Alibaba Cloud Apsara Stack Enterprise

Operations and Maintenance Guide -Cloud Essentials and Security

Version: 1911, Internal: V3.10.0

Issue: 20200317



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, consequent

ial, 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 intellectu al 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

Document conventions

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

Style	Description	Example
{} or {a b}	This format is used for a required value, where only one item can be selected.	<pre>switch {active stand}</pre>

Contents

Legal disclaimerI
Document conventionsI
1 Operations of basic platforms
1.1 Apsara Stack Operations (ASO)
1.1.1 Apsara Stack Operations overview
1.1.2 Log on to Apsara Stack Operations
1.1.3 Web page introduction
1.1.4 Operations and maintenance dashboard
1.1.5 Alert Monitoring
1.1.5.1 Overview
1.1.5.2 Alert events
1.1.5.3 Alert history 12
1.1.5.4 Alert configuration 12
1.1.5.4.1 Alert contacts 12
1.1.5.4.2 Alert contact groups13
1.1.5.4.3 Static parameter settings14
1.1.5.5 Alert overview15
1.1.5.6 Alert subscription and push16
1.1.5.7 Alert masking18
1.1.5.7.1 Add a masking rule18
1.1.5.7.2 Remove the masking21
1.1.6 Physical servers22
1.1.6.1 View the physical server information
1.1.6.2 Add a physical server27
1.1.6.3 Modify a physical server28
1.1.6.4 Export the physical server information
1.1.6.5 Delete a physical server
1.1.7 Inventory Management33
1.1.7.1 View the ECS inventory
1.1.7.2 View the SLB inventory
1.1.7.3 View the RDS inventory 43
1.1.7.4 View the OSS inventory
1.1.7.5 View the Tablestore inventory45
1.1.7.6 View the Log Service inventory
1.1.7.7 View the EBS inventory
1.1.7.8 View the NAS inventory48
1.1.7.9 View the HDFS inventory
1.1.8 Products
1.1.8.1 Product list
1.1.8.2 ISV access configurations

-1.1.5.2.1 Compare the 15V access information	
1.1.8.2.2 Modify the ISV access information	
1.1.8.2.3 Delete the ISV access information	
1.1.9 ITIL Management	
1.1.9.1 Overview	
1.1.9.2 Dashboard	
1.1.9.3 Services	
1.1.9.3.1 Basic functions	
1.1.9.3.1.1 Overview	
1.1.9.3.1.2 Manage requests	
1.1.9.3.1.3 Manage tasks	
1.1.9.3.2 Manage incidents	
1.1.9.3.2.1 Create an incident request	
1.1.9.3.2.2 Manage incident requests	
1.1.9.3.2.3 Manage incident tasks	
1.1.9.3.3 Manage problems	63
1.1.9.3.3.1 Create a problem request	63
1.1.9.3.3.2 Manage problem requests	64
1 1 9 3 3 3 Manage problem tasks	66
1 1 9 4 Version control	68
1 1 9 5 Configure process templates	69
1 1 9 6 Configure CAB or ECAB	
1.1.10 Configurations	73
11110 Comgutations	
1.1.10.1 Overview	73
1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product	73
1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 73
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration value of a modified configuration 	73 73 ntion 74
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration item	73 73 ntion 74 75
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 73 ation 74 75 75
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76 76
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76 76 77 77
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 attion 74 75 75 76 76 76 76 78 78
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76 76 76 78 78 78
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ition 74 75 75 76 76 76 76 78 78 78 78
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76 76 76 78 78 78
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ition 74 75 75 76 76 76 76 76 78 78 78 78 78 80 80 80 80
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76 76 76 78 78 78 78 80 80 80 81 81
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configuration	73 ntion 74 75 75 76 76 76 76 76 76 78 78 78 78 78 80 80 81 81
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product	73 ntion 74 74 75 76 76 76 76 76 76 76 78 78 78 78 78 78 78 78 78
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product 1.1.10.3 Restore the configuration value of a modified configurative item	73 73 ation 74 75 75 76 76 76 76 76 76 76 78
 1.1.10.1 Overview 1.1.10.2 Modify a configuration item of a product	73 ntion 73 74 75 75 76 76 76 76 76 76 76 78
 1.1.10.1 Overview	73 73 ation 74 75 75 76 76 76 76 76 76 78

1.1.12.2.5 Delete a subview	86
1.1.12.2.6 Delete a view	.87
1.1.12.3 Resource management	87
1.1.12.3.1 Network elements	87
1.1.12.3.1.1 Device management	88
1.1.12.3.1.2 Modify the device password	.92
1.1.12.3.1.3 Configuration comparison	93
1.1.12.3.2 Service Load Balancers	.94
1.1.12.3.2.1 View the cluster monitoring information	94
1.1.12.3.2.2 View the instance monitoring information	95
1.1.12.3.3 Collect IP addresses	95
1.1.12.3.4 IP address ranges	96
1.1.12.3.4.1 Import the planning file	96
1.1.12.3.4.2 Manually add the IP address pool information	97
1.1.12.3.4.3 Modify the IP address pool information	97
1.1.12.3.4.4 Export the IP address pool information	.98
1.1.12.3.4.5 Delete the IP address pool information	98
1.1.12.4 Alert management	98
1.1.12.4.1 View and process current alerts	98
1.1.12.4.2 View history alerts	99
1.1.12.4.3 Add a trap	99
1.1.12.4.4 View a trap	02
1.1.12.5 Network reconfiguration 1	103
1.1.12.5.1 Physical network integration1	103
1.1.12.5.2 ASW scale-up 1	105
1.1.12.6 Fault check	108
1.1.12.6.1 IP address conflict check1	108
1.1.12.6.2 Leased line discovery 1	108
1.1.12.6.3 Network inspection	10
1.1.12.6.4 Configuration baseline audit1	12
1.1.13 Full Stack Monitor 1	13
1.1.13.1 SLA	13
1.1.13.1.1 View the current state of a cloud product1	13
1.1.13.1.2 View the history data of a cloud product	13
1.1.13.1.3 View the availability of an instance	14
1.1.13.1.4 View the availability of a product1	14
1.1.13.2 Operations full link logs1	115
1.1.13.3 Correlation diagnosis and alarm1	16
1.1.13.3.1 Full stack correlation alert 1	16
1.1.13.3.2 Server1	17
1.1.13.3.3 Network device1	20
1.1.13.3.4 ECS1	20
1.1.13.3.5 RDS	22
1.1.13.3.6 SLB1	23
1.1.13.3.7 VPC 1	24

1.1.14 Storage Operation Center	. 125
1.1.14.1 Pangu	125
1.1.14.1.1 Pangu grail	125
1.1.14.1.2 Cluster information	127
1.1.14.1.3 Node information	. 130
1.1.14.1.4 Pangu operation	131
1.1.14.2 EBS	. 132
1.1.14.2.1 IO HANG fault analysis	. 132
1.1.14.2.2 Slow IO analysis	. 133
1.1.14.2.3 Inventory settings	135
1.1.15 Task Management	. 136
1.1.15.1 Overview	136
1.1.15.2 View the task overview	137
1.1.15.3 Create a task	. 138
1.1.15.4 View the execution status of a task	. 142
1.1.15.5 Start a task	143
1.1.15.6 Delete a task	144
1.1.15.7 Process tasks to be intervened	144
1.1.16 Log Management	145
1.1.16.1 Log configurations	. 145
1.1.16.1.1 Clear	. 145
1.1.16.1.1.1 Configure parameters for automatic log clearance	145
1.1.16.1.1.2 Configure the manual log clearance time	146
1111011112 Compare the manual log clearance inferior	
1.1.16.1.2 Project	147
1.1.16.1.2 Project 1.1.16.1.2.1 Add a project	147 147
1.1.16.1.2 Project1.1.16.1.2.1 Add a project1.1.16.1.2.2 Delete a project	147 147 . 148
1.1.16.1.2 Project 1.1.16.1.2.1 Add a project 1.1.16.1.2.2 Delete a project 1.1.16.1.3 Agent	147 147 . 148 148
1.1.16.1.2 Project 1.1.16.1.2.1 Add a project 1.1.16.1.2.2 Delete a project 1.1.16.1.3 Agent 1.1.16.1.3.1 Add an agent	147 147 . 148 148 148
1.1.16.1.2 Project 1.1.16.1.2.1 Add a project 1.1.16.1.2.2 Delete a project 1.1.16.1.3 Agent 1.1.16.1.3.1 Add an agent 1.1.16.1.3.2 Modify an agent	147 147 . 148 148 148 148 151
1.1.16.1.2 Project 1.1.16.1.2.1 Add a project 1.1.16.1.2.2 Delete a project 1.1.16.1.3 Agent 1.1.16.1.3.1 Add an agent 1.1.16.1.3.2 Modify an agent 1.1.16.1.3.3 Delete an agent	147 147 . 148 148 148 148 151 151
1.1.16.1.2 Project	147 147 148 148 148 148 151 151 151
1.1.16.1.2 Project	147 147 148 148 148 148 151 151 152 152
1.1.16.1.2 Project	147 147 148 148 148 151 151 151 152 152 152
 1.1.16.1.2 Project 1.1.16.1.2.1 Add a project 1.1.16.1.2.2 Delete a project 1.1.16.1.3 Agent 1.1.16.1.3.1 Add an agent 1.1.16.1.3.2 Modify an agent 1.1.16.1.3.3 Delete an agent 1.1.16.1.4 Bucket management 1.1.16.1.4.1 OSS configurations 1.1.16.1.4.2 NAS configurations 1.1.16.1.4.3 FTP configurations 	147 147 148 148 148 151 151 152 152 152 152 153
1.1.16.1.1.2 Project	147 147 148 148 148 151 151 152 152 152 153 154
 1.1.16.1.12 Configure the manual tog creatance children in the configuration in the configuration is a configuration in the configuration in the configuration in the configuration in the configuration is a configuration in the configuration in the configuration in the configuration is a configuration in the configure the	147 147 148 148 148 151 151 152 152 152 152 153 154 154
1.1.16.1.1.2 Project	147 147 148 148 148 151 151 152 152 152 153 154 154 154
1.1.16.1.1.2 Project	147 147 148 148 148 148 151 151 152 152 152 152 153 154 154 154 154
1.1.16.1.2 Project. 1.1.16.1.2.1 Add a project. 1.1.16.1.2.2 Delete a project. 1.1.16.1.2.2 Delete a project. 1.1.16.1.3.4 Add an agent. 1.1.16.1.3.1 Add an agent. 1.1.16.1.3.2 Modify an agent. 1.1.16.1.3.3 Delete an agent. 1.1.16.1.4.4 Bucket management. 1.1.16.1.4.1 OSS configurations. 1.1.16.1.4.2 NAS configurations. 1.1.16.1.4.3 FTP configurations. 1.1.16.3 Log export. 1.1.16.3.1 Export logs. 1.1.16.4 Log clearance.	147 147 148 148 148 148 151 151 152 152 152 154 154 154 155 156
1.1.16.1.2 Project. 1.1.16.1.2 Project. 1.1.16.1.2.1 Add a project. 1.1.16.1.2.2 Delete a project. 1.1.16.1.3 Agent. 1.1.16.1.3 Agent. 1.1.16.1.3.1 Add an agent. 1.1.16.1.3.2 Modify an agent. 1.1.16.1.3.3 Delete an agent. 1.1.16.1.4.4 Bucket management. 1.1.16.1.4.1 OSS configurations. 1.1.16.1.4.2 NAS configurations. 1.1.16.1.4.3 FTP configurations. 1.1.16.3 Log export. 1.1.16.3 Log export. 1.1.16.3.1 Export logs. 1.1.16.4 Log clearance. 1.1.16.4.1 Containers.	147 147 148 148 148 148 151 151 152 152 152 154 154 154 155 156 156
1.1.16.1.1 Contingure the manual rog creatine content of the manual rog creating creatine content of the	147 147 147 148 148 148 151 151 152 152 152 154 154 154 155 156 e
1.1.16.1.1.2 Project.1.1.16.1.2 Project.1.1.16.1.2.1 Add a project.1.1.16.1.2.2 Delete a project.1.1.16.1.3 Agent.1.1.16.1.3 Agent.1.1.16.1.3.1 Add an agent.1.1.16.1.3.2 Modify an agent.1.1.16.1.3.3 Delete an agent.1.1.16.1.4.3 Delete an agent.1.1.16.1.4 Bucket management.1.1.16.1.4.1 OSS configurations.1.1.16.1.4.2 NAS configurations.1.1.16.1.4.3 FTP configurations.1.1.16.1.4.3 Log export.1.1.16.3 Log export.1.1.16.3.1 Export logs.1.1.16.4.1 Containers.1.1.16.4.11 Obtain the watermark information of one or mor containers.	147 147 148 148 148 151 151 152 152 152 153 154 154 154 155 156 e 156
1.1.101.1.2 Consigned the matching creating of the matching of	147 147 147 148 148 148 151 152 152 152 152 154 154 154 156 156 e 156 156
1.1.101.1.2 Consigned the matching electronic constraints1.1.101.1.2 Project	147 147 147 148 148 148 151 151 152 152 152 153 154 154 155 156 e 156 156 157

	1.1.16.4.1.5 Clear container logs	158
	1.1.16.4.1.6 View clear records	
	1.1.16.4.2 Servers	
	1.1.16.4.2.1 Obtain the watermark information of one or	more
	servers	159
	1.1.16.4.2.2 Add a log clearance rule	160
	1.1.16.4.2.3 Modify a log clearance rule	161
	1.1.16.4.2.4 Delete a log clearance rule	
	1.1.16.4.2.5 Clear server logs	161
	1.1.16.4.2.6 View clear records	
	1.1.16.4.3 Import clearance rules of containers or servers	162
	1.1.16.4.4 Export clearance rules of containers or servers	
	1.1.17 System Management	163
	1.1.17.1 Overview	163
	1.1.17.2 Department management	
	1.1.17.3 Role management	164
	1.1.17.4 Logon policy management	165
	1.1.17.5 User management	166
	1.1.17.6 Two factor authentication	169
	1.1.17.7 Application whitelist	
	1.1.17.8 Server password management	
	1.1.17.9 Operation logs	176
	1.1.17.10 View the authorization information	177
	1.1.17.11 Menu settings	
	1.1.17.11.1 Add a level-1 menu	179
	1.1.17.11.2 Add a submenu	181
	1.1.17.11.3 Hide a menu	
	1.1.17.11.4 Modify a menu	184
	1.1.17.11.5 Delete a menu	184
]	1.2 Apsara Stack Doctor (ASD)	185
	1.2.1 Apsara Stack Doctor introduction	185
	1.2.2 Log on to Apsara Stack Doctor	187
	1.2.3 ASA	189
	1.2.3.1 RPM Check	189
	1.2.3.2 Virtual IP Check	190
	1.2.3.3 Volume Check	191
	1.2.3.4 NTP Check	192
	1.2.3.5 IP Conflict Check	193
	1.2.3.6 DNS Check	194
	1.2.3.7 IP Details	195
	1.2.3.8 Quota Check	195
	1.2.3.9 Error Diagnostics	
	1.2.3.10 Versions	197
	1.2.4 Support tools	
	1.2.4.1 Diagnose with the OS tool	197

1.2.4.2 Use Support Tools	198
1.2.4.3 Update Support Tools	200
1.2.4.4 Diagnose with the inspection tool	201
1.2.4.5 Upload script files for EDAS diagnostics	202
1.2.4.6 EDAS diagnostics	203
1.2.5 Service Availability	204
1.2.5.1 View Service Availability	204
1.2.5.2 View Control Service Availability	205
1.2.6 Monitoring	207
1.2.6.1 View alert templates	207
1.2.6.2 View alert information	208
1.2.6.3 View the alert status	208
1.3 Operation Access Manager (OAM)	209
1.3.1 OAM introduction	209
1.3.2 Instructions	210
1.3.3 Quick start	211
1.3.3.1 Log on to OAM	211
1.3.3.2 Create a group	213
1.3.3.3 Add group members	213
1.3.3.4 Add group roles	214
1.3.3.5 Create a role	215
1.3.3.6 Add inherited roles to a role	215
1.3.3.7 Add resources to a role	215
1.3.3.8 Add authorized users to a role	216
1.3.4 Manage groups	218
1.3.4.1 Modify the group information	218
1.3.4.2 View group role details	218
1.3.4.3 Delete a group	219
1.3.4.4 View authorized groups	219
1.3.5 Manage roles	220
1.3.5.1 Search for roles	220
1.3.5.2 Modify the role information	220
1.3.5.3 View the role inheritance tree	220
1.3.5.4 Transfer roles	221
1.3.5.5 Delete a role.	221
1.3.5.6 View authorized roles	222
1.3.5.7 View all roles	222
1.3.6 Search for resources	222
1.3.7 view the personal information	223
1.3.8 Appendix	223 222
1.3.6.1 Default role of OAM	223
1.3.0.1.1 Detault tote of UAM.	∠∠ð nt
Framework	пі 201
1 3 8 1 3 Default roles of Wahann-rula	··· 224
1.0.0.1.0 Detault 10169 of Webapp-1016	440

1.3.8.1.4 Default roles of the workflow console	227
1.3.8.1.5 Default role of Tianjimon	227
1.3.8.2 Permission lists of operations platforms	228
1.3.8.2.1 Permission list of Apsara Infrastructure Managemen	t
Framework	228
1.3.8.2.2 Permission list of Webapp-rule	238
1.3.8.2.3 Permission list of the workflow console	238
1.3.8.2.4 Permission list of Tianjimon	239
1.4 Apsara Infrastructure Management Framework	239
1.4.1 Old version	239
1.4.1.1 What is Apsara Infrastructure Management Framework?	239
1.4.1.1.1 Overview	239
1.4.1.1.2 Basic concepts	
1.4.1.2 Log on to Apsara Infrastructure Management Framework	242
1.4.1.3 Web page introduction	244
1.4.1.3.1 Introduction on the home page	. 244
1.4.1.3.2 Introduction on the left-side navigation pane	247
1.4.1.4 Cluster operations	249
1.4.1.4.1 View cluster configurations	249
1.4.1.4.2 View the cluster dashboard	251
1.4.1.4.3 View the cluster operation and maintenance center	256
1.4.1.4.4 View the service final status	260
1.4.1.4.5 View operation logs	262
1.4.1.5 Service operations	. 263
1.4.1.5.1 View the service list	263
1.4.1.5.2 View the service instance dashboard	264
1.4.1.5.3 View the server role dashboard	. 267
1.4.1.6 Machine operations	
1.4.1.6.1 View the machine dashboard	270
1.4.1.7 Monitoring center	273
1.4.1.7.1 Modify an alert rule	. 273
1.4.1.7.2 View the status of a monitoring instance	273
1.4.1.7.3 View the alert status	274
1.4.1.7.4 View alert rules	274
1.4.1.7.5 View the alert history	275
1.4.1.8 Tasks and deployment summary	276
1.4.1.8.1 View rolling tasks	276
1.4.1.8.2 View running tasks	278
1.4.1.8.3 View history tasks	279
1.4.1.8.4 View the deployment summary	. 279
1.4.1.9 Reports	282
1.4.1.9.1 View reports	282
1.4.1.9.2 Add a report to favorites	284
1.4.1.10 Appendix	284
1.4.1.10.1 Project component info report	284

1.4.1.10.2 IP list	285
1.4.1.10.3 Machine info report	285
1.4.1.10.4 Rolling info report	287
1.4.1.10.5 Machine RMA approval pending list	289
1.4.1.10.6 Registration vars of services	291
1.4.1.10.7 Virtual machine mappings	291
1.4.1.10.8 Service inspector report	291
1.4.1.10.9 Resource application report	292
1.4.1.10.10 Statuses of project components	293
1.4.1.10.11 Relationship of service dependency	295
1.4.1.10.12 Check report of network topology	295
1.4.1.10.13 Clone report of machines	296
1.4.1.10.14 Auto healing/install approval pending report	297
1.4.1.10.15 Machine power on or off statuses of clusters	297
1.4.2 New version	298
1.4.2.1 What is Apsara Infrastructure Management Framework?	299
1.4.2.1.1 Introduction	299
1.4.2.1.2 Basic concepts	300
1.4.2.2 Log on to Apsara Infrastructure Management Framework	302
1.4.2.3 Homepage introduction	304
1.4.2.4 Project operations	307
1.4.2.5 Cluster operations	308
1.4.2.5.1 View the cluster list	308
1.4.2.5.2 View the cluster details	310
1.4.2.5.3 View operation logs	313
1.4.2.6 Service operations	314
1.4.2.6.1 View the service list	314
1.4.2.6.2 View the server role details	315
1.4.2.7 Machine operations	316
1.4.2.8 Monitoring center	317
1.4.2.8.1 View the monitoring instance status	317
1.4.2.8.2 View the alert status	318
1.4.2.8.3 View alert rules	319
1.4.2.8.4 View the alert history	320
1.4.2.9 View tasks	322
1.4.2.10 Reports	322
1.4.2.10.1 View reports	322
1.4.2.10.2 Add a report to favorites	324
1.4.2.11 Tools	324
1.4.2.11.1 Machine tools	324
1.4.2.11.2 IDC shutdown	326
1.4.2.12 Appendix	.326
1.4.2.12.1 Project component info report	326
1.4.2.12.2 IP list	327
1.4.2.12.3 Machine info report	328

1.4.2.12.4 Rolling info report	
1.4.2.12.5 Machine RMA approval pending list	
1.4.2.12.6 Registration vars of services	333
1.4.2.12.7 Virtual machine mappings	
1.4.2.12.8 Service inspector report	
1.4.2.12.9 Resource application report	
1.4.2.12.10 Statuses of project components	335
1.4.2.12.11 Relationship of service dependency	337
1.4.2.12.12 Check report of network topology	
1.4.2.12.13 Clone report of machines	338
1.4.2.12.14 Auto healing/install approval pending report	
1.4.2.12.15 Machine power on or off statuses of clusters	339
1.5 Network operations	341
1.5.1 What is Apsara Network Intelligence?	
1.5.2 Log on to the Apsara Network Intelligence console	
1.5.3 Query information	
1.5.4 Manage cloud service instances	
1.5.5 Tunnel VIP	345
1.5.5.1 Create a Layer-4 listener VIP	
1.5.5.2 Query the tunnel VIP of a cloud service	
1.5.6 Create a Direct Any Tunnel VIP	
1.5.7 Leased line connection	
1.5.7.1 Overview	347
1.5.7.2 Manage an access point	348
1.5.7.3 Manage an access device	
1.5.7.4 Establish a leased line connection	351
1.5.7.5 Create a VBR	355
1.5.7.6 Create router interfaces	358
1.5.7.7 Create a routing table	
1.5.8 Manage Business Foundation System flows in a VPC	
1.5.9 Configure reverse access to cloud services	
2 Elastic Compute Service (ECS)	366
2.1 ECS overview	366
2.2 Log on to the Apsara Stack Operations console	
2.3 ECS operations and maintenance	369
2.3.1 Overview	369
2.3.2 VM	369
2.3.2.1 Overview	369
2.3.2.2 Search for VMs	
2.3.2.3 Start a VM	
2.3.2.4 Stop a VM	370
2.3.2.5 Restart a VM	371
2.3.2.6 Cold migration	
2.3.2.7 Reset a disk	373
2.3.3 Disks	

