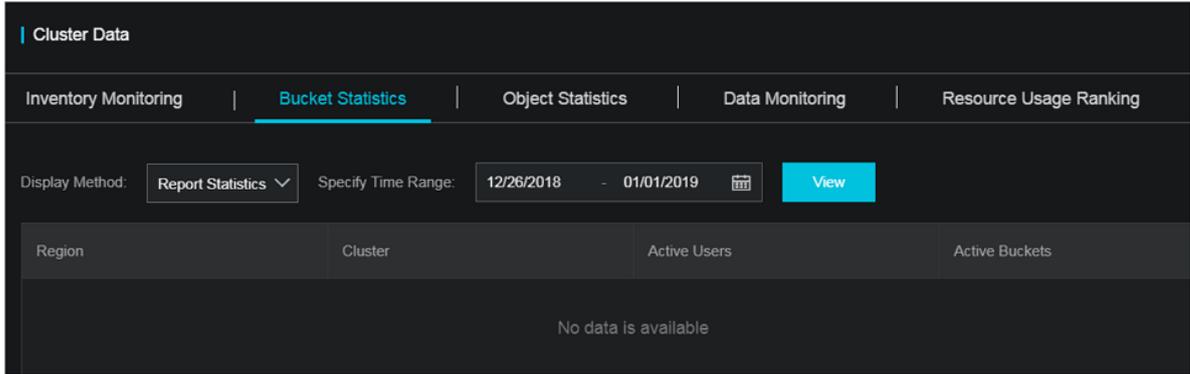
Context

You can collect statistics on the count of buckets by cluster.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. Choose **Business Data > Cluster Data** in the left-side navigation pane.
3. Click **Bucket Statistics**.
4. You can select one of the following display methods to view bucket statistics: **Report Statistics**, **Growth Trend**, or **Current Overall Statistics**.
5. If you select **Report Statistics** or **Growth Trend**, specify the query time range, and click **View**.

Figure 2-12: Bucket statistics



2.3.1.2.2.6 Object statistics

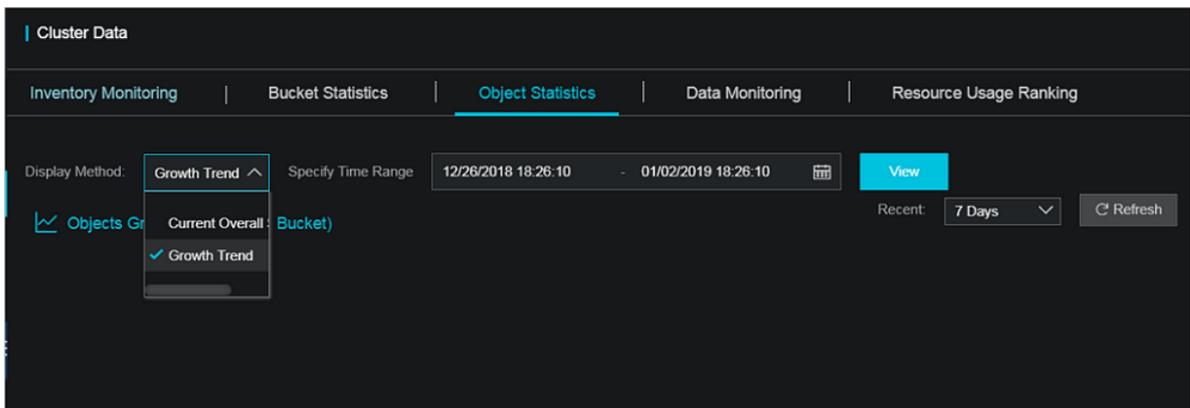
Context

You can view the statistics on the number and trends of objects by cluster.

The statistics shown are not in real-time. The data is collected at 1-hour intervals. You can select **Display Method** to query historic object counts.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. Choose **Business Data > Cluster Data** in the left-side navigation pane.
3. Click **Object Statistics**.
4. You can select **Current Overall Statistics** or **Growth Trend** to view object statistics.
5. If you select **Growth Trend**, specify the time range, and click **View**, as shown in the following figure.



2.3.2 Use of tools

2.3.2.1 Typical commands for tsar

- View help details of tsar.

Command: `tsar -help`

- View the nginx operation data of each minute from the past two days.

Command: `tsar -n 2 -i 1 -nginx`

`-n 2` indicates the data generated in the past two days. `-i 1` indicates one result record generated each minute.

- View the tsar load status and operation data of each minute from the past two days.

Command: `tsar --load -n 2 -i 1`

2.4 Table Store

2.4.1 Storage Operations and Maintenance System

2.4.1.1 Overview

Storage Operations and Maintenance System helps quickly locate problems during O&M and notifies users of the current running status of their services. Appropriate use of Storage Operations and Maintenance System can significantly improve O&M efficiency.

The domain name of Storage Operations and Maintenance System is in the format of "chiji.ots.{global:intranet-domain}."

The main modules of Storage Operations and Maintenance System are **User Data**, **Cluster Management**, **Inspection Center**, **Monitoring Center**, **System Management**, and **Platform Audit**. These modules provide comprehensive O&M functions to meet different requirements.

2.4.1.2 User data

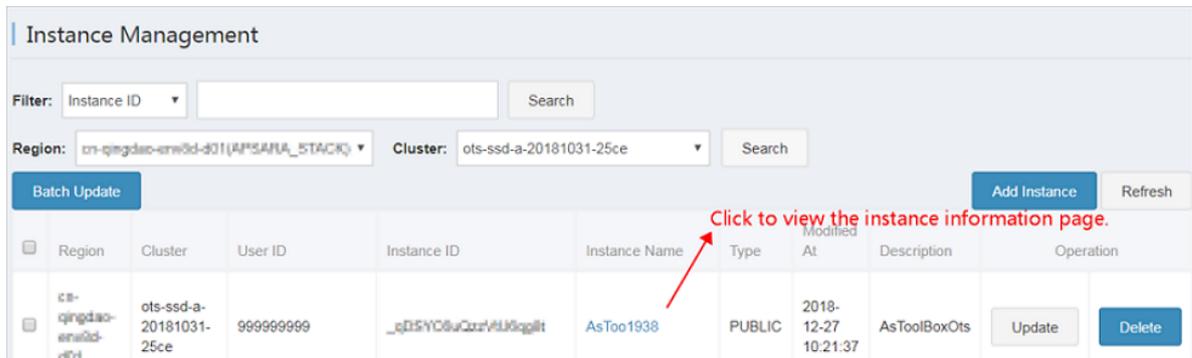
2.4.1.2.1 Manage instances

You can obtain instance details through the cluster instance list, specified query conditions, and instance meta information.

Function description

- Specify a region and a cluster name to obtain a list of instances.

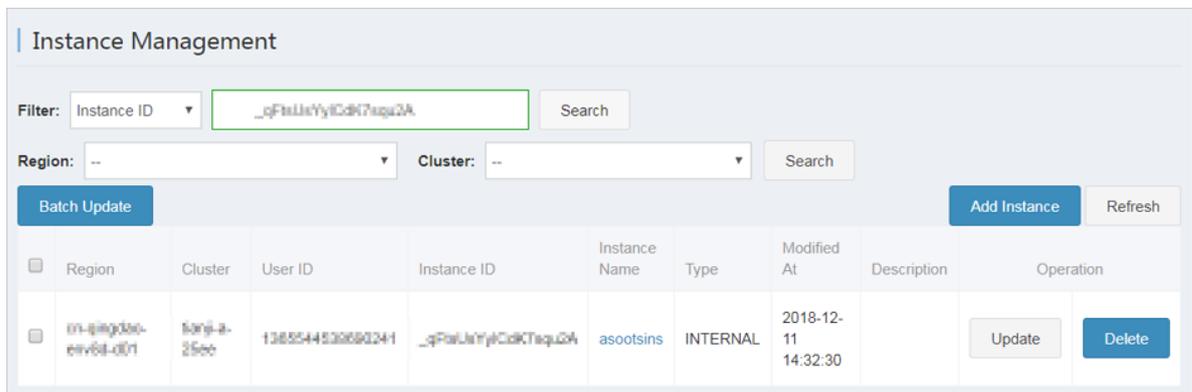
You can specify a region and a cluster to view the instances in the specified cluster, and the basic information of each instance.



On the Manage Instances page, you can:

- View the cluster instance list.
- View instance descriptions.
- View the links to details of instances by instance name.
- Update and delete an instance in the instance list.
- Search for instances based on specified conditions

This page allows you to search for instances of all clusters in all regions based on the specified search conditions.

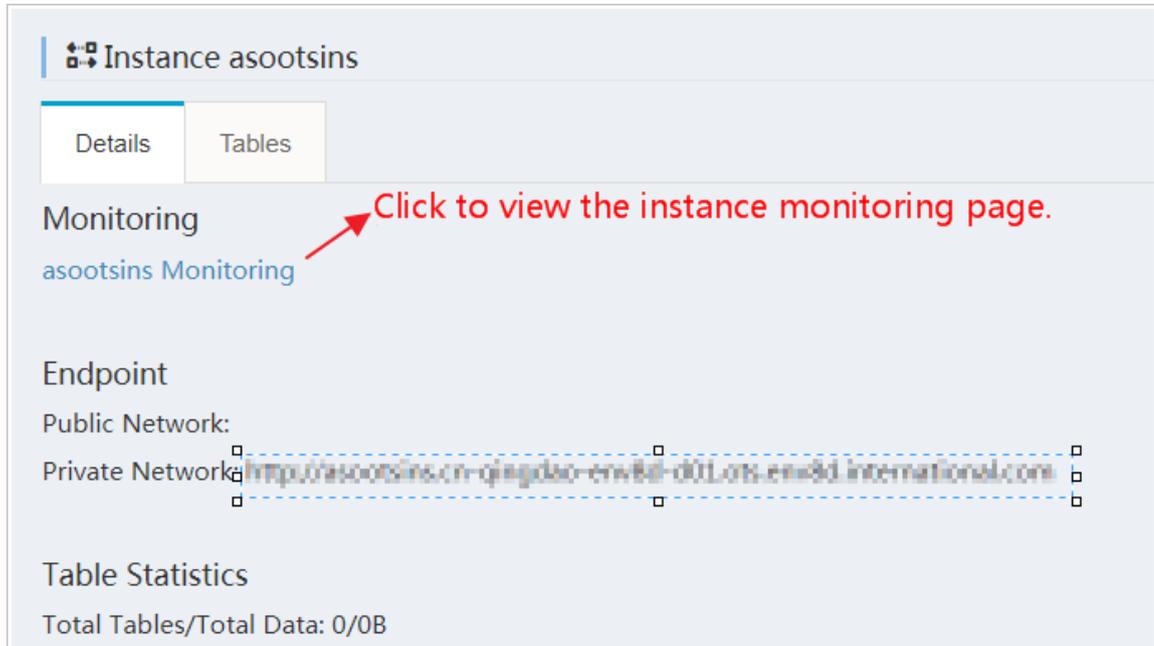


The available search conditions include:

- Instance ID
- Instance name
- User ID
- Apsara Stack account
- View instance details

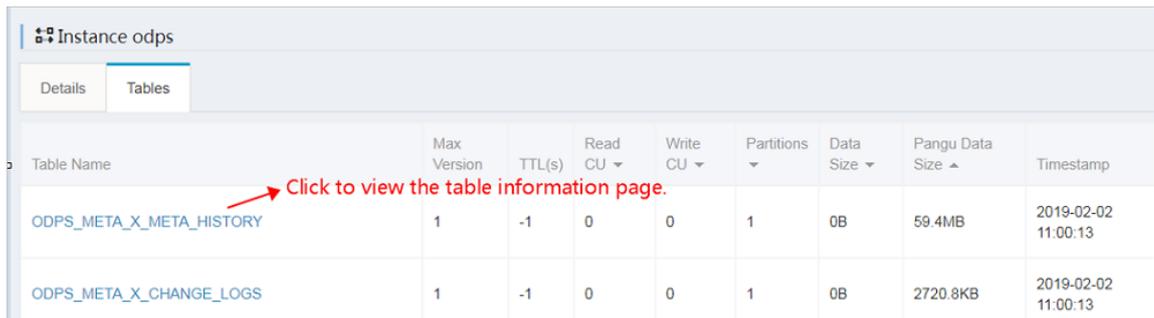
— Instance overview

Click instance name otssmoke96 to go to the **Details** tab page of the instance. This page provides detailed information about the instance, such as the instance monitoring link, intranet and Internet URLs, and statistics on tables in the instance.



— Table information

Click the **Table List** tab of the instance to view table information such as maxVersion, TTL, read CU, write CU, and timestamp.



- View table details

— Details

Click table name test_base_monitor in the table list to go to the table details page. On the Details tab page, you can view the link to the monitoring data for this table, as well as the summary information such as the number of partitions and table data size.

Table ODPS_META_X_META_HISTORY

Details | Partitions

Monitoring
[ODPS_META_X_META_HISTORY Monitoring](#)

Overview

Allow Read	true
Allow Write	true
Partitions	1
Table Data Size	0B
Pangu File Size	59.4MB

— Partition list

You can obtain the basic information of a partition, such as the partition ID and worker information. You can also specify query conditions to search for the partitions that meet your requirements.

Table ODPS_META_X_META_HISTORY

Details | Partitions

Search: Worker [] Search

ID	Partition ID	Start Key	End Key	Worker	Pangu File Size	Data Size	Youchao Files	Timestamp
1	1891d981-771c-45af-b239-84312b750ba9		\xfd\xfd\xfd\xfd...	a36f01001.cloud.f01.amtest10	59.4MB	0B	9	2019-02-02 11:00:13

The available query conditions include:

- Worker (see the value in the worker column for more information)
- Partition ID

2.4.1.3 Cluster management

2.4.1.3.1 Cluster information

You can obtain cluster information through cluster search, cluster usage, and top requests.

Function description

- Cluster list

Cluster Information				
Region: <input type="text" value="All"/>		OCM Cluster Synchronization		Refresh
Status	Cluster	Region	Storage Type	Operation
using	ots-hy-a-20181217-2e46	cn-qingdao-env8...	HYBRID	Delete
using	ots-ssd-a-20181031-25ce	cn-qingdao-env8...	SSD	Delete
using	tianji-a-25ee	cn-qingdao-env8...	HYBRID	Delete

Select All or specify a specific region to obtain a list of clusters. The functions are as follows:

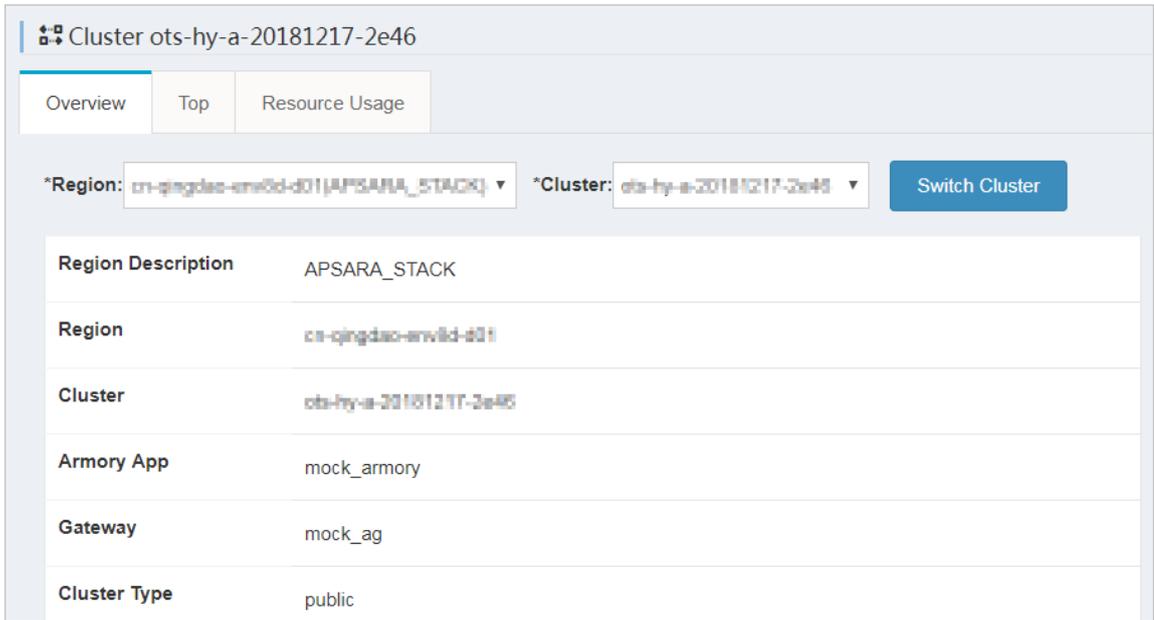
- OCM cluster synchronization: If you deploy an OCM service in each region of Table Store, the OCM service contains all cluster information of that region. This function synchronizes OCM clusters with their respective regions in Storage Operations and Maintenance System to obtain all clusters in the regions.
- Cluster deletion: You can use this function to remove a cluster from Storage Operations and Maintenance System after you confirm that the cluster is offline.

- Cluster details

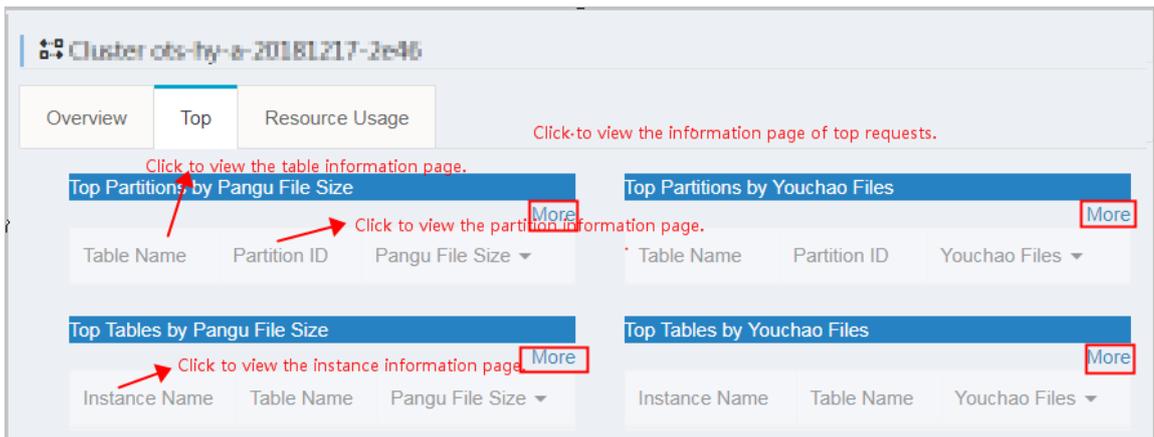
Cluster Information				
Region: <input type="text" value="All"/>		OCM Cluster Synchronization		Refresh
Status	Cluster	Region	Storage Type	Operation
using	ots-hy-a-20181217-2e46	cn-qingdao-env8...	HYBRID	Delete
using	ots-ssd-a-20181031-25ce	cn-qingdao-env8...	SSD	Delete
using	tianji-a-25ee	cn-qingdao-env8...	HYBRID	Delete

As shown in the preceding figure, you can click a cluster name to go to the cluster details page. You can view the following cluster details:

- Cluster overview: provides the basic information of a cluster.



— Top requests: provides top request information by partition and table.



— Usage: provides cluster usage details. Typically, the usage statistics collection task is automatically triggered in the backend at a specific interval. In special cases, you can click **Collect Data** to manually trigger the usage statistics collection task. After the usage statistics collection task is complete, refresh the page to display the latest usage statistics.

 **Note:**

The usage check result is either success or failure. In addition, you need to pay special attention to the cause of a usage check failure. (As shown in the following figure, the usage check failure is caused by the failure to obtain storage space.)

Cluster **ots-hy-a-20181217-2e46**

Overview | Top | **Resource Usage**

Click to manually collect resource usage information

Collected At: ~ Collect Data

Check Result :

Storage Resource Usage

Total Disk Size	Total File Size	Recycle Bin Size	Table Size	Free Space	Disk Usage Ratio (%)
					%

Gap Size | Hosts Total/Master/OTSServer/SqWorker | Hybrid Deployment | Cluster Type | Scale-out Requirement

///

OTSServer Resource Usage

Hosts	Failed Hosts	Avg/Max CPU Usage (%)	Increased CPU Cores	Avg/Max NetIn (MB/s)	Increased Hosts Due to Excessive NetIn	Avg/Max NetOut (MB/s)	Increased Hosts Due to Excessive NetOut
		/		/		/	

2.4.1.4 Inspection center

2.4.1.4.1 Abnormal usage

You can click Abnormal Usage in the left-side navigation pane to locate all cluster abnormalities and their causes.

Function description

Abnormal Resource Usage Collect Data

Cluster Name	Abnormal Resource Usage							
Date	Total Disk Size	Total File Size	Gap	Recycle Bin Size	Table Size	Free Space	Disk Usage Ratio (%)	Scale-out Requirement
2019-02-02	64.46TB	6.21TB	3.25TB	1.64TB	1.32TB	48.80TB	24.31%	3天增长, Reach Safe Level in -1Days, Growth Rate:-35.27GB/Days 3天增长, Reach Safe Level in -1Days, Growth Rate:-35.28GB/Days

You can click Abnormal Usage in the left-side navigation pane to inspect cluster abnormalities in all regions. Abnormalities are displayed in red, allowing you to quickly locate abnormal clusters.

Typically, the usage statistics collection task is automatically triggered in the backend at a specific interval. In special cases (such as a failure in backend task execution), you can click **Collect Data**

to manually trigger usage statistics collection. The triggering action is performed asynchronously. After you complete the usage statistics collection task, refresh the page to display the latest usage statistics.

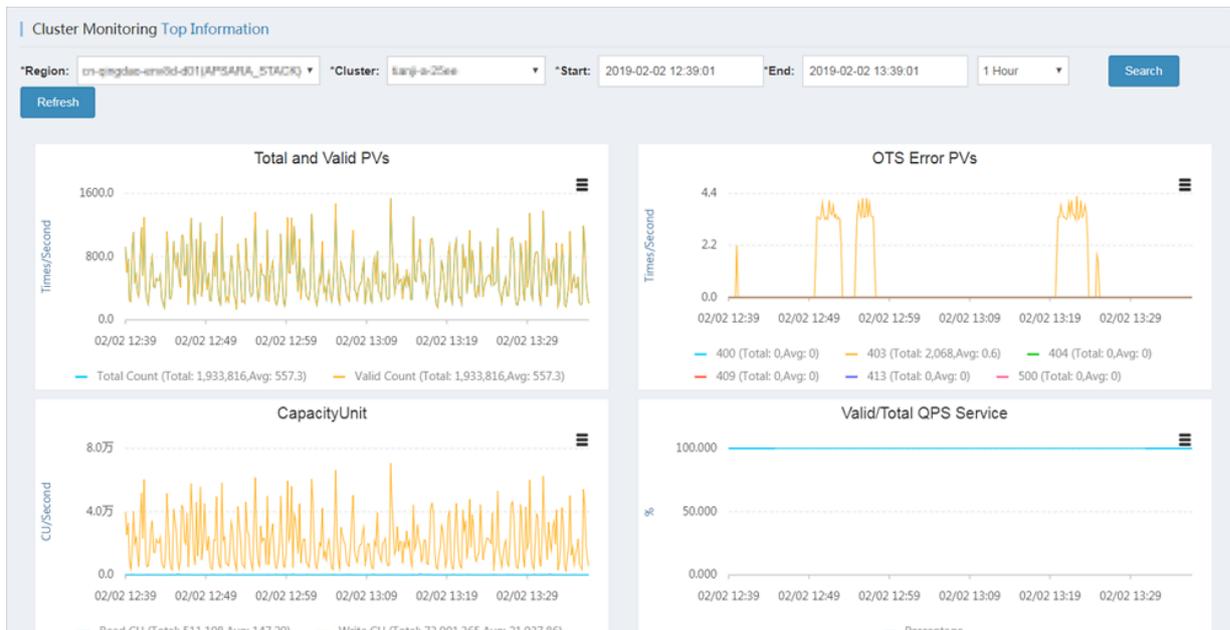
2.4.1.5 Monitoring center

2.4.1.5.1 Cluster monitoring

You can determine the service status of a cluster based on a series of metrics such as cluster-level monitoring information.

Function description

You can query the cluster service metrics in a specified time range, and determine whether a cluster service is healthy based on the monitored metrics in the following dimensions.



2.4.1.5.2 Application monitoring

By checking the instance-level and table-level monitoring metrics, you can determine whether a user's service is abnormal.

Function description

The following metrics generally reflect whether a service for a specified user is in a healthy state.



Note:

The Instance field is required, and the Table and Action fields are optional.



Function description

Four monitoring levels are supported for top requests: Instance, Instance-Action, Instance-Table, and Instance-Table-Action. You can view the top request details of a cluster based on 13 monitored metrics, such as the total number of requests and the total number of rows.

Topic	Total Requests	Total Rows	Total Failed Rows	Public Uplink	Public Downlink	Internal Uplink	Internal Downlink	Read CU	Write CU	Total Latency Max Avg	SQLWorker Latency Max Avg	HTTP Status	SQL Status
{instanceName=metric...	1,643,542	73,033,406	0	0B	0B	19.3GB	1308.2MB	245,919	73,070,441	614,911 us 13,686 us	613,801 us 12,844 us	{*200*:1643542}	{*0*:73175642}
{instanceName=odps, ...}	186,686	185,768	0	0B	0B	45.4MB	100.7MB	180,059	11,366	203,426 us 885 us	203,268 us 748 us	{*200*:186686}	{*0*:186686}

2.4.1.5.4 Request log search

You can search for a log based on a request ID to assist in problem investigation.

Function description

Query all log information about a request based on the request ID.

Host	Timestamp	File	Content
------	-----------	------	---------

2.4.1.6 System management

2.4.1.6.1 Manage tasks

You can maintain the backend tasks in Storage Operations and Maintenance System.

Function description

After the completion of Storage Operations and Maintenance System deployment in the Apsara Stack environment, the backend tasks of usage statistics collection are automatically built in Storage Operations and Maintenance System. You can perform the following operations on the backend tasks:

- View task details, and learn about the specific parameters and running time of each task.
- Enable or disable a task.



Note:

Disabled tasks will no longer run automatically.

- Execute a task immediately.

The following figure shows the task details page. Based on the monitoring rules, the task collects usage statistics at 2 am every day.

Monitoring Task Details	
Task ID	1
Task Name	collect_water_level
Task Script	
Task Script Parameter	
Remote HTTP Task URL	http://10.68.163.205/ots/apsarastack/v1/inner/httptask/run
Cluster	
Host Role	
Monitoring Rule	0 0 2 * * ?
Task Status	1
Alert Receiver Employee ID	
DingTalk Group Chat Robot Webhook	
Task Type	4
Alert Method	0
Task Result Format	0

2.4.1.6.2 View tasks

You can view the execution status of backend tasks and locate the causes of task exceptions.

The following figure shows the execution status of backend tasks in Storage Operations and Maintenance System. You can view the list of tasks, which either succeeded or failed.