2.3.3.1 Overview	
2.3.3.2 Search for disks	
2.3.3.3 View snapshots	
2.3.3.4 Mount a disk	
2.3.3.5 Detach a disk	375
2.3.3.6 Create a snapshot	
2.3.4 Snapshots	
2.3.4.1 Overview	
2.3.4.2 Search for snapshots	
2.3.4.3 Delete a snapshot	
2.3.4.4 Create an image	
2.3.5 Images	
2.3.5.1 Overview	
2.3.5.2 Search for images	379
2.3.6 Security groups	379
2.3.6.1 Overview	
2.3.6.2 Search for security groups	380
2.3.6.3 Add security group rules	
2.4 VM hot migration	
2.4.1 Overview	382
2.4.2 Limits on hot migration	
2.4.3 Complete hot migration on AG	383
2.4.4 Modify the position of the NC where the VM is located.	385
÷ -	
2.4.5 FAQ	
2.4.5 FAQ 2.5 Hot migration of disks	385
2.4.5 FAQ 2.5 Hot migration of disks 2.5.1 Overview	
2.4.5 FAQ 2.5 Hot migration of disks 2.5.1 Overview 2.5.2 Limits	
 2.4.5 FAQ 2.5 Hot migration of disks 2.5.1 Overview 2.5.2 Limits 2.5.3 O&M after hot migration 	
 2.4.5 FAQ 2.5 Hot migration of disks 2.5.1 Overview 2.5.2 Limits 2.5.3 O&M after hot migration 2.6 Upgrade solution 	
 2.4.5 FAQ 2.5 Hot migration of disks 2.5.1 Overview 2.5.2 Limits 2.5.3 O&M after hot migration 2.6 Upgrade solution 2.6.1 Overview 	385 388 388 388 388 389 389 389 389 389
 2.4.5 FAQ 2.5 Hot migration of disks 2.5.1 Overview 2.5.2 Limits 2.5.3 O&M after hot migration 2.6 Upgrade solution 2.6.1 Overview 2.6.2 Limits on GPU clusters 	385 388 388 388 388 389 389 389 389 389
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 388 389 389 389 389 389 389
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 388 389 389 389 389 389 389
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 389 389 389 389 389 389 389
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 389 389 389 389 389 389 390 390 390 390 391
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 389 389 389 389 389 389 390 390 390 390 391 402
 2.4.5 FAQ. 2.5 Hot migration of disks. 2.5.1 Overview. 2.5.2 Limits. 2.5.3 O&M after hot migration. 2.6 Upgrade solution. 2.6.1 Overview. 2.6.2 Limits on GPU clusters. 2.6.3 Limits on FPGA clusters. 2.7 Disk maintenance of an instance. 2.7.1 Overview. 2.7.2 Maintenance procedure. 2.7.3 Additional instructions. 2.8 Handle routine alarms. 	385 388 388 388 389 389 389 389 389 389 390 390 390 390 391 402 403
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 389 389 389 389 389 389 390 390 390 390 402 403 403
 2.4.5 FAQ 2.5 Hot migration of disks. 2.5.1 Overview. 2.5.2 Limits. 2.5.3 O&M after hot migration. 2.6 Upgrade solution. 2.6.1 Overview. 2.6.2 Limits on GPU clusters. 2.6.3 Limits on FPGA clusters. 2.7 Disk maintenance of an instance. 2.7.1 Overview. 2.7.2 Maintenance procedure. 2.7.3 Additional instructions. 2.8 Handle routine alarms. 2.8.1 Overview. 2.8.2 API proxy. 	385 388 388 388 389 389 389 389 389 389 389
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 389 389 389 389 389 389 390 390 390 390 390 402 403 403 403 405
 2.4.5 FAQ	385 388 388 388 389 389 389 389 389 389 390 390 390 390 390 391 402 403 403 405 405 405
 2.4.5 FAQ 2.5 Hot migration of disks	385 388 388 388 389 389 389 389 389 389 390 390 390 390 390 390 402 403 403 403 405 405 405 407
 2.4.5 FAQ. 2.5 Hot migration of disks. 2.5.1 Overview. 2.5.2 Limits. 2.5.3 O&M after hot migration. 2.6 Upgrade solution. 2.6.1 Overview. 2.6.2 Limits on GPU clusters. 2.6.3 Limits on FPGA clusters. 2.6.3 Limits on FPGA clusters. 2.7 Disk maintenance of an instance. 2.7.1 Overview. 2.7.2 Maintenance procedure. 2.7.3 Additional instructions. 2.8 Handle routine alarms. 2.8.1 Overview. 2.8.2 API proxy. 2.8.3 API Server. 2.8.4 RegionMaster. 2.8.5 RMS. 2.8.6 PYNC. 	385 388 388 388 389 389 389 389 389 389 390 390 390 390 390 390 390 402 403 403 403 405 405 405 406 407 408
 2.4.5 FAQ. 2.5 Hot migration of disks. 2.5.1 Overview. 2.5.2 Limits. 2.5.3 O&M after hot migration. 2.6 Upgrade solution. 2.6.1 Overview. 2.6.2 Limits on GPU clusters. 2.6.3 Limits on FPGA clusters. 2.7 Disk maintenance of an instance. 2.7.1 Overview. 2.7.2 Maintenance procedure. 2.7.3 Additional instructions. 2.8 Handle routine alarms. 2.8.1 Overview. 2.8.2 API proxy. 2.8.3 API Server. 2.8.4 RegionMaster. 2.8.5 RMS. 2.8.6 PYNC. 2.8.7 Zookeeper. 	385 388 388 388 389 389 389 389 389 389 390 390 390 390 390 390 402 403 403 403 405 405 405 405 406 407

2.8.9 Server groups	410
2.9 Inspection	411
2.9.1 Overview	411
2.9.2 Cluster basic health inspection	
2.9.2.1 Overview	411
2.9.2.2 Monitoring inspection	411
2.9.2.3 Inspection of basic software package versions	
2.9.2.4 Basic public resources inspection	411
2.9.3 Cluster resource inspection	
2.9.3.1 Overview	412
2.9.3.2 Cluster inventory inspection	412
2.9.3.3 VM inspection	
3 Auto Scaling (ESS)	416
3.1 Log on to the Apsara Stack Operations console	416
3.2 Product resources and services	
3.2.1 Application deployment	417
3.2.2 Troubleshooting	418
3.3 Inspection	419
3.3.1 Overview	419
3.3.2 Monitoring inspection	419
3.3.3 Basic software package version inspection	419
4 Object Storage Service (OSS)	
4.1 Log on to the Apsara Stack Operations console	
4.1 Log on to the Apsara Stack Operations console 4.2 OSS operations and maintenance	420 421
 4.1 Log on to the Apsara Stack Operations console 4.2 OSS operations and maintenance 4.2.1 User data 	420 421 421
 4.1 Log on to the Apsara Stack Operations console 4.2 OSS operations and maintenance 4.2.1 User data 4.2.1.1 Basic bucket information 	
 4.1 Log on to the Apsara Stack Operations console 4.2 OSS operations and maintenance	
 4.1 Log on to the Apsara Stack Operations console 4.2 OSS operations and maintenance	420 421 421 421 421 421 422 422
 4.1 Log on to the Apsara Stack Operations console 4.2 OSS operations and maintenance	420 421 421 421 421 421 422 423 423 425
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 421 421 422 423 423 425
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 421 422 423 423 425 425 425 426
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 421 422 423 423 425 425 425 426 427
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 421 422 423 423 425 425 425 425 426 427 428
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 423 425 425 425 425 426 427 428 432
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 421 422 423 423 425 425 425 425 426 427 428 432 433
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 423 423 425 425 425 425 425 426 427 428 432 433 433
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 422 423 425 425 425 425 426 427 428 432 433 433 433
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 423 425 425 425 425 425 426 427 428 432 433 433 433 433
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 423 423 425 425 425 425 425 426 427 428 432 433 433 433 433 433
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 422 423 425 425 425 425 426 427 428 432 433 433 433 433 433 434 434
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 423 425 425 425 425 425 425 426 427 428 432 433 433 433 433 433 433 434 434
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 422 423 425 425 425 425 426 427 428 432 433 433 433 433 433 434 434 434
 4.1 Log on to the Apsara Stack Operations console	420 421 421 421 422 423 423 425 425 425 425 425 426 427 428 432 433 433 433 433 433 433 434 434 434

5.1.4 Inspection center	440
5.1.4.1 Abnormal resource usage	440
5.1.5 Monitoring center	441
5.1.5.1 Cluster monitoring	441
5.1.5.2 Application monitoring	
5.1.5.3 Top requests	
5.1.5.4 Request log search	
5.1.6 System management	
5.1.6.1 Manage tasks	
5.1.6.2 View tasks	446
5.1.7 Platform audit	
5.1.7.1 Operation logs	446
5.2 Cluster environments	
5.3 System roles	448
5.4 Pre-partition a table	449
5.4.1 Pre-partitioning	
5.4.2 View partitions	
6 ApsaraDB for RDS	
6.1 Architecture	
6.1.1 System architecture	
6.1.1.1 Backup system	
6.1.1.2 Data migration system	
6.1.1.3 Monitoring system	
6.1.1.4 Control system	
6.1.1.5 Task scheduling system	454
6.2 Log on to the Apsara Stack Operations console	
6.3 Instance management	
6.4 Manage hosts	
6.5 Security maintenance	
6.5.1 Network security maintenance	459
6.5.2 Account password maintenance	
7 AnalyticDB for PostgreSOL	460
7.1 Overview	
7.2 Architecture	
7.3 Routine maintenance	
7.3.1 Check for data skew on a regular basis	
7.3.2 Execute VACUUM and ANALYZE statements	
7.4 Security maintenance	
7.4.1 Network security maintenance	
7.4.2 Account password maintenance	
8 KVStore for Redis	466
8 1 O&M tool	лаа Лаа
8.2 Architecture diagram	
8.2 Architecture magram	

8.3.1 Architecture	466
8.3.1.1 Backup system	467
8.3.1.2 Data migration system	467
8.3.1.3 Monitoring system	467
8.3.1.4 Control system	468
8.3.1.5 Task scheduling system	468
8.4 Log on to the Apsara Stack Operations console	468
8.5 Instance management	470
8.6 Host management	470
8.7 Security maintenance	471
8.7.1 Network security maintenance	471
8.7.2 Password maintenance	472
9 Apsara Stack Security	473
9.1 Log on to the Apsara Infrastructure Management Framework consol	e 473
9.2 Routine operations and maintenance of Server Guard	473
9.2.1 Check the service status	473
9.2.1.1 Check the client status	473
9.2.1.2 Check the status of Aegiserver	474
9.2.1.3 Check the Server Guard Update Service status	476
9.2.1.4 Check the Defender module status	476
9.2.2 Restart Server Guard	477
9.3 Routine operations and maintenance of Network Traffic Monitori	ng
System	479
9.3.1 Check the service status	479
9.3.1.1 Basic inspection	479
9.3.1.2 Advanced inspection	479
9.3.2 Common operations and maintenance	481
9.3.2.1 Restart the Network Traffic Monitoring System process	481
9.3.2.2 Uninstall Network Traffic Monitoring System	481
9.3.2.3 Disable TCP blocking	482
9.3.2.4 Enable TCPDump	482
9.4 Routine operations and maintenance of Anti-DDoS Service	483
9.4.1 Check the service status	483
9.4.1.1 Basic inspection	483
9.4.1.2 Advanced inspection	483
9.4.2 Common operations and maintenance	485
9.4.2.1 Restart Anti-DDoS Service	485
9.4.2.2 Troubleshoot common faults	486
9.5 Routine operations and maintenance of Threat Detection Service	491
9.5.1 Check the service status	491
9.5.1.1 Basic inspection	491
9.5.1.2 Advanced inspection	491
9.5.2 Restart TDS	493
9.6 Routine operations and maintenance of WAF	493
9.6.1 Check the service status	493

9.6.1.1 Basic inspection	493
9.6.1.2 Advanced inspection	494
9 7 Routine operations and maintenance of Sensitive Data Discovery	and
Protection	497
9.7.1 Check the service status	
9.7.1.1 Basic inspection	497
9.7.1.2 Advanced inspection: Check the status of the SddnSer	vice
corvico	498
9713 Advanced inspection. Check the status of the Sddpl	770 Nata
sorvico	Jata 500
9.7.1.4 Advanced inspection: Check the status of the SddnDrivi	300 1000
5.7.1.4 Advanced inspection. Check the status of the Sudpring	10ge 501
9715 Advanced inspection. Check the status of the Sddr	JUI
5.7.1.5 Auvanced inspection. Check the status of the Suup	503
0 7 9 Doctort SDDD	505
9.7.2 Restart SDDF	504
9.8 Routine operations and maintenance of Apsara Stack Security Cent	EI 500
9.8.1 Check service status	500
9.8.1.1 Basic inspection	
9.8.1.2 Advanced inspection	500
9.8.2 Restart the secure-console service	507
9.9 Routine operations and maintenance of secure-service	508
9.9.1 Check the service status	508
9.9.1.1 Basic inspection	
9.9.1.2 Advanced inspection: Check the secure-service status	508
9.9.1.3 Check the Dolphin service status	510
9.9.1.4 Check the data-sync service status	511
9.9.2 Restart secure-service	511
10 Apsara Stack DNS	514
10.1 Introduction to Apsara Stack DNS	514
10.2 Maintenance	514
10.2.1 View operational logs	514
10.2.2 Enable and disable a service	515
10.2.3 Data backup	515
10.3 DNS API	515
10.3.1 Manage the API system	515
10.3.2 Troubleshooting	518
10.4 DNS system	519
10.4.1 Check whether a server role is normal	519
10.4.2 Troubleshooting	521
10.4.3 Errors and exceptions	521
10.5 Log analysis	522
10.6 View and process data	522
-	

1 Operations of basic platforms

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 Operations & Maintenance (O&M) operations.

ASO has the following main functions:

· Operations and Maintenance Dashboard

The Operations and Maintenance Dashboard module displays the local version list, inventory overview, alert breakdown, and inventory curve of the cloud platform, which allows you to know the current usage of resources.

· Alert Monitoring

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

Resource Management

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

Inventory Management

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

Products

The Products module 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, reduce operation and maintenance costs, and finally enhance the user satisfaction.

Configurations

The Configurations module 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

The Offline Backup module is used to back up the key metadata of Apsara Stack.
The backed up metadata is used for the fast recovery of Apsara Stack faults.
NOC

The Network Operation Center (NOC) module provides the operations capabiliti es such as the visualization of network-wide monitoring, automated implementa tion, 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

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

- Storage Operation Center
 - The Storage Operation Center module provides the O&M operations of pangu and EBS.
- Task Management

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

Log Management

The Log Management module is used to access various business logs and allows you to search for, export, back up, and clear logs.

System Management

The System Management module consists of the user management, twofactor authentication, role management, department management, logon policy management, application whitelist, server password management, operation logs, authorization, and menu settings. As the module for centralize d management of accounts, roles, and permissions, the System Management module 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 access address in the format of http://region-id.aso.intranet-domainid.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

Log On	
<u>8</u>	Enter a user name
£	Enter the password
	Log On



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 log on to ASO.

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.

C-)	Apsara Stack Operation	cn-q	ingdao-env4b-d01 🗸				Unauthorized ⑦ English (US) 🗸 aliyuntest 🗸
88 (C)	5 perations and		Products		Inventory Overview	Alert Breakdown	1 2 3	4
G	Dashboard		Product	Version	sis and a second se		Remind: 30 Critical: 892 Major: 10	
@ W	6			20e12ff3-d1f7-4230-8e0d-3d1c9ad0e0 86j32477ee2-3b75-4dad-bdad-74eb6b 754760	SLB -			
କ				c8cbd0ef-4bbb-42a3-8734-bae444ad4 2dd			12856	
R			dnsProduct	11e29d30-9ff2-462a-a049-fc24c276383 4	RDS -			
v B	E			1d4aa38c-5847-4051-a83f-8cbcab8244 85	Une 201e 407a 2001e 507a 2001e			
٥			middleWare-staragent	87fe8e08-bf0c-4dca-a4bb-62febdd65c 42	Inventory Curve			
8 +			asto	4fc11ad8-971f-4e23-b727-7198054104 e8	80%			
			datahub	508704ea-ad12-4caf-93fc-d86e144b73 fb	60%			
ି ନ୍ତ				68231ccb-4673-40db-adaf-41b534f402 b8	40%			
				5a0d6ed6-aa22-47cb-a8c4-c0fe05bc0e 9a	0%, Jan 16, 2020 Jan 17, 2020 Jan 18, 1	2020	Jan 19, 2020	Jan 20, 2020
			dms-enterprise	a39f43a3-c835-4061-932d-909b7115b dd9	- SLS - SLB -	— ECS — RDS		

The description of each area is as follows.

Area		Description				
1	Authorization	Unauthorized : Click this button to go to the Authorization page.				
2	Help center	in the help center, you can view the alarm knowledge base and upload other documents related to operations.				
3	Language switching	English (US) : Select the language from the drop-down list to change the language of ASO.				
4	Information of the current logon user	aliyuntest Click this drop-down list to view the information of the current user, modify the password, and complete the logo settings and logon settings.				
5	Expand button	BEE: Move the pointer over this button to expand the left-side navigation pane.				
6	Left-side navigation pane	Click to select a specific Operations & Maintenance (O&M) operation.				

1.1.4 Operations and maintenance dashboard

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

Log on to Apsara Stack Operations. In the left-side navigation pane, click Operations and Maintenance.

The Dashboard page of Operations and Maintenance displays the current product version, inventory statistics, and alert statistics of the cloud platform. By viewing the dashboard, operations engineers can know the overall operating conditions of Apsara Stack products in time.



1.1.5 Alert Monitoring

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

1.1.5.1 Overview

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

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, select Alert Monitoring.

Enter the search content.	Search						
Basic							
Recovered	Total	Recovered	Total	Recovered	Total	Recovered	Total
0 ⊘	0 Q	0 📀	0 🗘	5,450 ⊘	5,454 <mark>0</mark>	1,449 📀	1,492 <mark>0</mark>
test							
Recovered	Total						
00	0 🗘						

- 3. Then, you can:
 - View the total number of alerts and the number of recovered alerts in the basic, critical, important, and minor monitoring metrics, and custom filters.



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

Search for alerts

Enter a keyword, such as cluster, product, service, severity, status, and monitoring metric name, in the search box at the top of the page and then click Search to search for the corresponding alert event.

Add a custom filter

Click **T**. On the displayed page, enter the filter name and then configure the

filter conditions.

After adding a custom filter, you can view the overview information that meets the filter conditions in Alert Monitoring.

Modify a custom filter

After adding a custom filter, you can click main as required to modify the filter

conditions and obtain the new filter results.

• Delete a custom filter

After adding a custom filter, you can click **market** as required to delete it if it is no

longer in use.

1.1.5.2 Alert events

The Alert Events module displays the information of all alerts generated by the system on different tabs. The alert information is aggregated by monitoring item

or product name. You can search for alerts based on filter conditions, such as monitoring metric type, product, service, severity, status, and time range when the alert is triggered, and then perform Operations & Maintenance (O&M) operations on the alerts.

Context

The Alert Events module displays the alert events on the following tabs:

- Hardware & System: Displays the alert information related to the hardware or system in the Apsara Stack environment.
- Base Modules: Displays the alert information related to the base products such as baseserviceAll, webappAll, middlewareAll, https-proxy, dns, dnsProduct, and minirds.
- Monitoring & Management: Displays the alert information related to the cloud monitoring and management products except the base modules and cloud products.
- Cloud Product: Displays the alert information related to the cloud products such as OSS, ECS, SLB, VPC, RDS, DataWorks, DTS, MaxCompute, yundun-advance, yundun-common, and Tablestore.
- Timeout Alert: Displays the information of all the timeout alerts.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Events.

Hardware & System Base Modules Monitoring & Management Cloud Product Timeout Alert									
By Monitoring Item V Monitoring Metric Type V	Product v	Service \lor Severity \lor	Status v	Start Date	e ~ End	Date 🛱	Enter the sear	ch content.	Search
±									
Monitoring Metric	Monitoring Type	Alert Details	Alerts	P1		P3	P4	P0	P5
postcheck_monitortianji_base-template	Event							746	
testimage_monitortianji_base-template	Event								
project_monitortianji_sub-template	Event							267	
ping_monitortianji_base-template	Event								
tianji_db_upgradetianji_base-template	Event								

- 3. Click the Hardware & System, Base Modules, Monitoring & Management, Cloud Product, or Timeout Alert tab and then you can:
 - Search for alerts

At the top of the page, you can search for alerts by Monitoring Metric Type, Product, Service, Severity, Status, Start Date, End Date, and search content.

- View alert sources
 - a. If the alert information is aggregated by Product Name on this page, click + at the left of the product name to display the monitoring metrics. If the alert information is aggregated by Monitoring Item on this page, skip this step.
 - b. Find the monitoring metric and severity to which the alerts you are about to view belong, and then click the number in the specific severity column.
 - c. Move the pointer over the alert source information in blue in the Alert Source column to view the alert source details.
- · View alert details
 - a. If the alert information is aggregated by Product Name on this page, click + at the left of the product name to display the monitoring metrics. If the alert information is aggregated by Monitoring Item on this page, skip this step.
 - b. Find the monitoring metric and severity to which the alerts you are about to view belong, and then click the number in the specific severity column.
 - c. Click the value in blue in the Alert Details column. On the displayed Alert Details page, you can view the alert information, such as the summary, reference, scope, and resolution.
- View the original alert information of an alert
 - a. If the alert information is aggregated by Product Name on this page, click + at the left of the product name to display the monitoring metrics. If the alert information is aggregated by Monitoring Item on this page, skip this step.
 - b. Find the monitoring metric and severity to which the alert you are about to view belongs, and then click the number in the specific severity column.
 - c. Click the number in blue in the Alerts column. The Alerts page appears.
 - d. Click Details in the Alert Information column to view the original alert information.
- Process an alert

Find the monitoring metric and severity to which the alert you are about to process belongs, and then click the number in the specific severity column.

Note:

If the alert information is aggregated by Product Name on this page, click + at the left of the product name to display the monitoring metrics.

 If an alert is being processed by operations engineers, click Actions > Process in the Actions column to set the alert status to In Process.

If multiple alerts are being processed by operations engineers, select these alerts and then click Process at the top of the page to process multiple alerts at a time.

- If the processing of an alert is finished, click Actions > Processed in the Actions column to set the alert status to Processed.

If the processing of multiple alerts is finished, select these alerts and then click Complete at the top of the page to complete multiple alerts at a time.

- To view the whole processing flow of an alert, click Actions > Alert Tracing in the Actions column.
- If an alert is considered as an incident when being processed, click Actions
 > Report to ITIL in the Actions column. Then, an incident request is created in the Information Technology Infrastructure Library (ITIL) to track the issue. For more information, see *Manage incidents*.

If multiple alerts are considered as incidents, select these alerts and then click Report to ITIL at the top of the page. Then, the system creates multiple incident requests in the ITIL to track the issues.

- View the recent monitoring data

Click Actions > Exploration in the Actions column at the right of an alert to view the trend chart of a recent monitoring metric of a product.

• Export a report

Click **mathematical at the top of the page to export the alert list.**

1.1.5.3 Alert history

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert History.
- 3. On the Alert History page, you can:
 - Search for alerts
 - At the top of the page, you can search for alerts by Monitoring Metric Type, Product, Service, Severity, Status, Start Date, End Date, and search content.
 - Export a list of alerts

Click **at the top of the page to export a list of history alerts.**

• View alert sources

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

• View alert details

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

 \cdot View the original alert information

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

1.1.5.4 Alert configuration

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

1.1.5.4.1 Alert contacts

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

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Configuration. You are on the Contacts tab by default.
- 3. Then, you can:
 - Search for alert contacts

Configure the corresponding product name, contact name, and phone number and then click Search. The alert contacts that meet the search conditions are displayed in the list.

• Add an alert contact

Click Add. On the displayed Add Contact page, complete the configurations and then click OK.

• Modify an alert contact

Find the alert contact to be modified and then click Modify in the Actions column. On the displayed Modify Contact page, modify the information and then click OK.

· Delete an alert contact

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

1.1.5.4.2 Alert contact groups

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Configuration.
- 3. Click the Contact Groups tab.

4. Then, you can:

• Search for an alert contact group

Enter the group name in the search box and then click Search. The alert contact group that meets the search condition is displayed in the list.

· Add an alert contact group

Click Add. On the displayed Add Contact Group page, enter the group name and select the contacts to add to the contact group. Then, click OK.

· Modify an alert contact group

Find the alert contact group to be modified and then click Modify in the Actions column. On the displayed Modify Contact Group page, modify the group name, description, contacts, and notification method. Then, click OK.

· Delete one or more alert contact groups

Find the alert contact group to be deleted and then click Delete in the Actions column. In the displayed dialog box, click OK.

Select multiple alert contact groups to be deleted and then click Delete All. In the displayed dialog box, click OK.

1.1.5.4.3 Static parameter settings

You can configure the static parameters related to alerts based on business needs. Currently, you can only configure the parameter related to timeout alerts.

Context

You cannot add new alert configurations in the current version. The system has a default parameter configuration for timeout alerts. You can modify the configuration as needed.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Configuration.
- 3. Click the Static Parameter Settings tab.
- 4. Optional: Enter the parameter name in the search box and then click Search to search for the static parameter configuration that meets the condition.
- 5. At the right of the static parameter to be modified, click Modify in the Actions column.
6. On the Modify Static Parameter page, modify the parameter name, parameter value, and description.

Configuration	Description
Parameter Name	Enter a parameter name related to the configuration
	•
Parameter Value	The default value is 5, indicating 5 days.
	After completing the configuration, the system
	displays the alert events that meet the condition
	according to this parameter value on the Timeout
	Alert tab of Alert Monitoring > Alert Events.
	For example, if the parameter value is 5, the system
	displays the alert events that exceed 5 days on the
	Timeout Alert tab of Alert Monitoring > Alert Events.
Description	Enter the description related to the configuration.

Modify Static Parameter \times
Parameter Name
Alarm Time Out
Parameter Code
ALARM_TIME_OUT
Parameter Value
5
Description
Alarms that exceed a specified number of days are classified as overdue, Unit: day

7. Then, click OK.

1.1.5.5 Alert overview

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

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose Alert Monitoring > Alert Overview. The Alert Overview page appears.



- The column chart displays the number of unsolved alerts in the last seven days
- The section at the bottom of the page displays the alert statistics in the current system by product.

1.1.5.6 Alert subscription and push

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Alert Monitoring > Subscribe/Push.

Subscribe	Pu	sh													
Add Chann	el														
Channel Name	Subscribed Language	Subscription Region	Filter Condition	Protocol	Push Interface Address	Port Number	URI	HTTP Method	Push Cycle (Minutes)	Pushed Alerts	Push Mode	Push Template	Custom JSON Fields	Push Switch	Actions
test	zh-CN	cn-qingdao- env4b-d01		http		80		POST			ALL	ANS			

- 3. On the Subscribe tab, click Add Channel.
- 4. On the Add Subscription page, complete the following configurations.

Configuration	Description
Channel Name	The name of the subscription channel.
Subscribed Language	Select Chinese or English.
Subscription Region	Select the region where the subscripti on is located.

Configuration	Description
Filter Condition	 Select a filter condition. Basic Critical Important Minor Custom filter
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 alerts pushed each time .
Push Mode	 Select one of the following methods: ALL: All of the alerts are pushed in each push cycle. TOP: Only alerts with high priority are pushed in each push cycle.
Push Template	 Select one of the following templates: ASO: The default template. ANS: Select this template to push alerts by DingTalk, SMS, or email. Currently, you can only configure one channel of this type. Note: A preset ANS template exists if the system already connects with the ANS product. To restore the initial configurations of the template with one click, click Reset.

Configuration	Description
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 alerts. If the switch is not turned on here, you can enable the push feature in the Push Switch column after configuring the subscription channel.

5. After completing the configurations, click OK.

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

6. 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 the ANS push channel, you must enter the mobile phone number, email address, and/or DingTalk to which alerts are pushed after clicking Test in the Actions column.

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

1.1.5.7 Alert masking

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

1.1.5.7.1 Add a masking rule

By adding a masking rule, you can mask alerts that you are not required to pay attention to.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Masking.
- 3. Click Add on the page.

Configuration	Description
Product	Optional. The product to which alerts to be masked belong.
Cluster	Optional. The cluster to which alerts to be masked belong.
Service	Optional. The name of the service to which alerts to be masked belong.
Alert Item	Optional. The alert name to be masked.
	Note: If the number of alerts is large, you may have to wait for a few minutes when selecting an alert item.
Monitoring Metric	Optional. The monitoring metric to which alerts to be masked belong.

4. On the Add page, complete the configurations to mask a certain type of alerts.

Configuration	Description
Alert Plan	Optional. The alert details of the alerts to be masked.
	Example:
	{"serverrole":"ecs-yaochi.ServiceTest#"," machine":"vm010012016074","level":"error"}
Severity	 Optional. Alerts are classified into the following levels: P0: indicates the cleared alerts, corresponding to alerts whose Alert Level is Restored in Monitoring > Alert History of Apsara Infrastructure Management Framework. P1: indicates the critical alerts, corresponding to alerts whose Alert Level is P1 in Monitoring > Alert History of Apsara Infrastructure Management Framework. P2: indicates the major alerts, corresponding to alerts whose Alert Level is P2 in Monitoring > Alert History of Apsara Infrastructure Management Framework. P2: indicates the major alerts, corresponding to alerts whose Alert Level is P2 in Monitoring > Alert History of Apsara Infrastructure Management Framework. P3: indicates the minor alerts, corresponding to alerts whose Alert Level is P3 in Monitoring > Alert History of Apsara Infrastructure Management Framework.
	 P4: indicates the remind alerts, corresponding to alerts whose Alert Level is P4 in Monitoring > Alert History of Apsara Infrastructure Management Framework
	 P5: indicates the system alerts, corresponding to alerts whose Alert Level is P5 in Monitoring > Alert History of Apsara Infrastructure Management Framework.

Add		×
Product		
Select		~
Cluster		
Select		~
Service		
Select		~
Alert Item		
Monitoring Metric		
Select		~
Alert Plan		
Enter data in JSON format.		
Severity		
Select		~
	ОК	Cancel

5. Then, click OK.

Result

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

In Alert Monitoring > Alert Events and Alert Monitoring > Alert History, you cannot view alerts that meet the conditions in the masking rule.

1.1.5.7.2 Remove the masking

You can remove the masking for masked alerts.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Masking.
- 3. Optional: Select a product, service, or alert item, and then click Search.

4. Find the alert masking rule and then click Delete in the Actions column to remove the masking.



5. In the displayed dialog box, click OK.

Result

After removing the masking, you can view alerts masked by the deleted masking rule in Alert Monitoring > Alert Events and Alert Monitoring > Alert History.

1.1.6 Physical servers

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

1.1.6.1 View the physical server information

You can view the physical server list and the details of physical servers in the system.

Product tab

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Resource Management > Physical Servers.

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

Resource Manag	Product	Server	Physical View of Devic	9				Physical Serve	s: 23) Servers with Aler	ts: (139) Alerts: <mark>512</mark>
Physical Servers		٩	Product 🗸	Enter a value Q						Ū
	+ cn-qingdao-env	4b-d01	Product	Hostname	Cluster	Group	IP Address	Host Server	Alerts	Operation
			middleWare-starag	ent vm010004029063	BasicCluster-A-2019102 8-eaea	StaragentInit		a56h11112.cloud.h12.a mtest72		
			middleWare-staraç	ent vm010004024188	BasicCluster-A-2019102 8-eaea	Staragent2_1_Controlle r		a56g13106.cloud.g14.a mtest72		
			middleWare-staraç	ent vm010004021246	BasicCluster-A-2019102 8-eaea	Staragent2Fs		a56g10104.cloud.g11.a mtest72		

- 3. On this tab, view the physical server information.
 - Expand the navigation tree on the left level by level to view the list of physical servers where a cluster of a product is located.
 - Enter the product name, cluster name, group name, or hostname in the search box in the upper-left corner to quickly locate the corresponding node.
 - In the search box on the right, search for physical servers by product, cluster, group, or hostname and view the details of a physical server.
 - Click Details in the Operation column at the right of a product to go to the Physical Server Details page. Then, view the basic information, monitoring details, and alert information of the physical server to which the product belongs.

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

The Monitoring Details tab displays the CPU usage, system load, disk usage , memory usage, network throughput, and disk I/O. When viewing the monitoring information, you can select the monitoring item in the upper-right corner of each monitoring graph and then select the time range to view the monitoring value in the specific time range.

In the upper-right corner of the CPU Usage, System Load, Disk Usage, Memory Usage, Network Throughput, and Disk IO sections, you can:

- Click the **mathematical button to view the monitoring graph in full screen.**
- Click the **w** button to download the monitoring graph to your local computer.
- Click the physical button to manually refresh the monitoring data.
- Click the production and then the button changes to green. The system

automatically refreshes the monitoring data every 10 seconds. To disable the auto refresh function, click the button again.

Server tab

1. Log on to Apsara Stack Operations.

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

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

- 3. Click the Server tab.
- 4. On this tab, view the physical server list.
 - Expand the navigation tree on the left by IDC > rack to view the physical server list in a rack.
 - Enter the rack name in the search box in the upper-left corner and then press Enter to search for and view the list of all the physical servers in the rack.

Product	Server Phys	sical View of Device				Physical Servers: 230 Servers	with Alerts: 139 Alerts: 612
Enter a va	alue Q	Hostname V Enter a value	٩				+ 🗉
amte:	st72 310	Hostname	Device Function	IP Address	SN	Alerts	Operation
	a56g10001.cloud.g10.amtest72	a56h11101.cloud.h12.amtest72	Worker				
	a56g10002.cloud.g10.amtest72 a56g10003.cloud.g10.amtest72	a56h11010.cloud.h11.amtest72	Worker				
	a56g10004.cloud.g10.amtest72 a56g10005.cloud.g10.amtest72	a58h11012.cloud.h11.amtest72	Worker				

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

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

The Monitoring Details tab displays the CPU usage, system load, disk usage, memory usage, network throughput, and disk I/O. When viewing the monitoring information, you can select the monitoring item in the upper-right corner of each monitoring graph and then select the time range to view the monitoring value in the specific time range.

In the upper-right corner of the CPU Usage, System Load, Disk Usage, Memory Usage, Network Throughput, and Disk IO sections, you can:

- Click the **main** button to view the monitoring graph in full screen.
- Click the **button to download the monitoring graph to your local**

computer.

- Click the **button to manually refresh the monitoring data**.
- Click the production and then the button changes to green. The system

automatically refreshes the monitoring data every 10 seconds. To disable the auto refresh function, click the button again.

Physical View of Device tab

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Resource Management > Physical Servers.

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

- 3. Click the Physical View of Device tab.
- 4. On the Physical View of Device tab, expand the navigation tree on the left by IDC
 > rack to view the corresponding rack information and the server information of a rack on the right.

Racks and servers use different colors to identify the alert condition of servers.

- Red indicates a critical alert.
- Orange indicates a moderate alert.
- Blue indicates the normal status.

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

Product Server Physical View of Device	Critical Alerts	Server	s 1
Rack ∨ Enter a value Q + (B) — 100%	G11	Normal Servers	🛛 🗖 Moderate Alerts 0 💼 Critical Alerts 1
antes/72		1	
G10		2	
		3	
		4	
		5	
		6 💶	
		7	
		8	
		9	
		10 🔳	
		11	•••
		12	
		13	
		14	•••
		15 🔳	
		16 🔳	
		17	
		18	

- 5. To view the details of a physical server, you can:
 - a) Find the physical server whose details you are about to view in the left-side navigation tree or rack graph on the right.
 - b) On the rack details page displayed on the right, click the color block of the server to view the basic information of the server.
 - c) Click Details in the Operation field of the basic information.



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

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

The Monitoring Details tab displays the CPU usage, system load, disk usage , memory usage, network throughput, and disk I/O. When viewing the monitoring information, you can select the monitoring item in the upper-right corner of each monitoring graph and then select the time range to view the monitoring value in the specific time range.

In the upper-right corner of the CPU Usage, System Load, Disk Usage, Memory Usage, Network Throughput, and Disk IO sections, you can:

- Click the **mail button to view the monitoring graph in full screen.**
- Click the work button to download the monitoring graph to your local

computer.

- Click the **button to manually refresh the monitoring data**.
- Click the production and then the button changes to green. The system

automatically refreshes the monitoring data every 10 seconds. To disable the auto refresh function, click the button again.

1.1.6.2 Add a physical server

Operations personnel can add the information of existing physical servers in the environment to Apsara Stack Operations (ASO).

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Resource Management > Physical Servers.

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

- 3. Click the Server tab or Physical View of Device tab.
- 4. Click the **button** in the upper-right corner of the Server tab or in the

upper-left corner of the rack graph on the Physical View of Device tab.

5. On the displayed Add Physical Server page, configure the physical server information.

Configuration	Description
Zone	The name of the region where the physical server to be added is located.
Data Center	The name of the data center where the physical server to be added is located.
Rack	The rack where the physical server to be added is located.
Room	The room where the physical server to be added is located.
Physical Server Name	The name of the physical server to be added.
Memory	The memory of the physical server to be added.
Disk Size	The disk size of the physical server to be added.
CPU Cores	The CPU cores of the physical server to be added.
Rack Group	The rack group to which the physical server to be added belongs.
Server Type	The server type of the physical server to be added.
Server Role	The function or purpose of the physical server to be added.
Serial Number	The serial number (SN) of the physical server to be added.
Operating System Template	The template used by the operating system of the physical server to be added.
IP Address	The IP address of the physical server to be added.

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

6. Click OK.

1.1.6.3 Modify a physical server

You can modify the physical server information in the system when the information is changed in the Apsara Stack environment.

Server tab

1. Log on to Apsara Stack Operations.

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

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

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

Server	Physical View of Device				Physical Servers: (23) Servers	with Alerts: (139) Alerts: (512)
٩	Hostname 🗸 Enter a value	٩				+ 🗉
	Hostname	Device Function	IP Address	SN	Alerts	Operation
	a56h11101.cloud.h12.amtest72	Worker				Details Modify Delete
	a58h11010.cloud.h11.amtest72	Worker				Details Modify Delete
	a56h11012.cloud.h11.amtest72	Worker				

6. On the displayed Modify Physical Server page, modify the physical server information.

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

7. Click OK.

Physical View of Device tab

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Resource Management > Physical Servers.

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

3. Click the Physical View of Device tab.

4. Expand the navigation tree on the left by IDC > rack to find the physical server to be modified.



In the upper-left corner, you can also select to search for the physical server to be modified by rack, hostname, IP address, device function, SN, or IDC.

- 5. On the rack details page on the right, click the color block of a server to view the basic information of the server.
- 6. Click Modify in the Operation field of the basic information.



7. On the displayed Modify Physical Server page, modify the physical server information.

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

8. Click OK.

1.1.6.4 Export the physical server information

You can export the information, namely the zone, hostname, disk size, CPU cores, data center information (data center, rack, room, and rack group), device function, serial number (SN), operating system template, IP address, and number of alerts, of all the physical servers in the system for offline review.

Product tab

1. Log on to Apsara Stack Operations.

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

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

Product	Server	Physical View of Device					Physical Server	s: 233) Servers with Alerts:	(139) Alerts: (512)
Enter a value	٩	Product V Enter :	i value Q						Ē
+ cn-qingdao-env	4b-d01	Product	Hostname	Cluster	Group	IP Address	Host Server	Alerts	Operation
		middleWare-staragent	vm010004029063	BasicCluster-A-2019102 8-eaea	StaragentInit		a56h11112.cloud.h12.a mtest72		
		middleWare-staragent	vm010004024188	BasicCluster-A-2019102 8-eaea	Staragent2_1_Controlle r		a58g13108.cloud.g14.a mtest72		

3. In the upper-right corner, click the putton to export the information of all

the physical servers from the dimension of products to your local computer.

Server tab or Physical View of Device tab

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Resource Management > Physical Servers.

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

- 3. Click the Server tab or the Physical View of Device tab.
- 4. Click the 🕕 button in the upper-right corner of the Server tab or at the top

of the Physical View of Device tab to export the information of all the physical servers from the dimension of servers to your local computer.

1.1.6.5 Delete a physical server

You can delete a physical server that does not require to be monitored based on business needs.

Server tab

1. Log on to Apsara Stack Operations.

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

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

- 3. Click the Server tab.
- 4. Optional: In the search box on the right, search for the physical server to be deleted by hostname, IP address, device function, or serial number (SN).
- 5. Find the physical server to be deleted and then click Delete in the Operation column.
- 6. In the displayed dialog box, click OK.

Physical View of Device tab

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Resource Management > Physical Servers.

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

- 3. Click the Physical View of Device tab.
- 4. Expand the navigation tree on the left by IDC > rack to find the physical server to be deleted.



In the upper-left corner, you can also select to search for the physical server to be deleted by rack, hostname, IP address, device function, SN, or IDC.

- 5. On the rack details page on the right, click the color block of a server to view the basic information of the server.
- 6. Click Delete in the Operation field of the basic information.



7. Click OK in the displayed dialog box.

1.1.7 Inventory Management

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

1.1.7.1 View the ECS inventory

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > ECS.

Note:	
You can click 👩	in the upper-right corner to configure the inventory

thresholds.



3. View the ECS inventory.

Where,

- The CPU Inventory Details(Core) and Memory Inventory Details(TB) sections display the used and available CPU (core) and memory (TB) of all ECS instance type families in the last five days.
- The ECS Instances Inventory Details section allows you to perform a paging query on the inventory details 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 the following table.

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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	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
	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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	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
	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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	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
	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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	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
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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	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
	ecs.c5. 4xlarge	None	16	32.0	8
	ecs.c5. 6xlarge	None	24	48.0	8
	ecs.c5. 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)
	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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.re5. 30xlarge	None	120	1980.0	15
	ecs.re5. 45xlarge	None	180	2970.0	15

1.1.7.2 View the SLB inventory

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > SLB.





3. View the SLB inventory.

Where,

- The section in the upper-left corner displays the used inventory and percentage of internal VIP and public VIP.
- The section in the upper-right corner displays the inbound and outbound network card traffic.
- The SLB Inventory Details section allows you to perform a paging query on the SLB inventory details by Type and Date.

1.1.7.3 View the RDS inventory

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > RDS.

Note: You can click in the upper-right corner to configure the inventory

thresholds of each engine.



3. View the RDS inventory.

Where,

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

1.1.7.4 View the OSS inventory

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

Procedure

1. Log on to Apsara Stack Operations.

et Inventory Detail:

E Search

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



3. View the OSS inventory.

Where,

- The Inventory Availability History(TB) section displays the available OSS inventory in the last five days.
- The Current Inventory Usage(TB) section displays the used OSS inventory and the corresponding percentage.
- The OSS Bucket Inventory Details section allows you to perform a paging query on the OSS inventory details by Date.

1.1.7.5 View the Tablestore inventory

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > OTS.





3. View the Tablestore inventory.

Where,

- The Inventory Availability History(TB) section displays the available Tablestore inventory in the last five days.
- The Current Inventory Usage(TB) section displays the used Tablestore inventory and the corresponding percentage.
- The OTS Bucket Inventory Details section allows you to perform a paging query on the Tablestore inventory details by Date.

1.1.7.6 View the Log Service inventory

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > SLS.

Note: You can click si in the upper-right corner to configure the inventory

thresholds and global quota.

sls-inner sls-public					ø		
History Inventory Records(TB) 70(TB)		Cur 3/	ment Quota Details(G) 11796875(TB)				
60(TB)		2	.9296875(TB)				
50(TB)			44140625(TB)				
40(TB)							
30(TB)		14	46484375(TB)				
20(TB)							
10(TB)			500(G)				
0 Jan 8, 2020	Jan 9, 2020	Jan 10, 2020					
Log Service Inventory Details							
Date Select a date mil Search							
Date	Region ID	Total(TB)	Used(TB)	Available(TB)	Usage (%)		

3. Click the sls-inner tab to view the inventory of base Log Service instances.

Where,

- The History Inventory Records(TB) section displays the available and total inventory of base Log Service instances in the last five days by using the line graph.
- The Current Quota Details(G) section displays the capacity consumed by each base Log Service instance.
- The Log Service Inventory Details section allows you to perform a paging query on the inventory details of base Log Service instances by Date.
- 4. Click the sls-public tab to view the inventory of Log Service instances you have applied for.
 - The Inventory Availability History(TB) section displays the available Log Service inventory in the last five days.
 - The Current Inventory Usage(TB) section displays the used Log Service inventory and the corresponding percentage.
 - The SLS Bucket Inventory Details section allows you to perform a paging query on the Log Service inventory details by Date.

1.1.7.7 View the EBS inventory

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

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose Inventory Management > EBS.

ECS-IO8-A-eb33	ECS-IO7-A-eb38	ECS-IO8-A-eb37					
Inventory Availability History(TB) 70(TB)				Current In	wentory Usage(TB)		
60(TB) 50(TB) 40(TB) 30(TB) 20(TB)						2% 1.06TB	
10(T8) 0 Jan 19, 2020 I EBS Bucket Inventory Deta	Jan 20, 2020 ils	Jan 21, 2020	Jan 22, 2020	Jan 23, 2020			
Select a date 🕅	Search						

3. If multiple ECS clusters exist in the environment, click the tab of the corresponding ECS cluster to view the EBS inventory.

Where,

- The Inventory Availability History(TB) section displays the available EBS inventory in the last five days.
- The Current Inventory Usage(TB) section displays the used EBS inventory and the corresponding percentage.
- The EBS Bucket Inventory Details section allows you to perform a paging query on the EBS inventory details by Date.

1.1.7.8 View the NAS inventory

By viewing the Network Attached Storage (NAS) inventory, you can know the current usage and surplus of NAS product resources to perform Operations & Maintenance (O&M) operations according to actual requirements.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > NAS.



3. View the NAS inventory.

Where,

- The Inventory Availability History(TB) section displays the available NAS inventory in the last five days.
- The Current Inventory Usage(TB) section displays the used NAS inventory and the corresponding percentage.
- The NAS Bucket Inventory Details section allows you to perform a paging query on the NAS inventory details by Date.

1.1.7.9 View the HDFS inventory

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Inventory Management > DFS.



3. View the HDFS inventory.

Where,

- The Inventory Availability History(TB) section displays the available HDFS inventory in the last five days.
- The Current Inventory Usage(TB) section displays the used HDFS inventory and the corresponding percentage.
- The DFS Bucket Inventory Details section allows you to perform a paging query on the HDFS inventory details by Date.

1.1.8 Products

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

1.1.8.1 Product list

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

Prerequisites

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

Context

After logging on to Apsara Stack Operations (ASO), you can view operations and maintenance icons of different products and different ISV icons on the Product List page based on your permissions. For example, a Tablestore operations engineer can only view the OTS Storage Operations and Maintenance System icon. Click OTS Storage Operations and Maintenance System to go to the Tablestore 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, the system can dynamically assign different permissions based on different roles.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Products > Product List.
- 3. On the Product List page, you can view operations and maintenance icons of different products and different ISV icons based on your permissions.

1.1.8.2 ISV access configurations

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

1.1.8.2.1 Configure the ISV access information

You can configure the ISV access information in the system based on business needs. Then, you can access the corresponding ISV page by clicking the icon on the Product List page.

- 1. Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Products > ISV Access Configurations.
- 3. Click Add on the page.
4. On the displayed Add page, configure the ISV access information.

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

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



5. Then, click Add.

Result

You can view the added ISV icon in Products > Product List. Click the icon and then you can be redirected to the corresponding page.

1.1.8.2.2 Modify the ISV access information

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Products > ISV Access Configurations.
- 3. Optional: In the search box on the page, enter the ISV name and then click Search. Fuzzy search is supported.
- 4. Find the ISV whose access information is to be modified. Click Modify in the Actions column.

l IS	ISV Access Configurations								
E	nter Name	Search Add	1						
N	ame	Кеу	lcon	Level-one Category	Level-two Category	Usage	Access Link	Description	Actions
te	st	test		ISV	ISV		http://example.com		Modify Delete

- 5. On the displayed Modify page, modify the name, key, icon, level-one category, level-two category, usage, access link, or description of the ISV.
- 6. Then, click Modify.

1.1.8.2.3 Delete the ISV access information

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Products > ISV Access Configurations.
- 3. Optional: In the search box on the page, enter the ISV name and then click Search. Fuzzy search is supported.
- 4. Find the ISV whose access information is to be deleted. Click Delete in the Actions column.
- 5. In the displayed dialog box, click OK.

Result

Then, the ISV information is not displayed in Products > Product List.

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, reduce operation and maintenance costs, and finally enhance the user satisfaction.

ITIL has the following functions:

 \cdot Dashboard

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

Services

The Services section 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

The Version Control section 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 Operations & 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 or ECAB.

1.1.9.2 Dashboard

The Dashboard module 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.
- 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.

The Services module 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.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the Request tab.
- 3. On the Request tab, you can:
 - · Create a request

Click ____ New and then select a request type. Complete the configurations

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*.

Requests are classified into three types based on the processing status.

- **P**: In processing, indicating the requests that are waiting to be processed.
- Closed, indicating the requests that have the whole process completed.

- **Recycle bin, indicating the recycled requests.**

• 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 are about 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, indicating tasks that are waiting to be processed by you.
- Task pool, indicating 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 under

• Solution: Processed by me, indicating the history tasks that have been processed by you. After you process the tasks under solution, they are displayed under solution.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the My Task tab.
- 3. On the My Task tab, you can:
 - 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 are about 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 alert information in Apsara Stack Operations (ASO). The alert module transfers the alert information to the ITIL module to generate the incident request based on the actual conditions, such as the alert level and the alert 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.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the Request tab.
- 3. Click _____ and then select Incident. Configure the incident request on the

Configuration	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.
Region	The region to which the request belongs.
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.

displayed page.

Configuration	Description	
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:	
	· Critical	
	• Major	
	• Minor	
	• Remind	
	· Cleared	
	• System	
Alarm Code	The alert ID.	
Summary	The summary of this request.	
Description	The detailed description about the request.	
Suggestion	Optional. The suggestion about the request processing.	

4. After completing the configurations, click Confirm.

1.1.9.3.2.2 Manage incident requests

After creating an incident request, you can change the priority of, comment on, suspend, resume, recycle, restore, and delete 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.
- 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 are about to manage, and then click Detail.

- 5. On the request details page, you can:
 - Change the priority

Click Change Priority at the top of the page. 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 on the incident request

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

• Suspend the incident request

Click Suspend at the top of the page. In the displayed dialog box, enter the Remarks. Perform this operation for incident requests that currently do not require to be processed.

• Resume the incident request

Click Resume at the top of the page. In the displayed dialog box, enter the Remarks. Perform this operation for suspended incident requests that require to be processed.

Recycle the incident request

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 (

Restore the incident request

Perform this operation for incident requests in the recycle bin (

Click Restore to restore the recycled incident request. After being restored,

Q

the incident request is in the in processing (

) list and restored to the

status before the request is recycled.

• Delete the incident request

Perform this operation for incident requests in the recycle bin (

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.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the My Task tab.
- 3. Click the [2] (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. In the task list, find the task that you are about to manage and then click Detail.

5. On the task details page, click Diagnose at the top of the page. In the displayed Diagnose dialog box, complete the configurations and then click OK.