Status	Name	Type	Started At	Ended At	Operation
Abnormal	collect_water_level	Remote HTTP	2019-02-02 06:00:00	2019-02-02 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-02-01 06:00:00	2019-02-01 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-31 06:00:00	2019-01-31 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-30 06:00:00	2019-01-30 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-29 06:00:00	2019-01-29 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-28 06:00:00	2019-01-28 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-27 06:00:00	2019-01-27 06:00:10	View All View Exceptions

For abnormal tasks, click **View All Tasks** or **View Abnormal Tasks** to view the specific cause of a task failure, as shown in the following figure.

collect_water_level task result

total 1 count, 0 execute success, /1 execute fail, 1 execute warning

Executelp	StartTime	EndTime	TaskResult	Warning	IsSuccess
HTTP	Feb 2, 2019 2:00:00 AM	Feb 2, 2019 2:00:10 AM	"env: APSARA_STACK, inner task collect water level fail: Trigger collect water level fail, cluster list: [ots-hy-a-20181217-2e46, ots-ssd-a-20181031-25ce]"	env: APSARA_STACK, inner task collect water level fail: Trigger collect water level fail, cluster list: [ots-hy-a-20181217-2e46, ots-ssd-a-20181031-25ce]	fail

2.4.1.7 Platform audit

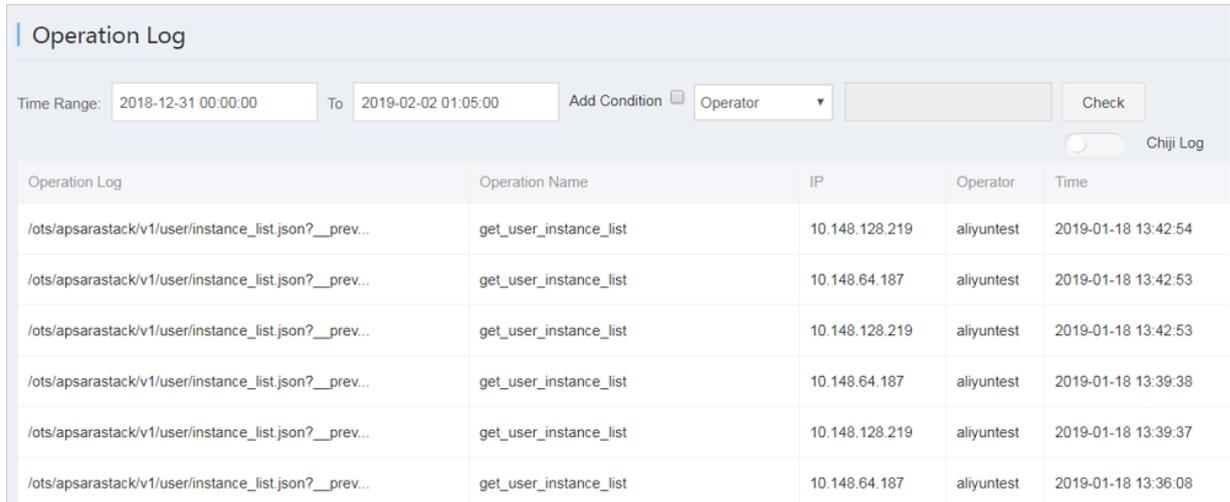
2.4.1.7.1 Operation logs

You can view the management and control operation logs of Storage Operations and Maintenance System.

Function description

The **Management Platform Operation Logs** page provides the operation logs of Storage Operations and Maintenance System. You can query audit records generated in a specified time

range and filter the records as required. This helps management personnel to quickly obtain information about the platform status.



The screenshot shows the 'Operation Log' interface. At the top, there is a search bar with 'Time Range' set from '2018-12-31 00:00:00' to '2019-02-02 01:05:00'. There is an 'Add Condition' button and a dropdown menu currently set to 'Operator'. A 'Check' button is also present. Below the search bar is a 'Chiji Log' toggle switch. The main part of the interface is a table with the following data:

Operation Log	Operation Name	IP	Operator	Time
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:42:54
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:42:53
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:42:53
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:39:38
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:39:37
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:36:08

2.4.2 Cluster environment description

Two environments are provided for Table Store: the internal environment for cloud services such as MaxCompute, Log Service, and StreamSQL, and the external environment deployed for users.

Some cloud services use both environments simultaneously. For example, metadata of StreamSQL is stored in the internal environment, but its dimension table data (user data) is stored in the external environment.

Table Store services include TableStoreOCM, TableStoreInner/TableStore, TableStorePortal, chiji, and TableStoreSqlInner/TableStoreSql.

- TableStoreOCM: the tool used to manage information about clusters, users, and instances
- TableStoreInner/TableStore: the Table Store data service node
- TableStorePortal: the backend of the Table Store O&M platform
- chiji: the Table Store O&M platform frequently used for fault location
- TableStoreSqlInner/TableStoreSql: the Table Store backend tool

2.4.3 System role description

- TableStoreOCM
 - OCMInit: the OCM initialization tool used to create tables and bind POP APIs
 - OCM: the service node of OCM
 - ServiceTest: the service test image of OCM
- TableStoreInner/TableStore

- InitCluster: the process of adding cluster information to OCM, including the domain name and type of the cluster, as well as the pre-configured Table Store account information
- LogSearchAgent: the Table Store log collection service node
- MeteringServer: the Table Store metering node (only available in Table Store)
- MonitorAgent: the data collection node of the Table Store Monitor system
- MonitorAgg: the data aggregation node of the Table Store Monitor system
- OTSAlertChecker: the Table Store alarm service module
- OTSFrontServer: the frontend server of Table Store, which can be Nginx, OTS Server, or Replication Server
- OTSServer: the OTS frontend server
- OTSTEngine: the Nginx service for OTS frontend servers
- PortalAgServer: the backend service for Storage Operations and Maintenance System
- ServiceTest: the test service that runs scheduled smoke tests
- SQLOnlineReplicationServer: the Table Store disaster recovery service
- SQLOnlineWorker: the application that was used to generate alarms but does not provide actual services now
- TableStoreAdmin: all O&M tools of Table Store, including the splitting and merging tools
- TableStorePortal
 - PortalApiServer: the backend service for Storage Operations and Maintenance System
- TableStoreSqlInner/TableStoreSql
 - Tools: the backend tool for Table Store, such as sqlonline_console
 - UpgradeSql: the backend hot upgrade tool for Table Store

2.4.4 Pre-partition a table

2.4.4.1 Pre-partitioning

After you create a table, Table Store creates a default data partition for this table. With business development, this data partition can split automatically depending on the data volume or data access load. The table with only one data partition may be unable to provide sufficient service capability during a stress test or data import. In this case, you need to complete pre-partitioning manually.

Pre-partitioning rules

You can estimate the number of partitions required based on the criteria of 10 GB data per partition. However, considering other factors such as the number of hosts and concurrent write operations by developers, we recommend that the number of partitions be no more than 256. If data can be written into the table evenly, you can partition the table equally based on the number of partitions required.



Note:

As data is written into the table, the system automatically splits the table to ensure sufficient partitions for increasing data.

Splitting methods

Pre-partitioning of a data table can be completed using `split_merge.py`, which can be obtained from `/apsara/TableStoreAdmin/split` on the host of `TableStoreAdmin` in `TableStoreInner`.

You can use any of the following methods to split a data table:

Specify splitting points

```
python2.7 split_merge.py split_table -p point1 point2 ...table name
```

Specify the number of partitions and partition key format

- Partition key of the int type

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_digit table name
```

- Partition key starting with lowercase MD5 code ([0-9, a-f])

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_hex_lower table name
```

- Partition key starting with uppercase MD5 code ([0-9, A-F])

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_hex_upper table name
```

- Base 64-coded partition key ([+0-9, A-Z, a-z])

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_base64 table name
```

- `--only_plan` : creates splitting points but does not split the table; `--force`: directly splits the table without manual confirmation.

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_digit --only_plan table name
```

Split an existing partition based on existing data volume

```
python2.7 split_merge.py split_partition -n PART_COUNT (number of
partitions) partition_id
```



Note:

The preceding methods can also be used to split a data table that already saves data.

2.4.4.2 View partitions

You can view data partitions in a data table on the control platform.

On the control platform, find the data table of a specific instance and click **Table Partitions** to view information about all data partitions in the data table, including the ID, range, worker, Pangu file size, and data size of each partition. The data size of a partition is the size of the user's original data (which may not be the real-time data because data is updated after files are merged in the background of the system). The Pangu file size is the compressed data size (the actual storage space required on the disk is three times the file size because three copies of data need to be saved).

2.5 Distributed File System (DFS)

2.5.1 Introduction

2.5.1.1 Apsara Distributed File System introduction

Apsara Distributed File System is the distributed file system component of Apsara. It integrates disks in low reliability PC servers into a whole and provides secure, stable, and easy-to-use file storage capability to external systems. Apsara Distributed File System services are provided by masters and chunkservers. Masters store and manage meta information, while chunkservers store and manage data. The common terms used in Apsara Distributed File System are listed in the following table.

Term	Description
Multi-master	<p>To prevent a single point of failure caused by a single master, multiple masters of Apsara Distributed File System are deployed . One master serves as the primary master to provide services externally, and other masters provide backup as secondary masters . Once the primary master fails to provide services externally, Apsara Distributed File System automatically switches services to one secondary master to guarantee the service availability.</p> <p>Because Apsara Distributed File System adopts the simplified Paxos algorithm for master election, we recommend that you configure an odd number of machines as masters. The recommended number of masters is five.</p>
CS	The Pangu Chunkserver used to store user data.
Supervisor	The module controlling the O&M of Apsara Distributed File System. The Supervisor module is deployed in Apsara Infrastructure Management Framework V0.16 and later. After Apsara Infrastructure Management Framework is deployed, changes in Apsara Distributed File System clusters, such as hot upgrades and CS enabling/disabling actions, can only be carried out after being approved on Supervisor

Term	Description
	. Supervisor serves as an interface between Apsara Distributed File System and Apsara Infrastructure Management Framework. It provides approval and alarm services, and is the data source of Pangu Portal.
Monitor	Monitor monitors hardware and environment statuses of machines and reports alarms to Apsara Infrastructure Management Framework when an exception occurs.
OpLog	As a service, Apsara Distributed File System keeps records of all operations in the form of operation logs (OpLogs). Once a machine restarts after fault occurrence, the data in its memory can be restored by replaying all OpLogs.
Checkpoint	To reduce the number of OpLogs, Apsara Distributed File System periodically dumps data from the memory. The dumped files are called checkpoints. After the machine restarts, the latest checkpoint is loaded into the memory. Then , instead of replaying all OpLogs, you only need to replay the OpLogs generated after the checkpoint to restore the data.

2.5.1.2 Overview of Apsara Infrastructure Management Framework

Apsara Infrastructure Management Framework is an automatic data center management system . It manages the hardware lifecycles and various static resources in the data center, such as programs, configurations, operating system images, and data. Apsara Infrastructure Management Framework provides a set of universal version management, deployment, and hot upgrade solutions for the applications and services of various Apsara and Alibaba Cloud products. It implements automatic operations and maintenance on its services in a large-scale distributed environment. Therefore, it greatly improves the operations and maintenance efficiency and system availability.

2.5.1.2.1 Basic concepts of Apsara Infrastructure Management Framework

Cluster

A logical set of physical machines that provide services.

Service

Software that provides specific functions in Apsara Infrastructure Management Framework. A cloud product is typically a service. The service name is globally unique. We recommend that you use a combination of lowercase letters with the BU name as a prefix, for example, aliyun.oss.

Each service corresponds to a service package, which is a standard tar.gz file. The directory structure of a service package must comply with the service package specifications of Apsara Infrastructure Management Framework.

A service can be deployed on a group of hardware servers, that is, a cluster, to provide the related service capabilities. For example, Apsara Distributed File System, Job Scheduler, and Apsara Name Service and Distributed Lock Synchronization System are all services.

Server role

A service can be divided into one or more server roles based on functions. A server role is an indivisible deployment unit and indicates a certain functional component of a service that runs on a hardware server. The deployment of a server role onto a server indicates that the server provides the corresponding function. Multiple server roles, for example, PanguMaster and TianjiClient, can be deployed on the same server.

We recommend that you use "Upper Camel Case" to name the server role, for example, PanguMaster. To support multiple tenants, the full name of a server role contains the service name prefix as the namespace, for example, pangu.PanguMaster.

Server role instance

Server role instance that is deployed in a cluster. A server role instance is expressed by "<ServerRoleName>#[instanceNO]", where "ServerRoleName" is the name of the server role, and "instanceNO" is the instance number. It can be a combination of letters and digits. A number of instances of different server roles can be deployed on one server in a cluster. For example, several versions of pang lib can be deployed in one cluster. Different instances of one server role are identified with a number sign (#) and a suffix, for example, PanguLib#56 and PanguLib#57.

Application

Process-level service component contained by a server role. Each application works independently. Application is the minimum unit for deployment and upgrade in Apsara Infrastructure Management Framework, and can be deployed on every server.

2.5.2 Configuration update

2.5.2.1 Overview

After configurations are modified on the portal of Apsara Infrastructure Management Framework, Apsara Infrastructure Management Framework pushes the modification to each machine where the server role resides. The process automatically detects changes on the configuration file and updates the configuration. Configuration update covers the startup parameters, flags, `apsara_log_conf.json`, and other customized configurations.

You can also update the configuration by setting the flag value (which takes effect immediately) in the memory using `puadmin` and modifying the configuration on Apsara Infrastructure Management Framework. You need to record the modification to make sure the modification remains effective after a process restart.



Note:

Configuration changes on startup parameters and some flags taking effect only after a restart are sent immediately, but do not take effect before a process restart. Therefore, before updating the configuration, make sure you understand whether the modification requires a process restart to take effect.

2.5.2.2 Operation method

Perform configurations using the portal of Apsara Infrastructure Management Framework

To modify the configuration of an Apsara Distributed File System service in running status, you can use the service configuration update function of Apsara Infrastructure Management Framework. Choose **O&M > Cluster Maintenance > Manage > Service Configuration Update** to go to the Service Configuration Update page.

Modify or add a configuration file and click Submit.

For more information about the directory structure of the configuration file, see [Configuration file structure](#).

Perform configurations using puadmin

Puadmin is included in the PanguTools server role. You can use puadmin to set and view the flags. Note: Flags configured through puadmin cannot remain persist and will become invalid after the corresponding process is restarted. For example:

Set a flag

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicateWindowSize 1572864 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicateWriteDataBlockSize 524288 -c
```

When modifying the master flag, you need to specify the volume using the `-v` parameter.

The following example uses the v2 volume. If the `-v` parameter is not specified, the default volume is configured. In the following example for obtaining a flag, the `-v` parameter is like the parameter in the flag setting.

```
/apsara/deploy/puadmin flag -set pangu_master_ReplicationTaskDestinationLengthLimit 6643777536 -v v2
```

Obtain a flag

```
/apsara/deploy/puadmin flag -get pangu_chunkserver_ReplicateWindowSize -c
```

```
/apsara/deploy/puadmin flag -get pangu_chunkserver_ReplicateWriteDataBlockSize -c
```

```
/apsara/deploy/puadmin flag -get pangu_master_ReplicationTaskDestinationLengthLimit -v v2
```

2.5.2.3 Configuration file structure

To modify configurations on the portal of Apsara Infrastructure Management Framework, you must understand the directory structure of the configuration files on the portal. Taking pangu as an example, the directory structure for configuration update is as follows:

The directory structure of configuration files of server roles pangu_chunkserver, pangu_supervisor, and pangu_master are the same.

To modify a flag, use the `user/pangu_*/pangu_*_flag.json` file. To modify the startup parameter, use the `user/pangu_*/pangu_*.json` file. To modify the log configuration file, use the `user/pangu_*/conf/apsara_log_conf.json` file.

**Note:**

The directory structure is of crucial significance and must be correct. An incorrect directory structure may result in unexpected problems.

The directory structure of `pangu_monitor` is slightly different from that of others. The configuration file under `conf` describes the metric items, rather than the log configuration. The configuration file `pangu_monitor_items.json` under `monitor` is differentiated based on the application line. For the MaxCompute application line, `cs_load_usage` is set to generate an alarm as warning when the value reaches 400 and an alarm as error when the value reaches 500. To modify the metric item values, select the corresponding product line based on the value of `CLUSTER_TYPE` in the tag. Copy the file to the portal of Apsara Infrastructure Management Framework, and modify the corresponding value.

Due to template differences, these files may not exist. In this case, add them manually. `xx.json` describes the startup parameters and `xx_flag.json` describes the flags.

**Note:**

Among all configurations, the key is a character string and the value of any type is enclosed in quotation marks. No space is allowed for the value.

2.5.2.4 Validation of configuration changes

After a task is submitted, Apsara Infrastructure Management Framework pushes the modified configuration file and calls the `update_config` script provided by the application. The `update_config` script computes and saves the configuration differences under the `config_update` folder. The `pangu_master`, `pangu_chunkserver`, `pangu_supervisor`, and `pangu_monitor` scripts periodically load the application configuration changes. Except for the startup parameters and flags that cannot take effect immediately, other configurations take effect immediately after pushing. For startup parameters and flags that cannot take effect immediately, execute the service upgrade process to make them take effect.

2.5.2.5 Overwriting relationship between configurations

The principle for configuration management is to save all configurations in the code, rather than making configuration changes directly on the cluster. Configurations are performed on the portal of Apsara Infrastructure Management Framework only when the code configuration cannot satisfy the requirement. The configuration on the portal of Apsara Infrastructure Management Framework eventually overwrites the configuration in the code.

2.5.2.6 Configuration validity checks

Configuration files are all JSON files. A schema is defined on Apsara Infrastructure Management Framework to check the syntax of JSON files. If you try to submit a JSON file with a syntax error, an error is returned and the submission fails. You need to modify the JSON file before submitting it again. This mechanism ensures all issued JSON files are legitimate. To check the validity:

Choose **O&M > Service O&M > Pangu > Service Instance**. Select a cluster and the associated schema.

In the following example, the cluster is associated with the schema `pangu_conf_check`. To ensure the correctness of the configuration syntax, use the schema to associate with the cluster of Apsara Distributed File System. If the schema is not associated, syntax check is not performed.

2.5.3 Apsara Distributed File System cluster O&M

2.5.3.1 Global flag settings of Apsara Distributed File System

For the convenience of modifying internal variables of a process, Apsara Distributed File System uses global flags to describe variables that can be modified externally. In addition, it allows users to set and view global flags using `puadmin`. `Puadmin` is an application of `PanguTools` and typically deployed on the AG. It is stored in `/apsara/deploy/puadmin`.

Usage

Syntax formats:

- `puadmin flag -set flag_name flag_value [option] [option]`
- `puadmin flag -get flag_name flag_value [option] [option]`

Examples:

- Set a flag based on the chunkserver:`/apsara/deploy/puadmin flag -set pangu_chunkserver_xxx 300 -c`

- Obtain a flag based on the chunkserver: `/apsara/deploy/puadmin flag -get pangu_chunkserver_xxx -c`
- Set a flag based on the master: `/apsara/deploy/puadmin flag -set pangu_master_xxx 100000 -m` or `/apsara/deploy/puadmin flag -set pangu_master_xxx 100000`
- Obtain a flag based on the master: `/apsara/deploy/puadmin flag -get pangu_master_xxx -m` or `/apsara/deploy/puadmin flag -get pangu_master_xxxx`

2.5.3.2 Operations on Apsara Distributed File System files

Apsara Distributed File System provides the pu tool for various file operations. For example:

- Create a folder named newdir: `/apsara/deploy/pu mkdir pangu://localcluster/newdir/`
- Create a folder named newdir: `/apsara/deploy/pu rmdir pangu://localcluster/newdir/`
- Upload the file named newfile to the newdir folder in Apsara Distributed File System: `/apsara/deploy/pu cp newfile pangu://localcluster/newdir/`
- Read the file named newfile from the newdir folder in Apsara Distributed File System to the local disk and name the file dstfile: `/apsara/deploy/pu get pangu://localcluster/newdir/newfile dstfile`
- Restore the newdir folder that has been deleted: `/apsara/deploy/pu restore pangu://localcluster/newdir/`



Note:

- Deleted files are stored in the recycle bin, that is, the `deleted` folder. The files and folders in this folder cannot be deleted by running the `pu rm` or `rmdir` command. To empty the recycle bin, run `/apsara/deploy/puadmin fs -crb -f`. If the GC function is disabled on Apsara Distributed File System, the preceding commands are invalid.
- You can run `./puadmin fs -quota pangu://localcluster/deleted/` to check the size of the `deleted` directory and the deleting process.

2.5.3.3 Common commands of puadmin

Puadmin is a common command line management tool of Apsara Distributed File System. It allows you to check the status of Apsara Distributed File System, modify its flags, and change its running status. Common commands are listed as follows:

- `puadmin cs -ls`: outputs the storage space statistics of Apsara Distributed File System.
- In the 0.16.1 version, the `-a` option is added to optimize the output information, as shown in the following figure.

```
The pangu disk status:
Total Disk Size:                75044 GB
Total Free Disk Size:           73663 GB
Total Pangu Usable Disk Size:   75044 GB
Total Pangu Usable Free Size:   73663 GB
Total File Size:                0 GB
Total User Reserved Size:       0 GB
Total User Used Size:           0 GB
Total Garbage Size:             0 GB
Total Abnormal Size:            0 GB
Redundancy Ratio:               3
TotalChunkNumber:4167           NonTempChunkNumber:4167           NonTempChunkDataSize:0 GB
```

- Check an abnormal chunkserver: `puadmin fs -abnchunk -t [none|onecopy|lessmin|lessmax]`
- Check which file an abnormal chunkserver belongs to: `puadmin fs -whois 5973023903449089`
- Delete a file: `pu rm /systest/pangu /rsld04281.et2sqa/RAFWriteRead/BlockSize_4096/16/0`
- Empty the recycle bin: `puadmin fs -crb -f`
- Check the election state: `puadmin gems`
- Check the version consistency: `puadmin env -gbi -c`
- Check a decommissioned chunkserver: `./puadmin cs -ls --puadmin_ShowDecommissionChunkserver=true |grep tcp`

2.5.3.4 GC functions of Apsara Distributed File System

On Apsara Distributed File System, the files that you delete are moved to the recycle bin first and then deleted some time later (one day by default, with the flag being `pangu_master_DelayTimeForFileGC`, in seconds). When the memory usage on `pangu_master` exceeds the threshold (85% of the total memory), Apsara Distributed File System automatically empties the recycle bin and the deleted files are cleared.

To avoid deleted files getting permanently cleared from the disk, you can set `pangu_master_ForceFileGCThreshold` to 100 to disable this function. If the function is disabled successfully, the following information is displayed: Pangu master in mode: `PANGU_MASTER_MODE_SAFE_OPEN & PANGU_MASTER_MODE_ALLOW_WRITE` (You can ignore the content after `&`.)

GC function of Apsara Distributed File System are as follows:

Disable GC

Command

```
/apsara/deploy/puadmin ms -stat --safe=on
```

Output

Master Address:

```
nuwa://localcluster/sys/pangu/master
```

Pangu master in mode:

```
PANGU_MASTER_MODE_SAFE_OPEN & PANGU_MASTER_MODE_ALLOW_WRITE
```

Output description

If PANGU_MASTER_MODE_SAFE_OPEN is included in the output, the GC is disabled successfully. In this case, deleted files are not automatically recycled by GC.

Enable GC

Command

```
/apsara/deploy/puadmin ms -stat --safe=off
```

Output

Master Address:

```
nuwa://localcluster/sys/pangu/master
```

Pangu master in mode:

```
PANGU_MASTER_MODE_SAFE_CLOSE & PANGU_MASTER_MODE_ALLOW_WRITE
```

Output description

If PANGU_MASTER_MODE_SAFE_CLOSE is included in the output, GC is enabled successfully. In this case, deleted files are automatically recycled. By default, GC is enabled on Apsara Distributed File System.

Check the GC status

Command

```
/apsara/deploy/puadmin ms -stat -M
```

Output

Master Address:

```
nuwa://localcluster/sys/pangu/master
```

```
Pangu master in mode:
```

```
PANGU_MASTER_MODE_SAFE_CLOSE & PANGU_MASTER_MODE_ALLOW_WRITE
```

Output description

If PANGU_MASTER_MODE_SAFE_CLOSE is included in the output, GC is enabled. If

PANGU_MASTER_MODE_SAFE_OPEN is included in the output, GC is disabled.

2.5.3.5 Cluster rebalance

Apsara Distributed File System implements even data storage during operation. However, because a newly added machine is less occupied after capacity expansion, data skew is caused. In addition, new machines suffer from heavy load. Therefore, we recommend that you enable the rebalance function on the background. Run the following commands to enable the rebalance function:

```
/apsara/deploy/puadmin rebalance -expand #Compute data distribution.
```

```
/apsara/deploy/puadmin rebalance -start #Start rebalance.
```

```
/apsara/deploy/puadmin rebalance -stat #Check for data migration completed by rebalance.
```

```
/apsara/deploy/puadmin rebalance -stop #Stop rebalance.
```



Note:

Rebalance needs to be manually stopped once it starts. We recommend that you stop rebalance when disk utilizations of the machines in the cluster are close to each other after rebalance is run for a period.

In the same time, to prevent the frontend reading and writing operations from being affected by rebalance on the background, the rebalance traffic is restricted. To speed up rebalance, you can set a greater traffic threshold for it using the pangu chunkserver flag in the unit of MB/s. The configuration covers the whole cluster. Run the following commands to increase and restore the traffic threshold:

Increase the traffic threshold:

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
inReadNetThroughput 300 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
axReadNetThroughput 300 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
inWriteNetThroughput 300 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
axWriteNetThroughput 300 -c
```

Restore the traffic threshold to the default value:

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
inReadNetThroughput 10 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
axReadNetThroughput 30 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
inWriteNetThroughput 10 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM
axWriteNetThroughput 30 -c
```

2.5.3.6 Directory quota operations

Set the quota

Command

Run the following command to restrict the directory quota on Apsara Distributed File System:

```
puadmin fs -quota <dir-name> --set=<entryCount,physicalSize,fileCount,
logicalSize>
```

You can set four quota values corresponding to the total number of subdirectories and files under the directory (entryCount), physical size of the files (physicalSize, disk space occupied by the files), number of files (fileCount), and logical size of the files (logicalSize).

The quaternary values can be positive integers or strings "default" and "unlimited." The string "default" indicates that the quota to which the entry corresponds remains unchanged. The string "unlimited" indicates that no quota is set for the entry.

Example

- `$puadmin fs -quota pangu://localcluster/testQuota/ --set=500,200,300, default`

Set a quota for testQuota. The maximum number of subdirectories and files under the directory is 500, the maximum physical size is 200, the maximum number of files is 300, and the maximum logical size remains unchanged.

- `$puadmin fs -quota pangu://localcluster/testQuota/ --set=500,200,300,unlimited`

Set a quota for testQuota. The maximum number of subdirectories and files under the directory is 500, the maximum physical size is 200, the maximum number of files is 300, and the maximum logical size is not restricted.



Note:

- Quota operations are directory specific. You need to add a forward slash (/) to the end of the directory to differentiate directories from files.
- The maximum number of files is irrelevant with replicas. The maximum physical file size is related to replicas. For example, (3,3) indicates that the file size is three times of the original file size.
- If the system contains RaidFile, its logical size is accurate. The physical size of a single file may have a deviation of one or several bytes due to rounding.

View the quota

Command

```
puadmin fs -quota <dir name> Or pu quota <dir name>
```

Obtain the directory quota, including the setting quota under the directory and the actual number of directories, number of files, and file size.

Example

```
$/apsara/deploy/puadmin fs -quota pangu://localcluster/apsara/
```

The output is as follows:

```
quota under pangu://localcluster/apsara/ EntryNumber  Limit:unlimited
      Used:3 FileNumber  Limit:unlimited  Used:2 FilePhysicalLength
Limit:unlimited  Used:626292888 FileLogicalLength  Limit:unlimited
Used:2087642962
```

Delete the quota

Command

```
puadmin fs -quota <dir> -r
```

Delete the quota restriction on the directory, that is, set the entry number, file number, file logic length, and file physical length to unlimited.