Configuration	Description
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. If No is selected, the system goes to the Resolve phase. 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	The information about the task.

6. The system goes to the Resolve phase after the Diagnose phase. After processing the incident offline, click Resolve at the top of the page. In the displayed Resolve dialog box, configure the resolved date and the handling steps. Then, click OK.

The Resolve phase consists of the incident troubleshooting and solving. ITIL only tracks this step in a standardized way and processes the log records.

7. The system goes to the Confirm phase after the Resolve phase. This phase reviews the processing result of the incident. Then, click Confirm at the top of the page. 8. 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.
- It is not solved. Analyze again: 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.
- It is not solved. Process again: 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 causes 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.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the Request tab.

3. Click _____ and then select Problem. Configure the problem request on the

Configuration	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.	
Region	The region to which the request belongs.	
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:	
	• Critical	
	• Major • Minor	
	· Remind	
	ClearedSystem	
Alarm Code	The alert ID.	
Summary	The summary of this request.	
Description	The detailed description about the request.	
Suggestion	Optional. The suggestion about the request processing.	

displayed page.

4. After completing the configurations, click Confirm.

1.1.9.3.3.2 Manage problem requests

After creating a problem request, you can change the priority of, comment on, suspend, resume, recycle, restore, and delete 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.
- 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 are about to manage, and then click Detail.
- 5. On the request details page, you can:
 - Change the priority

Click Change Priority at the top of the page. 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 on the problem request

Click Comment at the top of the page. In the displayed dialog box, enter the comment for this problem request. Perform this operation for collaborative scenarios. For example, users can comment on the problem request to share

- the information with each other and guide each other when they process the same problem.
- Suspend the problem request

Click Suspend at the top of the page. In the displayed dialog box, enter the Remarks. Perform this operation for problem requests that currently do not require to be processed.

Resume the problem request

Click Resume at the top of the page. In the displayed dialog box, enter the Remarks. Perform this operation for suspended problem requests that require to be processed.

Recycle the problem request

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 (

Restore the problem request

Perform this operation for problem requests in the recycle bin (

Click Restore to restore the recycled problem request. After being restored, the problem request is in the in processing (

status before the request is recycled.

Delete the problem request

Perform this operation for problem requests in the recycle bin (

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.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the My Task tab.
- 3. Click the (My Task) button.



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. In the task list, find the task that you are about to manage and then click Detail.
- 5. On the task details page, click Diagnose at the top of the page. In the displayed Diagnose dialog box, complete the configurations and then click OK.

Configuration	Description
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.
	If No is selected, the system goes to the Resolve phase.
	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.

Configuration	Description
Remarks	The information about the task.

6. The system goes to the Resolve phase after the Diagnose phase. After processing the problem offline, click Resolve at the top of the page. In the displayed Resolve dialog box, configure the resolved date and the handling steps. Then, click OK.

The Resolve phase consists of the problem troubleshooting and solving. ITIL only tracks this step in a standardized way and processes the log records.

- 7. The system goes to the Confirm phase after the Resolve phase. This phase reviews the processing result of the problem. Then, click Confirm.
- 8. 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.
- It is not solved. Analyze again: 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.
- It is not solved. Process again: 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

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

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose ITIL Management > Version Control.

Select a node in the tree structure or enter a name in the search box and then click the search icon. The version and cluster information is displayed on the right.



Before the search, click not to synchronize the information to Apsara Stack

Operations (ASO).

1.1.9.5 Configure process templates

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

Log on to Apsara Stack Operations. In the left-side navigation pane, choose ITIL

Management > Process Template Configuration. 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
- \cdot 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:

ID 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.

Factorial states and a set of the process. A process usually ends with archiving.

Resolve indicates the phase. A phase is usually composed of roles with

specific functions.

• **Each** 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.



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

• 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.
- 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.
- 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.

• 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/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 or 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 topic, 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.
- 2. In the left-side navigation pane, choose ITIL Management > CAB/ECAB Configuration.
- 3. Click the CAB Configurations tab.
- 4. Select one or more users on the left and then click to add them to the

list on the right.

Users in the list on the right are the current CAB configuration.

Note:

- You can use the search box in the upper-left corner to search for users. Fuzzy search is supported.
- You can select one or more users on the right and then click

cancel the configuration for the selected users.

to

«

1.1.10 Configurations

1.1.10.1 Overview

The Configurations module 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.

1.1.10.2 Modify a configuration item of a product

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Configurations > Configuration Items.

- 3. 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.

Click Get in the Actions column to load the actual data from the product to your local computer.

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 as follows:

- a. Click Add in the upper-right corner.
- b. 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.
- c. Click OK.

Then, this configuration item is displayed in the list. You can search for and 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 in the upper-right corner.

Note:

To import configuration items as a file, we recommend that you export a file before the import and then complete the configurations based on the format in the exported file.

1.1.10.3 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.

- **1.** Log on to Apsara Stack Operations.
- 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.
- 5. Click OK in the displayed dialog box to restore the configuration value of the configuration item.

1.1.10.4 Manage kernel configurations

You can add, modify, or delete a kernel configuration.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Configurations > Kernel Configurations.
- 3. On the Kernel Configurations page, you can:
 - 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 to be modified. 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 to be deleted. Click Delete in the Actions column. In the displayed dialog box, click OK.

1.1.10.5 Scan configurations

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

Prerequisites

Before the scan, make sure that the following conditions are met:

- The configurations to be scanned are added in the kernel configurations list. For more information about how to add a kernel configuration, see *Manage kernel configurations*.
- The hostname or IP address of the host to be configured is obtained from Apsara Infrastructure Management Framework.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 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 Get.

1.1.11 Offline Backup

The Offline Backup module is used to back up the key metadata of Apsara Stack. Currently, you can only back up the pangu metadata. The backed up metadata is used for the fast recovery of Apsara Stack faults.

1.1.11.1 Service configuration

The Service Configuration module consists of the backup service configuration and product management.

1.1.11.1.1 Configure the backup server

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

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Service Configuration > Backup Service Configuration.

4. On the Backup Service Configuration page, click Modify in the Actions column at the right of the backup server to configure the backup server information.

Configuration	Description
Backup Server IP Address	The IP address of the backup server.
	The backup server must meet the following
	requirements:
	• The backup server is an independent physical server.
	$\cdot $ The backup server is managed and controlled by
	Apsara Infrastructure Management Framework.
	• The backup server has its network connected with other servers in Apsara Stack.
	\cdot Apsara Distributed File System cannot be
	deployed on the server, at least cannot be
	deployed on its disk that stores the backup metadata.
Backup Server	The storage path of backup files on the backup
Monitoring Path	server.
	The backup service detects new backup files by
	monitoring the specified folder on the backup server
	and determines whether the backup is successful
	by comparing the MD5 values of the backup file and
	the original file.
Backup Retention	The actual time (in days) that backup files are stored . The backup file that exceeds the time is to be deleted
	ucicicu.

1.1.11.1.2 Add a backup product

The Product Management module allows you to add the backup product information. In the current version, you can only back up the pangu metadata.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Service Configuration > Product Management.
- 4. Click Add in the upper-right corner.

5. In the displayed Add Product dialog box, add the product information based on the following table and then click OK.

Configuration	Description
Product	Enter pangu here because you are about to back up the pangu data.
Backup Items	Enter the information based on the pangu information of the cloud product to be backed up in the format of backup product name_pangu. For example, ecs_pangu.
Backup Script	The backup script name. For example, metadata_b ackup.py.
Retry Times	Generally, enter 3.

The added product is displayed on the Backup Service > Backup Configuration page.

6. Generally, you are required to add multiple backup items by completing the preceding steps.

Then, you can click Modify or Delete in the Actions column to modify or delete a backup item.

1.1.11.2 Backup service

The Backup Service module consists of the backup configuration, backup details, and service status.

1.1.11.2.1 Backup configuration

After adding a product backup item, you are required to configure the backup in Apsara Stack Operations (ASO).

Prerequisites

Make sure that a product backup item is added. For more information about how to add a product backup item, see *Add a backup product*.

Context

The backup item is the minimum unit of backup. You can back up the metadata of different pangus, such as ecs pangu, rds pangu, and ots pangu, according to different situations of Apsara Stack.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Backup Service > Backup Configuration.

The left part of the Backup Configuration page displays the current backup configurations in a hierarchical tree structure. The root node is a product list and displays the backup products provided by the current backup system. Currently, only pangu metadata backup is provided.

4. Click a product backup item on the left and then configure the backup information on the right.

Configuration	Description
Product Cluster Location	The IP address of the actual transfer server.
Backup File Folder	A folder on the transfer server. You are only required to enter the value in the field, instead of manually creating the folder. For example, enter /apsarapangu/disk8/ pangu_master_bak /product name_pangu/bin.
Script Execution Folder	A folder on the transfer server. You are only required to enter the value in the field, instead of manually creating the folder. For example, enter /apsarapangu/disk8/ pangu_master_bak /product name_pangu/bin.
Script Parameters	You must enter the value in the format ofip=xxx .xxx.xxx, in which the IP address is any IP address of pangu master.
Backup Schedule	Enter 1 here, indicating that the backup is only performed once.
Backup Schedule Unit	Select Day, Hour, or Minute. Select Hour here, indicating that the backup is performed by the hour.
Time-out	Select the timeout. Enter 3600 minutes here.

- 5. Then, click Modify to complete the configurations and trigger the backup.
- 6. Follow the preceding steps to configure all the backup items.

1.1.11.2.2 View the backup details

You can view the backup details of each backup item in Apsara Stack Operations (ASO) during the backup.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Backup Service > Backup Details.
- 4. On the Backup Details page, enter the product and backup item, select the start date and end date, and then click Search.
- 5. View the backup details of a backup item, including the product, backup item, file name (file that requires to be backed up), start time, and state.

The state consists of four types: not started, in process, timeout, and error.

1.1.11.2.3 View the backup server status

You can view the memory, disk, and CPU usage of the current backup server before and after the backup.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Backup Service > Service Status.
- 4. On the Service Status page, view the memory, disk, and CPU usage of the current backup server.

1.1.11.3 View the backup status

The Service Status module allows you to view the status of the current backup service, including the backup product, completed backup items, timeout backup items, and failed backup items, and view the status of the current backup server.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Service Status > State.

- 4. On the State page, view the current backup status.
 - View the numbers of backup items that are in process, completed, timed out, and failed in the current system.
 - View the statuses of the latest backup items of the current product.

The backup status consists of the following types: success, not started, in process, timeout, and failure.

• View the status of the current backup server on the right, namely the memory , disk, and CPU usage.

1.1.12 NOC

1.1.12.1 Overview

The Network Operation Center (NOC) module 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 of network-wide 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.

1.1.12.2 Dashboard

1.1.12.2.1 View the dashboard

The Dashboard tab allows you to monitor the current devices, network, and traffic.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Dashboard tab to view the dashboard information.

Item	Description	
Device Management	Device Overview	The model distribution of used network devices.

Item	Description		
	Ports Usage	 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 all the device ports within a certain time range, of which the top 5 are displayed. 	
	Configuration Management	 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	Alerts	The total number of alerts generated by network devices.	
	Alerting Devices	The number of network devices that generate alerts and the total number of network devices.	
	Alarm Details	The details of the alert.	
Traffic Dashboard	SLB Overview	The bandwidth utilizatio n of SLB clusters.	

Item		Description
	XGW Overview	The bandwidth utilizatio n of XGW clusters.

Dashboard Network Topology Custom Vi	ew	
Device Management		
Device Overview		SLB Overview
Total Devices by Model Arccent	Ateris Atering Devices	Cluster Bandwidth UNication 30M 1H 0H 12H
H3C H	Alarm Details Time Device Name Alert Item Details	
Ports Usage		
Ports Utilization		
• DSW		XGW Overview
o ASW → 0% → 0%		Cluster Bandwidh Utilization
Error Packets by PortTOP5 $\label{eq:Time-Range: 1} \mbox{Time-Range: 1} \mbox{Month } \lor$		
Ranking Device Port Total Error Packets		

1.1.12.2.2 View the network topology

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Network Topology tab.
- 4. On the Network Topology tab, view the physical network topology of a physical data center.

You can select Standard Topology, Dynamic Topology, or Historical Topology as the Topology Type as needed.



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

- Green: The link works properly.
- Red: The link has an error.

• Grey: The link is inactive.

If the Topology Type is Standard Topology, the Refresh Alert switch is turned on by default. You can turn off the Refresh Alert switch, and then devices or link statuses in the topology are not updated after new alerts are triggered.

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

1.1.12.2.3 Create a view

You can create a custom view based on business needs to display the monitoring data and graph information you are interested in.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Custom View tab.
- 4. Click Create View.

Dashboard	Net	work Topology	Custom View				
Select a view.	~	01/01/2020 18:25:12	- 01/02/2020 00:25:12	8	Search	Create View	Delete View

5. In the displayed dialog box, enter the view name and the description, and then click OK.

The view name cannot be the same as an existing name. If the message A view with the same name already exists appears, you must change the view name to a unique one, and then click OK.

1.1.12.2.4 Add a subview

By default, no subviews exist in a view after you create the view. You can add subviews to the view as needed.

Prerequisites

Make sure a view is created. For more information about how to create a view, see

Create a view.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Custom View tab.
- 4. In the search box, select the view and then click Search.

Dashboard	Ι	Network Topology	Custom View				
test01		✓ Start Date	- End Date	Ē	Search	Create View	Delete View

- 5. Click the + button.
- 6. On the displayed page, select the device, monitoring object, monitoring metric, and monitoring submetric.

		×
Device	LSW-VM-G1-1.AMTEST92 V	
Monitoring	interface 🗸	
Object	LJ	
Monitoring	Ten-GigabitEthernet1/0/1	
Monitoring Submetrie	Select ^	
Judineuro	out_bps	
	in_bps	
	in_pps	
	out pps	

7. Then, click OK.

After the subview is added, the system automatically displays the subview on the view to which the subview belongs.



8. Optional: You can add other subviews as needed.

1.1.12.2.5 Delete a subview

Prerequisites

!) Notice:

Deleting a view will delete its subviews at the same time, so proceed with caution.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Custom View tab.
- 4. In the search box, select the view to which the subview to be deleted belongs and then click Search.
5. Click the x button in the upper-right corner of the subview to be deleted.



6. In the displayed dialog box, click OK.

1.1.12.2.6 Delete a view

You can delete a view that is no longer in use.

Prerequisites

!) Notice:

Deleting a view will delete its subviews at the same time, so proceed with caution.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Custom View tab.
- 4. In the search box, select the view to be deleted and then click Search.
- 5. Click Delete View at the top of the page.
- 6. In the displayed dialog box, click OK.

1.1.12.3 Resource management

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

1.1.12.3.1 Network elements

The Network Elements module displays the basic information and running status of physical network devices, and allows you to configure and manage physical

network devices, including device management, password management, and configuration comparison.

1.1.12.3.1.1 Device management

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

1.1.12.3.1.1.1 View the network monitoring information The Network Monitoring tab allows you to view the basic information, running status, and traffic monitoring of Apsara Stack physical network devices, and know the health status of devices in the whole network in time.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Network Monitoring tab under Device Management.
- 4. Then, you can:
 - View the basic information, ping status, and SNMP status of Apsara Stack physical network devices.



You can also click Export to CSV to export the network device information to your local computer as required.

If a problem exists in the business connectivity or gateway connectivity, the value in the Ping Status column or SNMP Status column changes from green to red. Then, the operations personnel are required to troubleshoot the problem.

- In the search box in the upper-right corner, enter the device name or IP address to search for the monitoring information of a specific device.
- $\cdot \,$ View the port information and alert information of a device.
 - a. Click a device name, or click View in the Details column at the right of a device.
 - b. Under Port, view the port list, port working status, and other link information of the device.
 - c. Under Alert Info, view the alert information of the device.

During the daily operations, you must pay close attention to the alert information list of the device. Normally, no data exists under Alert Info, indicating that the device works properly.

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

- View the traffic information of a device for a specific port and time range.
 - a. Click a device name, or click View in the Details column at the right of a device.
 - b. Find the port that you are about to view under Port, and then click View in the Details column.
 - c. Select a time range on the right and then click Search to view the traffic in the selected time range.

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

1.1.12.3.1.1.2 View logs

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

Context

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Syslogs tab under Device Management.
- 4. In the upper-right corner, select the name of the device that you are about to view from the drop-down list, and then select a time range. Click Search to view if the device generates system logs during the selected time range.

No search results exist if the device has a configuration exception or does not generate any logs during the selected time range.

- 5. Optional: You can filter the search results based on the log keyword.
- 6. Optional: Click Export to CSV in the upper-right corner to export the search results to your local computer.

1.1.12.3.1.1.3 Collection settings

The Collection Settings tab allows you to configure the collection interval of

physical network element devices and manage OOB network segments. 1.1.12.3.1.1.3.1 Configure the collection interval

Before collecting the network device information, you must configure the collection interval of network device information according to the business requirements.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Collection Settings tab under Device Management.
- 4. In the Collection Interval Settings section, configure the auto scan interval, device scan interval, port scan interval, and link scan interval.

If you have no special requirements, we recommend that you use the initial default value.

5. Click Submit.

Then, the system collects the device information based on your configuration.

1.1.12.3.1.1.3.2 Modify the collection interval

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Collection Settings tab under Device Management.
- 4. In the Collection Interval Settings section, modify the values.



To not save your modification before the submittal, click Reset in the upperright corner to reset the collection interval to the former version.

5. Click Submit.

One minute later, the modified collection interval of network device information is synchronized to the system.

1.1.12.3.1.1.3.3 Add an OOB network segment

If this is the first time to use the Network Elements function of Network Operation Center (NOC), you must add the device loopback IP address range planned by the current Apsara Stack network device, which is generally the IP address range of the netdev.loopback field in the IP address planning list.

Context

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

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

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Collection Settings tab under Device Management.
- 4. In the OOB Network Segments section, click Add Network Segment.
- 5. In the displayed dialog box, enter the IP address range containing the mask information, subnet mask, and select a data center.
- 6. Click Submit.

The initial data is synchronized to the system after the submittal.

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

1.1.12.3.1.1.3.4 View the OOB network segment information

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Collection Settings tab under Device Management.
- 4. In the OOB Network Segments section, click Refresh on the right.
- 5. In the list, view the network segment information of your managed device.

Note:

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

1.1.12.3.1.2 Modify the device password

You can modify the passwords of physical network devices as required.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Password Management tab.

4. Optional: Enter the name of the device whose password is to be modified in the search box of the Devices on Live Network section and then click Search.

To search for another device, click Reset to reset the configured search condition.

5. Select one or more devices and then click Add.

Then, the selected devices are displayed in the Target Devices section on the right.

Note:

To remove a device from the Target Devices section, click Manage > Delete in the Actions column at the right of the device. You can also click Clear in the upperright corner to remove all the devices in the Target Devices section.

- 6. The system must verify the old password before you modify it. Enter the Username and Old Password in the lower-right corner and then click Verify. You must verify the old password for all the devices in the Target Devices section.
- 7. After the verification is passed, modify the password for one or more devices as required.
 - Modify the password of a device

Click Manage > Set Username and Password in the Actions column at the right of a device. Enter the username and password in the displayed dialog box and then click OK.

 $\cdot\,$ Modify the passwords of all devices

Click Modify under the Target Devices section to modify the passwords of all the devices added to the Target Devices section.

1.1.12.3.1.3 Configuration comparison

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. Click the Config Comparison tab.

4. Optional: Enter the name of the device whose configurations you are about to compare in the Device Name search box and then click Search.

To search for another device, click Reset to reset the configured search condition.

5. Select the device and then click Compare Configuration.

After the comparison, click Refresh and then click Export Results to export the differences.

1.1.12.3.2 Service Load Balancers

The Service Load Balancers module displays the basic information, running status, and water level of network product Server Load Balancer by using cluster monitoring and instance monitoring.

1.1.12.3.2.1 View the cluster monitoring information

The Cluster Monitoring tab allows you to view the basic information, inbound limit (bit/s), outbound limit (bit/s), inbound limit (PPS), outbound limit (PPS), active connection limit, inactive connection limit, and water level of a single device node in a cluster.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Service Load Balancers.
- 3. Click the Cluster Monitoring tab.
- 4. Select the cluster that you are about to view from the drop-down list and then click Search.

The information of all device nodes in the cluster is displayed.

Cluster Monitoring	nstance Monitoring				
cn-qingdao-env4b-d01 🗸 🗸	Search				Search Q
Node IP Address	Status	Local IP Address	Site ID	LVS group ID	Details
	online				
	online				
<pre> Prev 1 Next > </pre>					Items per Page 10 🗸

5. Find a device node and then click View in the Details column.

6. On the Node Message page, view the basic information, inbound limit (bit/ s), outbound limit (bit/s), inbound limit (PPS), outbound limit (PPS), active connection limit, and inactive connection limit of the device node.

Node Message		
	Status : online	Local IP Address Range :
NIC : ourningu SITE ID : 1 Inhannel I imit (Bit(c) - 1049578	Active Connection Limit : 10000 Outbaund Limit (2014) - No Limit	Proxy Check Type : map Inactive Connection Limit : 0
Outbound Limit (PPS) : No Limit	Outoound Linnik (bills) : No Linnik	וווסטנווס בוחות (רריג) . 10000

1.1.12.3.2.2 View the instance monitoring information

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Service Load Balancers.
- 3. Click the Instance Monitoring tab.
- 4. Select the cluster where the instance that you are about to view is located from the drop-down list. Enter the lb-id or VIP address that you are about to search for in the field and then click Search.
- 5. In the search result, view the monitoring information of the instance.

Where,

- The first section is the basic information of the SLB instance, which allows operations engineers to troubleshoot problems and confirm the owner where a device belongs.
- The second section is the operating water level graph of the instance. Select a time range and then click Search or select 5MIN, 30MIN, 1H, or 6H in the Quick Query section to view the operating water level graph of the instance in a specific time range, including the detailed bps and pps.

1.1.12.3.3 Collect IP addresses

The system regularly collects the IP addresses of all the physical networks in the current Apsara Stack environment based on the configured collection interval. You

can search for the information of devices and ports to which a network segment or IP address belongs based on the network segment/IP address and subnet mask.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Collection.
- 3. Enter the network segment/IP address and subnet mask in the corresponding search boxes and then click Search.

If the network segment address you are searching for belongs to a network segment in the current Apsara Stack environment, the system displays the information of devices and ports to which the network segment address belongs.

Note:

If you enter an IP address in the search box and then click Search, the system calculates the corresponding network segment address based on the IP address and subnet mask.

1.1.12.3.4 IP address ranges

The IP Address Ranges module is used to manage the planning information in the Apsara Stack environment, including the network architecture and IP address planning. You can modify, import, and export the planning information.

1.1.12.3.4.1 Import the planning file

No data is imported when the system is initialized. You must import the planning file to obtain the IP address allocation information of the current Apsara Stack environment. You can also import a new planning file for a change in the environment.

Prerequisites

The IP address allocation list is obtained from Apsara Stack Deployment Planner.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Ranges.
- 3. Click Import in the upper-right corner.

- 4. In the displayed dialog box, click Browse and then select the IP address allocation list.
- 5. Click Import.

1.1.12.3.4.2 Manually add the IP address pool information

You can also manually add new IP address pool information to Apsara Stack Operations (ASO) for centralized management.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Ranges.
- 3. Click Add.
- 4. In the displayed dialog box, complete the IP address pool information.
- 5. Click Add.

1.1.12.3.4.3 Modify the IP address pool information

If an IP address range is changed, you can modify the IP address pool information.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Ranges.
- 3. Optional: On the IP Address Ranges page, configure the search conditions and then click Search.

Note:

To reset the search conditions, click Reset to clear your configurations with one click.

- 4. Find the IP address pool whose information you are about to modify and then click Manage > Edit in the Actions column.
- 5. In the displayed dialog box, modify the network architecture and IP address planning.
- 6. Then, click Edit.

1.1.12.3.4.4 Export the IP address pool information

You can export the IP address pool information to your local computer and then view the information offline.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Ranges.
- 3. Select the IP address pool whose information you are about to export and then click Export.

1.1.12.3.4.5 Delete the IP address pool information

You can delete the IP address pool information that is no longer in use.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Ranges.
- 3. Find the IP address pool whose information you are about to delete and then click Manage > Delete in the Actions column.

1.1.12.4 Alert management

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

1.1.12.4.1 View and process current alerts

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Alert Management > Alert Dashboard.
- 3. Click the Current Alerts tab.
- 4. Enter a keyword in the search box in the upper-right corner and then click Search.

Alerts that meet the search condition are displayed.

- 5. Optional: You can filter the search results by device name, device IP address, or alert name.
- 6. Click Details in the Details column at the right of an alert to view the detailed alert information.
- 7. Find the reason why the alert is triggered and then process the alert.
 - If the alert does not affect the system normal operation, you can click Ignore in the Actions column to ignore the alert.
 - If the alert is meaningless, you can click Delete in the Actions column to delete the alert.

After processing an alert, you can search for it on the History Alerts tab.

8. Optional: Click Export to CSV to export the alert information to your local computer.

1.1.12.4.2 View history alerts

You can view history alerts on the History Alerts tab.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Alert Management > Alert Dashboard.
- 3. Click the History Alerts tab.
- 4. Select Alert Source, Alerting IP Address, Alerting Device, Alert Name, Alert Item, or Alerting Instance from the drop-down list and then enter a keyword in the field. Select a time range and then click Search.

Alerts that meet the search conditions are displayed.

- 5. Click Details in the Details column at the right of an alert to view the detailed alert information.
- 6. Optional: Click Export to CSV to export the alert information to your local computer.

1.1.12.4.3 Add a trap

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Alert Management > Alert Settings.
- 3. On the Alert Settings page, click Configure Trap.
- 4. In the displayed Configure Trap dialog box, complete the configurations.

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

Configuration	Description	Example
Trap Name	The name of the alert event .	linkdown or BGPneighbor down. You can customize this value.
Trap OID	The OID of the alert event.	.1.3.6.1.4.1.25506.8.35.12.1 .12 Configure the value strictly according to the device document. You cannot customize this value.
Тгар Туре	The type of the alert event . Select a value from the drop-down list.	-
Trap Index	The index ID of the alert item.	This value is the KV information in the trap message, which is used to identify the alert object. Generally, this value can be an API name, protocol ID, or index ID. Configure the value strictly according to the device document. You cannot customize this value.

Configuration	Description	Example
Trap Msg	The message of the alert item.	This value is the KV information in the trap message, which is used to identify the alert data. Generally, this value can be the additional informatio n of the alert item, such as a system message or a message indicating the location of the state machine or the current status. Configure the value strictly according to the device document. You cannot customize this value.
Alert Type	Indicates whether this alert is of the fault type or the event type.	-

Configuration	Description	Example
Association	Indicates whether this alert has an event alert.	-
	If Fault is selected as	
	the Alert Type and this	
	alert has an association	
	alert, select Event Alert as	
	Association and then add	
	the trap of the association	
	alert.	

Configure Trap				🚭 Clear	×
Trap Name :	1		Alert Type:	● Fault ◯ Event	
Trap OID:			Association:	Event Alert None	
Trap Type:	Select ~				
Trap Index:		+ Submit	Trap Msg:		÷

5. Then, click Submit.

After the submittal, the system checks if the trap OID and trap name are the same as the existing ones. If not, the alert settings of the added trap are finished.

The system pays attention to the alert events of the configured trap OID and such alert events are displayed on the Current Alerts and History Alerts tabs of Alert Dashboard.

1.1.12.4.4 View a trap

You can view a trap configured in the current system.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Alert Management > Alert Settings.

3. Enter a keyword in the search box in the upper-right corner and then click Search.

Note:

After the search results are displayed, you can click Export to CSV in the upperright corner to export the trap information to your local computer.

- 4. Optional: You can filter the search results by trap name, trap type, or OID.
- 5. Find a trap and then move the pointer over Details in the Actions column to view the detailed trap information.

Dive:

If a trap is no longer in use, you can click Delete in the Actions column at the right of the trap.

1.1.12.5 Network reconfiguration

The Network Reconfiguration module allows you to automatically reconfigure the network of the data center in Apsara Stack Operations (ASO).

1.1.12.5.1 Physical network integration

The Physical Network Integration module allows network operations engineers to perform automated integration of physical networks in Apsara Stack Operations (ASO) by entering the integration parameters. Network Operation Center (NOC) automatically generates and issues the configurations to specific devices and then automatically performs the network integration test.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose NOC > Network Reconfiguration > Physical Network Integration.
- 3. Enter the project name and then click Create to create a project.

The network operations engineer must create a project file for this change to store the parameters related to the change. You can click Manage > Import in the History section to import the project information for later usage.

- 4. Click Save Project in the upper-right corner to save the project details.
- 5. Click Next.

- 6. Select a device.
 - a) In the Select Device step, enter a device name in the search box of the Devices on Live Network section and then click Search.

After adding a device, you can click Reset to clear the search condition and then search for and add another device.

b) Click Add at the right of the device required by this change to add it to the Target Device section on the right.

To remove the device from the Target Device section, click Manage > Delete at the right of the device. You can also click Manage > Set the username and password to modify the logon username and password of the device.

- c) Click Save Project in the upper-right corner to save the information of devices added to the Target Device section.
- 7. Click Next.
- 8. Configure the interface parameters.
 - a) In the Configure Interfaces step, click Edit.
 - b) Complete the parameter configurations and then click Add to add the interface to the list.

You can click Manage > Edit or Manage > Delete in the list to modify or delete the interface.

- c) Click Save Project in the upper-right corner to save the information.
- 9. Click Next.
- 10.Configure the route parameters.
 - a) In the Configure Routes step, click Edit.
 - b) Complete the parameter configurations and then click Add to add the route to the list.

You can click Manage > Edit or Manage > Delete in the list to modify or delete the route.

c) Click Save Project in the upper-right corner to save the information.

11.Click Next.

12.Configure the route policies.

- a) In the Configure Route Policies step, click Edit.
- b) Complete the parameter configurations and then click Add to add the route policy to the list.

You can click Manage > Edit or Manage > Delete in the list to modify or delete the route policy.

c) Click Save Project in the upper-right corner to save the information.

13.Click Next.

14Jn the Generate Integration Configurations step, click Generate to generate the integration configurations.

The system generates the integration configuration commands and rollback commands of all the devices with parameters configured.

Operations engineers can automatically generate the configurations of each device based on the configured parameters. After the generation, click View in the Actions column to view the corresponding commands on the left.

You can also click Export to export the file, which contains the configuration commands and rollback commands of detection devices, to your local computer.

1.1.12.5.2 ASW scale-up

You can automatically scale up ASW devices in Apsara Stack Operations (ASO) by using ASW scale-up. After network operations engineers enter the scaleup parameters, Network Operation Center (NOC) automatically generates the configuration and pushes the configuration to a specific device for automatic scaleup.

Prerequisites

Before scaling up ASW devices in ASO, you must plan the IP addresses and configure the ASW in *Apsara Stack Deployment Planner*.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Network Reconfiguration > ASW Scale-up.

- 3. Select devices to be implemented.
 - a) In the Select Device step, enter a device name in the search box of the Devices on Live Network section and then click Search.

After adding a device, you can click Reset to clear the search condition and then search for and add another device.

b) Click Add at the right of the device to be implemented for this change to add the device on live network to the Target Device list.

To remove a device, click Manage > Delete in the Target Device list. You can also modify the logon username and password of the device by clicking Manage > Set the username and password.

- 4. Click Next.
- 5. Disable the DSW ports.
 - a) In the Disable DSW Port step, click Port Settings at the right of the device to be implemented.
 - b) Disable the corresponding port and then click Implement.
 - c) In the displayed dialog box, click OK to run the script commands.
- 6. Click Next.
- 7. Configure the DSW ports.
 - a) In the Configure DSW Port step, click Edit at the right of the device to be implemented. The Interface Parameter Configuration list is displayed.
 - b) Select the Display Ports, enter the Port Description, IP Address, and Subnet Mask, and then click Add to add the interface parameter to the list.
 Then, you can click Manage > Edit or Manage > Delete to modify or delete the interface parameter.
 - c) After adding the interface parameter, click Implement at the right of the device.
 - d) In the displayed dialog box, click OK to run the script commands.

If an exception occurs after the implementation, you can click Back to roll back to the version before the implementation.

8. Click Next.

- 9. Configure the BGP.
 - a) In the Configure BGP step, click Edit at the right of the device to be implemented. The Interface Parameter Configuration list is displayed.
 - b) Enter the Group Name, Peer ASN, and Peer IP Address, and select the Local Port Name. Then, click Add to add the interface parameter to the list.
 Then, you can click Manage > Edit or Manage > Delete to modify or delete the interface parameter.
 - c) After adding the interface parameter, click Implement at the right of the device.
 - d) In the displayed dialog box, click OK to run the script commands.

If an exception occurs after the implementation, you can click Back to roll back to the version before the implementation.

10.Click Next.

11 In the Upload ASW Configurations step, upload the new ASW configuration.

12.Click Next.

13Enable the DSW ports.

- a) In the Enable DSW Port step, click Port Settings at the right of the device to be implemented.
- b) Enable the corresponding port and then click Implement.
- c) In the displayed dialog box, click OK to run the script commands.

14.Click Next.

15Perform the scale-up test.

- a) In the Test Scale-up step, click Select at the right of the device to be implemented. The route table is displayed.
- b) In the ASW IP Address search box, enter the IP address to be tested and then click Add to add it to the ASW Connectivity Test list.
- c) Click Test and then the system returns the test results.

1.1.12.6 Fault check

The Fault Check module consists of IP address conflict check, leased line discovery, and network inspection.

1.1.12.6.1 IP address conflict check

You can check if conflicted IP addresses exist in the current Apsara Stack environment by using IP address conflict check.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Fault Check > IP Address Conflict Check.

On the IP Address Conflict Check page, the system automatically checks if conflicted IP addresses exist in the current Apsara Stack environment. If yes, the conflicted IP addresses are displayed in the list. You can also view the port information, device name, and logon IP address to which each conflicted IP address belongs.

1.1.12.6.2 Leased line discovery

You can configure the leased line discovery of devices in Apsara Stack Operations (ASO) and implement it automatically. After network operations engineers configure the discovery parameters, Network Operation Center (NOC) automatically generates the discovery configuration, pushes the configuration to a specific device, and then automatically performs the discovery test.

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose NOC > Fault Check > Leased Line Discovery.

- 3. Select a discovery source.
 - a) In the Select Sources step, enter a device name in the search box of the Devices on Live Network section and then click Search.

After adding a device, you can click Reset to clear the search condition and then search for and add another device.

b) Click Add for Discovery at the right of the device to add a device on live network to the Devices for Discovery list on the right.

To remove a device from the Devices for Discovery list, click Manage > Delete in the list. You can also modify the logon username and password of the device by clicking Manage > Set the username and password.

- 4. Click Next.
- 5. Configure the discovery parameters.
 - a) In the Configure Parameters step, click Edit. The Configure Parameters list is displayed.
 - b) Enter the Link Name, Destination IP Address, Source IP, Discovery Interval, Discoveries, and Discovery Timeout, and then click Add to add the information to the list.

Then, you can click Manage > Edit or Manage > Delete to modify or delete the discovery parameter.

- 6. Click Next.
- 7. In the Generate Discovery Configuration step, click Generate to generate the discovery configuration commands and rollback commands of all devices with discovery parameters configured.

Then, click View in the Actions column to display the corresponding commands on the left.

You can also select one or more devices and then click Export to export the files containing configuration commands and rollback commands of discovery devices to your local computer.

- 8. Click Next.
- 9. In the Push Configuration step, click Push Configurations.
- 10In the displayed dialog box, click Continue to push the discovery configuration commands to the corresponding device.

Then, you can click View Logs to view the detailed pushed logs.

11.Click Next.

- 12.In the Start Discovery step, click Started at the right of a device for discovery to perform the leased line discovery test.
- 13.Then, click Next.
- 14Jn the Roll Back Discovery step, click Roll Back at the right of each device that you have performed the leased line test to roll back the corresponding NQA configuration in the device.

You can click View Logs to view the detailed rollback logs.

1.1.12.6.3 Network inspection

You can configure the inspection of network devices in Apsara Stack Operations (ASO) and implement it automatically for daily fault checking of network devices.

Context

Generally, the time interval of a network inspection is a week or a day.

Procedure

- 1. Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Fault Check > Network Inspection.
- 3. In the Create/Import Project step, enter the project name and then click Create to create a project.

Network operations engineers must create a project file for this inspection. Parameters related to the project are saved in the file and you can click Manage > Import in the History section to import the project information if needed.

- 4. Click Save Project in the upper-right corner to save the project details.
- 5. Click Next.

- 6. Select devices for inspection.
 - a) In the Select Device for Inspection step, enter a device name in the search box of the Devices on Live Network section and then click Search.

After adding a device, you can click Reset to clear the search condition and then search for and add another device.

b) Select one or more devices and then click Add for Inspection to add the devices to the Target Devices list on the right.

To remove a device from the Target Devices list, click Manage > Delete in the list. You can also modify the logon username and password of the device by clicking Manage > Set Username and Password.

c) Click Save Project in the upper-right corner to save the information of devices for inspection.

Note:

The system only saves the information of devices whose Status is Accessible in the Target Devices list.

- 7. Click Next.
- 8. Select check items.
 - a) In the Select Check Item step, select one or more check items on the left and then click Add for Inspection.

The added check items are displayed on the right.

To remove an added check item, click Delete in the Manage column at the right of the check item.

- b) Click Save Project in the upper-right corner to save the current information.
- 9. Click Next.
- 10In the Start Inspection step, click Check in the Action column at the right of each check item to create an inspection task.
- 11 After the inspection, click Refresh to refresh the inspection result.
- 12.Click Details in the Check Details column of each check item to view the inspection details of the check item.
- 13.Optional: You can also click Export Result to export all the information of check items to your local computer for offline analysis.

1.1.12.6.4 Configuration baseline audit

The Configuration Baseline Audit module allows you to compare the baseline configurations of devices with the current running configurations.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose NOC > Fault Check > Configuration Baseline Audit.
- 3. Select one or more devices in the device list and then click Audit. Then, the system starts to audit the baseline configurations of the selected devices.

The statuses during the audit process and the corresponding descriptions are as follows.

Status	Description
Pending	The initial status.
Auditing	The baseline configurations of the device are being audited in the background.
Pass	The baseline configurations of the device are the same as the running configurations.
Fail	The baseline configurations of the device are different from the running configurations.
Disconnected	The system fails to connect to the device.
No Data	The system fails to obtain the baseline configurations of the device.

- 4. After the audit is complete, click Refresh to update the audit results.
- 5. In the Actions column of the device, click View the result to show the audit result on the right.

1.1.13 Full Stack Monitor

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

1.1.13.1 SLA

The SLA module allows you to view the current state, history data, instance availability, and product 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.

1.1.13.1.1 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.
- 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 properly.
- Yellow: warning. The service has some latency, but can still work properly.
- Red: hitch. The service is temporarily interrupted and cannot work properly.
- 4. Find the product whose running state you are about to view. Click Check in the Operation column.
 - The Overall Availability section displays the availability of a product. You can view the availability by hour, day, or minute.
 - The Related Events section displays the current exception events. Click Show Details to view the event details.

1.1.13.1.2 View the history data of a cloud product

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

- **1.** Log on to Apsara Stack Operations.
- 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 is displayed on this page. Different colors represent different statuses:

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

1.1.13.1.3 View the availability of an instance

You can view the current instance availability ratio of a cloud product to know the instance damages.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 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, 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.

1.1.13.1.4 View the availability of a product

You can view the availability ratio of a cloud product to know its monthly availability index.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Full Stack Monitor > SLA.
- 3. Click the Availability of Product tab.
- 4. Select the Product and Time Range, and then click Search to view the availability ratio of the product.

For example, if the availability ratio of Elastic Compute Service (ECS) is 100.00%, it indicates that ECS runs properly this month, without any faults.

1.1.13.2 Operations full link logs

The Operations Full Link Logs module allows you to search for logs of ECS-, SLB-, and All in ECS-related applications.

Context

- Currently, you can search for logs of multiple product components, such as pop, openapi, pync, and opsapi, on the ECS tab.
- If each SLB service node properly enables the ilogtail reporting feature, you can search for logs of pop, slb-yaochi, and slb-control-master on the SLB tab.
- You can search for vm_adapter logs, all in ECS-Apsara Infrastructure Management Framework adaption layer logs, and all the other ECS operations logs on the All in ECS tab.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Full Stack Monitor > Operations Full Link Logs.
- 3. Click the ECS, SLB, or All in ECS tab.
- 4. Enter a keyword in the Query field. Select the time range in the Time field. Then, click Search.

Note:

You can enter any string in the Query field as the search condition, such as the instance ID, request ID, or the keyword "error".

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

6. 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.

- 7. Enter a keyword in the search box to search for the related information in the search results.
- 8. Optional: After the search, you can click Export Log to export the search results to your local computer.

1.1.13.3 Correlation diagnosis and alarm

The Correlation Diagnosis and Alarm module allows you to perform an aggregate query on the system alert events, and perform a correlation query on physical servers, network devices, ECS instances, RDS instances, SLB instances, and VPC instances.

1.1.13.3.1 Full stack correlation alert

The Full Stack Correlation Alert tab consists of two sections: full stack topology and full stack alert. The full stack topology section allows you to view the physical network topology in the current data center. The full stack alert section allows you to view the alert event list after the aggregation and the corresponding details.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Full Stack Monitor > Correlation Diagnosis and Alarm.
- 3. Click the Full Stack Correlation Alert tab.
- 4. Then, you can:
 - View the full stack topology.

Select the product that you are about to view from the drop-down list and then enter the corresponding instance ID in the field. Click Add to add multiple products and then click Determine to obtain the full stack topology.



Currently, you can only view the full stack topology of ECS instances, RDS instances, SLB instances, and NC servers.

In the topology, you can click the instance icon to obtain the instance information or click the network connection to obtain the connection information.

• View the full stack alerts.

By default, the Full Stack Alert section displays the alert events aggregated in the current system by using the correlation diagnosis.

Complete the following steps to view the full stack alerts of an instance in a specific time range.

- a. 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. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
- b. In the displayed alert list, click at the right of Alert Type and Alert Level to filter the alert results.
- c. Click Details at the right of an alert event.
- d. On the Detail page, you can view the details of the exception event related to the alert, including the alert basic information, associated event information, impacted instances in ECS, and impacted instances in RDS.

1.1.13.3.2 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 alert information of the network where a server is located, and the full stack correlation alert information.

- **1.** Log on to Apsara Stack Operations.
- 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 instance ID in the search box, select the time range, and then click Search. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.

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

View the uplink network topology of the host, which visually shows the alerts of network devices (blue indicates the normal status and red indicates the abnormal status).

You can click SERVER in the topology to view the performance data of the server, including the CPU utilization, TCP retransmission rate, NIC traffic, and packet loss statistics.

In the topology, click the connection between a server and a network device or the connection between two network devices to view the device port information. Click a port to view the water level graph of the port.

• Title Message

View the basic operating data for the operating system of the host.

NC Diagnostics Info

View the real-time diagnosis and alert information of the host.

- 🗾 indicates the diagnosis is passed.
- **F** indicates the detection does not obtain results.
- **Image indicates an exception at the warning level exists.**
- 📷 indicates a fatal exception exists.
- **maindicates the item is being diagnosed.**
- NC Retransmit Root Cause Location

Used to detect the packet loss on the NC server or in the transmission process from NC server to ASW. After the system detects the TCP retransmission, the

- backend diagnoses the server metrics and configurations. The analysis results are displayed after the diagnosis.
- Network Alert Info

View the alert information of the network devices that are included in the uplink network topology of the host.

• Full Stack Alert

View the list of aggregated alert events and the corresponding details.

1.1.13.3.3 Network device

You can use the network device IP address or network device name to search for and view the essential information, real-time diagnosis information, and full stack correlation alert information of a network device.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 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 device ID in the search box, select the time range, and then click Search. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
- 5. You can view the following information on this page.
 - Essential Information
 - View the basic information of the network device.
 - Diagnostic Information

View the real-time diagnosis and alert information of the network device.

Full Stack Alert

View the list of aggregated alert events of the network device.

1.1.13.3.4 ECS

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

- **1.** Log on to Apsara Stack Operations.
- 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. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
- 5. You can view the following information on this page.
 - · Topology

View the uplink network topology of the host to which the ECS instance belongs. The topology visually shows the alerts of network devices (blue indicates the normal status and red indicates the abnormal status).

• ECS Basic Info and HostNC Basic Info

View the basic information of the ECS instance and the host to which the ECS instance belongs.

• ECS Diagnosis Info and Host Nc Diagnostic Information

View the diagnosis and alert information of the ECS instance and the host to which the ECS instance belongs.

· AVS diagnosis and ECS-Alarm

View the AVS diagnosis information and exceptions of the virtual machine and NC server.

- The operating water level of the ECS instance, including the CPU utilization, disk I/O, and Internet/intranet inbound and outbound traffic.
- netdev

View 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

View 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 Alert Info

View the alert information of the network devices that are included in the uplink network topology of the host to which the ECS instance belongs.

• Full Stack Alert

View the aggregated alert events on the ECS instance and the uplink devices of the ECS instance.

1.1.13.3.5 RDS

You can use the RDS instance ID to search for and view the full stack information, availability diagnosis results, and full stack correlation alert information of an RDS instance.

- **1.** Log on to Apsara Stack Operations.
- 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. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
- 5. You can view the following information on this page.
 - · Topology

View the uplink network topology of the host to which the RDS instance belongs. The topology visually shows the alerts of network devices (blue indicates the normal status and red indicates the abnormal status).

Basic Info

View the basic information of the RDS instance, including the primary database IP address, secondary database IP address, SLB ID, and Proxy IP address.

Instance Performance data

View the performance and water level data of the RDS instance.

· Diagnosis Info

View the availability detection results of the RDS instance in the selected time range.

Network Alert Info

View the alert information of the network devices that are included in the uplink network topology of physical machines in the primary database.

Full Stack Alert

View the aggregated alert events on the RDS instance and the uplink devices of the RDS instance.

1.1.13.3.6 SLB

You can use the Server Load Balancer (SLB) instance ID to search for and view the deployment information of an SLB cluster, and the traffic diagnosis results and bandwidth chart of an SLB instance.

- 1. Log on to Apsara Stack Operations.
- 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. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
- 5. You can view the following information on this page: the topology of an SLB instance, the deployment information of an SLB cluster, the diagnosis information, the SLB bandwidth chart, and the full stack correlation alert information.

Where,

· Topology

View the uplink network topology of the host to which the SLB instance belongs. The topology visually shows the alerts of network devices (blue indicates the normal status and red indicates the abnormal status).

SLB Clusters

View the deployment information of the SLB cluster, namely the service name and host ID.

The Instance ID is the name of the physical machine to which the SLB subservice belongs. You can click the ID to go to the server page for a deep query.

· Diagnostics

View the availability detection results of the SLB instance in the selected time range.

• SLB Bandwidth Chart

View the bandwidth chart of the SLB instance in the selected time range.

Full Stack Alert

View the aggregated alert events on the SLB instance and the uplink devices of the SLB instance.

1.1.13.3.7 VPC

You can use the global_tunnel_id of the VPC leased line to search for the leased line traffic, or use the router interface ID to view the router interface information and the corresponding leased line traffic.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Full Stack Monitor > Correlation Diagnosis and Alarm.
- 3. Click the VPC tab.
- 4. The Topology section displays the topology of the XGW cluster. You can perform the following operations:
 - Enter the global_tunnel_id of the leased line in the search box, select the time range, and then click Search to view the leased line traffic. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
 - Enter the router interface ID, namely the instance ID of the router interface, in the search box, select the time range, and then click Search to view the router interface information and the leased line traffic. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.

1.1.14 Storage Operation Center

The Storage Operation Center module consists of pangu and EBS.

1.1.14.1 Pangu

The Pangu module displays the pangu grail, cluster information, node information, and pangu cluster status.

1.1.14.1.1 Pangu grail

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

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Storage Operation Center > Pangu > Pangu Grail.
- 3. Select the product that you are about to view from the Service drop-down list.

The Pangu Grail module displays the data overview, heatmap of health, and top 5 data of each accessed cloud product as of the current date.

• Overview

The Overview section 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

The Heatmap of Health section displays the health information of all the clusters in the selected product. Clusters in different health statuses are displayed in different colors.

Where,

- Green indicates the normal status.
- Yellow indicates a warning.
- Red indicates the abnormal status.
- Dark red indicates a fatal error.
- 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.

Move the pointer over the color block of each cluster to view the correspond ing service name, server name, and IP address.



· Data of Top 5 Services

The Data of Top 5 Services section displays the data of the top 5 unhealthiest clusters in the time range from zero o'clock to the current time in the current date for the selected product.

This section displays the top 5 clusters in terms of abnormal water levels, abnormal masters, abnormal disks, and abnormal chunk servers. Click the cluster name to go to the corresponding cluster information page.

~	✓ Data of Top 5 Services(Jan 8, 2020, 00:00 ↔ Jan 8, 2020, 20:31:00)								
		Service	Cluster Name	Abnormal Water Level	Health	Service	Cluster Name	Abnormal Masters	Health
		tianji		53.82		ecs			
		nas		47.39		ecs			
		ecs		17.49		sis			
		055				odps			
		ecs		6.05		055			
		Service	Cluster Name	Abnormal Disks	Health	Service	Cluster Name	Abnormal Chunk Servers	Health
		ecs				ots			
		ots				ecs			
		tianji				tianji			
		datahub				datahub			
		055				055			

1.1.14.1.2 Cluster information

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

Procedure

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

By default, data of the first cluster in the Cluster Name drop-down list is displayed.

3. Select the cluster that you are about to view from the Cluster Name drop-down list and then view the following information.