Example

```
$/apsara/deploy/puadmin fs -quota pangu://localcluster/apsara/ -r
```

The output is as follows:

```
quota under pangu://localcluster/apsara/ EntryNumber Limit:unlimited
  Used:3 FileNumber Limit:unlimited Used:2 FilePhysicalLength
Limit:unlimited Used:626292888 FileLogicalLength Limit:unlimited
Used:2087642962
```

2.5.3.7 Directory pin operations

You can pin a directory to prevent deletion or renaming operations on the directory by mistake.

Note: To pin a directory, you must pin its parent directory first. To unpin a directory, ensure that all its subdirectories are unpinned. Generally, you can pin a high-level directory. Directory pin operations are as follows:

Pin a directory

- **Command (To pin the full path, end the directory with a forward slash (/).)**

```
$puadmin fs -pin pangu://localcluster/apsara/
```

Output

```
Directory pinned!
```

- **Command (Pin a local directory.)**

```
$puadmin fs -pin /apsara/deploy/
```

Output

```
No Master Address specify, using nuwa://localcluster/sys/pangu/
master Directory pinned!
```

Unpin a directory

- **Command (An error is returned when a directory with a pinned subdirectory is unpinned .)**

```
$puadmin fs -unpin /apsara/
```

Output

```
No Master Address specify, using nuwa://localcluster/sys/pangu/
master -- Error: unpin:Sub dirs are pinned, please unpin all first.
Pinned dirs: deploy/
```

- **Command (Unpin the subdirectory first.)**

```
$puadmin fs -unpin /apsara/deploy/
```

Output

```
No Master Address specify, using nuwa://localcluster/sys/pangu/
master Directory unpinned!
```



Note:

An error is returned when you try to delete a pinned directory. For example, if you run `$pu rmdir -f -p /apsara/`, the following error message is returned: `error: Directory is pinned.`

Check whether a directory is pinned

Command

```
$pu dirmeta /apsara/
```

Output

```
pangu://localcluster//apsara/ Length      : 0 FileNumber      : 0
DirNumber      : 2 Pinned          : 1
```

Output description

If the value of Pinned is 1, the directory is pinned. If the value is 0, the directory is not pinned.

2.5.4 Master O&M

2.5.4.1 Overview

Apsara Distributed File System has started to use multiple masters since version 0.16. During the operations and maintenance, you can switch services from the primary master to a secondary master, view the status, replace masters, synchronize logs, and expand volumes. Currently, masters can compose a federation. The default volume is accessed by default. To access a specified volume, add the `-v` parameter.

For example, you can run the following command to list all volumes under a master:

```
/apsara/deploy/puadmin vol -ls
```

Output:

```
Name: v2 Name: PanguDefaultVolume
```

It indicates that the master has two volumes. If the `-v` parameter is not added, the command takes effect only on the PanguDefaultVolume volume. To perform operations on the v2 volume, add `-v v2` to the command. For example:

```
/apsara/deploy/puadmin lscs: List chunkservers of the default volume.
```

```
/apsara/deploy/puadmin lscs -v v2: Lists chunkservers of the v2 volume.
```

2.5.4.2 Primary master switchover

Usage scenario

A master group consists of multiple masters. Only one master is in primary state. When an exception occurs on the primary master, the PE assesses the situation and switches services from the primary master to a secondary master if necessary. Assume the primary master is A, the secondary masters are B and C, and services need to be switched from the primary master to B.

Command

```
/apsara/deploy/puadmin ms -sp B.tcpAddress
```



Note:

To obtain the TCP address of B, run `/apsara/deploy/puadmin gems`.

Operation result

- If the operation succeeds, `Switch primary form A.tcpAddress to B.tcpAddress succeed.` is displayed.
- If the operation fails, services are probably switched to C.
- You can run `/apsara/deploy/puadmin ms -elec` to check whether primary master switchover succeeds.

2.5.4.3 Status check for multiple masters

You can view the election status of multiple masters and the status of log synchronization among multiple masters.

2.5.4.3.1 View the election status

Command

```
/apsara/deploy/puadmin ms -elec
```

Output result

```
ElectMasterStatus : ELECT_MASTER_OVER_ELECTION PrimaryId      :
tcp://10.101.164.1:10260 PreferredWorkerid   : PrimaryLogId   :
882290350 TotalWokerNumber   : 3 ElectConsentNumber : 2 SyncConsen
tNumber   : 2 ElectSequence     : [b2bca55e-530d-49ad-96d6-f3e564aece
6d,1,1782091301] WorkerStatus   : tcp://10.101.164.10:10260 :
ELECT_WORKER_STATUS_SECONDARY tcp://10.101.164.12:10260 : ELECT_WORK
ER_STATUS_SECONDARY tcp://10.101.164.1:10260 : ELECT_WORKER_STATUS_
PRIMARY
```

Output result description

The current environment has one primary master. Its TCP address is tcp://10.101.164.1:10260. (See: tcp://10.101.164.1:10260 : ELECT_WORKER_STATUS_PRIMARY.)

The current environment has two secondary masters and their TCP addresses are

tcp://10.101.164.10:10260 and tcp://10.101.164.12:10260 respectively.

(See: tcp://10.101.164.10:10260: ELECT_WORKER_STATUS_SECONDARY and tcp://10.101.164.12:10260: ELECT_WORKER_STATUS_SECONDARY.)

2.5.4.3.2 View the status of log synchronization among multiple masters

Command

```
/apsara/deploy/puadmin ms -elec -s
```

Output result

```
PrimaryStatus : PRIMARY_STARTUP_SERVICE_STARTED PrimaryCurrentLogId
: 882290350 WorkerSyncStatus : tcp://10.101.164.10:10260[SyncedLogI
d:882290350, LastFailTime:1970-01-01 08:00:00, WorkerType: NORMAL,
LogGap:0] tcp://10.101.164.12:10260[SyncedLogId:882290350, LastFailTi
me:1970-01-01 08:00:00, WorkerType: NORMAL, LogGap:0]
```

Output result description

The status log ID of the master is 882290350, a monotonically increasing sequence.

The statuses of secondary masters are NORMAL and VIRTUAL.

A master in VIRTUAL status can only serve as a secondary master and receive logs synchronized from the primary master. It has no influence on the service.

2.5.4.4 Rules for generating .cpt and .log files

The .cpt and .log files generated on Apsara Distributed File System masters are named in the following formats: `pangu_master_op.logId.cpt` and `pangu_master_op.logId.log`. In the file name, `logId` indicates the memory state when the .cpt file is created and the .log file state. Each .cpt file is followed by a .log file with the same `logId`. The .log file cannot be empty. Otherwise, the master cannot use the .cpt and .log file pair to restore the memory.

For example, the following files exist:

```
/apsarapangu/pangu_master_op.673284360.cpt
```

```
/apsarapangu/pangu_master_op.673284360.log
```

```
/apsarapangu/pangu_master_op.673285365.log
```

```
/apsarapangu/pangu_master_op.673517452.cpt
```

```
/apsarapangu/pangu_master_op.673517452.log
```

`pangu_master_op.673517452.cpt` and `pangu_master_op.673517452.log` are the latest .cpt and .log file pair.

2.5.4.5 Master replacement

You can replace a master.

2.5.4.5.1 Procedure

Apsara Distributed File System has started to use multiple volumes since version 0.15. Each master belongs to a volume.

To replace a master, choose **O&M > Cluster O&M > Cluster O&M > Target Cluster > Manage > Change Machine** to go to the Change Machine page.

The machines included in `PanguMaster#` are specified by the machine group `PanguMater#Volume_*`. To replace a master, you only need to modify the machine group `PanguMater#Volume_*`.

For example, there are two volumes. To replace the master in the default volume, you only need to modify the `PanguMaster#Volume_` machine group.



Note:

- The newly added machine must have no master oplogs under `/apsarapangu` and `/apsarapangu/backup`, and its `/apsarapangu/conf/` folder must be empty.
- Only one master can be replaced at a time. Check that the rolling operation is completed on Apsara Infrastructure Management Framework and puadmin confirms the replacement before you start the next replacement.
- To ensure data security, do not replace three masters consecutively. If you need to replace three masters, perform the following steps before each replacement:
 1. Run the `/apsara/deploy/pu touch pangu://localcluster/xxx` command to generate one oplog.
 2. Send the `/apsara/deploy/puadmin ms -dump` command to three masters to enable every master to generate a checkpoint. Confirm that the new checkpoint has been generated on each master before going to the next step. By default, the checkpoint of a master is under the `/apsarapangu/` directory and in the file name format similar to `pangu_master_op.370748130.cpt` (with different numbers). Check that the file generation time is later than the command running time.
 3. Run the `/apsara/deploy/pu rm pangu://localcluster/xxx` command.
 4. Modify configurations on the portal of Apsara Infrastructure Management Framework and then replace the master.
- If the master is damaged, you can replace it with a new one using the master replacement process.
- If a master in the master group is damaged but you want to replace another master, the undamaged master cannot be replaced.
- Machine replacement is not allowed across machine groups.

For example, assume there are two volumes in an environment: `PanguMater#Volume_1` : [1, 2, 3] and `PanguMater#Volume_v2` : [4, 5, 6].

The server roles are `PanguMaster#`: `PanguMater#Volume_1` | `PanguMater#Volume_v2`.

In this case, replacing [1, 2, 3] and [4, 5, 6] with [1, 5, 3] and [4, 2, 6] results in volume chaos. Therefore, it is not allowed.

On some application lines, the supervisor and the `pangu_master` are deployed on one machine. Consequently, if you put a machine replaced by another out of service directly, the supervisor may get offline as well. If all supervisors are offline, no supervisor approves the follow-up operations and maintenance, and the operations and maintenance will get stuck. Therefore, you need

to modify the machine group of supervisors during master replacement to prevent putting the supervisor out of service by mistake.

**Note:**

Replacement of `pangu_master` is a high-risk operation. Before one replacement is completed, do not start another replacement. A replacement is done when:

- The rolling operation is completed on Apsara Infrastructure Management Framework.
- On the portal of Apsara Infrastructure Management Framework, the number of in-service/out-of-service machines in the corresponding cluster is 0.
- Use `/apsara/deploy/puamin gems` on the AG to confirm that the new in-service machine works normally and the out-of-service machine cannot be found.

2.5.4.5.2 Procedure for replacing a master

The master replacement process involves two steps, one for getting the new master into service and one for getting the existing master out of service. The supervisor is not informed when a new master is put into service. The supervisor only receives the application request for getting a master out of service. The supervisor then sends an out-of-service command to the master. The primary master deletes the existing master and returns OK to the supervisor only when a new master joins the cluster and logs are synchronized to the cluster. After the supervisor receives the OK message, it approves the replacement Apsara Infrastructure Management Framework. Then, Apsara Infrastructure Management Framework gets the existing master out of service.

If the existing master is the primary master, the supervisor switches services from the primary master to a secondary master before sending the out-of-service command. If the new master is not in service, the out-of-service task for the existing master remains unapproved until the task times out.

The supervisor approves the task after the new master gets into service and completes log synchronization. The log synchronization time depends on the size of the existing `.cpt` file, the difference between the `.log` file and the `.cpt` file, and the pressure on the existing master. At present, the `.cpt` file copy speed between masters is restricted to 50 MB/s. The `.cpt` file synchronization speed between masters is 35000 entries per second. If the `.cpt` file is 20 GB in size and the `.log` file contains 10 million entries, the new master takes 410 ($20 \times 1024/50$) seconds to synchronize the `.cpt` file. If the pressure on the existing master is 15000 entries per second, it takes 500 [$1000/(3.5 - 1.5)$] seconds to synchronize 10 million log entries. The total time is 910 seconds, that is, 15 minutes.

2.5.4.5.3 Manual replacement procedure

When automatic replacement fails, you can manually replace the new master after it starts properly.

Before manual replacement, confirm with the self-service team that the supervisor is unable to approve the replacement. Assume that the IP address of the master to be brought into service is new, the IP address of the master to be put out of service is old, and the new master has started properly.

If it does not start properly, use Apsara Infrastructure Management Framework to deploy and start a new master. If the new master cannot be started, contact the support team of Apsara Infrastructure Management Framework.

After the new master is started properly, perform the following operations: Perform the following steps on the AG:

1. Run the `puadmin gss` command and check whether the new master is included in the command output, to determine whether the new master has been added to the master group. If it is included in the command output, go to step 2. If it is not included in the command output, run the `/apsara/deploy/puadmin ms -elec --role --add=tcp://new:10260 -virtual yes` command to add the new master to the master group and set the status of the new master to virtual.
2. Run the `puadmin gss` command to view the synchronization status, as shown in the following figure.

```

$./puadmin gss
PrimaryStatus : PRIMARY STARTUP_SERVICE_STARTED
PrimaryCurrentLogId : 5068945
WorkerSyncStatus :
  tcp://100.81.240.140:10260 [SyncedLogId:5068945, LastFailTime:]
  tcp://100.81.240.143:10260 [SyncedLogId:5068945, LastFailTime:]

```

In the preceding figure, the red rectangle indicates the ID of the primary master, while the brown rectangle indicates the log ID of the new master. If the difference of the two IDs is less than 10000, it can be inferred that log synchronization is completed.

3. After synchronization, run the `/apsara/deploy/puadmin ms -elec --role --add=tcp://old:10260 -virtual yes` command to set the status of the old master to virtual.
4. Run the `/apsara/deploy/puadmin ms -elec --role --add=tcp://new:10260 -virtual no` command to set the status of the new master to normal.

5. Run the `/apsara/deploy/puadmin ms -elec --role --rm=tcp://old:10260` command to delete the old master.

2.5.4.6 Manually synchronize logs between the primary master and a secondary master

Usage scenario

If automatic log synchronization fails due to a large log quantity difference resulting from a fault or mis-operation, you need to manually synchronize logs between the primary master and a secondary master.

Prerequisites

- This operation is performed without stopping the Apsara service.
- Multiple masters of Apsara Distributed File System are running.

Procedure

1. Run the following command to make the primary master generate a new checkpoint file.

```
/apsara/deploy/puadmin ms -dump
```

2. Copy the new `.cpt` and `.log` files generated on the primary master to the `/apsarapangu/` directory of the secondary master.
3. Ensure that the latest `.cpt` file and all the subsequent `.log` files are copied. Restart the `pangu` process on the secondary master.



Note:

By default, the `/apsara/deploy/puadmin ms -dump` command is sent to the primary master and a successful sending message is displayed. Wait until the files are generated.

2.5.4.7 Rename a chunkserver online

Usage scenario

If the chunkserver name configured on a machine is incorrect during cluster deployment, you need to rename the chunkserver without service interruption. In this case, you can use `puadmin` to change the chunkserver name recorded in the master memory.

For example, you can change the name of the chunkserver on `tcp://10.101.164.7:10260` to `rs1d04271.et2sqa_new`.

Command (for renaming a chunkserver)

```
/apsara/deploy/puadmin ms -modify --csm -n tcp://10.101.164.7:10260,
rs1d04271.et2sqa_new
```

Expected output

Modify chunkserver name success.

Command (for result confirmation)

```
/apsara/deploy/puadmin cs -ls --service=tcp://10.101.164.1:10260 | grep
tcp://10.101.164.7:10260
```

Expected output

```
31. NORMAL (80000/129228) (ttl= 20) tcp://10.101.164.7:10260 rs1d04271.
et2sqa_new SendBuffer : 0(KB)
```

**Note:**

If there are multiple masters, confirm that the new chunkserver name meets expectations and the name has been changed on every master. After the change, a checkpoint is generated on every master.

Command (for a master to generate a checkpoint)

```
/apsara/deploy/puadmin ms -dump --Server=tcp://10.101.164.1:10260
```

Expected output

Start to generate checkpoint now.

2.5.4.8 Multi-master tools

Command	Description
/apsara/deploy/puadmin ms -elec	Views the election status.
/apsara/deploy/puadmin ms -elec -m tcp://10.138.26.24:10260	Views the ElectWorker status.
/apsara/deploy/puadmin ms -elec -s	Views the status of oplog synchronization among multiple masters.
/apsara/deploy/puadmin ms -sp tcp.Address	Switches over the primary master.
/apsara/deploy/puadmin ms -elec -M tcpAddress	Views all the ElectWorker metrics and their values of Apsara Distributed File System.

Command	Description
<code>/apsara/deploy/puadmin ms -elec -l</code>	Obtains the latest oplog.

2.5.5 Chunkserver O&M

2.5.5.1 Set the chunkserver status

You can manually set the chunkserver status. The available states include: NORMAL, READONLY, and SHUTDOWN. This topic describes how to set the chunkserver status to SHUTDOWN.

Command

```
/apsara/deploy/puadmin cs -stat tcp://10.101.164.7:10260 --set=SHUTDOWN
```

Expected output

```
set chunkserver status: SHUTDOWN
```

Setting result check

Run the `/apsara/deploy/puadmin lscs | grep tcp://10.101.164.7` command to check the setting result.



Note:

After the command is executed, the check takes effect only after the master detects that the chunkserver enters the SHUTDOWN state. This process takes less than one minute.

Expected output

```
31. SHUTDOWN (NA) (ttl= 20) tcp://10.101.164.7:10260 rs1d04271.
et2sqa_new SendBuffer : 0(KB), Backup: Doing
```

2.5.5.2 Set the disk status

You can manually set the chunkserver disk status. The available states include OK, SHUTDOWN, and ERROR. This topic describes how to set the disk status to DISK_ERROR.

Command

```
/apsara/deploy/puadmin cs -stat tcp://10.101.164.7:10260 -d 1 --set=
ERROR
```

Expected output

```
set disk status success. After set : DISK_ERROR
```

Setting result check

Run the `/apsara/deploy/puadmin cs -stat tcp://10.101.164.7:10260 -d 1` command to check the setting result.

Expected output

```
DiskId:1 DiskStatus:DISK_ERROR
```



Note:

If a cluster involves both the storage and journal scenarios, do not change the status of the disks and SSDCache from ERROR to OK. Instead, use the function for automatically bringing empty disks into service, which is enabled by default.

If the status of a disk is changed from ERROR to OK in the preceding scenario, the following output is generated:

```
[admin@e18g06550 /apsarapangu/disk5/zhousu/dailycores/core_10.2017011117]
$/apsara/deploy/puadmin cs -stat tcp://100.81.240.125:10260 -d 1 --set=ok
Set disk status failed. Disk status: DISK_ERROR
--
Error: stat: For disk status, only OK/SHUTDOWN/ERROR are supported. READONLY is not supported now.
For chunkserver status, only NORMAL/READONLY/SHUTDOWN are supported.
```

2.5.5.3 Chunkserver scale-out and scale-in

2.5.5.3.1 Procedure

To perform chunkserver-based capacity expansion/reduction, choose **O&M > Cluster O&M > Target Cluster > Management > Change Machine** to go to the Change Machine page.

Change the machine list of PanguChunkserver#. To expand the capacity, add entries to the list.

To reduce the capacity, remove entries from the list. Then, submit the change.



Note:

The added machines must belong to the current cluster. The machines in the default cluster can be used by this cluster only after they are added to this cluster via cluster capacity expansion.

2.5.5.3.2 Capacity expansion procedure

You can expand the capacity by adding machines to the machine list of the chunkserver server role. The supervisor does not need to participate in chunkserver-based capacity expansion. The machines should belong to the target cluster. If not, add them to the cluster first via cluster capacity expansion. Apsara Infrastructure Management Framework deploys and starts the chunkserver on the machine.

**Note:**

- The started chunkserver should be registered with Apsara Name Service and Distributed Lock Synchronization System. The NuwaConfig server role and NuwaLib server role should be deployed on the machine before capacity expansion, to ensure correct Nuwa configuration. If the added machine is used by another cluster and contains data of the cluster, the chunkserver without the NuwaConfig server role and NuwaLib server role may read dirty data and therefore register with Apsara Name Service and Distributed Lock Synchronization System of another cluster.
- The new chunkserver should be a machine without `/apsarapangu/pangu_chunkserver_op*`, which is chunkserver metadata. In the test environment, a new chunkserver contains chunkserver metadata. As a result, the chunkserver enters its original state after being started. In this case, clear the metadata before expanding the capacity with the chunkserver.
- If there are multiple zones, add the new chunkserver to the zones after capacity expansion. Otherwise, the chunkserver cannot be used.

2.5.5.3.3 Capacity reduction procedure

Same as chunkserver-based capacity expansion, chunkserver-based capacity reduction is performed by modifying the chunkserver server role list. You can reduce the capacity by removing chunkservers that are not needed from the list.

The capacity of Apsara Infrastructure Management Framework can be reduced only with the supervisor's approval. After receiving a chunkserver disconnection task, the supervisor sends the SHUTDOWN command to the corresponding chunkserver, waits for chunkserver data replication to complete, and then sends the approve command to Apsara Infrastructure Management Framework. After receiving the approve command, Apsara Infrastructure Management Framework stops the process of the chunkserver. The supervisor keeps sending the decommission command to the master to delete the chunkserver from the cluster. If the chunkserver process stops and

the chunkserver enters the DISCONNECTED state, the decommission command is executed correctly. The chunkserver is removed from the cluster.

The number of chunkservers that can be removed at a time depends on a flag of the supervisor. The flag is named `pangu_supervisor_MaxConcurrentChunkserverShutdownCount` and has the default value 1, indicating that one chunkserver can be removed at a time. To remove multiple chunkservers at a time, change the value.

The period from the time when a chunkserver process is stopped to the time when the chunkserver enters the DISCONNECTED state depends on two flags of `pangu_master`, namely, `pangu_chunkserver_normal_ttl` and `pangu_chunkserver_disconnecting_ttl`. The default value of the first flag is 4, while the default value of the second is 16, so the total period is $(4 + 16) \times 15/60 = 5$ minutes. Set the first flag on ODPS to 40, so that the total period changes to 14 minutes. This means that a chunkserver is completely out of service at least 5 minutes or 14 minutes after data replication.

Chunkserver disconnection duration

The supervisor performs approval only after chunkserver data replication is completed. The data replication time depends on the number of nodes in a cluster and the amount of data on the chunkserver to be disconnected. Currently, the maximum data replication speed is limited to 30 MB/s. Take the AY44B typical to OSS as an example. Most of the chunkservers are approaching their maximum storage capacity. The disk capacity totals 44 TB, the cluster has 500 chunkservers, and data replication is performed at 30 MB/s. When all the chunkservers replicate data, the estimated time required to disconnect one chunkserver is $44 \times 1024 \times 1024 / (30 \times 500) = 3075$ seconds, about 50 minutes. If the chunkservers have high workload and cannot replicate data at the maximum rate, the time can be lengthened to as many as 150 minutes (at the minimum data replication speed). If most of the data are RAID files (8+3), the time is multiplied by 8 and increased to 400 to 1200 minutes. Therefore, the estimated time of data replication on the layer of Apsara Distributed File System for disconnecting a chunkserver under AY44B is 50 to 150 minutes (without RAID files), or 400 to 1200 minutes (with RAID files (8+3)). If the current state of the chunkserver is DISCONNECTED, the supervisor directly approves the request.

**Note:**

Before disconnecting a chunkserver, add the chunkserver to the blacklist of Job Scheduler. Otherwise, when TempFile is used, Job Scheduler may distribute instances to the chunkserver that has been shut down, which results in a failure to write TempFile.

When a chunkserver is disconnected, the supervisor automatically adds the chunkserver to the blacklist of Apsara Distributed File System, so that no more new chunk will be distributed to the chunkserver. You can run the following commands to set and view the blacklist:

- Set the blacklist: `/apsara/deploy/puadmin upgrade -util -sbcs "tcp://x.x.x.x:10260, tcp://x.x.x.y:10260"`
- View the blacklist: `/apsara/deploy/puadmin upgrade -util -gbc`

2.5.5.3.4 Disable manual chunkserver-based capacity reduction using puadmin

To reduce the capacity in Apsara Infrastructure Management Framework, use the standard capacity reduction procedure instead of manually running puadmin to shut down the chunkserver. After data replication is completed, run the puadmin command to decommission the chunkserver from the cluster. Failing to do so results in inconsistency between the chunkserver list on Apsara Infrastructure Management Framework and that on pangu_master. For example, the number of chunkservers recorded on pangu_master is 100, while the number recorded in Apsara Infrastructure Management Framework is 120. In this case, when the chunkservers not existing in pangu_master are used, the supervisor considers that the chunkservers do not belong to the current cluster and reject requests consequently to ensure security.

Therefore, do not manually decommission chunkservers using puadmin. Use the standard capacity reduction procedure instead.

2.5.6 Cluster status tracking

2.5.6.1 View the cluster status

After initiating a task on Apsara Infrastructure Management Framework, you can view the status of the task on the portal of Apsara Infrastructure Management Framework. Choose **Upgrade Task > Run Task** to view the list of running tasks. You can select the related cluster to view details.

To view historical tasks, choose **Upgrade Task > Run Task**.

You can view details by double-clicking an entry.

2.5.6.2 Submission history

You can view the history of operations on a cluster on Apsara Infrastructure Management Framework. Choose **Tianji O&M > Manage Cluster**, select the related cluster, and choose **Monitor > Operation Log**.

On the page, you can perform the following operations:

- View the historical operations on the cluster.
- Compare any two submissions.
- View difference details and the content of each change.

2.5.7 FAQs

2.5.7.1 Identify the machine where pangu_supervisor runs and the log location

On the AG, run the `/apsara/deploy/nuwa_console --address=nuwa://localcluster/sys/pangu/supervisor` command to obtain the IP address of the machine where the supervisor runs.

The log path of pangu_supervisor is `/apsara/pangu_supervisor/log/pangu_supervisor.LOG`.

2.5.7.2 Adjust the flag of pangu_supervisor

If you need to adjust the flag of the supervisor on the portal of Apsara Infrastructure Management Framework during operations and maintenance, wait until the current rolling ends. If you need to effect the change immediately, use puadmin to set the flag. The flag set using puadmin becomes invalid upon restart.

It should be noted that, puadmin automatically increases the port number by 1 due to a historical issue. If the default listener port number of the supervisor is 10263, change it to 10262 when using puadmin to adjust the flag of the supervisor.

To view the flag, run the following command: `/apsara/deploy/puadmin flag -get pangu_supervisor_MaxTolerantFailedChunkserver -s tcp://supervisor_ip_addr:10262`.

To set the flag value, run the following command: `/apsara/deploy/puadmin flag -set pangu_supervisor_MaxTolerantFailedChunkserver 100 -s tcp://supervisor_ip_addr:10262`.

2.5.7.3 Supervisor approval problems

2.5.7.3.1 Supervisor approval prerequisites

Before obtaining an approval from the supervisor, ensure the cluster and master group are normal, all the chunkservers are in NORMAL state, and inter-cluster data replication is not in process.

The checklist is as follows:

- `CLUSTER_TYPE` in `tag.conf` is incorrect, which results that the supervisor cannot be started.
- If the supervisor cannot be started, when you manually run `/apsara/pangu_supervisor/start`, noticeable error information is generated.
- If the master group has no primary master, check whether Apsara Name Service and Distributed Lock Synchronization System is available.
- If the master logs are not synchronized, the number of logs of the primary master differs greatly from that of the secondary master. The default number is 100.
- If the `.cpt` file has not been created by masters for a long time, the logid contained in the `.cpt` file name differs greatly from the current logid. The default value is 1048576.
- The number of abnormal chunkservers in the cluster exceeds the limit (10 for ODPS and 2 for others) specified by `pangu_supervisor_MaxTolerantFailedChunkserver`. This value affects hot upgrades, but does not affect chunkserver disconnection. Before an upgrade, add the abnormal chunkservers to the blacklist of the supervisor. The blacklisted chunkservers are not included in the statistics on abnormal chunkservers. To add them to the blacklist, run the following command: `/apsara/deploy/puadmin flag -set pangu_supervisor_ChunkserverBlackList "csip1,csip2" -s tcp://supervisor_ip_addr:10262`.