Note:

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

· Overview

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



Alarm Monitor

•

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

✓ Alarm Mor	V Alarm Monitor					
			Alarm Log			
Warning:0						
Fuzzy Search:		٩				
Level			Desc			

• Replica

Displays the replica information of the selected cluster.

• Run Chart of Clusters

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

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



Issue: 20200317

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



Rack Information

Contains Servers in Rack and Storage. Where,

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

✓ Rack Information				
		Servers in Rack		
3.				
2.5				
2-				
1.5 -				
1-				
05-				
0.5-				
V-1	a56g13		a56h11	

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



1.1.14.1.3 Node information

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Storage Operation Center > Pangu > Node Information.

By default, data, namely the master information and chunk server information, of the first cluster in the Cluster Name drop-down list is displayed.

3. Select the cluster that you are about to view from the Cluster Name drop-down list and then view the following information.



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

• 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.

Clus	ter Name: ECS-IO7-A-eb38	\checkmark	
`	Master Info		
	Server		Role
			SECONDARY
			SECONDARY
			PRIMARY

· 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.

∨ Chu	✓ Churk Server Info							
Total:	Total-5 Normal-5 Disconnected-0							
Fuzzy	Fuzzy Search: Enter a keyword Refresh							
	Server		DiskBroken Disks/Disks	SSDCacheBroken Disks/Disks	Status	Backup	Storage (TB)	Usage(%)
+	a56g13210.cloud.h14.amtest 72		0/10	0/10	NORMAL		13.8476	23.9800%
+	a56h11108.cloud.h12.amtest 72		0/10	0/10	NORMAL		13.8476	26.3800%
+	a56g13211.cloud.h14.amtest 72		0/10	0/10	NORMAL		13.8476	24.1900%
+	a56h11210.cloud.h13.amtest 72		0/10	0/10	NORMAL		13.8476	28.6300%
+	a56g13110.cloud.g14.amtest 72		0/10	0/10	NORMAL		13.8476	24.1000%

1.1.14.1.4 Pangu operation

The Pangu Operation module allows you to view the pangu cluster status.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Storage Operation Center > Pangu > Pangu Operation.
- 3. Select a service from the Service drop-down list to view the pangu cluster status of this service.

Clusters in different statuses are in different colors. Where,

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

Service:	ecs	~	
	ECS-IO8-A-eb33	ECS-IO8-A-eb37	ECS-IO7-A-eb38

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

1.1.14.2 EBS

The EBS module provides the following functions: IO HANG fault analysis, Slow IO analysis, and inventory settings.

1.1.14.2.1 IO HANG fault analysis

The IO HANG Fault Analysis module allows you to view the affected virtual machine (VM) list, VM cluster statistics, and device cluster statistics.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Storage Operation Center > EBS > IO HANG.

By default, the system displays the affected VM list, VM cluster statistics, and device cluster statistics in the last 24 hours.

- 3. Select the time range (One Hour, Three Hours, Six Hours, One Day, or customize the time range) that you are about to view and then click Search. View the following information:
 - Affected VM List

The Affected VM List section displays the IO HANG start time and recovery time of all the VMs, and the cluster name and userId to which these VMs belong.

To view the information of a cluster, user, or VM, enter the cluster name, userId, or VM name in the search box to perform a fuzzy search.

• VM Cluster Statistics

The VM Cluster Statistics section displays the number of affected VMs in a cluster.

To view the VM statistics of a cluster, enter the cluster name in the search box to perform a fuzzy search.

Device Cluster Statistics

The Device Cluster Statistics section displays the number of affected devices in a cluster.

To view the device statistics of a cluster, enter the cluster name in the search box to perform a fuzzy search.

1.1.14.2.2 Slow IO analysis

The Slow IO Analysis module allows you to view the Slow IO list, top 10 NCs, cluster statistics, top 5 cluster statistics, and reasons.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Storage Operation Center > EBS > Slow IO.

By default, the system displays the Slow IO list, top 10 NCs, cluster statistics, top 5 cluster statistics, and reasons in the last 24 hours.

- 3. Select the time range (One Hour, Three Hours, Six Hours, One Day, or customize the time range) that you are about to view and then click Search. View the following information:
 - Slow IO List

The Slow IO List section displays the following Slow IO-related data: cluster name, NC, virtual machine, device_id, storage_type, start time, recovery time, number of Slow IO, and reason.

To view the information of a cluster, NC, or block device, you can enter the cluster name, NC IP address, or device_id in the search box to perform a fuzzy search.

You can also sort by Cluster Name, NC, Virtual Machine, device_id, storage_ty pe, Start Time, Recovery Time, Number of Slow IO, and Reason as needed.
Top Ten NC

The system displays the information of top 10 NCs by using a graph and a list. Where,

- The Graphic Analysis section displays the proportion for the number of Slow IO in each cluster of the top 10 NCs by using a pie chart.
- The Top Ten NC section displays the nc_ip, cluster name, slow_io, percentage, and major_reason of the top 10 NCs with the most Slow IO by using a list.

To view the information of a cluster or NC, enter the NC IP address or cluster name in the search box to perform a fuzzy search.

You can also sort by nc_ip, Cluster Name, slow_io, and major_reason as needed.

Cluster Statistics

The Cluster Statistics section displays the cluster name, number of devices, number of Slow IO, percentage, and major reason of a cluster with Slow IO.

To view the information of a cluster, enter the cluster name in the search box to perform a fuzzy search.

You can also sort by Cluster Name, Number of Device, Number of Slow IO, and Major Reason as needed.

• Top Five Cluster Statistics

The system displays the statistics of top 5 clusters by using a graph and a list.

Where,

- The Top Five Cluster Statistics section displays the cluster name, number of devices, number of Slow IO, percentage, and major problem of the top 5 clusters with the most Slow IO by using a list.

To view the information of a cluster, enter the cluster name in the search box to perform a fuzzy search.

You can also sort by Top Five Cluster, Number of Device, Number of Slow IO , and Major Problem as needed.

- The Graphic Analysis section displays the proportion for the number of Slow IO in each of the top 5 clusters by using a pie chart.
- Reason

The system displays the reason statistics by using a graph and a list.

Where,

- The Reason section displays the number of Slow IO from the dimension of reasons.

To view the information of a reason, enter the reason information in the search box to perform a fuzzy search.

You can also sort by Reason and Number of Slow IO as needed.

- The Graphic Analysis section displays the proportion of reasons by using a pie chart.

1.1.14.2.3 Inventory settings

The Inventory Settings module allows you to view the sales status of a cluster, configure the oversold ratio of a cluster, and configure whether a cluster is on sale.

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose Storage Operation Center > EBS > Inventory Settings.

By default, the system displays the data, namely the cluster name, oversold ratio, and sales status, of all the clusters in the current environment.

Inventory Infor	Inventory Information				
ECS-IO8-A-020d Oversold, Ratio:3.0% io8 On Sale	ECS-IO8-A-01b1 Oversold Ratio:3.0% io8 On Sale	ECS-IO7-A-01b8 Oversold Ratio:2.5% io7 On Sale	ECS-IQ7River-A- Ulc7 Oversold Ratio:2.5% io7 On Sale		
ECS-IO8-A-020d					
Adjust Setting Oversell Ratio(%) :	3.0		Confirm		
Adjustment of sales status :					

- 3. Complete the following configurations:
 - Select a cluster. Enter a number in the Adjust Setting Oversell Ratio(%) field, and then click Confirm to configure the oversold ratio of the cluster.
 - Select a cluster. Turn on or off the Adjustment of sales status switch to configure whether the cluster is on sale.

1.1.15 Task Management

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

1.1.15.1 Overview

The Task Management module has the following functions:

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

• Supports uploading the .tar package as the script.

1.1.15.2 View the task overview

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Task Management > Task Overview.

Dashboard				Tasks To Be Intervened				
Pending for Intervention	Running	Failed		Task Name	Task Description	Start Tim	e A	ctions
		20	00	test01	test	Dec 30, 2	019, 10:59:45 D	
Create Task			Create Task					
Running Status in Last 7 Days								
1								
0.8				Running Tasks(Running tim	ne more than 1 day)			
0.6				Task Name	Task Description	Target Group	Start Time	Running Duration
0.4								
0.2								
0 December 30 December 29	December 28 December 2	7 December 26 December	25 December 24					

- 3. On the Task Overview page, you can:
 - In the Dashboard section, view the number of tasks in the Pending for Intervention, Running, Failed, or Completed status in the system.

Click the status or number to view the task list of the corresponding status.

 $\cdot\,$ In the Create Task section, click Create Task to create an operations task.

For more information about how to create a task, see *Create a task*.

- If a task has a breakpoint and runs to the breakpoint, the task stops and waits for manual confirmation. You can view and process tasks to be intervened in the Tasks To Be Intervened section.
- In the Running Status in Last 7 Days section, view the running trend of tasks and whether tasks are successful in the last seven days.
- In the Running Tasks section, view tasks running in the last 24 hours.

1.1.15.3 Create a task

You can create daily changes as tasks to run on the cloud platform.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Task Management > Task Management.
- 3. Click Create.
- 4. In the displayed dialog box, configure the task information.

Configuration	Description
Task Name	The name of the operations task.
Task Description	The description of the operations task.
Target Group	The task target. You can configure the target group in the following ways:
	 Select from the drop-down list by product > cluster > service > server role > virtual machine (VM) or physical machine. Enter the VM or physical machine in the field and then press Enter. You can enter multiple VMs or
	 physical machines in sequence. Click the button next to Target Group. In the displayed dialog box, enter the target group, with one VM or physical machine in one line, and then click OK.

Configuration	Description
Execution Batch	Optional. This option appears after you enter the target group.
	If the Execution Batch is not selected, Target Group
	is displayed in the Target Group column in the task
	list of the Task Management > Task Management
	page. If you select the Execution Batch, Batch
	Execution Policy is displayed in the Target Group
	column.
	You can select the following options as the Execution
	Batch.
	• Default Order
	If the number of machines is equal to or less
	than 10, the machines are allocated to different
	batches by default, with one machine in batch
	1, one machine in batch 2, two machines in
	batch 3, three machines in batch 4, and the other
	machines in batch 5. You can adjust the batch for
	machines as needed.
	If the number of machines is more than 10, the
	machines are allocated to different batches
	by default, with one machine in batch 1, three
	machines in batch 2, five machines in batch 3,
	N/3-1 (an integer) machines in batch 4, N/3-1 (
	an integer) machines in batch 5, until all of the
	machines are allocated. Where, N is the total
	number of servers in the cluster. You can adjust
	the batch for machines as needed.
	• Single-Machine Order: By default, each batch has one machine. You can adjust the batch for machines as needed.

Configuration	Description
Execution Method	If you select the Execution Batch, the Execution Method can only be Manual Execution and cannot be selected.
	If the Execution Batch is not selected, you can select one of the following execution methods:
	 Manual Execution: You must manually start the task. With this option selected, you must click Start in the Actions column to run the task after the task is created. Scheduled Execution: Select the execution time.
	The task automatically runs when the time is reached.
	Regular Execution: Select the interval and execution times to run the task. The task runs again if the execution condition is met.
	• Advanced: Configure the command to run the task periodically.

Configuration	Description
Add Script	Click Add Script. Select one or more .tar packages to upload the script file. After the upload, you can delete and reupload the script.
	After uploading the script, if Manual Execution is selected as the Execution Method, you must select whether to turn on the Intervention Required switch. If the switch is turned on, the task stops and waits for manual intervention after the script runs.

Create Task						×
* Task Name				Task Description		
test				test		
* Target Group 🚄						
a50 🗸	vm010004024198 \times	vm010004028255 \times	vn	n010004021104 ×	vm010004028003 \times	
	vm010004020034 \times	vm010004029054 \times	vn	n010004021096 ×	vm010004024236 \times	
Execution Batch 🕐 🔵 🔵 Default Order 👁) 🔿 Single-Machine	Order 🐵				
* Execution Method						
Manual Execution	~					
+Add Script						
Supported Extension	: .tar					
						Create

5. Then, click Create.

Result

The created task is displayed in the task list.

1.1.15.4 View the execution status of a task After a task runs, you can view the execution status of the task.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Task Management > Task Management.
- 3. Optional: Enter the task name, select the task status, and configure the start time and end time of the task. Then, click Query to search for the task.
- 4. Find the task that you are about to view and then click Target Group or Batch Execution Policy in the Target Group column.



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

Tasks					
Task Name Task S	tatus V Start Date - En	d Date 🕅 Query	Create		•
Task Name	Task Description	Time	Task Status	Target Group	Actions
		End Time :Nov 22, 2019, 14:09:49			
baoxun22	dds	Created At :Nov 15, 2010, 11:23:17 Start Time :Nov 22, 2019, 11:39:59 End Time :Nov 22, 2019, 11:40:37			
4		Created At :Nov 15, 2019, 10:51:18 Start Time :Nov 15, 2019, 10:51:22 End Time :Nov 15, 2019, 10:52:10			
test1		Created At :Nov 11, 2019, 14:43:49 Start Time :Nov 11, 2019, 14:43:52 End Time :Nov 11, 2019, 14:44:04			

5. In the displayed dialog box, view the task execution status based on the machine color. Click a machine to view the execution result of the task.

Batch Execution Policy			Successful	🛑 Failed	Not Executed	😑 Unreachable	×
Batch1	Batch2	Batch3		Batch4			
vm010004024196	vm010004028255	vm010004021104			vm010004020	034	
		vm010004028003			vm010004029	054	
					vm010004021	096	
Batch5							
vm010004024236							
						Go	se

1.1.15.5 Start a task

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Task Management > Task Management.
- 3. Optional: Enter the task name, select the task status, and configure the start time and end time of the task. Then, click Query to search for the task.
- 4. Find the task that you are about to start and then click Start in the Actions column.
- 5. In the displayed dialog box, select the batches to start and then click Start.

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

If the task has the Intervention Required switch turned on, you must intervene the task after clicking Start. The Task Status changes to Pending for Intervention and you can only continue to run the task by clicking Continue in the Actions column.

Tasks								
Task Name Task Status V Start Date - End Date 577 Create								
Task Name	Task Description	Time	Task Status	Target Group	Actions			
test03		Created At Dec 30, 2010, 14:34:17 Start Time :Dec 30, 2010, 14:30:47 End Time :Dec 30, 2010, 14:40:08			Modify Start Delete			
test02		Created At Dec 30, 2019, 11:03:32 Start Time :Dec 30, 2019, 14:43:14 End Time :Dec 30, 2019, 14:43:40						
test01	test	Created At :Dec 30, 2010, 10:50:45 Start Time :Dec 30, 2010, 14:20:38 End Time :	OPending for Intervention					

1.1.15.6 Delete a task

For better management, you can delete a task that is no longer in use.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Task Management > Task Management.
- 3. Optional: Enter the task name, select the task status, and configure the start time and end time of the task. Then, click Query to search for the task.
- 4. Find the task to be deleted and then click Delete in the Actions column.
- 5. Click OK in the displayed dialog box.

1.1.15.7 Process tasks to be intervened

If a task has a breakpoint and runs to the breakpoint, the task stops and waits for manual confirmation. The task can only continue to run after the manual confirmation.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Task Management > Task Overview.

3. In the Tasks To Be Intervened section, find the task to be intervened and then click Details in the Actions column.

Tasks To Be Intervened							
Task Name	Task Description	Start Time	Actions				
test01	test	Dec 30, 2019, 10:59:45	Details				

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

1.1.16 Log Management

The Log Management module is used to access various business logs and allows you to search for, export, back up, and clear logs.

1.1.16.1 Log configurations

Before managing logs, you must complete the configurations for log clearance, projects, agents, and buckets.

1.1.16.1.1 Clear

You can configure the parameters for automatic log clearance and the manual log clearance time on the Clear tab.

1.1.16.1.1.1 Configure parameters for automatic log clearance

To avoid logs filling up the disk space, the system supports automatically clearing logs. You can configure the parameters for automatic log clearance based on business needs.

Procedure

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

By default, you are on the Clear tab.

Clear Project Agent	Bucket Management		
Log Saved At	Storage Limit	Clear Cycle	Actions
7Days	80%	2Days	

3. In the Actions column, click Modify.

- ConfigurationDescriptionLog Saved AtThe system clears logs saved before the configured
time.Clear CycleThe system automatically clears logs according to
this cycle.Storage LimitIf the disk usage of the service cluster that stores
logs exceeds the configured limit, the system clears
logs by day, from oldest to latest, until the disk usage
is below the configured limit.
- 4. On the displayed page, complete the following configurations.

Modify Configurations	
Log Saved At	
7 Days	
Clear Cycle	
2 Days	
Storage Limit	
80 %	

5. Click OK.

1.1.16.1.1.2 Configure the manual log clearance time

The system allows you to manually clear logs. You can manually clear logs of the configured time range based on your requirements.

Procedure

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

By default, you are on the Clear tab.

Clear Project Agent Bucket Man	agement		
Log Saved At	Storage Limit	Clear Cycle	Actions
7Days	80%	2Days	

3. In the Actions column, click Clear Manually.

4. On the displayed page, configure the start time and end time of logs you are about to clear and then click OK.

Then, the system immediately clears logs of the configured time range.

1.1.16.1.2 Project

You can add or delete projects on the Project tab.

1.1.16.1.2.1 Add a project

You must add a project to configure the relationship among the project, product, and InstanceId.

Context

Logs accessed to the Log Management module are named in the format of {InstanceId}-yyyy.MM.dd. yyyy.MM.dd is the date when logs are accessed. For example, 2019.09.10. You can only search for logs by project and product after configuring the relationship among the project, product, and InstanceId.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Project tab.



4. Optional: Click Download Template to download the project template and complete the project information based on the template.



Before adding a project, you must download the project template. If a project file that meets the requirements already exists in your local computer, skip this step.

5. Click Add.

6. On the displayed page, click Select File. Select the project file from your local computer and then click Open to add the project.

Then, you can view the information of the added project in the project list.

At the right of the project, click Show in the Actions column to view the product and instance ID of this project.

1.1.16.1.2.2 Delete a project

You can delete one or more projects.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Project tab.
- 4. Delete one or more projects based on business needs.
 - Select the project name from the Project Name drop-down list and then click Search. Find the project to be deleted and then click Delete in the Actions column. In the displayed dialog box, click OK to delete a single project.
 - Select multiple projects to be deleted and then click Delete In Batch. In the displayed dialog box, click OK to delete multiple projects at a time.

1.1.16.1.3 Agent

You can configure the paths and format of logs to be accessed on the Agent tab.

1.1.16.1.3.1 Add an agent

Before using the Log Management module to collect logs, you must add an agent to configure the paths and format for logs to be accessed.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Agent tab.
- 4. Click Add.

5. On the displayed page, complete the configurations.

Configuration	Description
Product	Select the product whose logs you are about to collect.
Cluster	Select the cluster whose logs you are about to collect .
Service	Select the service whose logs you are about to collect
Service Role	Select the server role whose logs you are about to collect.
Path	The path used to store logs. You can enter at most three paths, separated by commas (,).

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

Configuration	Description
Deploy	Used to determine the collection rules of logs, which corresponds to the log format of the business system. Currently, json and csv are supported.
	Where,
	 json: instance and time are required properties. The property value of instance is the instanceId, and the property value of time is the log time.
	For example,
	<pre>{"instance":"i-uw905d8ny00drzx9****"," memory":"475136.0","write_disk_rate ":"11400.0","tx_intra":"47.0","cpu":" 0.114014113987","write_iops":"1.0"," memory_usage":"0.1616926321","time":" 2018-11-18 00:00:00","rx_intra":"7.0"," flow_intra":"54.0","read_disk_rate":"0.0 ","read_iops":"0.0"}</pre>
	 csv: Commas (,) are used to separate properties and the first property is instanceId.
	For example,
	uw905d8ny00drzx9****,2019-08-15 00:00:07, 15.75.128.85,6405,0,0,15.74.181.5

Des dant	
Product	
Select Product	~
Cluster	
Select Cluster	~
Service	
Select Service	~
Service Role	
Select Service	~
Path(Multiple paths are supported. You can split hree paths are supported.)	the paths by commas (,). A maximum of
Deploy 🕕	
Calaat	

6. Click OK.

Then, the system collects logs according to the configured paths and rules.

1.1.16.1.3.2 Modify an agent

After logs are accessed to an agent, you can modify the paths and format of accessed logs based on business needs.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Agent tab.
- 4. Optional: At the top of the page, select the product, cluster, service, and server role, and then click Search to search for agents that meet the conditions.
- 5. Find the agent to be modified and then click Modify in the Actions column.

Clear Project Agent Bud	ket Management			
Product Cluster Service Select Product V Select Cluster Select 1	Service Role Select Service Role Select Service Role Second	Add		
Product	Cluster	Service	Service Role	Actions
aso	asoCluster-A-20191028-eb04	aso-inventoryMgr	Inventory#	

6. On the displayed page, modify the paths of log collection and the log format, and then click OK.

1.1.16.1.3.3 Delete an agent

You can delete an agent that is no longer in use based on business needs.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Agent tab.
- 4. Optional: At the top of the page, select the product, cluster, service, and server role, and then click Search to search for agents that meet the conditions.
- 5. Find the agent to be deleted and then click Delete in the Actions column.
- 6. Click OK in the displayed dialog box.

1.1.16.1.4 Bucket management

You can configure the information of the backup server which is used to back up logs on the Bucket Management tab.

1.1.16.1.4.1 OSS configurations

You can specify the storage path for log backup by configuring the server information of Object Storage Service (OSS).

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Bucket Management tab.
- 4. Click the OSS Configurations sub-tab.
- 5. In the Actions column, click Modify.
- 6. On the displayed page, modify the configurations.

Configuration	Description
Endpoint	The OSS endpoint. For more information about how to obtain the endpoint, see the OSS Developer Guide .
Bucket	The bucket name of OSS.
AccessKey ID	The username used to access the OSS server, which generally corresponds to the AccessKey ID of OSS. For more information about how to obtain the AccessKey ID, see the OSS Developer Guide .
AccessKey Secret	The key used to access the OSS server, which generally corresponds to the AccessKey Secret of OSS. For more information about how to obtain the AccessKey Secret, see the OSS Developer Guide .
Path	The path on the OSS server, which is used to store the log backup file.

7. Then, click OK.

1.1.16.1.4.2 NAS configurations

You can specify the storage path for log backup by configuring the server information of Network Attached Storage (NAS).

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Log Management > Deploy.
- 3. Click the Bucket Management tab.
- 4. Click the NAS Configurations sub-tab.
- 5. In the Actions column, click Modify.
- 6. On the displayed page, modify the configurations.

Configuration	Description	
Endpoint	The NAS endpoint. For more information about how to obtain the endpoint, see the NAS Developer Guide .	
Path	The path on the NAS server, which is used to store the log backup file.	

7. Then, click OK.

1.1.16.1.4.3 FTP configurations

You can specify the storage path for log backup by configuring the FTP server information.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Deploy
- 3. Click the Bucket Management tab.
- 4. Click the FTP Configurations sub-tab.
- 5. In the Actions column, click Modify.
- 6. On the displayed page, modify the configurations.

Configuration	Description
FTP Domain Name	The access address of the FTP server.
Port Number	The port number used to access the FTP server.
Username	The username used to access the FTP server.
Password	The password used to access the FTP server.
Path	The path on the FTP server, which is used to store the log backup file.

7. Then, click OK.

1.1.16.2 Display logs

You can view logs of the Apsara Stack environment on the Log Display page.

Prerequisites

Before viewing logs, you must make sure that:

- You have configured the relationship among the project, product, and InstanceId. For more information, see *Add a project*.
- You have configured logs to access an agent. For more information, see *Add an agent*.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Display.
- 3. At the top of the page, you can configure the project name, product, index, start time and end time of logs, and log keyword to search for logs that meet the conditions.

After the search, if you cannot view all the log contents in the list, click Details in the Actions column.

1.1.16.3 Log export

The Log Export module allows you to export logs and monitor backup tasks.

1.1.16.3.1 Export logs

You can export or back up logs accessed to Apsara Stack Operations (ASO) to other storage servers. Currently, you can back up logs to an Object Storage Service (OSS) server, Network Attached Storage (NAS) server, or FTP server.

Prerequisites

Before backing up a log file, make sure that:

- You have configured the backup server of OSS, NAS, or FTP. For more information, see *OSS configurations*, *NAS configurations*, and *FTP configurations*.
- Confirm with the deployment personnel that the network of the log management server is connected to the network of the backup server of OSS, NAS, or FTP.

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose Log Management > Log Export.

By default, you are on the Export Log tab.

- 3. Optional: At the top of the page, configure the project name, product, index, and start date and end date of logs, and then click Search to search for logs that meet the conditions.
- 4. Then, you can:
 - Select one or more indexes to be backed up, and then click Back Up to OSS to back up logs to the specified directory of the OSS server.
 - Select one or more indexes to be backed up, and then click Back Up to NAS to back up logs to the specified directory of the NAS server.

Note:

After you click Back Up to NAS, the system backs up logs to the storage path of the server where the log management service is located if the NAS backup server information is not configured in the system.

- Select one or more indexes to be backed up, and then click Back Up to FTP to back up logs to the specified directory of the FTP server.
- Select the index to be exported and then click Download in the Actions column to download the corresponding logs to your local computer.

Result

After the backup, you can view the execution result of the backup task on the Tasks tab of the Log Management > Log Export page.

1.1.16.3.2 View tasks

After backing up logs, you can view the execution result of the backup task on the Tasks tab.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Export.
- 3. Click the Tasks tab.
- 4. Optional: Configure the task status, and the start time and end time of the task, and then click Search to search for the task that meets the conditions.

5. Find the task that you are about to view and then click Index List to view the index names included in the task.

1.1.16.4 Log clearance

The Log Clearance module allows you to clear logs in a specific log file of containers or servers specified in the system.

1.1.16.4.1 Containers

You can obtain the real-time watermark information of containers, and add clearance rules and clear logs in containers in time according to the watermark information.

1.1.16.4.1.1 Obtain the watermark information of one or more containers

You can obtain the watermark information of a container to know the disk usage in the container.

Procedure

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

Contai	ners Servers							
Product Select P	roduct ∨ Select Clus	ter V Service	Service Role	✓ Search	Add	Obtain Watermarks (Clear Logs Import	Export
	Product	Cluster	Service	Service Role	Path	Maximum Disk Usage	Current Disk Usage ↓	Clear Method

- 3. Click the Containers tab.
- 4. Then, you can:
 - At the top of the page, select the product, cluster, service, and server role, and then click Search. In the search results, select the container whose watermark information you are about to obtain. Click Obtain Watermark in the Actions column to obtain the watermark information of a single container.
 - Select multiple containers and then click Obtain Watermarks to obtain the watermark information of multiple containers.

1.1.16.4.1.2 Add a log clearance rule

You can add a clearance rule for a specific log file in the container as needed.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers tab.
- 4. At the top of the page, click Add.
- 5. On the displayed page, complete the configurations.

For more information, see the following table.

Configuration	Description
Product	Select the product to which the container whose logs are to be cleared belongs.
Cluster	Select the cluster to which the container whose logs are to be cleared belongs.
Service	Select the service to which the container whose logs are to be cleared belongs.
Service Role	Select the server role to which the container whose logs are to be cleared belongs.
Path	The path used to store the log files to be cleared. To clear a specific log file, enter the full path name, such as /tmp/test/test.log. To clear all of the log files under a path, you can use the wildcard, such as /tmp /test/*.log.
Maximum Disk Usage	If the actual disk usage exceeds the configured value, the value in the Current Disk Usage column is displayed in red.
Clear Method	 Select the method to clear logs. Delete: Directly deletes the log file. Clear: Clears the log contents without deleting the log file.

6. Click OK.

1.1.16.4.1.3 Modify a log clearance rule

You can adjust a configured log clearance rule in time based on business needs.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers tab.
- 4. Optional: At the top of the page, select the product, cluster, service, and server role, and then click Search to search for clearance rules that meet the conditions.
- 5. Find the clearance rule to be modified and then click Modify in the Actions column.
- 6. On the displayed page, modify the path, maximum disk usage, and clear method.
- 7. Click OK.

1.1.16.4.1.4 Delete a log clearance rule

You can delete a log clearance rule that is no longer in use based on business needs.

Procedure

- 1. Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers tab.
- 4. Optional: At the top of the page, select the product, cluster, service, and server role, and then click Search to search for clearance rules that meet the conditions.
- 5. Find the clearance rule to be deleted and then click Delete in the Actions column.
- 6. In the displayed dialog box, click OK.

1.1.16.4.1.5 Clear container logs

After configuring a log clearance rule, you can clear logs in a container in time based on the watermark information of the container.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers tab.
- 4. Then, you can:
 - At the top of the page, select the product, cluster, service, and server role, and then click Search. In the search results, find the container whose logs are to be cleared. Click Clear in the Actions column to clear logs of a single container.
 - Select multiple containers and then click Clear Logs to clear logs of multiple containers.
1.1.16.4.1.6 View clear records

After clearing logs, you can view the detailed clear records.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers tab.
- 4. Optional: At the top of the page, select the product, cluster, service, and server role, and then click Search to search for clearance rules that meet the conditions.
- 5. Find the clearance rule whose clear records you are about to view. Click View Clear Records in the Actions column to view the detailed clear records.

1.1.16.4.2 Servers

You can obtain the real-time watermark information of servers, and add clearance rules and clear logs in servers in time according to the watermark information.

1.1.16.4.2.1 Obtain the watermark information of one or more servers

You can obtain the watermark information of a server to know the disk usage in the server.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Servers tab.
- 4. Then, you can:
 - At the top of the page, configure the product, cluster, server, and IP address, and then click Search. In the search results, select the server whose watermark information you are about to obtain. Click Obtain Watermark in the Actions column to obtain the watermark information of a single server.
 - Select multiple servers and then click Obtain Watermarks to obtain the watermark information of multiple servers.

1.1.16.4.2.2 Add a log clearance rule

You can add a clearance rule for a specific log file in the server as needed.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Servers tab.
- 4. Click Add.
- 5. On the displayed Add Server Log page, complete the configurations.

For more information, see the following table.

Configuration	Description			
Product	Select the product to which the server whose logs are to be cleared belongs.			
Cluster	Select the cluster to which the server whose logs are to be cleared belongs.			
Server	Select the machine name of the server whose logs are to be cleared.			
IP	The IP address of the server whose logs are to be cleared.			
Path	The path used to store the log files to be cleared. To clear a specific log file, enter the full path name, such as /tmp/test/test.log. To clear all of the log files under a path, you can use the wildcard, such as /tmp /test/*.log.			
Maximum Disk Usage	If the actual disk usage exceeds the configured value, the value in the Current Disk Usage column is displayed in red.			
Clear Method	 Select the method to clear logs. Delete: Directly deletes the log file. Clear: Clears the log contents without deleting the log file. 			

6. Click OK.

1.1.16.4.2.3 Modify a log clearance rule

You can adjust a configured log clearance rule in time based on business needs.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Servers tab.
- 4. Optional: At the top of the page, configure the product, cluster, server, and IP address, and then click Search to search for clearance rules that meet the conditions.
- 5. Find the clearance rule to be modified and then click Modify in the Actions column.
- 6. On the displayed page, modify the path, maximum disk usage, and clear method.
- 7. Click OK.

1.1.16.4.2.4 Delete a log clearance rule

You can delete a log clearance rule that is no longer in use based on business needs.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Servers tab.
- 4. Optional: At the top of the page, configure the product, cluster, server, and IP address, and then click Search to search for clearance rules that meet the conditions.
- 5. Find the clearance rule to be deleted and then click Delete in the Actions column.
- 6. In the displayed dialog box, click OK.

1.1.16.4.2.5 Clear server logs

After configuring a log clearance rule, you can clear logs in a server in time based on the watermark information of the server.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Servers tab.

- 4. Then, you can:
 - At the top of the page, configure the product, cluster, server, and IP address, and then click Search. In the search results, find the server whose logs are to be cleared. Click Clear in the Actions column to clear logs of a single server.
 - Select multiple servers and then click Clear Logs to clear logs of multiple servers.

1.1.16.4.2.6 View clear records

After clearing logs, you can view the detailed clear records.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Servers tab.
- 4. Optional: At the top of the page, configure the product, cluster, server, and IP address, and then click Search to search for clearance rules that meet the conditions.
- 5. Find the clearance rule whose clear records you are about to view. Click View Clear Records in the Actions column to view the detailed clear records.

1.1.16.4.3 Import clearance rules of containers or servers If log clearance rules are configured in your local computer, you can import the clearance rules of multiple containers or servers at a time.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers or Servers tab.
- 4. Click Import.
- 5. Select the .xls or .xlsx file to be imported and then click Open to import multiple log clearance rules at a time.

1.1.16.4.4 Export clearance rules of containers or servers You can export the log clearance rules of multiple containers or servers at a time based on business needs.

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose Log Management > Log Clearance.
- 3. Click the Containers or Servers tab.
- 4. Select the clearance rules of containers or servers to be exported and then click Export.

1.1.17 System Management

1.1.17.1 Overview

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

1.1.17.2 Department management

Department management allows you to create, modify, delete, and search for 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. Department s are displayed in a hierarchy and you can create sub-departments under each level of departments.

Procedure

1. Log on to Apsara Stack Operations.

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. On this page, you can:

· 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 catalog tree 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

(!) Notice:

Before deleting a department, make sure that no user exists in the department. Otherwise, the department cannot be deleted.

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

1.1.17.3 Role management

You can add custom roles in 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 Manager (OAM) system and cannot be modified or deleted by users. The user-created roles can be modified and deleted.

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Roles.

- 3. On the Role Management page, you can:
 - Search for roles

Note:

To search for roles in ASO, you must have the ASO security officer role or system administrator role.

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:

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

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

• Modify a role

Note:

To modify a role in ASO, you must have the ASO security officer role.

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

U Notice:

Before deleting a role, make sure that the role is not bound to any user. Otherwise, the role cannot be deleted.

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

1.1.17.4 Logon policy management

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

Context

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

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Logon Policies.
- 3. On the Logon Policy Management page, you can:
 - · 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 IP addresses prohibited for logon. 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.

1.1.17.5 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

User management provides different 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 auditor officer. The permissions of these three roles are as follows:

Inotice:

To guarantee the system security, you must modify the password of these three default users as soon as possible.

- The system administrator can view, modify, delete, and add the information in the Operations and Maintenance Dashboard, Alert Monitoring, Resource Management, Inventory Management, Configurations, Offline Backup, Help Center, and Application Whitelist modules, and view the users, roles, departments, logon policies, and server passwords in the System Management module.
- The security officer can view, modify, delete, and add the users, roles, departments, logon policies, and server passwords in the System Management module.
- The auditor officer can read and write Apsara Stack Operations (ASO) system logs

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Users. Click the Users tab.
- 3. On the Users tab, you can:
 - · Search for users



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

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

• Add a user



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

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.

The added user is displayed in the user list. The Primary Key Value of the user is used to call the application API. In other words, the primary key value is used for authentication if other applications need to call the applications in ASO.

• Modify a user



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

Find the user to be modified, and then click Modify in the Actions column. In the 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.



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 in next to the logon username and then select Personal Information. The appeared Personal Information dialog box displays the personal information of the current user.

· Add a custom logo

In the upper-right corner, click in next to the logon username and then select Logo Settings. In the displayed Custom Settings dialog box, click to upload the

custom system logo image and system name image and then click Upload.

Logon settings

In the upper-right corner, click 🔽 next to the logon username and then select

Logon Settings. In the displayed Logon Settings dialog box, configure the logon timeout, multiple-terminal logon settings, maximum allowed password retries, account validity, and logon policy. Then, click Save.

1.1.17.6 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

This authentication method uses the password and mobile phone to provide double protection for accounts. You can obtain the logon key after configurin g 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 instructio ns if you select this authentication method. The third-party manufacturer provides the USB key hardware and the service that the backend authenticates and verifies the certificates. The USB key hardware 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 thirdparty 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 the 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 authentication 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 card number of a user. Therefore, bind the name and ID card number with a user account when you configure the authentication method in ASO.

Both USB key authentication and PKI authentication depend on the authentication server provided by the third-party manufacturer to verify the encrypted informatio n 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 Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Two Factor Authentication.
- 3. On the Two Factor Authentication page, you can:
 - 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 are about to enable the Google two-factor authentication, and then click Create Key in the Actions column. After the key is created, you can click Show Key to display the key 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:

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.
- USB key authentication
 - a. Select USB Key Authentication as the Current Authentication Method.
 - b. In the Authentication Server Configuration section, 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 the server list. Click Test to test the connectivity of the authentication server.

- c. In the User List section, click Add User. The added user is displayed in the user list.
- d. Find the user that you are about 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 controls to automatically enter the serial number. If the serial number fails to be entered, you must enter it manually. The serial number of USB key authentication is written in the USB key hardware. Therefore, you must insert the USB key, install the drive and browser controls, and then read the serial number by calling the browser controls.

- e. Then, click Enable Authentication in the Actions column.
- PKI authentication
 - a. Select PKI Authentication as the Current Authentication Method.
 - b. In the Authentication Server Configuration section, 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 the server list. Click Test to test the connectivity of the authentication server.
 - c. In the User List section, 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 are about to enable the PKI authentication, and then click Bind in the Actions column. Enter the full name and ID card number of the user to bind the user account with the name and ID card number.
 - e. Then, click Enable Authentication in the Actions column.
- No authentication

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

1.1.17.7 Application whitelist

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

Context

All the Apsara Stack Operations (ASO) services are accessed based on Operation Access Manager (OAM) permission management. Therefore, if your account does not have the corresponding role, your access requests are rejected. The applicatio n 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 permission s consist of read-only and read/write. The configured value is the logon user permission.

The application whitelist is managed by the system administrator. You can access this page after logging on as a system 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 Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Application Whitelist.
- 3. On the Application Whitelist page, you can:
 - 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

In the Permission drop-down list, modify the permission of the service to Read/Write or Read-only.

• Delete a whitelist

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

1.1.17.8 Server password management

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

Context

Server password management 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 passwords 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, or IP address.

Procedure

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Server Password.

The Password Management tab displays the passwords of all the servers in the Apsara Stack environment.

3. On this tab, you can:

• Search for servers

On the Password Management tab, configure the product, hostname, or IP address, and then click Search to search for specific servers.

- Show passwords
 - a. On the Password Management tab, find a server.
 - b. Click Show in the Password column, and then the system displays the host password in plain text, which turns into cipher text after 10 seconds. Alternatively, directly click Hide to display the cipher text.
- Update passwords
 - a. On the Password Management tab, find a server.
 - b. Click Update Password in the Actions column.
 - c. In the displayed Update Password dialog box, enter the Password and Confirm Password, and then click OK.

Then, the server password is updated.

- Update multiple passwords at a time
 - a. On the Password Management tab, select multiple servers.
 - b. Click Batch Update.
 - c. Enter the Password and Confirm Password, and then click OK.

Then, the passwords of the selected servers are updated.

- Configure the password expiration period
 - a. On the Password Management tab, select one or more servers.
 - b. Click Configuration.
 - c. 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 after the configuration and will be updated again after an expiration period.

 $\cdot \,$ View the history of server password updates

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

- Show history passwords of servers
 - a. On the History Password tab, find a server.
 - b. Click Show in the Password column, and then the system displays the host password in plain text, which turns into cipher text after 10 seconds. Alternatively, directly click Hide to display the cipher text.
- View and modify the password configuration policy

Click the Configuration tab. View the metadata, including 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 automatically updated by 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 dialog box, enter the Initial Password, Password Length, and Retry Times, and then click OK.

1.1.17.9 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

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

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose System Management > Operation Logs.

- 3. On the Log Management page, you can:
 - 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 and then click OK in the displayed dialog box.

Export logs

Click to export the logs of the current page.

1.1.17.10 View the authorization information

The Authorization page allows customers, field engineers, or operations engineers to quickly view the service with an authorization problem and then troubleshoot the problem.

Prerequisites

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

If you are not an administrator-level user, a message indicating that you do not have sufficient permissions is displayed when you access this page.

Procedure

1. Log on to Apsara Stack Operations.

- 2. In the left-side navigation pane, choose System Management > Authorization.
- 3. View the basic information and specifications on the Authorization Details page.

Note:

For formal authorization, you must enter the authorization code to view the authorization information. Obtain the authorization code in the authorization

letter attached by the project contract or contact the business manager (CBM) of your project to obtain the authorization code.

You can view the authorization information, including the authorization version , customer information, authorization type, Elastic Compute Service (ECS) instance ID, cloud platform version, authorization creation time, and the authorization information and authorization specifications of a service, in the current Apsara Stack environment on this page.

See the detailed authorization information and the corresponding description in the following table.

Authorization informatio n	Description
Authorization Version	 You can use the BP number in the version to associate with a project or contract. Where, TRIAL in the version indicates that the authorization is a trial one. The trial authorization is valid within 90 days from the date of deployment. FORMAL in the version indicates that the authorization is a formal one. The authorization information of the service comes from the signed contract.
Authorization Type	Indicates the current authorization type and authorization status.
Customer information	Includes the customer name, customer ID, and customer user ID.
ECS Instance ID	The ECS instance ID in the Deployment Planner of the field environment.
Cloud Platform Version	The Apsara Stack version of the current cloud platform.
Authorization Created At	The start time of the authorization.

Authorization informatio n	Description
Authorization informatio n of a service	 Includes the service name, service content, authorization mode, service authorizations, subscription start time, subscription expiration time , and real-time authorization status. If the following information appears in the Authorization Status column of a service: RENEW Service Expired Indicates that the customer must renew the subscription as soon as possible. Otherwise, the field operations services, including ticket processing, are to be terminated. Specification Quota Exceeded Indicates that the specifications deployed in the field for a service have exceeded the quota signed in the contract, and the customer must scale up the service as soon as possible.
Authorization specificat ions of a service	Includes the service name, and the name, unit , current number, and ceiling of authorization specifications.

1.1.17.11 Menu settings

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

1.1.17.11.1 Add a level-1 menu

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

- **1.** Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, choose System Management > Menu Settings.
- 3. Click Add.

4. On the displayed page, complete the configurations for the level-1 menu you are about to add.

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

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

Configuration	Description
Deletable	Whether this level-1 menu can be deleted after being added. Turn on or off the switch to configure whether the menu can be deleted. By default, the menu can be deleted. The setting cannot be modified after being configured.

Add Level-1 Men	u)
Menu Icon		_
	~	
• Menu Name		1
Menu Name(繁體)]
Menu Name(English)]
• Menu Order		
	1 -	
• Show	 Deletable(Yes) 	

5. Click OK.

Result

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

1.1.17.11.2 Add a submenu

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

Procedure

1. Log on to Apsara Stack Operations.

2. In the left-side navigation pane, choose System Management > Menu Settings.

3. Add a level-2 menu

- a) Find the level-1 menu to which you are about to add a level-2 menu, and then click Add in the Actions column.
- b) On the displayed page, complete the configurations for the submenu you are about to add.

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

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

Operations and Maintenance Guide - Cloud Essentials and Security / 1 Operations of basic platforms

Configuration	Description	
Parent Menu	The parent menu of this menu.	
Add Submenu	×	
• Menu Name		
Menu Name(繁體)		
Menu Name(English)		
L		
Menu Order		
 Show 	Deletable(Yes)	
Link Address		
	0	
Parent Menu		

c) Click OK.

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

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



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

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

1.1.17.11.3 Hide a menu

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

Prerequisites

I) Notice:

You cannot hide the System Management menu and its submenus.

Procedure

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

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

• Hide a level-2 or level-3 menu

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

1.1.17.11.4 Modify a menu

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

Procedure

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

1.1.17.11.5 Delete a menu

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

Prerequisites

I Notice:

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

Procedure

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

1.2 Apsara Stack Doctor (ASD)

1.2.1 Apsara Stack Doctor introduction

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

Features

- Provides data filtering, analysis, and processing for O&M data consumers.
- Provides encapsulation, orchestration, and permission 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 by using web-based 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, as shown in Figure 1-2: Levels of

release granularity.

Figure 1-2: 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 to form 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. 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)

A process-level service component contained by a server role. Applications can be grouped into three types:

- 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 file.

1.2.2 Log on to Apsara Stack Doctor

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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

Log On	
<u> </u>	Enter a user name
ß	Enter the password
	Log On



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 log on to ASO.
- 5. In the left-side navigation pane, click Products.
- 6. In the Basic O&M region, click ASD.

1.2.3 ASA

Apsara Stack Assistant (ASA) is a tool provided to help you improve the efficiency in testing, operating, maintaining, and releasing cloud products in Apsara Stack while ensuring version stability. ASA has also retained the inspection, scanning, and version tracking capabilities of Apsara Stack V2.

1.2.3.1 RPM Check

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

Procedure

1. Log on to Apsara Stack Doctor.

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

Host	Status
test_tianji_machine180	unavailable
ecsapigatewaylitetageb0a	unavailable
a36f04114.cloud.f05.amtest61	normal
vm010148064142	normal
vm010148064143	normal
vm010148064141	normal
vm010148064146	normal
a36f07206.cloud.f09.amtest61	normal
vm010148064026	normal
vm010148064023	normal

Table 1-1: Description of parameters on the RPM Check page

Parameter	Description
Host	The name of a host.
Status	 The status of a machine. Valid values: normal: indicates that the machine is operating normally. unavailable: indicates that the machine is not operating normally or unavailable.

1.2.3.2 Virtual IP Check

The Virtual IP Check module allows you to obtain the virtual IP addresses that are incorrectly bound to IP addresses of backend services.

Procedure

1. Log on to Apsara Stack Doctor.

2. In the left-side navigation pane, choose ASA > Virtual IP Check.

	Virtual IP Address	Virtual Port	Port	Backend IP Address	Cluster	Service	Server Role	Status
		9090	21069		ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal
		9090	21069		ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal
		9090	21069		ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal
ĺ		9090	21069		ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal

Table 1-2: Parameters on the Virtual IP Check page

r
The virtual IP address.
The port corresponding to a virtual IP address.
The port corresponding to the IP address of a backend service.
The IP address of a backend service.
The cluster to which the IP address of a backend service belongs.
The service to which the IP address of a backend service belongs.
The server role to which the IP address of a backend service belongs.
The health status, indicating whether the binding between the virtual IP address and the IP address of the backend service is normal.
 address is correctly bound to the IP address of the backend service. abnormal: indicates that the virtual IP address is not bound to the

1.2.3.3 Volume Check

The Volume Check module allows you to view the volume details of Docker hosts.

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

Container ID	Container Name	Host IP Address	Path	Disk Quota	Total Partition Space	Partition Space Used	Directory Space Used
2edb931eb098	bcc- api.Controllercontroller.1558621857		/opt/backup_minirds	{"/":"40g"}	20G	1.1G	4.0K
2edb931eb098	bcc- api.Controllercontroller.1558621857		/apsarapangu/disk8	{"/":"40g"}	45G	5.3G	4.0K
2edb931eb098	bcc- api.Controllercontroller.1558621857		/apsarapangu	{"/":"40g"}	45G	5.3G	16K

Table 1-3: Parameters on the Volume Check page

Parameter	Description
Container ID	The unique ID of a Docker container.
Container Name	The name of a Docker container.
Host IP Address	The IP address of a Docker host. Typically, a Docker virtual machine can be either a physical host or virtual host.
Path	The disk partition mount point of a Docker volume.
Disk Quota	The quota of a disk.
Total Partition Space	The total space of a mount point calculated by running the df command.
Partition Space Used	The space used by a mount point directory.
Directory Space Used	The total space of a mount point calculated by running the du command.

1.2.3.4 NTP Check

The NTP Check module allows you to check whether the system time of all machines, including Docker virtual machines and physical machines, is synchronized with the NTP time. If not, the time offset is reported in milliseconds.

1. Log on to Apsara Stack Doctor.

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

Host	Time Offset
a36f04114.cloud.f05.amtest61	0
vm010148064142	0
vm010148064143	0
vm010148064141	0
vm010148064146	0

Table 1-4: Parameters on the NTP Check page

Parameter	Description
Host	The name of a host.
Time Offset	The time offset. Unit: milliseconds.

1.2.3.5 IP Conflict Check

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

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose ASA > IP Conflict Check.

		Virtual Host	Туре	Server Role	Physical Host	IP

Table 1-5: Parameters or	the IP Conflic	t Check page
--------------------------	----------------	--------------

Parameter	Description
IP	A conflicting IP address.
Physical Host	The name of the physical host with the conflicting IP address.

Parameter	Description
Server Role	The server role that requests the resource.
Туре	The IP address type. Valid values: docker, vm, and physical.
Virtual Host	The hostname of the Docker virtual machine.

1.2.3.6 DNS Check

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.

Domain	Virtual IP Address	Owner	IP

Table 1-6: 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 that is bound to the domain name.
1.2.3.7 IP Details

The IP Details module allows you to check the details of all IP addresses in the current environment, including the IP addresses of physical machines, Docker machines, and virtual machines, as well as virtual IP addresses.

Procedure

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

IP	Virtual Host	Туре	Physical Host	Server Role
		vip		Server Role Information
		vip		Server Role Information
		vip		Server Role Information
		vip		Server Role Information
		vip		Server Role Information

Table 1-7: Parameters on the IP Details page

Parameter	Description	
IP	The IP address of a resource.	
Virtual HostThe name of a virtual machine.		
Туре	The resource type. Valid values: • physical • docker • vm	
Physical Host	The name of a physical host.	
Server Role	The server role that requests the resource.	