Note:

Separate multiple chunkserver IP addresses with a comma (,) and do not include tcp addresses or port numbers.

- The cluster contains abnchunks. You can run `/apsara/deploy/puadmin fs - abnchunk -t lessmin` to check abnchunks. The logs of `pangu_supervisor` contain the `HasAbnormalChunkserverForMasterVolume` information.
- After the supervisor approves the masters during migration, masters of Apsara Distributed File System 0.15.3 are terminated by Apsara Infrastructure Management Framework. It takes one or several hours for the master of Apsara Distributed File System 0.16 to start. Although the master process exists, it stays in the DISCONNECTED state for a long time. This problem

occurs when the disk speed is low and `pangu_OperationLogSyncDisk` is set to true. Before an upgrade, set this value to false in the template to accelerate startup.

- If the number of replica tasks of the cluster exceeds the limit specified by `pangu_supervisor_DefaultOngoingReplicaForUpgrade`, approval congestion occurs. To query the tasks, run the following command: `/apsara/deploy/puadmin rep -stat`.
- On some machines, a hostname may vary across volumes, while the corresponding IP address does not. In this case, you can run `puadmin lscs -v volumename | grep tcp | awk '{print $6 $7}' | sort` to check whether a hostname varies across volumes.

If yes, perform the following steps:

1. Run `curl "127.0.0.1:7070/api/v3/column/m.id?m.ip=10.101.162.176"` to view the hostname of the IP address in Apsara Infrastructure Management Framework.
 2. Refer to [Rename a chunkserver online](#) and change the hostnames for consistency across volumes.
- Some server roles of `pangu_master` are not in GOOD state. To prevent the problem that dual masters are unavailable after a task is stopped and then restarted, the supervisor checks the version and status of the `pangu_master` server role before approval. If it is not in GOOD state, the supervisor rejects the request. In most cases, `pangu_monitor` generates an error-level alarm, which results that the server role enters the PROBATION state. The supervisor has a flag that can be used to skip checking the server role status. You can use the flag if you are certain about the server role status. To use the flag, set `pangu_supervisor_EnableCheckServerRoleState` to false.

2.5.7.3.2 Failure to approve master replacement

The possible causes are as follows:

- The new master does not start.
- The logs of the new master have not been synchronized.

2.5.7.3.3 Unable to approve chunkserver disconnection

If the prerequisites have been met, the possible cause is that the number of chunkservers is smaller than the mincopy value and data replication fails. The solution is as follows:

1. Identify the files that have not been copied: If disconnection of a chunkserver fails within a long time, access `/apsara/pangu_chunkserver/log` on the chunkserver that is in the SHUTDOWN doing state, open `pangu_chunkserver.LOG`, and check whether the value of

the `safeRemove` field is true or false. If it is false, search the previous log for the file ID of the file related to the disconnection failure.

2. Check the `mincopy` value: On the AG, run `/apsara/deploy/puadmin gfi FileId` and search for the `minCopy` value of the file. If the value is greater than the current number of chunkservers, data cannot be replicated.
3. Run `/apsara/deploy/puadmin whois FileId` to convert the file ID into a file name.
4. Run `/apsara/deploy/pu setreplica filename 3 5` to set `mincopy` and `maxcopy` to 3 and 5 respectively.

2.5.7.4 Accelerate chunkserver disconnection

By default, chunkservers are disconnected in serial. To accelerate disconnection, you can change the value of `pangu_supervisor_MaxConcurrentChunkserverShutdownCount`. For more information, see [Adjust the flag of pangu_supervisor](#). The value indicates the number of chunkservers that can be disconnected at the same time. If it is set to 100, a maximum of 100 chunkservers can be disconnected at the same time.

2.5.7.5 Rolling failure during a hot upgrade of Apsara Distributed File System

If you initiate a hot upgrade but final rolling fails despite it has the supervisor approval, the most possible cause is that an ERROR alarm is generated by `pangu_monitor` during probation. The portal of Apsara Infrastructure Management Framework provides an error machine list. You can search `/apsara/pangu_monitor/monitor/log/monitor.LOG` on the corresponding machine for the alarm by `level.*error`.

2.5.7.6 Manually replace binary files in emergency

2.5.7.6.1 Manual overwrite

Apsara Infrastructure Management Framework protects its files and automatically fixes the files when they are changed. This makes it difficult to manually replace binary files. Apsara Infrastructure Management Framework provides the overwriting method to address this problem. The following example uses `pangu_supervisor` to describe the procedure:

1. Copy of the deployment directory of `pangu_supervisor`. Assume that `buildid` is 2861. Run the following commands:

```
cp /cloud/app/pangu/PanguSupervisor#/pangu_supervisor/2861
/cloud/app/pangu/PanguSupervisor#/pangu_supervisor/overwrite -rf
```

2. Under `service_manage`, create a configuration file using the following naming rule: `service name.server role name.app name`. The full path is as follows: `/cloud/data/tianji/TianjiClient#/service_manager/overwrite.d/pangu.PanguSupervisor#.pangu_supervisor`.
3. The file content is as follows:

```
{ "service_name": "pangu", "sr_name": "PanguSupervisor#", "app_name": "pangu_supervisor", "against_work_dir": "/cloud/app/pangu/PanguSupervisor#/pangu_supervisor/2861", "work_dir": "/cloud/app/pangu/PanguSupervisor#/pangu_supervisor/overwrite", "expired_time": 9456814356 }
```

The parameters are described as follows:

- `against_work_dir` is the normal working directory.
 - `work_dir` is a temporary working directory for replacement. The files in this directory are not affected by the downloader of Apsara Infrastructure Management Framework and are not automatically fixed by Apsara Infrastructure Management Framework.
 - `expired_time` is the time-out interval. Set the time correctly so that overwriting can take effect. The time should be the number of seconds from 1970-1-1 till now.
4. After the file is saved, the secondary master restarts the supervisor. If you want to overwrite files on masters or chunkservers, perform overwriting on the masters or chunkservers individually or in batches. If you perform overwriting on them in batches, some processes are restarted at the same time, which may interrupt services. `/apsara/pangu_supervisor` is the overwritten directory. Apsara Infrastructure Management Framework does not fix the files changed in the directory. You can directly replace the target binary files.
 5. After overwriting, the updated configuration is delivered, and the secondary master starts processes normally. After a version upgrade, the overwriting loses effect.

2.5.7.6.2 Use the overwrite tool of Apsara Infrastructure Management Framework

To simplify the overwriting procedure, Apsara Infrastructure Management Framework provides an overwrite tool. The following example describes how to use this tool.

Assume that the version to be overwritten is 2899 and the validity period is 9999 days. Perform the following three steps:

1. Copy the directory.
2. Replace the binary files in the copied directory.

3. Perform overwriting.

The following example uses `pangu_chunkserver` to describe the procedure:

1. Under `/apsara`, run `ls -l` to confirm that the chunkserver deployment directory is `/cloud/app/pangu/PanguChunkserver#/pangu_chunkserver/2899`. Run the following command to copy the directory: `cp /cloud/app/pangu/PanguChunkserver#/pangu_chunkserver/2899 /cloud/app/pangu/PanguChunkserver#/pangu_chunkserver/2899.overwrite -rf`.
2. Replace the binary files.
3. Run the following command to perform overwriting: `/cloud/tool/tianji/overwrite add pangu PanguChunkserver# pangu_chunkserver 2899 9999`.



Note:

The chunkserver is restarted during this step and switched to the overwritten chunkserver.

4. Under `/apsara`, run `ls -l` to confirm whether overwriting is successful. If the output ends with `.overwrite`, overwriting is successful.
5. To manually cancel overwriting, run the following command: `/cloud/tool/tianji/overwrite remove pangu PanguChunkserver# pangu_chunkserver`. After the command is run, the overwriting loses effect.

2.6 ApsaraDB for RDS

2.6.1 Service architecture

2.6.1.1 System architecture

2.6.1.1.1 Backup System

Database backup can be initiated at any time. RDS can restore a database to any point in time based on the backup policy, making the data more traceable.

Automatic backup

RDS provides various types of backup. MySQL instances support physical backup and logical backup. PostgreSQL and PPAS instances support full backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can initiate a temporary backup as required. The backup files are retained for seven days.

Log management

For MySQL instances, RDS automatically generates binlogs and provides a download function for local incremental backup.

For PostgreSQL and PPAS instances, RDS automatically performs full physical backup.

Data tracing

RDS uses the backup files and logs to generate a temporary instance from any time point within seven days. Data can be restored after you verify the data has no errors.

The operation of creating a temporary instance does not affect the running of the current instance.

Each RDS instance can create only one temporary instance at a time. Temporary instances are valid for 48 hours. They can be triggered at most 10 times in a day.

2.6.1.1.2 Monitoring system

RDS provides multi-dimensional monitoring services across the physical, network, and application layers to ensure business availability.

Performance monitoring

RDS provides nearly 20 metric items for system performance monitoring, such as disk capacity, IOPS, connections, CPU utilization, network traffic, TPS, QPS, and cache hit rate. You can obtain the running status information of instances over the past year.

SQL auditing

The system records the SQL statements and related information sent to RDS instances, such as the connection IP address, database name, access account, execution time, and number of records returned. You can use SQL auditing to locate problems and check instance security.

Threshold alarms

RDS provides alarm SMS notifications in the event of exceptions in instance status or performance.

These exceptions can include instance locking, disk capacity, IOPS, connections, and CPU. You can configure alarm thresholds and up to 50 alarm contacts (of which five are effective at a time).

When an instance exceeds the threshold, an SMS notification is sent to the alarm contacts.

Web operation logs

The system logs all modifying operations in the RDS console for administrators to check. These logs are retained for a maximum of 30 days.

2.6.1.1.3 Control system

After a host or instance crashes, the RDS high-availability (HA) component checks for the exception and fails services over within 30 seconds. This guarantees that applications run normally and RDS is highly available.

The HA service uses the Detection, Repair, and Notice modules to ensure the availability of data link services. It also processes internal database exceptions.

HA policies

HA policies use a combination of service priorities and data replication modes to meet your business requirements.

There are two service priority levels:

- **Recovery Time Objective (RTO):** The database prioritizes restoring services to maximize availability time. Use the RTO policy if you require longer database uptime.
- **Recovery Point Objective (RPO):** The database prioritizes data reliability to minimize data loss. If you require high data consistency, use the RPO policy.

There are three data replication modes:

- **Asynchronous replication (Async):** When an application initiates an update request in the form of an add, delete, or modify operation, the primary node responds to the application immediately after completing the operation. Then, the primary node replicates data to the secondary node asynchronously. This means that the operation of the primary database is not affected if the secondary node is unavailable. However, data inconsistencies may occur between the primary and secondary nodes if the primary node is unavailable.
- **Forced synchronous replication (Sync):** When an application initiates an update request in the form of an add, delete, or modify operation, the primary node replicates data to the secondary node immediately after completing the operation. Then, the primary node waits for the secondary node to return a success message before it responds to the application. Because the primary node replicates data to the secondary node synchronously, unavailability of the secondary node affects the operation on the primary node. However, unavailability of the primary node does not cause data inconsistency.

- Semi-synchronous replication (Semi-Sync): Data is typically replicated in Sync mode. If an exception occurs (such as unavailability of the secondary node or a network exception between the two nodes) when the primary node tries to replicate data to the secondary node, the primary node suspends response to the application. The primary node will restore response to the application until the Sync replication times out and degrades to Async. If the application is allowed to update data during these conditions, data will be inconsistent when the primary node becomes unavailable. When data replication between the two nodes resumes because the secondary node or network connection is recovered, the data replication mode changes from Async to Sync.

You can select different combination modes of service priorities and data replication modes to improve availability based on your own business characteristics.

2.6.1.1.4 Task scheduling system

You can use the RDS console or APIs to create and delete instances or switch instances between the internal and public networks. All instance operations are scheduled, traced, and displayed as tasks.

Resource

The Resource module allocates and integrates lower-level RDS resources to enable and migrate instances. For example, when you use the RDS console or an API to create an instance, the Resource module calculates which physical server is most suitable to carry traffic. This module also allocates and integrates lower-level resources required to migrate RDS instances. As instances are created, deleted, and migrated, the Resource module calculates the fragmentation of resources, and periodically integrates resource fragments to handle traffic spikes.

2.6.2 RDS O&M overview

Apsara Stack Operations Console provides the following RDS O&M features:

- Instance management: allows you to view instance details, instance logs, and user information.
- Host management: allows you to view and manage hosts.

2.6.3 Log on to Apsara Stack Operations

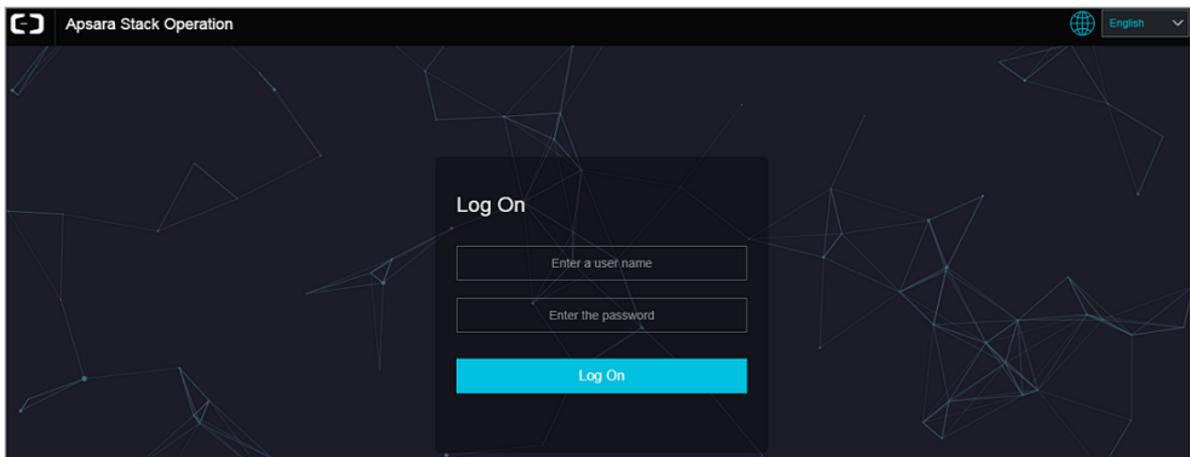
Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 2-13: Log on to ASO



Note:

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.

2.6.4 Instance management

Instance management allows you to view the instance details, instance logs, and user information.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. On the **Instance Management** tab page of **RDS**, you can view the following information:

- Instances

On the **Instance Management** tab page, view the instances under the account, as shown in [Figure 2-14: Instances](#).

Figure 2-14: Instances

The screenshot shows the 'Instance Management' tab in a dark-themed interface. At the top, there are two tabs: 'Instance Management' (selected) and 'Host Management'. Below the tabs is a search bar with the placeholder text 'Please enter' and a magnifying glass icon. To the right of the search bar is a blue button labeled 'Show Statistics'. The main content is a table with the following columns: Instance Name, Instance Status, Database Type, Instance Usage Type, and Actions. There are four rows of data, all with a status of 'CREATING' and 'Read-Only' usage type. The 'Actions' column contains links for 'User Information | Log Details | Error Details | Error Logs'.

Instance Name	Instance Status	Database Type	Instance Usage Type	Actions
...	CREATING	Redis	Read-Only	User Information Log Details Error Details Error Logs
...	CREATING	Redis	Read-Only	User Information Log Details Error Details Error Logs
...	CREATING	Redis	Read-Only	User Information Log Details Error Details Error Logs
...	CREATING	Redis	Read-Only	User Information Log Details Error Details Error Logs

- Instance details

Click the ID of an instance to view instance details, as shown in [Figure 2-15: Instance details](#).

Figure 2-15: Instance details

The screenshot shows the 'Instance Details' page in a dark-themed interface. The page is organized into several sections: 'Instance Information', 'Region Information', and 'Instance Specifications'. Each section contains key-value pairs for various parameters.

Instance Information	
Instance ID: 1536	Instance Name: ...
Instance Type: x	Instance Link Type: ns
Database Type: Redis	Database Version: 2.8
Status: CREATING	Instance Lock Mode: 0
Region Information	
Data Center: cn-qingdao-env@d-01	Cluster Name: ...
Instance Specifications	
CPU (cores): 3	Maximum Connections: 10000
Maximum Storage: 20	Maximum Memory: 1024
Maximum QPS:	Maximum IOPS:

- **User information**

Click **User Information** in the **Actions** column, as shown in [Figure 2-16: User information](#).

Figure 2-16: User information

The screenshot shows a 'User Information' dashboard with a table of instance metrics. The table has the following columns: Instance Name, Instance Status, Database Type, Instance Usage Type, CPU Utilization, IOPS Utilization, Disk Utilization, and Connections Utilization. All instances listed are in the 'CREATING' status, using 'Redis' as the database type and '常规实例' (Standard Instance) as the usage type. Each row shows utilization bars for CPU, IOPS, Disk, and Connections, all at 0%.

Instance Name	Instance Status	Database Type	Instance Usage Type	CPU Utilization	IOPS Utilization	Disk Utilization	Connections Utilization
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%
[Redacted]	CREATING	Redis	常规实例	0%	0%	0%	0%

- **Log details**

a. Click **Log Details** in the **Actions** column, as shown in [Figure 2-17: Log details](#).

b. You can select **Start Time** and **End Time** to view log details from a specified period of time.

Figure 2-17: Log details

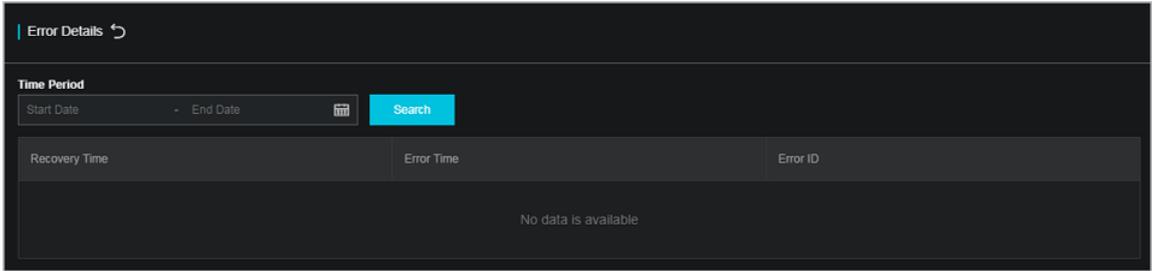
The screenshot shows a 'Log Details' dashboard. At the top, there is a 'Time Period' section with 'Start Date' and 'End Date' input fields, a calendar icon, and a 'Search' button. Below this is a table with the following columns: Slow Query Statement ID, IP Address of Host Linking Database, Database Name, Query Statement ID, Query Statement, Execution Time, Lockout Time, Parsed Rows, Returned Rows, Execution Start Time, and User Name. The table currently displays 'No data is available'.

Slow Query Statement ID	IP Address of Host Linking Database	Database Name	Query Statement ID	Query Statement	Execution Time	Lockout Time	Parsed Rows	Returned Rows	Execution Start Time	User Name
No data is available										

- **Error details**

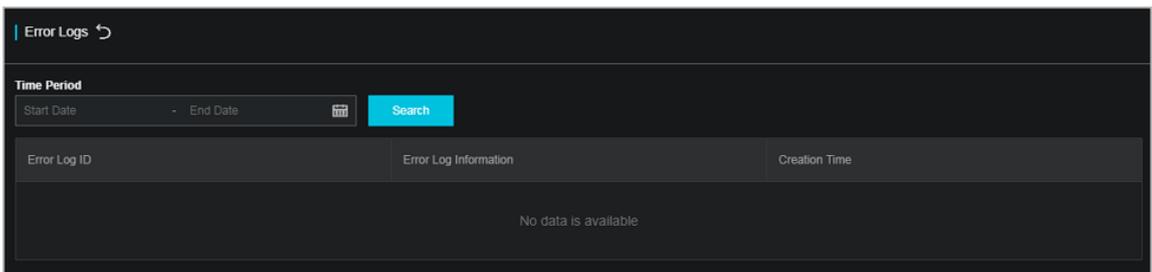
a. Click **Error Details** in the **Actions** column, as shown in [Figure 2-18: Error details](#).

b. You can select the **Start Time** and **End Time** to view error details during a specified period of time.

Figure 2-18: Error details

- **Error logs**

- Click **Error Logs** in the **Actions** column, as shown in [Figure 2-19: Error logs](#).
- You can select the **Start Time** and **End Time** to view error logs during a specified period of time.

Figure 2-19: Error logs

- On the **Instance Management** tab page, click **Show Statistics** to view the instance information by version and region, as shown in the following figure.

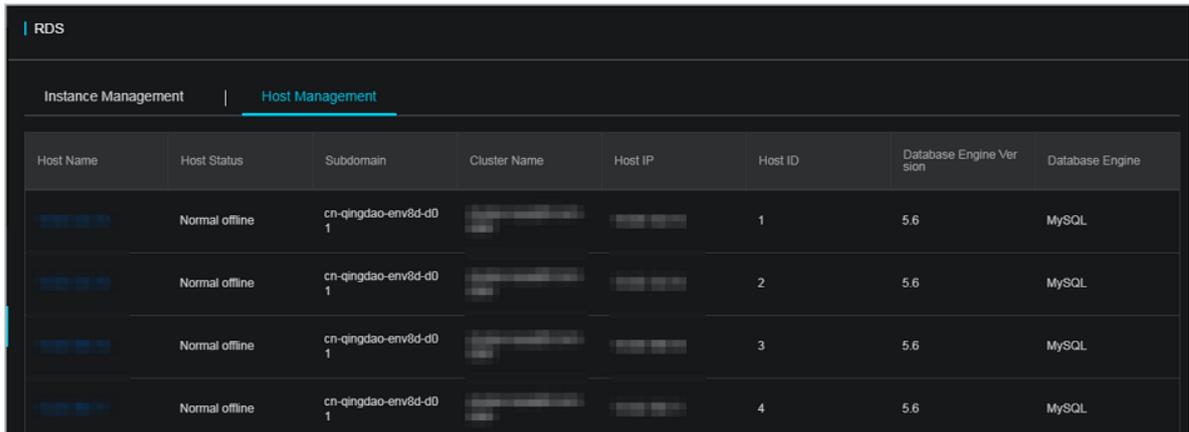


2.6.5 Host management

Host management allows you to view and manage hosts.

Procedure

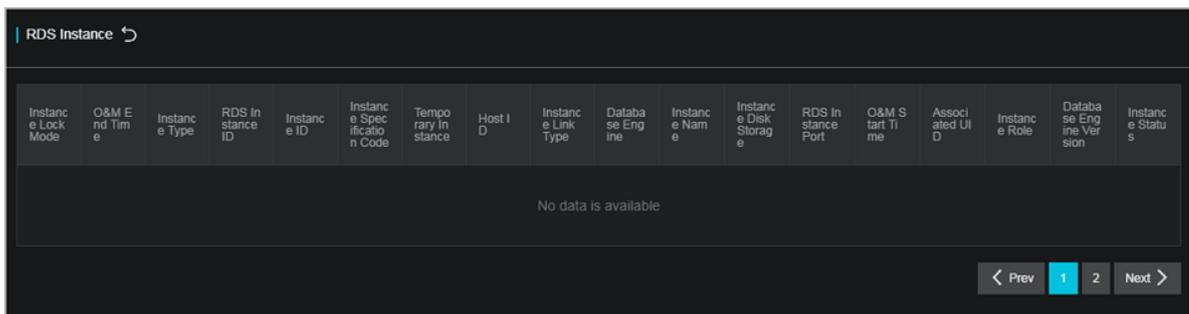
1. [Log on to Apsara Stack Operations](#).
2. On the **Host Management** tab page of **RDS**, you can view all host information.



The screenshot shows the RDS Host Management interface. It features a navigation bar with 'Instance Management' and 'Host Management' (the active tab). Below the navigation bar is a table with the following columns: Host Name, Host Status, Subdomain, Cluster Name, Host IP, Host ID, Database Engine Version, and Database Engine. The table contains four rows of data, all with a status of 'Normal offline' and a database engine of 'MySQL'.

Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine
[Redacted]	Normal offline	cn-qingdao-env8d-d01	[Redacted]	[Redacted]	1	5.6	MySQL
[Redacted]	Normal offline	cn-qingdao-env8d-d01	[Redacted]	[Redacted]	2	5.6	MySQL
[Redacted]	Normal offline	cn-qingdao-env8d-d01	[Redacted]	[Redacted]	3	5.6	MySQL
[Redacted]	Normal offline	cn-qingdao-env8d-d01	[Redacted]	[Redacted]	4	5.6	MySQL

3. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.



The screenshot shows the RDS Instance page. It features a navigation bar with 'RDS Instance' and a back arrow. Below the navigation bar is a table with the following columns: Instance Lock Mode, O&M End Time, Instance Type, RDS Instance ID, Instance ID, Instance Specification Code, Temporary Instance, Host ID, Instance Link Type, Database Engine, Instance Name, Instance Disk Storage, RDS Instance Port, O&M Start Time, Associated UID, Instance Role, Database Engine Version, and Instance Status. The table is currently empty, displaying 'No data is available' in the center. At the bottom right, there are navigation buttons: '< Prev', '1', '2', and 'Next >'.

Instance Lock Mode	O&M End Time	Instance Type	RDS Instance ID	Instance ID	Instance Specification Code	Temporary Instance	Host ID	Instance Link Type	Database Engine	Instance Name	Instance Disk Storage	RDS Instance Port	O&M Start Time	Associated UID	Instance Role	Database Engine Version	Instance Status
No data is available																	

2.6.6 Security maintenance

2.6.6.1 Network security maintenance

Network security maintenance includes device security and network security.

Device security

Check network devices, and enable security management protocols and configurations of devices.

Check for up-to-date versions of network device software and update to more secure versions in a timely manner.

For more information about the security maintenance method, see the product document for each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network situations to detect Internet and intranet traffic and defend the network against abnormal behavior and attacks.

2.6.6.2 Account password maintenance

Account passwords include the RDS system password and device password.

To ensure account security, periodically change the system and device passwords, and use passwords with high complexity.

2.7 KVStore for Redis

2.7.1 O&M tools

Apsara Stack Operations Console provides the following KVStore for Redis O&M features:

- Instance management: allows you to view the details, logs, and user information of instances.
- Host management: allows you to view and manage hosts.

2.7.2 Service architecture

2.7.2.1 System architecture

2.7.2.1.1 Backup system

Automatic backup

KVStore for Redis supports full backup. You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can initiate a temporary backup as required. The backup files are retained for seven days.

2.7.2.1.2 Data migration system

Data migration to or from KVStore for Redis

KVStore for Redis provides a professional tool and a migration wizard to help you migrate data to or from KVStore for Redis.

Backup file download

KVStore for Redis retains backup files for seven days. During this period, you can log on to the KVStore for Redis console to download the backup files.

2.7.2.1.3 Monitoring system

Performance monitoring

KVStore for Redis provides multiple metrics for system performance monitoring, such as disk capacity, memory usage, connections, CPU utilization, network traffic, QPS, and number of requests. You can obtain the running status information of instances over the past year.

Threshold alarms

KVStore for Redis provides alarm SMS notifications in the event of exceptions in instance status or performance.

These exceptions can include instance locking, insufficient disk capacity, abnormal IOPS, abnormal number of connections, and over high CPU utilization. You can configure alarm thresholds and up to 50 alarm contacts (of which only five are effective at a time). When a threshold is exceeded in an instance, an SMS notification is sent to the alarm contacts.

Web operation logs

The system logs all modification operations in the KVStore for Redis console for administrators to check. These logs are retained for a maximum of 30 days.

2.7.2.1.4 Control system

If a host or an instance crashes, the KVStore for Redis high-availability (HA) component fails services over within 30 seconds after the exception is detected. This guarantees that applications run properly and KVStore for Redis is highly available.

2.7.2.1.5 Task scheduling system

You can use the KVStore for Redis console or APIs to create or delete instances or switch instances between the intranet and Internet. All instance operations are scheduled, traced, and displayed as tasks.

2.7.3 Log on to Apsara Stack Operations

Prerequisites

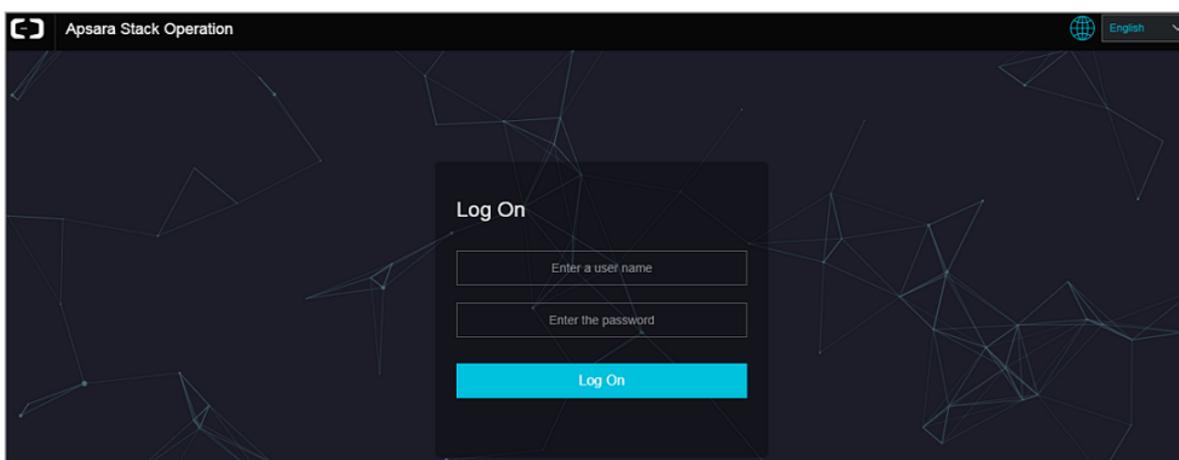
- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.

- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 2-20: Log on to ASO



Note:

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.

2.7.4 Instance management

Instance management allows you to view the details, logs, and user information of instances.

Procedure

1. On the **Instance Management** tab page of **RDS**, you can view the following information:

- **Instances**

On the **Instance Management** tab page, view the instances under the account.

- **Instance details**

Click the ID of an instance to view instance details.

- **User information**

Click **User Information** in the **Actions** column.

- **Log details**

a. Click **Log Details** in the **Actions** column.

b. You can set **Start Time** and **End Time** to view log details within the specified period of time.

- **Error details**

a. Click **Error Details** in the **Actions** column.

b. You can set the **Start Time** and **End Time** to view error details within the specified period of time.

- **Error logs**

a. Click **Error Logs** in the **Actions** column.

b. You can set the **Start Time** and **End Time** to view error logs within the specified period of time.

2. On the **Instance Management** tab page, click **Show Statistics** to view the instance information by version or region.

2.7.5 Host management

Host management allows you to view and manage hosts.

Procedure

1. On the **Host Management** tab page of **RDS**, you can view all host information.

Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine
...
...
...
...
...
...
...

2. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.

Instance Lock Mode	O&M End Time	Instance Type	RDS Instance ID	Instance ID	Instance Specific Code	Tempo Instance	Host ID	Instance Link Type	Datab Engine	Instance Name	Instance Disk Storage	RDS Instance Port	O&M Start Time	Instance Role	Datab Engine Version	Instance Status
0
0
0
0
0

2.7.6 Security maintenance

2.7.6.1 Network security maintenance

Network security maintenance is aimed at ensuring device security and network security.

Device security

Check network devices, and enable security management protocols and configurations of devices.

Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner.

For more information about the security maintenance method, see the product document of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check Internet and intranet traffic and defend the network against abnormal behaviors and attacks.

2.7.6.2 Account password maintenance

Account passwords include the KVStore for Redis system and device passwords.

To ensure account security, change the system and device passwords periodically, and use passwords that meet the complexity requirements.

2.8 ApsaraDB for MongoDB

2.8.1 Service architecture

2.8.1.1 System architecture

2.8.1.1.1 Backup system

Automatic backup

ApsaraDB for MongoDB supports both physical backup and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can initiate a temporary backup as required. The backup files are retained for seven days.

Log management

ApsaraDB for MongoDB generates operation logs and allows you to download them. You can use the operation logs for local incremental backup.

Data backtracking

ApsaraDB for MongoDB can use backup files and logs to generate a temporary instance for any time point within the past seven days. After verifying that the data in the temporary instance is correct, you can use the temporary instance to restore data to the specified time point.

Creating a temporary instance does not affect the running of the current instance.

Only one temporary instance can be created for each ApsaraDB for MongoDB instance at a time. A temporary instance is valid for 48 hours. You can create a maximum of 10 temporary instances for an ApsaraDB for MongoDB instance each day.

2.8.1.1.2 Data migration system

Database replication between instances

ApsaraDB for MongoDB allows you to easily migrate databases from one instance to another.

Data migration to or from ApsaraDB for MongoDB

ApsaraDB for MongoDB provides a professional tool and a migration wizard to help you migrate data to or from ApsaraDB for MongoDB.

Backup file download

ApsaraDB for MongoDB retains backup files for seven days. During this period, you can log on to the ApsaraDB for MongoDB console to download the backup files.

2.8.1.1.3 Monitoring system

Performance monitoring

ApsaraDB for MongoDB provides nearly 20 metrics for system performance monitoring, such as disk capacity, IOPS, connections, CPU utilization, network traffic, TPS, QPS, and cache hit rate. You can obtain the running status information of instances over the past year.

SQL auditing

The system records the SQL statements and related information sent to ApsaraDB for MongoDB instances, such as the connection IP address, database name, access account, execution time, and the number of records returned. You can use SQL auditing to locate problems and check instance security.

Threshold alarms

ApsaraDB for MongoDB provides alarm SMS notifications in the event of exceptions in instance status or performance.

These exceptions can include instance locking, insufficient disk capacity, abnormal IOPS, abnormal number of connections, and over high CPU utilization. You can configure alarm thresholds and up to 50 alarm contacts (of which only five are effective at a time). When a threshold is exceeded in an instance, an SMS notification is sent to the alarm contacts.

Web operation logs

The system logs all modification operations in the ApsaraDB for MongoDB console for administrators to check. These logs are retained for a maximum of 30 days.

2.8.1.1.4 Control system

If a host or an instance crashes, the ApsaraDB for MongoDB high-availability (HA) component fails services over within 30 seconds after the exception is detected. This guarantees that applications run properly and ApsaraDB for MongoDB is highly available.

2.8.1.1.5 Task scheduling system

You can use the ApsaraDB for MongoDB console or APIs to create or delete instances or switch instances between the intranet and Internet. All instance operations are scheduled, traced, and displayed as tasks.

2.8.2 ApsaraDB for MongoDB O&M overview

Apsara Stack Operations Console provides the following O&M features for ApsaraDB for MongoDB:

- Instance management: allows you to view instance details, instance logs, and user information.
- Host management: allows you to view and manage hosts.

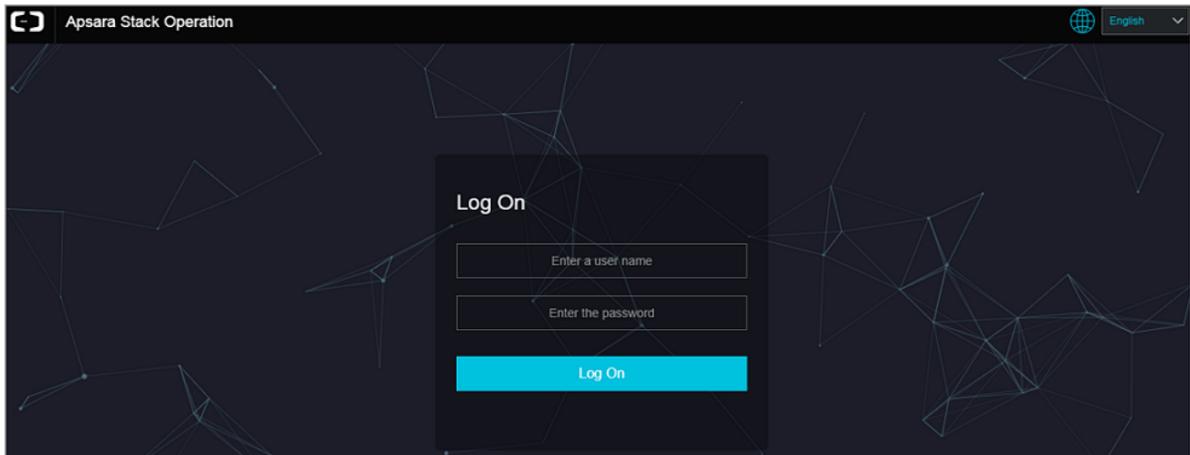
2.8.3 Log on to Apsara Stack Operations

Prerequisites

- ASO address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.
2. Enter the ASO access address `http://region-id.aso.intranet-domain-id.com` in the address bar and then press Enter.

Figure 2-21: Log on to ASO**Note:**

You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the password of your username as instructed when you log on to ASO for the first time. To improve security, the password must be 10-20 characters long and can contain English uppercase/lowercase letters (A-Z or a-z), numbers (0-9), and special characters, such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%), which meets the minimum complexity requirements.
4. Click **Log On** to go to the **Apsara Stack Operations** page.

2.8.4 Instance management

Instance management allows you to view instance details, instance logs, and user information.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. On the **Instance Management** tab of the **RDS** page, you can view the following information:

- Instances

On the **Instance Management** tab, view the instances under the account.

- Instance details

Click the ID of an instance to view instance details.

- **User information**

Click **User Information** in the **Actions** column to view user information.

- **Log details**

a. Click **Log Details** in the **Actions** column to view log details.

b. You can set **Start Time** and **End Time** to view log details in a specified period of time.

- **Error details**

a. Click **Error Details** in the **Actions** column to view error details.

b. You can set **Start Time** and **End Time** to view error details in a specified period of time.

- **Error logs**

a. Click **Error Logs** in the **Actions** column to view error logs.

b. You can set **Start Time** and **End Time** to view error logs in a specified period of time.

3. On the **Instance Management** tab, click **Show Statistics** to view instance information by version and region.

2.8.5 Host management

Host management allows you to view and manage hosts.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. On the **Host Management** tab of the **RDS** page, view information about all hosts.
3. Click a host name to go to the **RDS Instance** page. On this page, you can view all instances on this host.

2.8.6 Security maintenance

2.8.6.1 Network security maintenance

Network security maintenance is aimed at ensuring device security and network security.

Device security

Check network devices, and enable security management protocols and configurations of devices.

Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner.

For more information about the security maintenance method, see the product document of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check Internet and intranet traffic and defend the network against abnormal behaviors and attacks.

2.8.6.2 Account password maintenance

Account passwords include the ApsaraDB for MongoDB system and device passwords.

To ensure account security, change the system and device passwords periodically, and use passwords that meet the complexity requirements.

2.9 Apsara Stack Security

2.9.1 Log on to the Apsara Infrastructure Management Framework console

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

Prerequisites

Before logging on to the console, obtain the URL of Apsara Stack Operations Console, and the username and password to log on to the console from your system administrator.

Procedure

1. In the address bar of a browser, enter `https://ASOP URL`, and press Enter.
2. On the logon page, enter the username and password, and then click **Log On**.
3. In the left-side navigation pane, select **Products**.

4. In the product list, click **Apsara Infrastructure Management Framework** to go to the Apsara Infrastructure Management Framework console.

2.9.2 Routine operations and maintenance of Server Guard

2.9.2.1 Check the service status

2.9.2.1.1 Check the client status

Check the following status information about the Server Guard client to verify that the client is running properly:

Client logs

Client logs are stored in the data directory under the directory of the Server Guard process file, for example, `/usr/local/aegis/aegis_client/aegis_xx_xx/data`.

Client logs are saved by day, for example, `data. 1` to `data. 7`.

Client's online status

Run the following command to check the client's online status:

```
ps -aux | grep AliYunDun
```

Network connectivity

Run the following command to check whether the client has set up a TCP connection with the server:

```
netstat -tunpe | grep AliYunDun
```

Client UUID

Open the client log file `data.x` and check the character string following `Currentuid Ret.` This character string is the UUID of the current client.

Client processes

The Server Guard client has three resident processes: `AliYunDun`, `AliYunDunUpdate`, and `AliHids`.

When the client runs properly, all of the three processes run normally.



Note:

On a Windows OS client, the AliYunDun and AliYunDunUpdate processes exist in the form of services. The service names are Server Guard Detect Service and Server Guard Update Service, respectively.

2.9.2.1.2 Check the status of Aegiserver

Context

To check the running status of Aegiserver, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server of Aegiserver.
2. Run the following command to find the Aegiserver image ID:

```
docker ps -a |grep aegiserver
```

The following message is displayed:

```
b9e59994df41
reg.docker.example.com/aqs/aegiserverlite@sha256:f9d292f54c58646b672a
8533a0d78fba534d26d376a194034e8840c70d9aa0b3 "/bin/bash /startApp." 2
hours ago Up 2 hours 80/tcp, 7001/tcp, 8005/tcp, 8009/tcp yundun-aegis
.Aegiserverlite__.aegiserverlite. 1484712802
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep aegiserver
```

The following message is displayed:

```
root 153 0.6 25.8 2983812 1084588 ? S1 12:13 1:01 /opt/taobao/java
/bin/java -Djava.util.logging.config.file=/home/admin/aegiserver
lite/.default/conf/logging.properties -Djava.util.logging.manager
=org.apache.juli.ClassLoaderLogManager -server -Xms2g -Xmx2g -XX:
PermSize=96m -XX:MaxPermSize=384m -Xmn1g -XX:+UseConcMarkSweepGC -XX
:+UseCMSCompactAtFullCollection -XX:CMSMaxAbortablePrecleanTime=5000
-XX:+CMSClassUnloadingEnabled -XX:+UseCMSInitiatingOccupancyOnly -XX
:CMSInitiatingOccupancyFraction=80 -XX:+HeapDumpOnOutOfMemoryError -
XX:HeapDumpPath=/home/admin/logs/java.hprof -verbose:gc -Xloggc:/home
/admin/logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -Djava
.awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun
.net.client.defaultReadTimeout=30000 -XX:+DisableExplicitGC -Dfile.
encoding=UTF-8 -Ddruid.filters=mergeStat -Ddruid.useGloalDataSourceSt
at=true -Dproject.name=aegiserverlite -Dcatalina.vendor=alibaba -Djava
.security.egd=file:/dev/./urandom -Dlog4j.defaultInitOverride=true -
Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_EQUALS_IN_VALUE=true -
Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_HTTP_SEPARATORS_IN_V0=
true -Djava.endorsed.dirs=/opt/taobao/tomcat/endorsed -classpath /opt/
taobao/tomcat/bin/bootstrap.jar:/opt/taobao/tomcat/bin/tomcat-juli.jar
-Dcatalina.logs=/home/admin/aegiserverlite/.default/logs -Dcatalina.
base=/home/admin/aegiserverlite/.default -Dcatalina.home=/opt/taobao/
```

```
tomcat -Djava.io.tmpdir=/home/admin/aegiserverlite/.default/temp org.apache.catalina.startup.Bootstrap -Djboss.server.home.dir=/home/admin/aegiserverlite/.default -Djboss.server.home.url=file:/home/admin/aegiserverlite/.default start
```

5. Run the following command to perform the health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

6. View related logs.

- **Protocol logs:** View logs about upstream and downstream protocol messages between the server and client in `/home/admin/aegiserver/logs/AEGIS_MESSAGE.log`.
- **Operation logs:** View abnormal stack information during operation in `/home/admin/aegiserver/logs/aegis-default.log`.
- **Offline logs:** View the logs about client disconnection caused by time-out in `/home/admin/aegiserver/logs/AEGIS_OFFLINE_MESSAGE.log`.

2.9.2.1.3 Check the Server Guard Update Service status

Context

To check the status of Server Guard Update Service, follow the following steps:

Procedure

1. Run the `ssh host IP address` command to log on to the server of Aegiserver.
2. Run the following command to find the Aegiserver image ID:

```
docker ps -a |grep aegiserver
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep aegisupdate
```

5. Run the following command to perform the health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

2.9.2.1.4 Check the Defender module status

Context

To check the status of the Defender module of Server Guard, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Defender module of Server Guard.

2. Run the following command to find the image ID of the Defender module of Server Guard:

```
docker ps -a |grep defender
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep defender
```

5. Run the following command to perform health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

2.9.2.2 Restart Server Guard

Context

To restart Server Guard when a fault occurs, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts Server Guard.
2. Run the following command to find the image ID of Server Guard:

```
docker ps -a |grep application name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Restart related services.

- Restart the Server Guard client service.
 - For a server running a Windows OS, go to the service manager, locate *Server Guard Detect Service*, and restart this service.
 - For a server running a Linux OS, use either of the following methods to restart the Server Guard client service:
 - Run the `service aegis restart` command to restart the service.

- Run the `killall AliYunDun` command as the root user to stop the current process, and then restart the `/usr/local/aegis/aegis_client/aegis_xx_xx/AliYunDun` process.
- Restart the Aegiserver service.
 - a. Run the following command to view the Java process ID:

```
ps aux |grep aegiserver
```
 - b. Run the following command to stop the current process:

```
kill -9 process
```
 - c. Run the following command to restart the process:

```
sudo -u admin /home/admin/aegiserver/bin/jbossctl restart
```
 - d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/checkpreload.htm
```
- Restart Server Guard Update Service:
 - a. Run the following command to view the Java process ID:

```
ps aux |grep aegisupdate
```
 - b. Run the following command to stop the current process:

```
kill -9 process
```
 - c. Run the following command to restart the process:

```
sudo -u admin /home/admin/aegisupdate/bin/jbossctl restart
```
 - d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/checkpreload.htm
```
- Restart the Defender service of Server Guard.
 - a. Run the following command to view the Java process ID:

```
ps aux |grep secure-service
```
 - b. Run the following command to stop the current process:

```
kill -9 process
```
 - c. Run the following command to restart the process:

```
sudo -u admin /home/admin/secure-service/bin/jbossctl restart
```

- d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/checkpreload.htm
```

2.9.3 Routine operations and maintenance of Network Traffic Monitoring System

2.9.3.1 Check the service status

2.9.3.1.1 Basic inspection

During the basic inspection of Network Traffic Monitoring System, check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations > Project Operations** on the page that appears, enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
3. Select **BasicCluster**.
4. Check whether `yundun-beaver-advance` has reached the final status in **Service Instances List**.

2.9.3.1.2 Advanced inspection

During the advanced inspection feature of Network Traffic Monitoring System, check the status and features of the service.

Procedure

To inspect the service status, follow the following steps:

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Log on to two physical machines of Network Traffic Monitoring System, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **BasicCluster**.
 - d) Select `yundun-beaver-advance` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.

- e) Select **BeaverAdvance#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **Server Information**, and use TerminalService to log on to two physical machines of Network Traffic Monitoring System, respectively.
3. Check the log status of Network Traffic Monitoring System.
Run `sudo cat /var/log/messages`. If any record is returned, the logs are normal.
 4. Check the status of the mirrored traffic.
Run `sudo cat /proc/ixgbe_debug_info`. If the **speed is not 0** in the second-to-last row of the output, the mirrored traffic is normal.
 5. Check the configuration of the protected IP CIDR block.
Run `tail -f /dev/shm/banff-2018-xx.log`. In the command, *xx* indicates the month. For example, the log file for May in 2018 is named *banff-2018-05.log*. The IP CIDR block in the output should be the classic network SLB/EIP CIDR block (for CSW non-standard access, configure the VPC CIDR block).
 6. Check the network connectivity between Network Traffic Monitoring System and the VM.
Run `ping VMIP` to check the network connectivity. In the command, *VMIP* is a real IP address that falls in the CIDR block of the previous step.
 7. Check the tcp_decode process status.
Run `ps -ef | grep tcp_decode`. If any record is returned, the tcp_decode process is normal.
 8. Check the configuration of the traffic scrubbing server.
Run `cat /home/admin/beaver-dj-schedule/conf/dj.conf` and check whether the IP address specified in the unmarked configuration item `aliguard_smart` is the DNS VIP of the domain name `aliguard.${global:internet-domain}`.
 9. Check the following typical logs:
 - DDoS alert logs
Run the `grep -A 10 -B 10 LIDS /var/log/messages` command to view the DDoS alert logs.
 - TCP blocking command logs
Run the `grep add_to_blacklist.htm /var/log/messages` command to view the TCP blocking command logs.

- Outbound attack logs

Run the `grep zombie_new /var/log/messages` command to view the outbound attack logs.

2.9.3.2 Common operations and maintenance

2.9.3.2.1 Restart the Network Traffic Monitoring System process

Context

To restart the Network Traffic Monitoring System process, follow the following steps:

Procedure

1. Log on to the physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to restart the Network Traffic Monitoring System process:

```
rm -rf /dev/shm/drv_setup_path
```

2.9.3.2.2 Uninstall Network Traffic Monitoring System

Context

To uninstall Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to uninstall Network Traffic Monitoring System:

```
bash /opt/beaver/bin/uninstall.sh
```

2.9.3.2.3 Disable TCP blocking

Context

To disable TCP blocking for Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Open the `/beaver_client.sh` file on each server of Network Traffic Monitoring System, and add a number sign (#) to the start of the `./tcp_reset` line to comment out the line.

4. Run the following command on each server of Network Traffic Monitoring System to disable TCP blocking:

```
killall tcp_reset
```

2.9.3.2.4 Enable TCPDump

Context

To enable TCPDump for Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to enable TCPDump:

```
echo 1 > /proc/ixgbe_debug_dispatch
```



Note:

When TCPDump is enabled, the performance of Network Traffic Monitoring System may be affected. We recommend that you run the following command to disable TCPDump after packet capture is complete.

```
echo 0 > /proc/ixgbe_debug_dispatch
```

2.9.4 Routine operations and maintenance of Anti-DDoS Service

2.9.4.1 Check the service status

2.9.4.1.1 Basic inspection

The basic inspection of Anti-DDoS Service checks whether the service has reached the final status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console, and choose **Operations > Project Operations**. Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
2. Select **AliguardCluster**.
3. Check whether `yundun-aliguard` has reached the final status in **Service Instances List**.

2.9.4.1.2 Advanced inspection

The advanced inspection of Anti-DDoS Service checks the status and features of the service.

Procedure

To check the running status of Anti-DDoS Service, follow the following steps:

1. Log on to two physical machines of Anti-DDoS Service, respectively.
 - a) Log on to the Apsara Infrastructure Management Framework console, and choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **AliguardCluster**.
 - d) Select **yundun-aliguard** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **AliguardConsole#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and use TerminalService to log on to two physical machines of Anti-DDoS Service, respectively.
2. Check the deployment status of Anti-DDoS Service.

Run `/home/admin/aliguard/target/AliguardDefender/bin/aliguard_defender_check`, and check the output result.



Note:

If a server of Anti-DDoS Service has just restarted, wait for three to five minutes before running the script to check the deployment status.

- If the message `aliguard status check OK!` appears, Anti-DDoS Service has been correctly deployed and the service status is normal, as shown in [Figure 2-22: Check the status of Anti-DDoS Service](#).

Figure 2-22: Check the status of Anti-DDoS Service

```

1 [root@10.1.1.1111].cloud.111.111 /home/admin]
2 #aliguard_defender_check
3 myfwd
4 aliguard_log
5 netframe
6 route_monitor
7 neigh_monitor
8 aliguard_monitor
9 bgpd
10 rsyslogd
11 aliguard status check OK!

```

- If the error message shown in [Figure 2-23: Reinjection route error message](#) appears, the reinjection route is faulty.

Figure 2-23: Reinjection route error message

```

1 Error: route status error, we need two default routes to reinject the net flow!
2 Error: route error, can't get to the target ip.

```

Troubleshooting: The reinjection route is a default route generated by Anti-DDoS Service and is redirected to the interface through which the ISW is bound to the VPN in the next hop. If any problem occurs, check whether this route has been generated by Anti-DDoS Service. If this route has been generated, check whether the ISW has forwarded this route to downstream devices.

- If the error message shown in [Figure 2-24: BGP routing error message](#) appears, the BGP protocol (for traffic routing) is faulty.

Figure 2-24: BGP routing error message

```

1 Error: bgp status error!

```

Troubleshooting: If BGP routing is faulty, troubleshoot the problem as follows:

- a. Use the ISW to check whether the BGP neighbor is in the normal status.
- b. Check whether the BGP route of the ISW contains a 32-bit attacked IP address of which the route is redirected to Anti-DDoS Service in the next hop.

- c. Check whether the route policy in the BGP configuration of the ISW is correctly configured.
 - If the problem is caused by none of the above reasons, the core process is faulty. Contact Alibaba Cloud technical support.
3. Check the status of the NICs or optical modules of Anti-DDoS Service.

**Note:**

Anti-DDoS Service has special requirements on optical modules. Only optical modules equipped with Intel X520 or Intel 82599 NICs can be used.

Run `lspci | grep Eth`. If the command output contains four Intel 82599 NICs, the NICs are standard.

```
[root@cloud.am54 /root]
#lspci -v | grep Eth
02:00.0 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
02:00.1 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
04:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
      Subsystem: Intel Corporation Ethernet Server Adapter X520-2
04:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
      Subsystem: Intel Corporation Ethernet Server Adapter X520-2
81:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
      Subsystem: Intel Corporation Ethernet Server Adapter X520-2
81:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
      Subsystem: Intel Corporation Ethernet Server Adapter X520-2
```

2.9.4.2 Common operations and maintenance

2.9.4.2.1 Restart Anti-DDoS Service

Context

To restart Anti-DDoS Service when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts Anti-DDoS Service.
2. Run the following command to stop Anti-DDoS Service:

```
/home/admin/aliguard/target/AliguardDefender/bin/aliguard stop
```

**Note:**

If the `ERROR: Module net_msg is in use` message is displayed, run the command again later. If Anti-DDoS Service cannot be stopped after several attempts, restart the server of Anti-DDoS Service.

3. Run the following command to restart Anti-DDoS Service:

```
/home/admin/aliguard/target/AliguardDefender/bin/aliguard start
```

4. Run the service status check command five minutes after Anti-DDoS Service is restarted.

2.9.4.2.2 Troubleshoot common faults

Context

When an error occurs in Anti-DDoS Service, follow the following troubleshooting steps:

Procedure

1. Restart Anti-DDoS Service.

- If Anti-DDoS Service is in the normal status after being restarted but an error message is returned during the health check performed later, non-standard NICs or optical modules are used. To check whether standard NICs or optical modules are used, see [Check the status of the NICs or optical modules of Anti-DDoS Service](#). If non-standard NICs or optical modules are used, change the NICs or optical modules.
- If Anti-DDoS Service is in an unusual status after being restarted, go to the next step.

2. View the `aliguard_dynamic_config` file.

Carefully check whether each configuration item in the file is exactly the same as that in the plan.

**Note:**

Ensure that the AS number specified in `aliguard local` is 65515 and that the BGP password is correct.

3. Check the wiring and switch configuration.

**Note:**

If any incorrect configuration is found, the current fault is caused by incorrect wiring or switch IP address configuration, rather than incorrect deployment of Anti-DDoS Service. In this case, contact the network engineer.

Assume that the Anti-DDoS Service configurations to be checked are listed in the following figure, among which the server IP address is 10.1.4.12. To check whether the four ports of Anti-DDoS Service can ping the ports of the switch, follow the following steps:

Figure 2-25: Anti-DDoS Service configuration example

aliguard_host_ip	port	aliguard_port_ip	csr_port_ip
10.1.4.12	T0	10.1.0.34	10.1.0.33
10.1.4.12	T1	10.1.0.38	10.1.0.37
10.1.4.12	T2	10.1.0.50	10.1.0.49
10.1.4.12	T3	10.1.0.54	10.1.0.53
10.1.4.28	T0	10.1.0.42	10.1.0.41
10.1.4.28	T1	10.1.0.46	10.1.0.45
10.1.4.28	T2	10.1.0.58	10.1.0.57
10.1.4.28	T3	10.1.0.62	10.1.0.61

- a. Run the following commands to check the NIC PCI IDs of Anti-DDoS Service:

```
cd /sys/bus/pci/drivers/igb_uio
```

```
ls
```

Record the PCI IDs of the four NICs, for example, 0000:01:00.0, 0000:01:00.1, 0000:82:00.0, and 0000:82:00.1.

- b. Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard stop` command to stop Anti-DDoS Service.
- c. In the `/sys/bus/pci/drivers/igb_uio` directory, unbind the four NICs recorded in the first step from the `igb_uio` driver, as shown in [Figure 2-26: Unbind NICs](#).

Figure 2-26: Unbind NICs

```
1 echo "0000:01:00.0" >> unbind
2 echo "0000:01:00.1" >> unbind
3 echo "0000:82:00.0" >> unbind
4 echo "0000:82:00.1" >> unbind
```

- d. In the `/sys/bus/pci/drivers/ixgbe` directory, bind the four NICs to the `ixgbe` driver for Linux, as shown in [Figure 2-27: Bind NICs](#).

Figure 2-27: Bind NICs

```

1 echo "0000:01:00.0" >> bind
2 echo "0000:01:00.1" >> bind
3 echo "0000:82:00.0" >> bind
4 echo "0000:82:00.1" >> bind

```

- e. Set Anti-DDoS Service IP addresses for the NICs.

The local server IP address is 10.1.4.12, and the NIC IP addresses are set to 10.1.0.34, 10.1.0.38, 10.1.0.50, and 10.1.0.54, as shown in [Figure 2-25: Anti-DDoS Service configuration example](#).

- A.** Run the `ifconfig-a` command to display all NICs, and run the `ethtool -i` command to view the PCI ID of each NIC. Find the four NICs of which the IDs are the same as those recorded in the first step, for example, eth0, eth1, eth2, and eth3.
- B.** Run the following commands to move these NICs to the top of the queue:

```
ifconfig eth0 up
```

```
ifconfig eth1 up
```

```
ifconfig eth2 up
```

```
ifconfig eth3 up
```

- C.** Set Anti-DDoS Service IP addresses for the NICs. Run the following commands to set Anti-DDoS Service IP addresses for the NICs based on their PCI IDs in an ascending order:

```
ifconfig eth0 10.1.0.34 netmask 255.255.255.252
```

```
ifconfig eth1 10.1.0.38 netmask 255.255.255.252
```

```
ifconfig eth2 10.1.0.50 netmask 255.255.255.252
```

```
ifconfig eth3 10.1.0.54 netmask 255.255.255.252
```

- f.** Try to ping the peer IP addresses configured. If the peer IP addresses cannot be pinged, the switch configuration or wiring is incorrect.

```
ping 10.1.0.33
```

```
ping 10.1.0.37
```

```
ping 10.1.0.49
```

```
ping 10.1.0.53
```

- g. If these four IP addresses can all be pinged, you can directly start Anti-DDoS Service without unbinding the NICs.

Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard start` command to start Anti-DDoS Service.

After Anti-DDoS Service has been started for a while, run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard_rule -v 0.0.0.0 -d drop_icmp` command to disable the `drop_icmp` policy.

- h. Ping the peer IP addresses again.

```
ping 10.1.0.33
```

```
ping 10.1.0.37
```

```
ping 10.1.0.49
```

```
ping 10.1.0.53
```

If the peer IP addresses cannot be pinged, non-standard NICs or optical modules are used or the configuration is incorrect.

4. If these four peer IP addresses can be pinged after Anti-DDoS Service is started but an error is reported during a status check of Anti-DDoS Service, contact Alibaba Cloud technical support.

2.9.5 Routine operations and maintenance of the vulnerability analysis service

2.9.5.1 Check service status

2.9.5.1.1 Basic inspection

The basic inspection of the vulnerability analysis service checks whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations** > **Project Operations** on the page that appears, enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
3. Select **SasCluster**.

4. Check whether `yundun-cactus` has reached the final status in **Service Instances List**.

2.9.5.1.2 Advanced inspection: Checks the status of the Cactus -batch service

This topic describes how to check the running status of the Cactus-batch service for vulnerability analysis.

Procedure

To check the running status of the Cactus-batch service, follow the following steps:

1. [Log on to the Apsara Infrastructure Management Framework console](#).
2. Log on to two physical machines for vulnerability analysis, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **SasCluster**.
 - d) Select **yundun-cactus** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **CactusBatch#** from **service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and click Terminal to log on to two physical machines for vulnerability analysis, respectively.
3. Log on to two Cactus-batch Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep batch | awk '{print $1}') bash.
```

4. Check the Cactus-batch process status.

```
Run ps aux | grep java | grep cactus-batch. If any record is returned, the Cactus-batch process is normal.
```

5. Check the health status.

```
Run curl 127.0.0.1:7001/check.htm. If OK is returned, the service is normal.
```

6. Check the Beaver_server process status.

```
Run ps -afe | grep tcp_save_disk. If any record is returned, the Beaver_server process is normal.
```

7. Check the health status.

Run `ss -plnt | grep 8181`. If port 8181 is contained in the returned content, the service is normal.

8. View related logs.

- View the Tomcat logs in `/home/admin/cactus-batch/logs/jboss_stdout.log`.
- View the system logs of the Cactus-batch service in `/home/admin/logs/batch.log`.
- View the traffic data provided by Network Traffic Monitoring System in the files under the `/home/admin/beaver_logs` directory.

2.9.5.1.3 Advanced inspection: Check the status of the Cactus-keeper service

This topic describes how to check the running status of the Cactus-keeper service for vulnerability analysis.

Procedure

To check the running status of the Cactus-keeper service, follow the following steps:

1. [Log on to the Apsara Infrastructure Management Framework console.](#)

2. Log on to two physical machines for vulnerability analysis, respectively.

a) Choose **Operations > Project Operations**.

b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.

c) Select **SasCluster**.

d) Select **yundun-cactus** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.

e) Select **CactusBatch#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.

f) View **server Information**, and click Terminal to log on to two physical machines for vulnerability analysis, respectively.

3. Check the Tomcat process status.

Run `ps -afe | grep tomcat | grep -v grep`. If any record is returned, the Tomcat process is normal.

4. Check the scan engine status.

Run `ps -afe | grep heimdall | grep java | grep -v grep`. If at least one record is returned, the scan engine is normal.

5. Check the health status.

Run `curl 127.0.0.1:7001/check.htm`. If OK is returned, the service is normal.

6. Check the network connection between Cactus-keeper and the VM.

Run `ping VMIP`. For more information about `VMIP`, see [Advanced inspection of Network Traffic Monitoring System](#).

7. View related logs.

- View the system logs of the Cactus-keeper service in `/home/admin/logs/keeper.log`.
- View the scan engine logs in `/home/admin/logs/heimdall_log4j.log.N*`.



Note:

Each scan engine has a log file, and N* indicates the engine number.

2.9.5.2 Restart the vulnerability analysis service

2.9.5.2.1 Restart Cactus-batch

Context

To restart the Cactus-batch service, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Cactus-batch service.
2. Run the following command to find the image ID of the Cactus-batch service:

```
docker ps -a |grep service name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to stop the current service process:

```
kill -9 $(ps -ef | grep java | grep cactus-batch | grep -v grep | awk '{print $2}')
```

5. Run the following command to restart the Cactus-batch service:

```
/home/admin/cactus-batch/bin/jbossctl restart
```

6. Run the following command to check whether the service has been successfully restarted:

```
curl 127.0.0.1:7001/check.htm
```

If OK is returned, the service is normal.

2.9.5.2.2 Restart the Beaver_server service

Context

To restart the Beaver_server service, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Beaver_server service.
2. Run the following command to find the image ID of the Beaver_server service:

```
docker ps -a |grep service name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following commands to start or stop the Beaver_server service:

```
cd /home/admin/beaver_server
```

```
sh ./app.sh start|stop
```

2.9.5.2.3 Restart Cactus-keeper

Context

To restart the Cactus-keeper service, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Cactus-keeper service.
2. Run the following command to find the image ID of the Cactus-keeper service:

```
docker ps -a |grep service name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following commands to start or stop the scan engine process for vulnerability analysis:

```
cd /home/admin/cactus-keeper/bin
```

```
sh ./jbossctl start|stop
```

2.9.6 Routine operations and maintenance of Threat Detection Service

2.9.6.1 Check the service status

2.9.6.1.1 Basic inspection

During the basic inspection of Threat Detection Service (TDS), check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations > Project Operations**. Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
3. Select **BasicCluster**.
4. Check whether `yundun-sas` has reached the final status in **Service Instances List**.

2.9.6.1.2 Advanced inspection

The advanced inspection of TDS checks the status and features of the service.

Procedure

To check the TDS running status, follow the following steps:

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Log on to two TDS physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **BasicCluster**.
 - d) Select `yundun-sas` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select `SasApp#` from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and use TerminalService to log on to two TDS physical machines, respectively.
3. Log on to two TDS Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep sas | awk '{print $1}') bash.
```

4. Check the Java process status.

Run `ps aux | grep sas`. If any record is returned, the process is normal.

5. Check the health status.

Run `curl 127.0.0.1:3008/check.htm`. If OK is returned, the service is normal.

6. View related logs.

- View all logs in `/home/admin/sas/logs/sas-default.log`, including metaq message logs, execution logs of scheduled tasks, and error logs. Typically, you can locate TDS faults based on these logs.
- View the info logs generated when TDS is running in `/home/admin/sas/logs/common-default.log`.
- View the TDS error logs in `/home/admin/sas/logs/common-error.log`.
- View the logs about metaq messages received by TDS in `/home/admin/sas/logs/SAS_LOG.log`.



Note:

Asset verification has been performed on messages in this log file, and the number of messages in this log file is less than that in the `sas-default.log` file.

- View the logs generated when the alert contact sends an alert notification in `/home/admin/sas/logs/notify.log`.

2.9.6.2 Restart TDS

Context

To restart TDS when a fault occurs, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts TDS.

2. Run the following command to find the image ID of TDS:

```
docker ps -a | grep sas
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to locate the Java process:

```
ps aux | grep sas
```

5. Run the following command to stop the current process:

```
kill -9 process
```

6. Run the following command to restart the process:

```
sudo -u admin /home/admin/sas/bin/jbossctl restart
```

7. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/check.htm
```

2.9.7 Routine operations and maintenance of WAF

2.9.7.1 Check service status

2.9.7.1.1 Basic inspection

During the basic inspection of Web Application Firewall (WAF), check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations > Project Operations**. Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
3. Select **WafCluster**.
4. Check whether `yundun-waf` has reached the final status in **Service Instances List**.

2.9.7.1.2 Advanced inspection: Check the Tengine service status

This topic describes how to check the running status of the Tengine service of WAF.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Log on to two WAF physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **WafCluster**.
 - d) Select `yundun-waf` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **Tengine#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.

- f) View **server Information**, and click Terminal to log on to two WAF physical machines, respectively.
3. Log on to two Tengine Docker containers, respectively.
- Run `sudo docker exec -it $(sudo docker ps | grep tengine | awk '{print $1}') bash`.
4. Check the Nginx process status.
- Run `ps aux | grep nginx`. If any record is returned, the Nginx process is normal.
5. Check the service health status.
- Run `curl http://127.0.0.1/get_waf_status -H host:status.waf.example.com`. If the response is "success", the service is normal.
6. View related logs.
- **View the error logs in** `/opt/taobao/tengine/logs/error.log`.
 - **View the access logs in** `/opt/taobao/tengine/logs/access.log`.
 - **View the rule pulling logs in** `/home/admin/aliwaf/logs/waf_agent.log`.

2.9.7.1.3 Advanced inspection: Check the status of the tmd_server service

This topic describes how to check the running status of the tmd_server service of WAF.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console](#).
2. Log on to two WAF physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **WafCluster**.
 - d) Select **yundun-waf** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **Engine#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and use TerminalService to log on to two WAF physical machines, respectively.
3. Log on to two tmd-server Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep tmd-server | awk '{print $1}') bash.
```

4. Check the tmd-server process status.

Run `ps aux | grep tmd_server | grep -v grep`. If any record is returned, the tmd-server process is normal.

5. Check whether port 9002 of the tmd_server service is enabled.

Run `netstat -ano | grep 9002`. If any record is returned, port 9002 is enabled.

6. Check the health status.

Run `curl -v -m 10 -s 127.0.0.1:9002/copy_request`. If `HTTP/1.1 200 OK` is returned, the service is normal.

7. View related logs.

- View the **error logs** in `/home/admin/tmdserver/4/logs/error.log`.

2.9.7.1.4 Advanced inspection: Check the status of the gf_server service

This topic describes how to check the running status of the gf_server service of WAF.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console](#).

2. Log on to two WAF physical machines, respectively.

a) Choose **Operations > Project Operations**.

b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.

c) Select **WafCluster**.

d) Select **yundun-waf** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.

e) Select **Engine#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.

f) View **Server Information**, and use TerminalService to log on to two WAF physical machines, respectively.

3. Log on to two gf-server Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep gf-server | awk '{print $1}') bash.
```

4. Check the gf-server process status.

Run `ps aux | grep gf_server | grep -v grep`. If any record is returned, the gf-server process is normal.

5. Check whether port 9002 of the gf-server process is enabled.

Run `netstat -ano | grep 9002`. If multiple records are returned, port 9002 is enabled.

6. Check the health status.

Run `curl http://127.0.0.1:8002/status.taobao`. If `OK` is returned, the service is normal.

2.9.7.1.5 Check the etcd service status

This topic describes how to check the running status of the etcd service of WAF.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console](#).
2. Log on to two WAF physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **WafCluster**.
 - d) Select `yundun-waf` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **Engine#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and click **Terminal** to log on to two WAF physical machines, respectively.
3. Run the following command to go to the Docker container of the etcd service:

```
sudo docker exec -it $(sudo docker ps | grep etcd | awk '{print $1}')
```

```
bash
```

4. Run the following command to check whether the etcd process is normal:

```
ps aux | grep etcd | grep -v grep
```

If two records are returned, the etcd process is normal.

5. Run the following command to check whether port 4001 of the etcd process is enabled:

```
netstat -ano | grep 4001
```

If multiple records are returned, port 4001 is enabled.

6. Run the following command to perform the health check:

```
curl http://127.0.0.1:4001/v2/keys
```

2.9.7.1.6 Advanced inspection: Checks the status of the logcenter service

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Log on to a WAF physical machine.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **WafCluster**.
 - d) Select **yundun-waf** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **Engine#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and click Terminal to log on to a WAF physical machine.
3. Check the health status.

```
Run curl waflogcenter.${global:intranet-domain}:8108.
```

- If the message `Welcome to CWaf LogCenter!` is returned, the service is normal.
- If the status code `302` is returned, the service is unavailable.

2.9.7.2 Restart WAF

Context

To restart WAF when a fault occurs, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts WAF.
2. Run the following command to find the image ID of the service:

```
sudo docker ps -a |grep service name
```

3. Restart related services.
 - Restart the Tengine service.
 - a. Run the following command to restart the Tengine service:

```
sudo docker restart [imageid]
```

- b.** Run the following command to check whether the Tengine process has been successfully restarted:

```
curl http://127.0.0.1/get_waf_status -H host:status.waf.example.com
```

If the response is "success", the Tengine process has successfully restarted.

- Restart the tmd-server service.

- a.** Run the following command to restart the tmd-server service:

```
sudo docker restart [imageid]
```

- b.** Run the following command to check whether the tmd-server process has been successfully restarted:

```
ps aux | grep tmd_server | grep -v grep
```

If two records are returned, the tmd-server process has successfully restarted.

- c.** Run the following command to check whether port 9002 of the tmd-server process is enabled:

```
netstat -ano | grep 9002
```

If multiple records are returned, port 9002 is enabled.

- d.** Run the following command to check whether the tmd-server service has been successfully restarted:

```
curl -v -m 10 -s 127.0.0.1:9002/copy_request
```

- Restart the gf-server service.

- a.** Run the following command to restart the gf-server service:

```
sudo docker restart [imageid]
```

- b.** Run the following command to check whether the gf-server process has been successfully restarted:

```
ps aux | grep gf_server | grep -v grep
```

If two records are returned, the gf-server process has successfully restarted.

- c.** Run the following command to check whether port 8002 of the gf-server process is enabled:

```
netstat -ano | grep 8002
```

If multiple records are returned, port 8002 is enabled.

- d. Run the following command to check whether the gf-server service has been successfully restarted:

```
curl http://127.0.0.1:8002/status.taobao
```

- Restart the etcd service.

- a. Run the following command to restart the etcd service:

```
sudo docker restart [imageid]
```

- b. Run the following command to check whether the etcd process has been successfully restarted:

```
ps aux | grep etcd | grep -v grep
```

If two records are returned, the etcd process has been successfully restarted.

- c. Run the following command to check whether port 4001 of the etcd process is enabled:

```
netstat -ano | grep 4001
```

If multiple records are returned, port 4001 is enabled.

- d. Run the following command to check whether the etcd service has been successfully restarted:

```
curl http://127.0.0.1:4001/v2/keys
```

2.9.8 Routine operations and maintenance of Security Audit

2.9.8.1 Check service status

2.9.8.1.1 Basic inspection

During the basic inspection of Auditlog, check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations** > **Project Operations**. Enter `yundun-advance`, and click Details to go to the Cluster Operations page.
3. Select **BasicCluster**.
4. Check whether `yundun-security-auditlog` has reached the final status in **Service Instances List**.

2.9.8.1.2 Advanced inspection: check the status of the security-auditlog-app service

This topic describes how to check the status of the security-auditlog-app service.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)

To inspect the service status, follow the following steps:

2. Log on to two Security Audit physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **BasicCluster**.
 - d) Select **yundun-security-auditlog** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **SecurityAuditlogApp#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and click **Terminal** to log on to two Security Audit physical machines, respectively.
3. Log on to two auditlog-app Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep auditlog-app | awk '{print $1}') bash.
```

4. Check the security-auditlog process status.

```
Run ps aux | grep java | grep security-auditlog. If any record is returned, the security-auditlog process is normal.
```

5. Check the health status.

```
Run curl 127.0.0.1:3001/check.htm. If the response is "success", the service is normal.
```

6. View related logs.

- View the Tomcat logs in `/home/admin/security-auditlog/logs/jboss_stdout.log.`
- View the audit logs in `/home/admin/security-auditlog/logs/audit-exec.log.`
- View the business error logs in `/home/admin/security-auditlog/logs/biz-error.log.`

- View the check error logs in `/home/admin/security-auditlog/logs/check-error.log`.
- View the task scheduling logs in `/home/admin/security-auditlog/logs/job-exec.log`.
- View the remote service call logs in `/home/admin/security-auditlog/logs/remote-exec.log`.
- View the service call logs in `/home/admin/security-auditlog/logs/service-exec.log`.
- View the system error logs in `/home/admin/security-auditlog/logs/system-error.log`.
- View the download task logs in `/home/admin/security-auditlog/logs/task-exec.log`.
- View other logs in `/home/admin/security-auditlog/logs/main.log`.

2.9.8.1.3 Advanced inspection: Check the security-auditlog-syslog service status

This topic describes how to check the running status of the security-auditlog-syslog service.

Procedure

To inspect the service status, follow the following steps:

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Log on to two Security Audit physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **BasicCluster**.
 - d) Select **yundun-security-auditlog** from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select **SecurityAuditlogApp#** from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **Server Information**, and click **Terminal** to log on to two Security Audit physical machines, respectively.
3. Log on to two auditlog-syslog Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep auditlog-syslog | awk '{print $1}') bash.
```

4. Check the syslog-ng process status.

Run `ps aux | grep syslog-ng | grep -v grep`. If two records are returned, the syslog-ng process is normal.

5. Check the status of port 2514 of the syslog-ng process.

Run `netstat -ano | grep 2514`. If multiple records are returned, port 2514 of the syslog-ng process is normal.

6. Check the ilogtail process status.

Run `ps aux | grep ilogtail | grep -v grep`. If two records are returned, the ilogtail process is normal.

7. View related logs.

- View the ilogtail logs in `/usr/local/ilogtail/ilogtail.LOG`.
- View the syslog-ng logs in `/var/log/messages`.

2.9.8.2 Restart Security Audit

Context

To restart Security Audit when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server of Auditlog.
2. Run the following command to find the image ID of Auditlog:

```
docker ps -a | grep service name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Restart related services.

- Restart the security-auditlog-app service.
 - a. Run the following command to stop the current application process:

```
kill -9 $(ps -ef | grep java | grep security-auditlog | grep -v  
grep | awk '{print $2}')
```

- b. Run the following command to restart the application process:

```
/home/admin/security-auditlog/bin/jbossctl restart
```

- c. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/check.htm
```

If the response is "success", the service is normal.

- Restart the security-auditlog-syslog service.

- a. Run the following command to restart the syslog-ng process:

```
service syslog-ng restart
```

- b. Run the following command to check whether the syslog-ng process is normal:

```
ps aux | grep syslog-ng | grep -v grep
```

If two records are returned, the syslog-ng process is normal.

- c. Run the following command to check whether port 2514 of the syslog-ng process is enabled:

```
netstat -ano | grep 2514
```

If multiple records are returned, port 2514 of the syslog-ng process is enabled.

- d. Run the following command to restart the ilogtaild process:

```
/etc/init.d/ilogtaild stop
```

```
/etc/init.d/ilogtaild start
```

- e. Run the following command to check whether the ilogtail process is normal:

```
ps aux | grep ilogtail | grep -v grep
```

If two records are returned, the ilogtail process is normal.

2.9.9 Routine operations and maintenance of Apsara Stack Security Center

2.9.9.1 Check service status

2.9.9.1.1 Basic inspection

During the basic inspection of Apsara Stack Security Center, check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)

2. Choose **Operations > Project Operations**. Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
3. Select **BasicCluster**.
4. Check whether `yundun-secureconsole` has reached the final status in **Service Instances List**.

2.9.9.1.2 Advanced inspection

Check the running status of Apsara Stack Security Center.

Context

To check the running status of Apsara Stack Security Center, follow the following steps:

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console](#).
2. Log on to two physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **BasicCluster**.
 - d) Select `yundun-secureconsole` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select `SecureConsoleApp#` from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **server Information**, and use TerminalService to log on to two physical machines, respectively.
3. Log on to two secure-console Docker containers, respectively.


```
Run sudo docker exec -it $(sudo docker ps | grep secureconsole | awk '{print $1}') bash.
```
4. Check the console progress status.


```
Run ps aux | grep console. If any record is returned, the console progress is normal.
```
5. Check the health status.


```
Run curl 127.0.0.1:3014/check.htm. If OK is returned, the service is normal.
```
6. View related logs.
 - View the Tomcat logs in `/home/admin/console/logs/jboss_stdout.log`.

2.9.9.2 Restart the secure-console service

Context

To restart the secure-console service when an error occurs, follow the following steps :

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the secure-console service.
2. Run the following command to find the image ID of the secure-console service:

```
sudo docker ps -a |grep console
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to locate the Java process:

```
ps aux |grep console
```

5. Run the following command to stop the current process:

```
kill -9 process
```

6. Run the following command to restart the process:

```
sudo -u admin /home/admin/console/bin/jbossctl restart
```

7. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/check.htm
```

2.9.10 Routine operations and maintenance of secure-service

2.9.10.1 Check the service status

2.9.10.1.1 Basic inspection

During the basic inspection of secure-service, check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations** > **Project Operations** on the page that appears, enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
3. Select **BasicCluster**.

4. Check whether `yundun-secureservice` has reached the final status in **Service Instances List**.

2.9.10.1.2 Advanced inspection: Check the secure-service status

This topic describes how to check the secure-service running status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console](#).
2. Log on to two physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
 - c) Select **BasicCluster**.
 - d) Select `yundun-secureservice` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select `SecureServiceApp#` from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **Server Information**, and click Terminal to log on to two physical machines, respectively.
3. Log on to two secure-service Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep secureservice | awk '{print $1}') bash.
```

4. Check the secure-service process status.

Run `ps aux | grep secure-service`. If any record is returned, the secure-service process is normal.

5. Check the health status.

Run `curl 127.0.0.1:3010`. If OK is returned, the service is normal.

6. Run the following command to go to the Docker container:

```
sudo docker exec -it [imageId] /bin/bash
```

7. View related logs.

- View the Server Guard logs in `/home/admin/secure-service/logs/aegis-info.log`.
- View the error logs in `/home/admin/secure-service/logs/Error`.

- View the vulnerability analysis and scanning logs in `/home/admin/secure-service/logs/leakage-info.log`.
- View the cloud intelligence logs in `/home/admin/secure-service/logs/threat-info.log`.
- View the web attack logs in `/home/admin/secure-service/logs/web-info.log`.

2.9.10.1.3 Check the Dolphin service status

Context

To check the running status of the Dolphin service, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Dolphin service.

2. Run the following command to find the image ID of the Dolphin service:

```
sudo docker ps -a |grep dolphin
```

3. Run the following command to go to the Docker container:

```
sudo docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep dolphin
```

5. Run the following command to perform the health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

6. View related logs.

- View the info logs generated when the Dolphin service is running in `/home/admin/dolphin/logs/common-default.log`.
- View the Dolphin service error logs in `/home/admin/dolphin/logs/common-error.log`.
- View the metaq messages received by the Dolphin service in `/home/admin/dolphin/logs/dolphin-message-consumer.log`.

**Note:**

Currently, only Threat Detection Service (TDS) sends messages to the Dolphin service.

- View the metaq messages sent by the Dolphin service in `/home/admin/dolphin/logs/dolphin-message-producer.log`.

**Note:**

Currently, the Dolphin service sends messages only to TDS.

2.9.10.1.4 Check the data-sync service status

Context

To check the running status of the data-sync service, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the data-sync service.
2. Run the following command to find the image ID of the data-sync service:

```
sudo docker ps -a |grep data-sync
```

3. Run the following command to go to the Docker container:

```
sudo docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep data-sync
```

5. Run the following command to perform health check:

```
curl 127.0.0.1:7001/check_health
```

If OK is returned, the service is normal.

6. View related logs.

View the data-sync service logs in `data-sync.log`.

2.9.10.2 Restart secure-service

Context

To restart secure-service when a fault occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server of the service.
2. Run the following command to find the image ID of the service:

```
docker ps -a |grep application name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Restart related services.

- Restart secure-service.

a. Run the following command to view the Java process ID:

```
ps aux |grep secure-service
```

b. Run the following command to stop the current process:

```
kill -9 process
```

c. Run the following command to restart the process:

```
sudo -u admin /home/admin/secure-service/bin/jbossctl restart
```

d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001
```

- Restart the Dolphin service.

a. Run the following command to view the Java process ID:

```
ps aux |grep dolphin
```

b. Run the following command to stop the current process:

```
kill -9 process
```

c. Run the following command to restart the process:

```
sudo -u admin /home/admin/dolphin/bin/jbossctl restart
```

d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/checkpreload.htm
```

- Restart the data-sync service.

a. Run the following command to view the Java process ID:

```
ps aux |grep data-sync
```

b. Run the following command to stop the current process:

```
kill -9 process
```

c. Run the following command to restart the process:

```
sudo -u admin /home/admin/data-sync/bin/jbossctl restart
```

- d. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/check_health
```

2.10 Key Management Service (KMS)

2.10.1 Operations and maintenance of KMS components

2.10.1.1 Overview

KMS is deployed and managed from the Apsara Infrastructure Management Framework console. On the **Server Operations** page of the Apsara Infrastructure Management Framework console, you can access a host where KMS is deployed.

2.10.1.2 KMS_HOST

Determine whether the service role functions properly

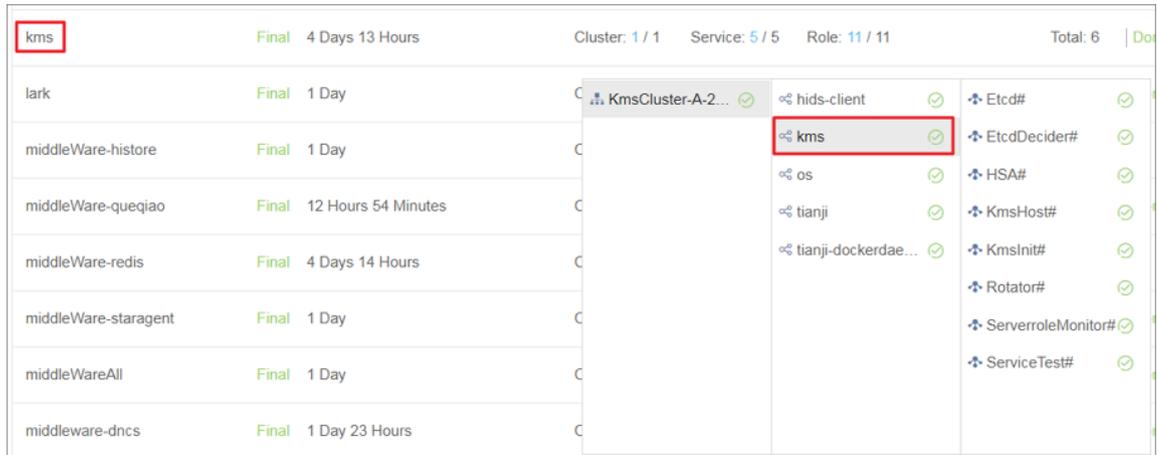
1. In the Apsara Infrastructure Management Framework console, check whether the KMS_HOST service role has been deployed.

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Tasks > Deployment Summary**. The **Deployment Summary** page is displayed.
- c. Click **Deployment Details**.
- d. On the **Deployment Details** page, locate kms.
- e. In the top navigation bar, choose **TasksDeployment Summary**. On the Deployment Summary page that appears, click **Deployment Details** to view the deployment status of the KMS_HOST service role, as shown in [Figure 2-28: Check whether the KMS_HOST service role has been deployed](#).

If a tick next to **KMS_HOST#** is green, the KMS_HOST service role has been deployed.

Figure 2-28: Check whether the KMS_HOST service role has been deployed

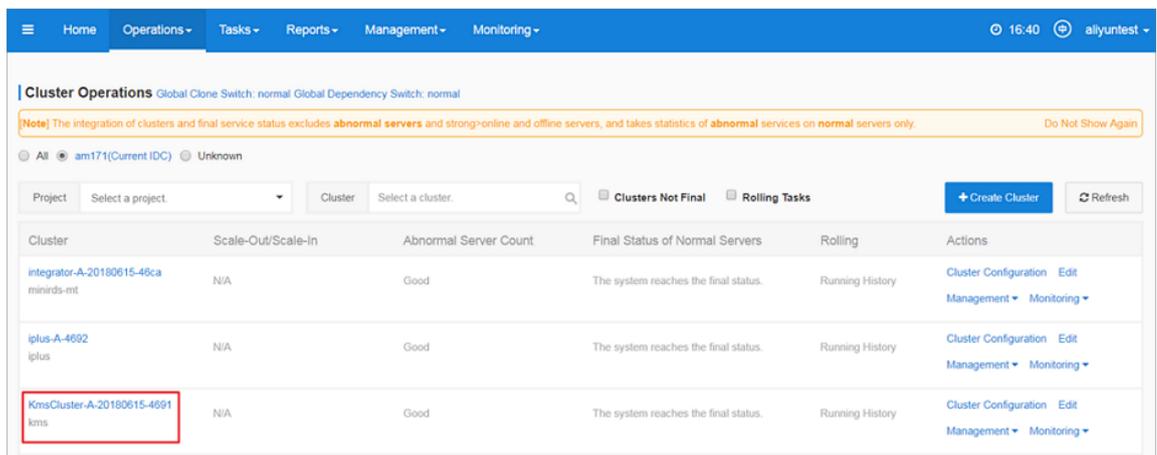


2. Obtain the IP addresses of the hosts where the KMS_HOST service role has been deployed.

Follow these steps:

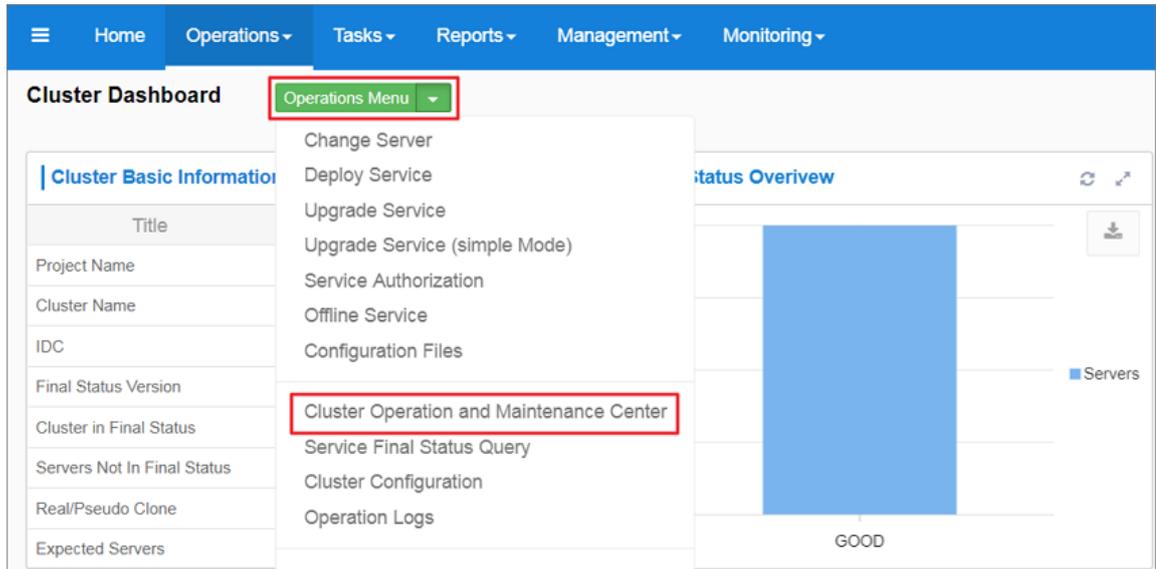
- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Operations** > **Cluster Operations**. On the Cluster Operations page that appears, search for the corresponding cluster, as shown in [Figure 2-29: Search for clusters](#).

Figure 2-29: Search for clusters



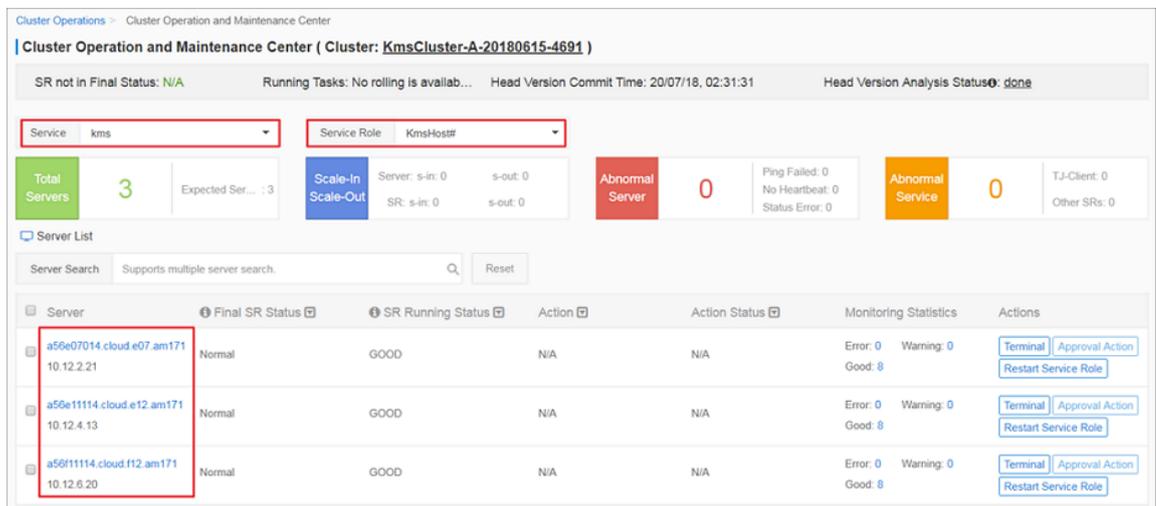
- c. Click a specified cluster URL to go to the **Cluster Dashboard** page.
- d. On the **Cluster Dashboard** page, choose **Operations Menu** > **Cluster Operation and Maintenance Center**, as shown in [Figure 2-30: Cluster Dashboard](#).

Figure 2-30: Cluster Dashboard



- e. On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of all hosts where the KMS_HOST service role has been deployed, as shown in [Figure 2-31: View IP addresses of hosts](#).

Figure 2-31: View IP addresses of hosts



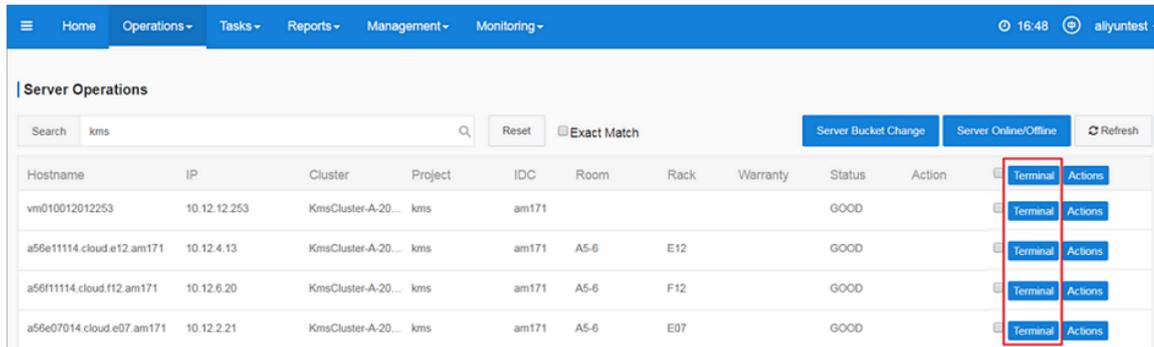
- 3. Access the KMS server and run the `curl http://ip:5555/status.html` command to check whether **success** is returned.

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.

- b. In the top navigation bar, choose **Operations > Server Operations**. On the Server Operations page that appears, locate a relevant host, as shown in [Figure 2-32: Search for hosts](#).

Figure 2-32: Search for hosts

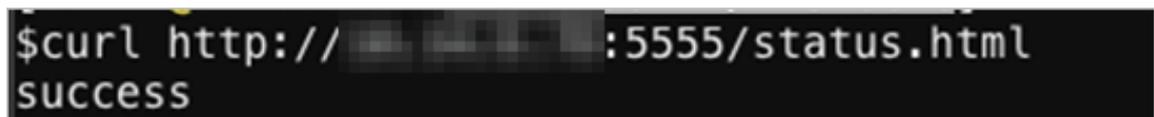


The screenshot shows the 'Server Operations' interface. At the top, there is a navigation bar with 'Home', 'Operations', 'Tasks', 'Reports', 'Management', and 'Monitoring'. Below this, a search bar contains the text 'kms'. To the right of the search bar are buttons for 'Reset', 'Exact Match', 'Server Bucket Change', 'Server Online/Offline', and 'Refresh'. Below the search bar is a table with columns: Hostname, IP, Cluster, Project, IDC, Room, Rack, Warranty, Status, Action, Terminal, and Actions. The table contains four rows of host data. The 'Terminal' and 'Actions' buttons for each row are highlighted with a red box.

Hostname	IP	Cluster	Project	IDC	Room	Rack	Warranty	Status	Action	Terminal	Actions
vm010012012253	10.12.12.253	KmsCluster-A-20...	kms	am171				GOOD		Terminal	Actions
a56e11114.cloud.e12.am171	10.12.4.13	KmsCluster-A-20...	kms	am171	A5-6	E12		GOOD		Terminal	Actions
a56f11114.cloud.f12.am171	10.12.6.20	KmsCluster-A-20...	kms	am171	A5-6	F12		GOOD		Terminal	Actions
a56e07014.cloud.e07.am171	10.12.2.21	KmsCluster-A-20...	kms	am171	A5-6	E07		GOOD		Terminal	Actions

- c. Select a host and click **Terminal** to log on to the host through a terminal session.
- d. Run the `curl http://ip:5555/status.html` command to check whether **success** is returned, as shown in [Figure 2-33: Enter a command](#). Verify all hosts where the KMS_HOST service role has been deployed based on the previous procedure. IP indicates the IP addresses of the hosts that are obtained in the previous step.

Figure 2-33: Enter a command



```
$ curl http://[redacted]:5555/status.html
success
```

Locate and determine exceptions

1. View log entries in `/cloud/log/kms/KmsHost#/kms_host`.
2. Check whether the KMS_HOST service role functions properly. If the KMS_HOST service role exits soon after it is started, check `debug.log` to locate the cause.
3. If the KMS_HOST service role runs properly with faulty functions, view `status.log` to locate the cause.

Possible exceptions and errors

- `xxx selfCheck error`



Note:

xxx indicates a dependent service.

1. Check whether the corresponding dependency configurations are correct. You can use `debug.log` to locate the configurations.
 2. Check whether xxx runs properly.
- `exit code 1`

Locate the cause of an unexpected exit based on `debug.log`.

2.10.1.3 HSA

Determine whether the service role is normal

1. In the Apsara Infrastructure Management Framework console, check whether the HSA service role has been deployed.

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Tasks > Deployment Summary**. The **Deployment Summary** page is displayed.
- c. Click **Deployment Details**.
- d. On the **Deployment Details** page, locate kms.
- e. In the top navigation bar, choose **TasksDeployment Summary**. On the Deployment Summary page that appears, click **Deployment Details** to view the deployment status of the HSA service role, as shown in [Figure 2-34: Check whether the HSA service role has been deployed](#).

If the tick next to `HSA#` is green, the HSA service role has been deployed.

Figure 2-34: Check whether the HSA service role has been deployed

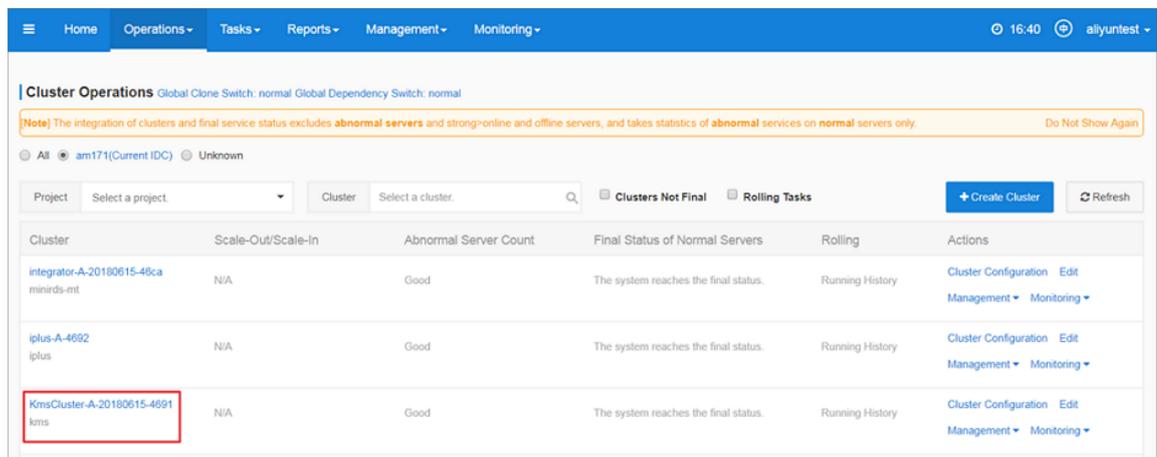
Service Name	Status	Duration	Cluster	Service	Role	Total
kms	Final	4 Days 13 Hours	1 / 1	5 / 5	11 / 11	6
lark	Final	1 Day				
middleWare-histore	Final	1 Day				
middleWare-queqiao	Final	12 Hours 54 Minutes				
middleWare-redis	Final	4 Days 14 Hours				
middleWare-staragent	Final	1 Day				
middleWareAll	Final	1 Day				
middleware-dnccs	Final	1 Day 23 Hours				

2. Obtain the IP addresses of the hosts where the KMS_HOST service role has been deployed.

Follow these steps:

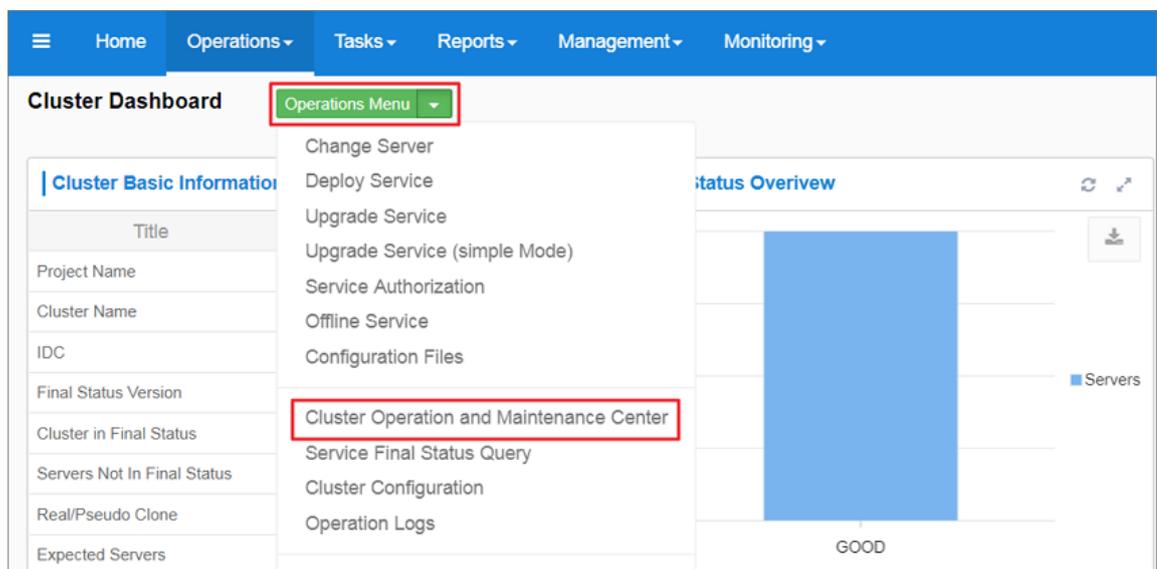
- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page that appears, search for the corresponding cluster, as shown in [Figure 2-35: Search for clusters](#).

Figure 2-35: Search for clusters



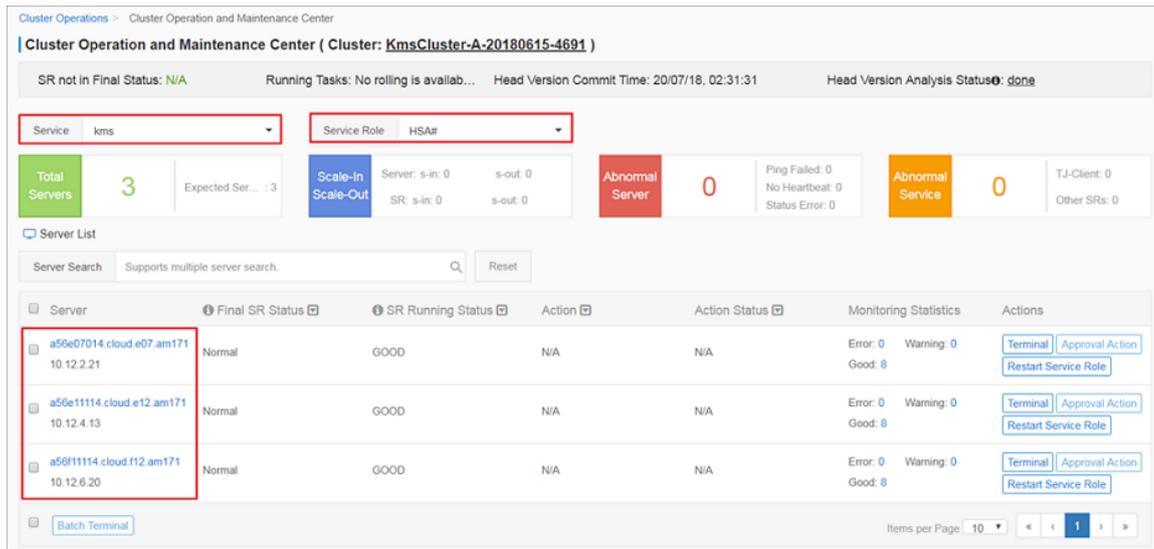
- c. Click a specified cluster URL to go to the **Cluster Dashboard** page.
- d. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**, as shown in [Figure 2-36: Cluster Dashboard](#).

Figure 2-36: Cluster Dashboard



- e. On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of the hosts where the HSA service role has been deployed, as shown in [Figure 2-37: Obtain the IP addresses of the hosts where the HSA service role has been deployed](#).

Figure 2-37: Obtain the IP addresses of the hosts where the HSA service role has been deployed

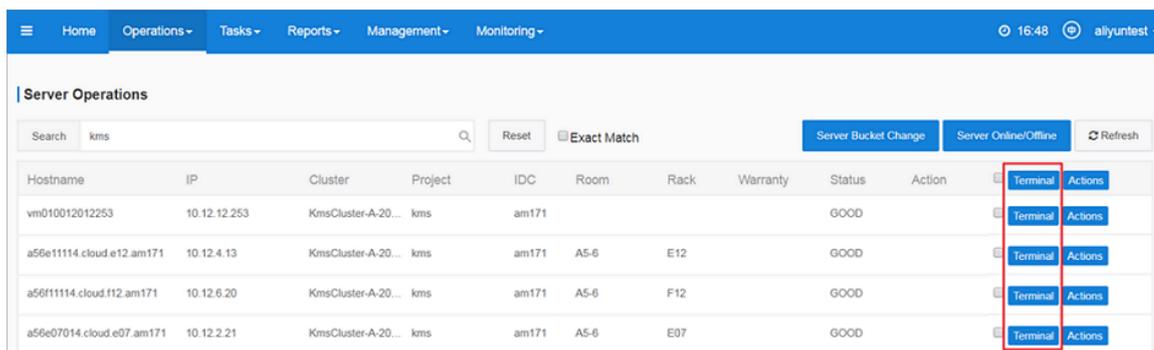


3. Access the KMS server and run the `curl http://ip:8081/status.html` command to determine whether **success** is returned.

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Operations > Server Operations**. On the Server Operations page that appears, locate a relevant host, as shown in [Figure 2-38: Search for hosts](#).

Figure 2-38: Search for hosts



- c. Select a host and click **Terminal** to log on to the host through a terminal session.

- d. Run the `curl http://ip:8081/status.html` command to verify the hosts where the HSA service role is deployed and determine whether `success` is returned, as shown in [Figure 2-39: Enter a command](#).

IP indicates the IP addresses obtained in the previous step. It is the IP address of a host where the HSA service role has been deployed.

Figure 2-39: Enter a command

```
$curl http://[redacted]:8081/status.html
success
```

Locate and determine exceptions

1. View log entries in `/cloud/log/kms/HSA#/hsa`.
2. Check whether the HSA service role runs properly. If the HSA service role exits soon after it is started, view `debug.log` to locate the cause.
3. If the HSA service role runs properly with faulty functions, view `status.log` to locate the cause.

Possible exceptions and errors

Error: `exit code 1`

View `debug.log` to locate the cause of the exceptional exit.

Typical possible causes include:

- The Etcd service role is not properly started.
- The Etcd service role is properly started but has no valid data.



Note:

This error may also occur when the secondary cluster encounters data synchronization exceptions during disaster recovery.

2.10.1.4 Etcd

Determine whether the service role functions properly

In the Apsara Infrastructure Management Framework console, determine whether Etcd and EtcdDecider service roles have been deployed.

Follow these steps:

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the top navigation bar, choose **Tasks > Deployment Summary**. The **Deployment Summary** page is displayed.
3. Click **Deployment Details**.
4. On the **Deployment Details** page, locate kms.
5. In the top navigation bar, choose **TasksDeployment Summary**. On the Deployment Summary page that appears, click **Deployment Details** to view the deployment status of Etcd and EtcdDecider service roles, as shown in *Figure 2-40: Check whether Etcd and EtcdDecider service roles have been deployed*.

If the ticks next to **Etcd#** and **EtcdDecider#** are green, they have been deployed.

Figure 2-40: Check whether Etcd and EtcdDecider service roles have been deployed

Service Name	Status	Duration	Cluster	Service	Role	Total
kms	Final	4 Days 13 Hours	1 / 1	5 / 5	11 / 11	6
lark	Final	1 Day	KmsCluster-A-2...	hids-client	Etcd#	
middleWare-histore	Final	1 Day		kms	EtcdDecider#	
middleWare-queqiao	Final	12 Hours 54 Minutes		os	HSA#	
middleWare-redis	Final	4 Days 14 Hours		tianji	KmsHost#	
middleWare-staragent	Final	1 Day		tianji-dockerdae...	KmsInit#	
middleWareAll	Final	1 Day			Rotator#	
middleWare-dnscs	Final	1 Day 23 Hours			ServerroleMonitor#	
					ServiceTest#	

Locate and determine exceptions

1. Etcd service role exceptions are complex. Log entries in `/cloud/log/kms/Etcd#/etcd` only provide some information.
2. To view the rest of the information, you must view the log entries in `/cloud/log/kms/EtcdDecider#/decider` for comprehensive analysis.

Possible exceptions and errors

Possible exceptions and errors are as follows:

- The startup parameters of the Etcd service role are not properly calculated for some special reasons.

To locate the specific reasons, retain the on-site log entries and contact Customer Services for troubleshooting.

Solution: Locate normal startup parameters in historic log entries of `debug.log` and manually start the Etcd service role.

- The abnormal Decider service role during service upgrades causes Etcd service role errors.

Typically, this error occurs when a rolling task exists. You can locate the cause in `debug.log` of the Decider service role.

- The data directory of the Etcd service role is missing and the Etcd service role fails to be started.

Solution: Scale in the abnormal Etcd node through the Apsara Infrastructure Management Framework console, and then restore it back to the original capacity.

2.10.1.5 Rotator

2.10.1.5.1 Primary IDC

The status of the Rotator service role is special. Even if the Apsara Infrastructure Management Framework console shows that the service role has been deployed, the Rotator service role is not necessarily working properly.

Rotator service role exceptions do not have any adverse impact on the API logic of KMS.

Typically, you must use log entries to locate the cause of a fault for the Rotator service role only after unexpected results are found. For example, the data on RDS does not match expectations.

Determine whether the Rotator service role is enabled in the primary IDC mode

View `current_idc_master` in `/cloud/log/kms/Rotator#/rotator/debug.log`, as shown in [Figure 2-41: View the IDC mode](#). Determine whether the Rotator service role (in the primary IDC) is enabled in the primary IDC mode.

Figure 2-41: View the IDC mode

```
[2017-10-16 13:07:50.458588] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] PKIVersion:pssl
[2017-10-16 13:07:50.497312] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] current_idc_master: true
[2017-10-16 13:07:50.553460] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] CurrentClients:map[a27d05007.ccloud.d05.ew9-5:0xc4206d6720 a27d08007.ccloud.d08.ew9-5:0xc420647da0 a27d11007.ccloud.d11.ew9-5:0xc4206d7980]
```

If the `current_idc_master` value indicates `true`, the Rotator service role is enabled in the primary IDC mode. If the `current_idc_master` value indicates `false`, the Rotator service role is enabled in the secondary IDC mode.

Check whether the Rotator service role is in the working state

The Rotator service role of the primary IDC is deployed to all nodes in the distributed lock mode. In this mode, only one node is in the working state and the other nodes are all in the standby state.

View `/cloud/log/kms/Rotator#/rotator/status.log` to determine whether the Rotator service role is in the working state.

As shown in [Figure 2-42: Working state](#) and [Figure 2-43: Standby state](#):

- ExecuteWorker: The node is in the working state.
- TryLock: The node is in the standby state.

Figure 2-42: Working state

```
[2017-10-23 16:51:51.554310] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d05007.cloud.d05.ew9-5 RotatorState:ExecuteWorker
[2017-10-23 16:52:51.554415] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d05007.cloud.d05.ew9-5 RotatorState:ExecuteWorker
```

Figure 2-43: Standby state

```
11007.cloud.d11.ew9-5 RotatorState:TryLock
[2017-10-17 18:35:20.618575] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d11007.cloud.d11.ew9-5 RotatorState:TryLock
[2017-10-17 18:36:11.867967] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d11007.cloud.d11.ew9-5 RotatorState:TryLock
[2017-10-17 18:36:20.620963] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d11007.cloud.d11.ew9-5 RotatorState:TryLock
```

Possible exceptions and errors

- Abnormal RDS database access. The statistic collection and key deletion tasks cannot be run properly.
- Abnormal HSA service role. The key update task cannot be run properly.
- Abnormal Log Service. The metering task cannot be run properly.
- Abnormal Etcd service role. The distributed lock is unavailable, and tasks cannot be run.
- If one of the tasks on the Rotator service role is abnormal, the Rotator service role may be unable to be deployed in the Apsara Infrastructure Management Framework console.

2.10.1.5.2 Secondary IDC

The Rotator service role of the secondary IDC is deployed to all nodes. Each node is in the working state. The work scope of each node is idempotent to those of other nodes within a certain time range.

Determine whether the Rotator service role is enabled in the secondary IDC mode

View current `idc master` in `/cloud/log/kms/Rotator#/rotator/debug.log`, as shown in [Figure 2-44: View the IDC mode](#). Determine whether the Rotator service role (in the secondary IDC) is enabled in the secondary IDC mode.

Figure 2-44: View the IDC mode

```
[2017-10-21 16:34:34.412535] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] PKIVersion:psl
[2017-10-21 16:34:34.446620] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] current idc master: false
```

If the `current idc master` value indicates `false`, the Rotator service role is enabled in the secondary IDC mode. If the `current idc master` value indicates `true`, the Rotator service role is enabled in the primary IDC mode.

Possible exceptions and errors

- Abnormal primary IDC networks. The Etcd service role of the primary IDC is inaccessible.
- Abnormal Etcd service role of the primary IDC. The Etcd service role of the primary IDC is inaccessible.
- Abnormal Etcd service role of the secondary IDC. Data write fails.
- Incorrect Etcd service role information of the primary IDC. An error occurs during data synchronization.



Note:

Rotator service role exceptions of the secondary IDC severely affect the KMS in the primary IDC. Handle this exception immediately.

2.10.2 Log analysis

2.10.2.1 Overview

Logtail is a log collection client provided by Log Service to facilitate your access to logs. After installing Logtail on a host that has KMS deployed, you can monitor a specified log. The newly written log entries are automatically uploaded to a specified log library.

Logtail is used to transmit the logs of KMS to Log Service. Then the portal or API of Log Service analyzes the logs. If Log Service has no portals, you have to log on to the hosts that have KMS deployed individually and check the hosts one by one.

2.10.2.2 Request IDs

After sending a request to KMS, you will receive a response from KMS. The response contains a request ID.

Request IDs are used in the following scenarios:

- Go to `/cloud/log/kms/KmsHost#/kms_host/audit.log` to view the KMS audit log.

You can use the `request_id` value to view the audit log information of the current access.

- For log entries whose `expected_code` values are not 200, you can view error information during debugging based on the request ID.

Path to the local log: `/cloud/log/kms/KmsHost#/kms_host/debug.log`



Note:

`/cloud/log/kms/KmsHost#/kms_host/debug.log` and `audit.log` are stored in the same host.

- If you need all details of a request, you can view detailed information in the trace log.

Path to the local log: `/cloud/log/kms/KmsHost#/kms_host/debug.log`



Note:

`/cloud/log/kms/KmsHost#/kms_host/debug.log` and `audit.log` are stored in the same host.

- You can use the request ID to associate cryptology-relevant APIs with the trace log of HSA.

Path to the local log: `/cloud/log/kms/HSA#/hsa/trace.log`



Note:

`/cloud/log/kms/KmsHost#/kms_host/trace.log` and `audit.log` may be stored in different hosts.

- You can also retrieve logs based on other information.

You can retrieve audit logs of KMS based on other information. However, you still need the request ID to associate the audit logs with other logs.

2.10.2.3 Common KMS errors

2.10.2.3.1 Overview

KMS has two HTTP status codes in audit.log: `expected_code` and `status_code`.

Typically, the expected code and status code of an error are the same. (`expected_code = status_code`). However, there are exceptions.

`status_code` is the HTTP status code that is actually returned to a user.

2.10.2.3.2 Error codes 4xx

Error codes 4xx indicate expected errors in KMS. For example, error code 403 indicates a user authentication request failure and error code 400 indicates that the parameters entered by a user are incorrect.

You can use the request ID to view detailed error information during debugging.

2.10.2.3.3 Error code 500

Typically, if the status code of an error is 500, the expected code of the error is also 500.

Errors of this type are not expected by KMS. Typically, they are severe errors and must be fixed immediately.

A dependent service probably encounters unexpected errors. We recommend that you contact Customer Services to troubleshoot the problem.

You can use the request ID to view detailed error information during debugging.

2.10.2.3.4 Error code 503

Error code 503 occurs in the following scenarios:

- The expected code is not 503 but the status code is 503.

Possible causes:

- The client-side user interrupted the connection in advance.
- The client has timed out because the response of the KMS server is too slow.

You can use the request ID in the trace log to determine whether the server has timed out and identify the modules that have timed out.

- `expected_code=status_code=503`

An expected error occurred in a dependent service of the KMS. This is caused by an unstable dependent service.

You can use the request ID to view detailed error information during debugging. We recommend that you contact Customer Services to troubleshoot the problem.

2.10.2.3.5 Dependent service degradation

KMS caches dependent service data to local memory. If a dependent service is unavailable, KMS uses the obsolete data cached to the local device to continue providing services.

In this scenario, the status code in the audit log of KMS is 200, but an additional debug log entry is generated.

In this scenario, some users can access KMS (data cached), but other users may encounter a 503 error (data not cached).

2.10.3 View and process internal data

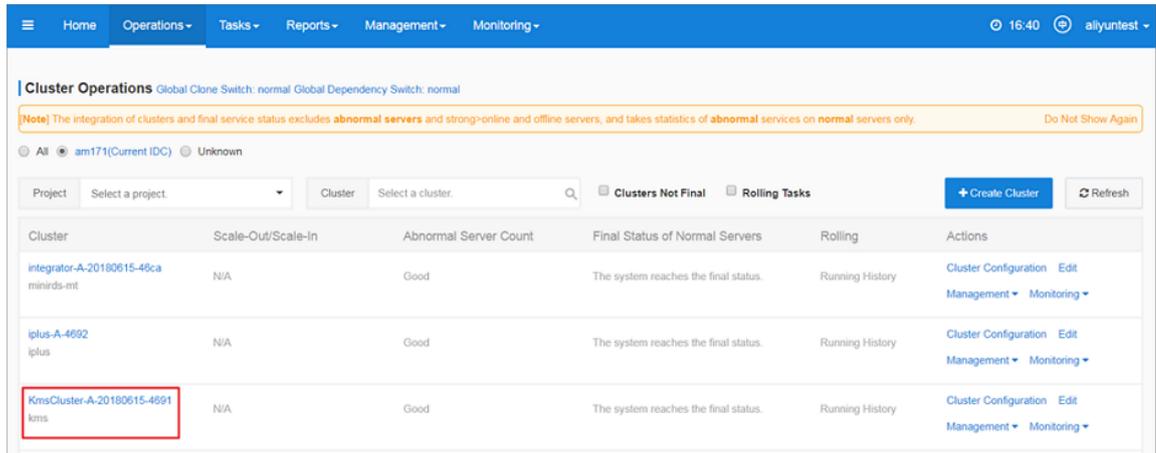
View updates to CMKs

1. Log on to the kmsdata database from an Apsara Stack server that has a MySQL client installed.

To obtain connection information of the kmsdata database, perform the following operations:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page that appears, search for the corresponding cluster, as shown in [Figure 2-45: Search for clusters](#).

Figure 2-45: Search for clusters



- c. Click a specified cluster URL to go to the **Cluster Dashboard** page.
- d. Scroll down the page and locate **Cluster Resource**.
- e. In **Cluster Resource**, locate the kmsdata database, as shown in [Figure 2-46: Cluster resources](#).

Figure 2-46: Cluster resources

The screenshot shows the 'Cluster Resource' table. The table has columns: Service, server..., app, name, type, status, error..., param..., result, res, reproc..., reproc..., reproc..., refer_v... The 'kmsdata' database is highlighted with a red box in the 'name' column.

Service	server...	app	name	type	status	error...	param...	result	res	reproc...	reproc...	reproc...	refer_v...
kms	kms.KmsInit#	db_init	kmsdata	db	done		{"minirds_p...	{"passwd": "...	ddec793d5...				{"#587c17...
kms	kms.KmsInit#	db_init	kms_servic...	accesskey	done		{"name": "k...	{"name": "k...	610f802a4...				{"#587c17...
kms	kms.KmsInit#	db_init	kms-internet	vip	done		{"check_ty...	{"nc_list": "...	4f59db9f05...				{"#587c17...
kms	kms.KmsInit#	db_init	kms-intranet	vip	done		{"check_ty...	{"nc_list": "...	3e4761db5...	done			{"nc_list": "...
kms	kms.KmsInit#	db_init	kms-intranet	dns	done		{"domain": ...	{"ip": "[*10...	4ec409337...				{"#587c17...
kms	kms.KmsInit#	db_init	kms-internet	dns	done		{"domain": ...	{"ip": "[*42...	3fa990dcde...				{"#587c17...

- f. In the **result** column corresponding to the cluster, right-click and choose **Show More** from the shortcut menu.

In the **Details** message that appears, you can view the connection information of the kmsdata database.

2. Enter the `select MIN(dk_version) from ekt_tbl;` SQL statement, and view the dk_version information, as shown in [Figure 2-47: View dk_version information](#).



Note:

dk_version indicates the date of the current day.

**Note:**

max_count indicates the upper limit. user_id indicates the PK of an Apsara Stack tenant account.

2.11 Log Service

2.11.1 Components operations and maintenance

Log Service is deployed, operated, and maintained using Apsara Infrastructure Management Framework. You can click **Machine Operation and Maintenance** to log on to the machine on which Log Service resides.

The Log Service console is the interface on which you operate Log Service and also the portal of all operations on Log Service. It is compliant with the standard deployment mode of Java application specified by Alibaba Cloud. Each console instance contains an Nginx server and a Jetty container.

Command portal

1. Log on to the Apsara Infrastructure Management Framework console.
2. From the top navigation bar, choose **Operations > Service Operations**.
3. Click **Management** in the Actions column corresponding to `sls-backend-server`. On the displayed page, click the service instance name to go to the **Service Instance Dashboard** page.
4. Locate the `webServer#` service role in **Service Role List** and click **Details** to view the machine.
5. Click Terminal in the Actions column corresponding to a machine, log on to the machine, and open the file directories to find the Nginx and Jetty services.

Commands

- Reboot the Nginx server:

```
sudo /etc/init.d/nginx restart
```

- Reboot the Jetty container:

```
sudo /etc/init.d/jetty restart
```

Directory structure

- Web application root directory: `/alidata/www/`
- Console application war directory: `/alidata/www/wwwroot/sls-console-aliyun-com/`

- Service application war directory: `/alidata/www/wwroot/sls-service-aliyun-com/`
- Static resources directoty: `/alidata/www/wwroot/static/`

Configuration files

- Nginx configuration: `/etc/nginx/conf.d/sls-console-aliyun-com.conf.console`
- Console configuration: `/alidata/www/wwroot/sls-console-aliyun-com/WEB-INF/classes/config/web.properties`
- Service configuration: `/alidata/www/wwroot/sls-service-aliyun-com/WEB-INF/classes/config/sls.properties`

Application logs

- Nginx log: `/apsara/nginx/logs/sls_console.log`
- Log root directory: `/alidata/www/logs/`
- Console application log: `/alidata/www/logs/java/sls/`
- Service application log: `/alidata/www/logs/java/sls-service/`
- Jetty log: `/usr/share/jetty/log/`

2.11.2 O&M and troubleshooting

2.11.2.1 Nginx

Error log: `/apsara/nginx/log/error.log`

Error	Solution
Bind Address Failed	Check the <code>/etc/init.d/nginx.conf</code> listener port.
open() ... failed	Check whether the corresponding item of the static resource file exists.

2.11.2.2 Console

Error log: `/alidata/www/logs/java/sls/error.log`

Error	Solution
SLS SDK Exception	Normal. No action is required.
Create Bean Failed	Check the dubbo settings in the Console configurations.

2.11.2.3 Service

Error log: `/alidata/www/logs/java/sls-service/applog/error.log`

Error	Solution
Create Bean Failed	Check the dubbo settings in the Service configurations.
Invoke failed	Check the scmg settings in the Service configurations.

2.12 Domain Name System (DNS)

2.12.1 Introduction to Apsara Stack DNS

This section describes Apsara Stack DNS and the features of its modules.

Database management system

The database management system compares versions in the baseline configurations with the versions in the database to manage database versions. This allows you to validate the database version in each update.

API system

The API system determines the business logic of all calls, and manages all data and tasks. This system is written in Java.

DNS resolution system

The DNS resolution system consists of BIND and Agent. Agent receives and processes task information passed from the API system, parses the tasks into commands, and then sends the commands to the BIND system.

2.12.2 Maintenance

2.12.2.1 View running logs

During operation and maintenance, if you need to view logs to troubleshoot errors, you can query logs stored in certain locations in different systems.

API logs: The running logs are stored in `/home/admin/gdns/logs/`. You can query the logs as needed.

Agent logs: The running logs are stored in `/var/log/dns/`. Each log contains log entries of a day.

BIND logs: The running logs are stored in `/var/named/chroot/var/log/` on the DNS server.

2.12.2.2 Enable and disable a service

To restart the API service, log on to the API server as an admin and run the `/home/admin/gdns/bin/appctl.sh restart` script. To ensure the continuity of services, we recommend that you do not restart the API services on all servers at the same time. This script takes the following command-line arguments: **start**, **stop**, and **restart**.

The DNS service provides services using the anycast IP address. Therefore, to restart the DNS service, run the `service ospfd stop` command to remove the corresponding OSPF routes to prevent requests from being routed to the DNS server, and then run the `service named stop` command to stop the DNS service.

To start the DNS service, you must first start the DNS service and then enable OSPF by running the following commands: `service named start` and `service ospfd start`.

To enable the Agent service, run the `/usr/local/AgentService/agent -s start` command. If you receive a message that the PID file already exists, delete the `/var/dns/dns.pid` file and run the command again.

To disable the Agent service, run the `/usr/local/AgentService/agent -s stop` command.

2.12.2.3 Data backup

If you need to back up data before updating the service, copy the `/var/named/` and `/etc/named/` directories to a backup location. When you need to restore your data, copy the backup back to the original directories. Do not trigger automatic update during the data restoration process. Otherwise, data inconsistency may occur.

2.12.3 DNS API

2.12.3.1 Manage the API system

You can manage the API system using Apsara Infrastructure Management Framework. Click **Server Operations** in the Apsara Infrastructure Management Framework console to quickly log on to the server of the API system.

Context

To determine whether the service roles are correct, follow these steps:

Procedure

1. In the Apsara Infrastructure Management Framework console, check whether the API is in its final status.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Choose **Tasks > Deployment Summary** to go to the **Deployment Summary** page.
 - c) Click **Deployment Details**.
 - d) On the **Deployment Details** page, locate the `dnsProduct` project.
 - e) Find the `dnsServerRole#` service role, and click **Details** in the Deployment Progress column to check whether the role is in its final status.

If a green check mark is displayed after `dnsServerRole#`, then `dnsServerRole#` is in its final status.
2. Obtain the IP addresses of the servers where the API services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Choose **Operations > Cluster Operations** to locate the specified cluster.
 - c) Click a cluster URL to go to the **Cluster Dashboard** page.
 - d) On the **Cluster Dashboard** page, select **Cluster Operation and Maintenance Center > in the drop-down Operations Menu**.
 - e) On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of servers where they API services are deployed.
3. Log on to the DNS API server, and run the `curl http://localhost/checkpreload.htm` command to check whether success is returned.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Choose **Operations > Server Operations** to locate the specified server.
 - c) Select a server and click **Terminal** to log on to the server.
 - d) Run the `curl http://localhost/checkpreload.htm` command to verify the server where the API services are deployed to confirm whether success is returned.

2.12.3.2 Troubleshooting

Procedure

1. View logs stored in `/home/admin/gdns/logs/`.
2. Check whether the API is working normally. If an error occurs with the API and the operation crashes when you call the API, check the log to troubleshoot errors.
3. If the API is working, but its functions do not work as expected, check the `application.log` file.

2.12.4 DNS system

2.12.4.1 Check whether the service role is correct

Procedure

1. In the console of Apsara Infrastructure Management Framework, check whether the Apsara Stack DNS system is in its final status.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Choose **Tasks > Deployment Summary** to go to the **Deployment Summary** page.
 - c) Click **Deployment Details**.
 - d) Click **Deployment Details**, and locate `bindServerRole#`.
 - e) Find the `bindServerRole#` service role, and click **Details** in the **Deployment Progress** column to check whether the role is in its final status.
2. Obtain IP addresses of the servers that are deployed with DNS services.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Choose **Operations > Cluster Operations** to locate the specified cluster.
 - c) Click a cluster URL to go to the Cluster Dashboard page.
 - d) On the Cluster Dashboard page, select **Cluster Operation and Maintenance Center > in the drop-down Operations Menu**.
 - e) On the Cluster Operation and Maintenance Center page, view and obtain all server IP addresses that `bindServerRole#` provides services for.
3. Log on to the DNS server, and check whether the status code returned by `python /bind/hello/check_health.py|echo $?` is 0.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Choose **Operations > Server Operations** to locate the specified server.
 - c) Select a server and click **Terminal** to log on to the server.
 - d) Run the `python /bind/hello/check_health.py|echo $?` command to verify the server that deploys `bindServerRole#` and determine whether the returned status code is 0.

2.12.4.2 Troubleshooting

Procedure

1. Check the BIND running log in `/var/named/chroot/var/log/`, and determine whether errors have occurred.

2. Check the Agent running log in `/var/log/dns/`, and determine whether errors have occurred.
3. Run the `named-checkconf` command to check whether errors have occurred in the configuration file.

2.12.4.3 Errors and exceptions

Error: exit code 1

Run the health check script to locate the cause of this error.

Common causes include:

- Named is not working.
- Agent is not working.
- Ospf is not working, or anycast and public IP address cannot be advertised because of a network information retrieval error.
- An error occurred while running the task.

2.12.5 Log analysis

Query log entries by RequestId

After you send a request, you will receive a response that contains the RequestId. The RequestId can be used in the following scenarios:

1. Query tasks in the database to obtain the RequestId.
2. Use the RequestId to query execution result and error message for the current request in the API system log.
3. Use the RequestId to query results stored in the log of `bindServerRole#`, and verify the results with information from multiple systems.

2.12.6 View and process data

Context

You can view task records and execution results.

Procedure

1. Log on to the API server to view database connection information.
2. Run the `use genesisdns` command of MySQL to log on to the database, and then run the `select * from task` statement to query all tasks and their status and progress.

2.13 API Gateway

2.13.1 API Gateway introduction

This topic describes Apsara Stack API Gateway and the features of its modules.

API Gateway console

The API Gateway console is used to configure and manage your APIs and related policies. With the API management system, you can query, update, edit, and delete APIs. You can also create , associate, disassociate, and delete API management policies. API Gateway also provides a full range of API lifecycle management functions, including creating, testing, publishing, and unpublishing APIs. It improves API management and iteration efficiency. All your data will eventually be used as the API metadata for API Gateway.

API Gateway

API Gateway is a complete API hosting service. It helps you use APIs to provide capabilities , services, and data to your partners. API Gateway is initialized based on the API metadata generated by the API management system, and ultimately acts as the agent to send API requests . API Gateway provides a range of mechanisms to enhance security and reduce risks arising from APIs. These mechanisms include attack prevention, replay prevention, request encryption, identity authentication, permission management, and throttling.

2.13.2 Routine maintenance

2.13.2.1 View operation logs

During O&M, if you need to view logs to troubleshoot errors, you can query logs stored in relevant locations on different systems.

API Gateway OpenAPI logs: The operation log files are stored in the `/alidata/www/logs/java/cloudapi-openapi/` directory. You can query the files as required.

API Gateway logs: The operation log files are stored in the `/alidata/logs/` directory. Each log contains log entries of a day. You can query the files as required.

2.13.2.2 Enable and disable API services

You can log on to API servers as an administrator and run the `sudo /etc/init.d/jetty restart` command to restart API services. To ensure the continuity of services, we recommend that you do not restart the API services on all servers at the same time.

You can run the `sudo sh /home/admin/stop.sh` command to stop API services, and run the `sudo sh /home/admin/start.sh` command to start API services.

2.13.3 API Gateway console O&M

2.13.3.1 System O&M

You can use Apsara Infrastructure Management Framework to operate and maintain the API Gateway console. To log on to the server in which the API Gateway console resides, choose **Operations > Server Operations** in Apsara Infrastructure Management Framework.

Procedure

1. In the Apsara Infrastructure Management Framework console, check whether API Gateway reaches the final status.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose **Tasks > Deployment Summary**.
 - c) On the Deployment Summary page that appears, click **Deployment Details**.
 - d) On the **Deployment Details** page, locate the apigateway project.
 - e) Click **Details** in the Deployment Progress column corresponding to the apigateway project. Check whether the server roles of ApigatwayLite# are in the final status. If a green check mark is displayed after dnsServerRole#, dnsServerRole# is in the final status.
2. Obtain the IP addresses of servers where the API services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Click the C tab in the left-side navigation pane.
 - c) Select apigateway from the project drop-down list.
 - d) Hover over the vertical dots next to the cluster, and choose Dashboard from the shortcut menu. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**.
 - e) On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of servers where the API services are deployed.
3. Log on to the DNS API server. Run the `curl http://localhost:18080/cloudapi-openapi/check_health` command, and check whether a successful command output is displayed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Locate the APIgatewayOpenapi# servers.

- c) Click Terminal in the Actions column corresponding to one of the servers to log on to the server.
- d) Run the `curl http://localhost:18080/cloudapi-openapi/check_health` command on servers where the API services are deployed, and check whether the command output is Ok.

2.13.3.2 Troubleshooting

Procedure

1. View relevant logs in the `/usr/share/jetty/logs/` directory.
2. Check whether the system is operating normally. If the system is not operating normally (that is, it quits soon after startup), check the logs to troubleshoot errors.
3. If the system is operating normally but does not function properly, view `stderrout.log`.

2.13.4 API Gateway O&M

2.13.4.1 System O&M

You can use Apsara Infrastructure Management Framework to operate and maintain API Gateway. To log on to the server in which the API Gateway console resides, choose Operations > Server Operations in the Apsara Infrastructure Management Framework console.

Procedure

1. In the Apsara Infrastructure Management Framework console, check whether API Gateway reaches the final status.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose **Tasks > Deployment Summary**.
 - c) On the Deployment Summary page that appears, click **Deployment Details**.
 - d) On the **Deployment Details** page, locate the apigateway project.
 - e) Click **Details** in the Deployment Progress column corresponding to the apigateway project. Check whether the server roles of ApigatwayLite# are in the final status. If a green check mark is displayed after dnsServerRole#, dnsServerRole# is in the final status.
2. Obtain the IP addresses of servers where the API services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Click the C tab in the left-side navigation pane.
 - c) Select apigateway from the project drop-down list.

- d) Hover over the vertical dots next to the cluster, and choose Dashboard from the shortcut menu. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**.
 - e) On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of servers where the API services are deployed.
3. Log on to the DNS API server. Run the `curl http://localhost/gateway_status` command, and check whether a successful command output is displayed.
- a) Log on to the Apsara Infrastructure Management Framework console.
 - b) Locate the `ecsapigatewaylitetag#` servers.
 - c) Click Terminal in the Actions column corresponding to one of the servers to log on to the server.
 - a) Run the `curl http://localhost/gateway_status` command on servers where the API services are deployed, and check whether the command output is `I'm fine, thank you, and you?`.

2.13.4.2 Troubleshooting

Procedure

1. View relevant logs in the `/alidata/logs/` directory.
2. Check whether the system is operating normally. If the system is not operating normally (that is, it quits soon after startup), check the logs to troubleshoot errors.
3. If the system is operating normally but does not function properly, view `system.log`.

2.13.5 Log analysis

You can perform log analysis based on the ID of an individual API request.

After you send a request, you will receive a response that contains the request ID from API Gateway.

You can use the request ID to perform the following operations:

- All API Gateway logs are uploaded to Log Service, where you can view the request ID.
- You can use the request ID to query the response to or error message for the current request in the API system logs.