3. Move the pointer over Server Role Information in the Server Role column to view server role details.

1.2.3.8 Quota Check

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

1. Log on to Apsara Stack Doctor.

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

Memory	CPU Disk			
Container ID	Container Name	Container Memory	Hostname	Host Memory
cb88159341f2a	dtdream-dtcenter.Uimuim.1559285092	4294967296	a36f04015.cloud.f04.amtest61	540732784640
c86de87d8d79c	vm010148065213	8643411968	a36f04015.cloud.f04.amtest61	540732784640
3eeee420a444c	asrbr-heimdallr.Heimdallrheimdallr.1559108650	4294967296	a36f04015.cloud.f04.amtest61	540732784640
773a7a37a2f71	drds-console.DrdsManagerdrds-manager.1558419453	8589934592	a36f04015.cloud.f04.amtest61	540732784640

- 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.

1.2.3.9 Error Diagnostics

Context

The Error Diagnostics page consists of the following tabs:

- Resource Errors: displays resource errors.
- Error with Self: displays internal errors.
- Error with Dependency: displays dependency errors.
- Normal: displays resources with no errors.

- 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.

1.2.3.10 Versions

The Versions module allows you to obtain version information and upgrade information of all services 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 to view information related to service versions, such as the IDC, service, and version.
 - Click the Server Role Versions tab to view information related to server role versions, such as the IDC, service, version, server role, and type.
 - Click the Version Tree tab to view information related to version trees.

1.2.4 Support tools

1.2.4.1 Diagnose with the OS tool

The OS tool allows you to perform OS diagnostics on physical machines in Apsara Stack.

Context

Metrics that can be diagnosed by the OS tool include: disk file metadata usage , memory usage, process status, time synchronization, kernel faults, high-risk operations, system loads, fstab files, read-only file systems, kdump services, kdump configurations, conman configurations, domain name resolution, disk I/O loads , 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.

Search	by physical machine name	Search Get Phy	sical Machine List		Run Diagnostic Script
	Physical Machine Name	Health Score	Host Address	Script Execution Status	Actions
	a36f01207.cloud.f03.amtest95			Not Executed	
	a36f04106.cloud.f05.amtest95			Not Executed	
	a36f01161.cloud.f02.amtest95			Not Executed	
	a36f12006.cloud.f12.amtest95			Not Executed	
	a36f01103.cloud.f02.amtest95			Not Executed	

- 3. Click Get Physical Machine List to obtain a list of all the physical machines in the system.
- 4. Optional: In the search bar, enter the name of a physical machine and click Search. This physical machine is displayed in the section below the search bar.
- 5. Select the physical machine and click Run Diagnostic Script in the upper-right corner.
- 6. When Script Execution Status changes from Not Executed to Diagnostic Result Decompression Finished, you can view the health score of the physical machine in the Health Score column.
- 7. After the diagnostics are complete, click View Report in the Actions column to view the diagnostic result.
- 8. Optional: For more information, click Download Report in the Actions column.

1.2.4.2 Use Support Tools

Support Tools allows you to diagnose some services and export diagnostic reports.

Procedure

1. Log on to Apsara Stack Doctor.

2. In the left-side navigation pane, choose Support Tools > Support Tools.

3. Optional: Select the target service, enter the host name or IP address, and click Search. The search results appear in the section below.

Diagnostic item	Description
Apsara Distributed File System Diagnostics	Collects and analyzes the running status of Apsara Distributed File System and its dependent services and environments, and provides diagnostic reports in case of exceptions.
ecs_vmdisk_usage_V3	Checks the ECS disk usage.
oss_used_summary	Checks the usage of OSS resources.
ots_examine	 Checks the following information: NTP Consistency of the Table Store versions Chunkserver status of Apsara Distributed File System Status of Apsara Name Service and Distributed Lock Synchronization System SQL status SQL partition and distribution Service availability of DNS Service availability of SLB Service availability of RDS Service availability of OTS Cluster Management (OCM) Service availability of Red Hat Package Manager (RPM) databases
ecs_error_log	Collects ECS logs.
ots_used_summary	Checks the usage of Table Store resources.
docker	Collects and analyzes data from Docker hosts, and generates reports based on the data.
ecs_diagnostor_v3	Collect the logs of end-to-end ECS links.
OS	 Collects and analyzes system logs, including the following operations: Collects information about the OS, network, disk, and hardware. Diagnoses and analyze system logs. Generates reports.

The following table lists the supported diagnostic items.

Diagnostic item	Description
oss_examine	Diagnoses OSS.

4. Find the row that contains the target machine and click Run Diagnostics in the Actions column corresponding to the target machine.



Alternatively, you can select the target service and click Search. In the search results, select multiple machines and click Run Diagnostics for batch diagnostics.

When Diagnostics Execution Status changes from Running to Succeeded, the diagnostics are completed.

Product:	pangu v Se	earch by hostname or IP add	Iress. 🛞 Sea	Run Diagnostics		Version: beta20190513 Upload File
	HostName	ClusterName	IP Address	Diagnostics Execution Status	Executed At	Actions
	a36f04013.cloud.f04.amtest61	ECS-IO7River-A-6ffe			May 22, 2019, 15:49:49	
	a36f04011.cloud.f04.amtest61	ECS-IO7River-A-6ffe			May 22, 2019, 15:49:49	
	a36f01109.cloud.f02.amtest61	ECS-IO7River-A-6ffe			May 22, 2019, 15:49:49	
	a36f04210.cloud.f06.amtest61	ECS-IO7River-A-6ffe			May 22, 2019, 15:49:48	

- 5. After the diagnostics are complete, click View Report in the Actions column to view the diagnostic result.
- 6. Optional: After the diagnostics are complete, click Download Report in the Actions column to download the diagnostic results to your local machine.

1.2.4.3 Update Support Tools

When the Support Tools toolkit has updates, you can update it to the latest version by uploading files.

- **1.** Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Support Tools > Support Tools.
- 3. In the upper-right corner of the page, click Upload File.

4. Select the toolkit file to upload, enter the verification code, and click Upload File. Contact level-2 support engineers to obtain the verification code.

	Upload File	Х		
File: J. Click to	select the file			
Verification code: Enter your verification code.				
	ation code.			
	Upload File			

1.2.4.4 Diagnose with the inspection tool

You can use the inspection tool to diagnose and inspect products such as Apsara File Storage NAS, Block Storage, and Apsara Name Service and Distributed Lock Synchronization System.

- **1.** Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Support Tools > Inspection Tool.
- 3. Click Get Gateway List to obtain all gateways in the system.

4. Select the target product from the Product drop-down list and click Search. The search result appears in the section below.

Apsara Stack Doctor (ASD) supports diagnostics for products such as Apsara File Storage NAS, Block Storage, and Apsara Name Service and Distributed Lock Synchronization System.

• Diagnostics of Apsara File Storage NAS

ASD allows you to collect Apsara File Storage NAS information, including disk status, key-value (KV) status, KV server spacing, version, recycle bin, memory , and TCP.

• Diagnostics of Block Storage

ASD allows you to collect the utilization information about storage clusters.

• Diagnostics of Apsara Name Service and Distributed Lock Synchronization System

ASD allows you to check the following information about this product:

- The health status of the end-to-end service link.
- The disk space of the product.
- Whether the nuwazk log is properly stored.
- Whether the nuwaproxy log is properly stored.
- 5. You can select multiple machines and click Run Diagnostics to perform batch diagnosis. Alternatively, you can select only one machine and click Run Diagnostics in the Actions column corresponding to the machine.

Product:	nas	✓ Search	Get Gateway List	Run Diagnostics	5	
	Admin Gateway	IP	Diagnostics Execution Status		Executed At	Actions
	vm010148128107		Tunnel Error			

6. After the diagnostics are complete, you can click Download Inspection Log in the Actions column corresponding to the machine to download the diagnostic results to your local machine.

1.2.4.5 Upload script files for EDAS diagnostics

Before the diagnostics, you can unload script files to be executed for server roles.

Procedure

1. Log on to Apsara Stack Doctor.

- 2. In the left-side navigation pane, choose Support Tools > EDAS Diagnostics.
- 3. In the upper-right corner of the page, click Upload Diagnostic Script.
- 4. Select the product, service, and server role.

If the server role has script files, the script files will be displayed in the Existing Scripts field. You can click the name of a script file to view details.

	Uplo	ad Diagnostic S	Script		
Product:	edas	~	ļ		
service ·	odae hef	~	1		
Service.	Cuas-IISI		ļ		
Server Role :	Hsflnit	~	J		
Existing Scripts:					
Script File:	上 Choose File				
		Upload Diagno	ostic Script		

- 5. Click Choose File. In the dialog box that appears, select the script file to be uploaded. Click Open to add the script file to be uploaded.
- 6. Click Upload Diagnostic Script.

1.2.4.6 EDAS diagnostics

The EDAS diagnostics tool allows you to inspect EDAS.

Prerequisites

Before the diagnosis, make sure that the server role to be diagnosed has an executable script file. If not, you need to upload the script file to be executed for the server role. For more information about how to upload the script file, see *Upload script files for EDAS diagnostics*.

- **1.** Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Support Tools > EDAS Diagnostics.

- 3. Optional: Select one or more services from the Service drop-down list and click Refresh. The filtered services appear in the section below.
- 4. Find the server role to diagnose, and click Run Diagnostics in the Actions column corresponding to the server role.



You can select multiple server roles at a time from the filtered services and click Run Diagnostics. In the dialog box that appears, click OK to run diagnostics.

When Diagnostic Status changes from Diagnosing to Diagnostics Succeeded, the tasks are completed.

Product:	edas v	Service: Select an ite	m. V Refresh	Run Diagnostics	Upload Diagnostic Script
	Service	Server Role	Diagnostic Status	Cause of Failure	Actions
	edas-edasService	EdasServer			
	edas-edasService	CaiFs		No configuration snapshot json	
	edas-edasService	EagleeyeConsole	Not Run		
	edas-edasService	EdasEam	Not Run		Run Diagnostics Download Report

5. After the tasks are completed, you can click Download Report in the Actions column corresponding to the server role to download the original diagnostic information.

1.2.5 Service Availability

1.2.5.1 View Service Availability

Service Availability allows you to view the availability statuses of cloud services in Apsara Stack.

Context

It is used to verify the continuity of these cloud services.

During the hot upgrade of a service, you can use Service Availability to check whether the upgrade causes a service interruption, helping you detect and solve problems in a timely manner.

Procedure

1. Log on to Apsara Stack Doctor.

2. In the left-side navigation pane, choose Service Availability > Service Availability.

3. In the search bar, select the service you want to view and click Search to view its service status.

Service status	Description
Pending	The service availability inspection is not enabled for this service.
UNKNOW	The service availability status of the service is unknown.
ERROR	The service availability status of the service is abnormal.
ОК	The service availability status of the service is normal.

The following table describes the service statuses.

Service A	Vailability						
Product:	All		Search				
Produc				Product Status	Exception Message	Checkpoint Time	

1.2.5.2 View Control Service Availability

The Control Service Availability page displays the statistics of the global environment, product response times, and product QPS.

Procedure

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Service Availability > Control Service Availability.
- 3. View the following information:
 - Global statistics

Global Statistics displays the environment information of all control gateways, including global queries per second (QPS), global response time statistics, and error details.

On the Global Statistics tab, select Last 1 Hour, Last 2 Hours, Last 24 Hours, or Select Time from the Time drop-down list and select HTTP status code. Click Update to view the information of the global environment within the specified time range.

HTTP status code	Description
200	The request is successful. It is generally used for GET and POST requests.
400	The syntax of the request from the client is incorrect, which cannot be understood by the server.
403	The server understands the request from the client but refuses to execute it.
404	The server cannot find the resource based on the request from the client.
500	The request cannot be completed because the server has an internal error.
503	The server is temporarily unable to process the request from the client.
201	Created. The request is successful, and a new resource is created.
204	No content. The server has processed the request but does not return any content.
409	A conflict occurs when the server processes the request.
202	Accepted. The request has been accepted but has not been processed.
405	The method specified in the request from the client is forbidden.

The following table describes the HTTP status codes.

Product response time statistics

Product Response Time Statistics displays the latency of each service from a specified period of time. You can view product response time statistics to identify whether exceptions have occurred in a service API based on the number of responses within a specified period of time.

On the Product Response Time Statistics tab, select Last 1 Hour, Last 2 Hours, Last 24 Hours, or Select Time from the Time drop-down list, Product to be queried, and HTTP Status Code. Click Update to view the average latency of a service within a specified period of time.

Product QPS statistics

Product QPS statistics displays the requests of each service within a specified period of time. You can view product QPS statistics to identify whether exceptions have occurred in the service status based on the number of requests within a specified period of time.

On the Product QPS Statistics tab, select Last 1 Hour, Last 2 Hours, Last 24 Hours, or Select Time from the Time drop-down list, Product to be queried, and HTTP Status Code. Click Update to view the latency of a service from a specified period of time.

1.2.6 Monitoring

The Monitoring module allows you to view alert templates, alerts, and alert status in the system.

1.2.6.1 View alert templates

Alert templates are used to configure alert monitoring settings. You can filter alert template content by service and template.

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Monitoring > Monitoring Templates.
- 3. In the search bar, select a service and a template, and click Search.
- 4. View the alert template content in the search result.



1.2.6.2 View alert information

During routine O&M, you can view alert information to obtain up-to-date information about services. When a service fails, you can filter out the alert information that you need based on the service, cluster name, and alert name to quickly resolve the failure.

Procedure

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Monitoring > Alerts.
- 3. Perform the following operations:
 - To view all alerts in the system, click Search without selecting any filters.

Service:	Select an iter	n		Cluster:	Enter a cluster na	me	\otimes	Alert:	Enter	an alert nam	e	\otimes	Search
A	lert Name	Metric		Alert Rule	Monitoring Dimension	Subject		Data Sourc	e	Enabled	Monitor Interval	Ale	rt Status
private.0 2019102 vice_tjm able_ala _ba	C_dnsCluster-A- 21-165c_dnsSer on-service_avail rmdnsService se-template	dnsService_tim on-check_servic e_available_clu ster_serverrole	{"name" on-servi rm","terr	:"dnsService_tjm ce_available_ala nplate":"base-terr plate"}	["{ \'cluster \":\"\$\$CLUSTER \$\$\", \"serverrole \":\"dnsService.b indSererRole#\" }"]	private. service not available ala rm	sou proj	urceType:ME ect:tjm_dnsS	TRIC	true	60	INSUFFI	CIENT_DATA
private.(2019102 vice_time e_alarm as	C_dnsCluster-A- 21-165c_dnsSer on-max_open_fil dnsService_b e-template	dnsService_tim on-check_log_k eyword_max_o pen_file_serverr ole	("name" on-max m","tem	:"dnsService_tjm <_open_file_alar plate":"base-tem plate"}	["{ \"cluster \":\"\$\$CLUSTER \$\$\", \"serverrole \":\"dnsService.b indServerRole# \" \"]	private. 文件打 开数到达最大值 Alarm-02.659.0 002.00006	sou proj	urceType:ME ect:tjm_dnsS	TRIC	true	60	INSUFFI	CIENT_DATA

• In the search bar, select a service, enter a cluster name and an alert name, and then click Search to view information about an alert.

1.2.6.3 View the alert status

After alerts are triggered, you can view the status of all alerts in the system.

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Monitoring > Alert Status.

- 3. Perform the following operations:
 - To view the status of all alerts in the system, click Search without selecting any filters.

Service: Select an item V	Cluster: Enter a cluste	er name 🛞 Alert:	Enter an alert name	Alert Status: Select a status:	atus V Period: Star	t Time End Time	Search
Alert Name	Status Last Updated At	Last Alert Time	Server Role	First Alert Time	Alert Rule	Monitoring Dimension	Alert Level
private.testimage_monitor_alarm_tia nji_base-template	Nov 29, 2019, 11:47:48	Nov 29, 2019, 10:37:10	drds-console.ServiceTes t#	Nov 27, 2019, 10:35:39	{"name":"testimage_monitor_alarm","t emplate":"base-lemplate"}	serverrole=drds-console.ServiceTest#, machine=vm010148065201,level=erro r	
private.testimage_monitor_alarmtia nji_base-template	Nov 29, 2019, 11:47:48	Nov 29, 2019, 10:37:10	gpdb-yaochi.ServiceTest #	Nov 27, 2019, 10:35:39	{"name":"testimage_monitor_alarm","t emplate":"base-template"}	serverrole=gpdb-yaochi.ServiceTest#, machine=vm010148065163,level=erro r	

• In the search bar, select a service, enter a cluster name and an alert name, and select a status and a time range. Then, click Search to view the status of an alert.

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.





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 has two types of subjects: users and groups.

user

Administrators and operators of operations systems.

group

A collection of users.

role

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

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

RoleHierarchy

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

RoleCell

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

resource

The description of an authorized object. For more information about resources of operations platforms, see *Permission lists of operations platforms*.

ActionSet

The description of authorized actions. An ActionSet can contain multiple actions. For more information about actions of operations platforms, see *Permission lists of operations platforms*.

1 1 5

available authorizations

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

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

Note:

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

1.3.3 Quick start

This topic describes how to add and assign roles quickly.

1.3.3.1 Log on to OAM

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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

Log On					
<u>8</u>	Enter a user name				
£	Enter the password				
Log On					



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 log on to ASO.
- 5. In the left-side navigation pane, select Products.
- 6. Click OAM under Apsara Stack O&M.

1.3.3.2 Create a group

Create a user group for centralized management.

Procedure

- 1. Log on to 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 Description.
- 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.

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group and then click Manage in the Actions column.
- 4. Click Add Member in the Group Member section.

5. Select the search mode, enter the corresponding information, and then click Details. The user details are displayed.

Three search modes are available:

- **RAM User Account: Search for the user in the format of** *RAM username@primary* account ID.
- Account Primary Key: Search for the user by using the unique ID of the user's cloud account.
- Logon Account Name: 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 in the Actions column at the right of the member.

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*.
- \cdot You are the owner of the group and the role.

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group and then click Manage in the Actions 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 remove a role from the group, click Remove in the Actions column at the right of the role in the Role List section.

1.3.3.5 Create a role

Procedure

- **1.** Log on to 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 Description, 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 owned roles, see Search for roles.

Procedure

1. Log on to OAM.

- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role and then click Manage in the Actions column.
- 4. Click the Inherited Role tab.
- 5. Click Add Role. 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.

1. Log on to OAM.

- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role and then click Manage in the Actions column.
- 4. Click the Resource List tab.
- 5. Click Add Resource.
- 6. Complete the configurations. For more information, see *Table 1-8: Configurations*.

Configuration item	Description
BID	The deployment region ID.
Product	The cloud product to be added, for example, rds.
	Note: The cloud product name must be lowercase. For example, enter rds, instead of RDS.
Resource Path	For more information about resources of cloud products and operations platforms, see <i>Permission lists of operations</i> <i>platforms</i> .
Actions	An ActionSet, which can contain multiple actions. For more information about actions of operations platforms, see <i>Permission lists of operations platforms</i> .
Available Authorizations	The maximum number of authorizations in the cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.
Description	The description of the resource.

Table 1-8: Configurations

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 OAM.
- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role and then click Manage in the Actions column.
- 4. Click the Authorized Users tab.
- 5. Click Add User.
- 6. Select the search mode and enter the corresponding information.

Four search modes are available:

- RAM User Account: search in the format of RAM username@primary account ID.
- Account Primary Key: search by using the unique ID of the user's cloud account.
- Logon Account Name: 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 Authorized Users 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 Authorized Users tab.

1.3.4 Manage groups

Group Management allows you to view, modify, or delete groups.

1.3.4.1 Modify the group information

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

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group and then click Manage in the Actions column.
- 4. Click Modify in the upper-right corner.
- 5. In the displayed Modify Group dialog box, modify the Group Name and Description.
- 6. Click Confirm.

1.3.4.2 View group role details

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

Prerequisites

A role is added to the group.

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group and then click Manage in the Actions column.
- 4. In the Role List section, click Details at the right of a role.

- 5. On the Role Information page, you can:
 - Click the Inherited Role tab to view the information about the inherited roles.
 To view the detailed information of an inherited role, click Details in the

Actions column 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 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 Actions column.

1.3.4.4 View authorized groups

You can view the groups to which you are added on the Authorized Groups page.

Context

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

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Authorized Groups.
- 3. On the Authorized Groups page, view the name, owner, description, and modified time of the group to which you belong.

1.3.5 Manage roles

Role Management allows you to view, modify, transfer, or delete roles.

1.3.5.1 Search for roles

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

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Optional: Enter the role name.
- 4. Click Search to search for roles that meet the search condition.



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 OAM.

- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role and then click Manage in the Actions column.
- 4. Click Modify in the upper-right corner.
- 5. In the displayed Modify Role dialog box, modify the Role Name, Description, 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 OAM.

- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role and then click Manage in the Actions 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 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 Details. The user details or group details are displayed.

Four search modes are available:

- RAM User Account: search in the format of RAM username@primary account ID.
- Account Primary Key: search by using the unique ID of the user's cloud account.
- Logon Account Name: 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 OAM.

- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. At the right of the role to be deleted and then click Delete.

1.3.5.6 View authorized roles

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

Procedure

1. Log on to OAM.

- 2. In the left-side navigation pane, choose Role Management > Authorized Roles.
- 3. On the Authorized Roles page, you can view the name, owner, description, modified time, and expiration time of the role assigned to you.

Click Details at the right of a role to view the inherited roles, resources, and inheritance tree information of the 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 a role, click Details to view the inherited roles, resources, and inheritance tree information of the 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 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 a role, click Details in the Actions column to view the inherited roles, resources, and inheritance tree information of the role.

1.3.7 View the personal information

You can view the personal information of the current user and test the permissions on the Personal Information page.

Procedure

1. Log on to OAM.

- 2. In the left-side navigation pane, select Personal Information.
- 3. In the Basic Information section, you can view the username, user type, created time, AccessKey ID, and AccessKey Secret of the current user.

Note:

Click Show or Hide to show or hide the AccessKey Secret.

- 4. In the Test Permission section, test if the current user has a certain permission.
 - a) Enter the resource information in the Resource field.



Use the English input method when entering values in the Resource and Action fields.

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

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

1.3.8.1.1 Default role of OAM

This topic describes the default role of Operation Access Manager (OAM) and the corresponding available authorizations.

Role name	Role descriptio n	Resource	Actions	Available authorizations
Super administrator	An administra tor with root permissions	*•*	*	10

1.3.8.1.2 Default roles of Apsara Infrastructure Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations.

Role name	Role descriptio	Resource	Actions	Available
	n			authorizations
Tianji_Project read-only	Has the read-only permission to Apsara Infrastructure Management Framework projects, which allows you to view the configurations and statuses of all projects and clusters	*:tianji: projects	["read"]	0
Tianji_Project administrator	Has all the permission s 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 descriptio n	Resource	Actions	Available authorizations
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 permission s 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 permission s to Apsara Infrastructure Management Framework data centers, which allows you to view and modify the data center information	*:tianji:idcs	["*"]	0

Role name	Role descriptio n	Resource	Actions	Available authorizations
Tianji administrator	Has all the permission s to Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastructure Management Framework configurations	*:tianji	["*"]	0

1.3.8.1.3 Default roles of Webapp-rule

This topic describes the default roles of Webapp-rule and the corresponding available authorizations.

Role name	Role descriptio n	Resource	Actions	Available authorizations
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 Default roles of the workflow console This topic describes the default roles of the workflow console and the corresponding available authorizations.

Role name	Role descriptio n	Resource	Actions	Available authorizations
grandcanal. ADMIN	The workflow console administrator, who can query the workflow and activity details, and retry, roll back , terminate, and restart a workflow	26842: grandcanal	["write" ," read"]	0
grandcanal. Reader	Has the read-only permission to the workflow console and can only perform the read operation	26842: grandcanal	["read"]	0

1.3.8.1.5 Default role of Tianjimon

This topic describes the default role of Tianjimon and the corresponding available authorizations.

Role name	Role descriptio n	Resource	Actions	Available authorizations
Tianjimon operations	Has all Tianjimon permission s, which allows you to perform basic monitoring and operations	26842: tianjimon:*	["*"]	0

1.3.8.2 Permission lists of operations platforms This topic describes the permissions of operations platforms.

1.3.8.2.1 Permission list of Apsara Infrastructure Management Framework

This topic describes the permissions of Apsara Infrastructure Management Framework.

Resource	Action	Description
*:tianji:services:[sname]:tjmontemplates:[tmplname]	delete	DeleteServiceTjmonTmpl
*:tianji:services:[sname]:tjmontemplates:[tmplname]	write	PutServiceTjmonTmpl
*:tianji:services:[sname]: templates:[tmplname]	write	PutServiceConfTmpl
*:tianji:services:[sname]: templates:[tmplname]	delete	DeleteServiceConfTmpl
*:tianji:services:[sname]: serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTj monTmpl
*:tianji:services:[sname]: serviceinstances:[siname]:tssessions	terminal	CreateTsSessionByService
*:tianji:services:[sname]: serviceinstances:[siname]:template	write	SetServiceInstanceTmpl
*:tianji:services:[sname]: serviceinstances:[siname]:template	delete	DeleteServiceInstanc eTmpl
*:tianji:services:[sname]: serviceinstances:[siname]:template	read	GetServiceInstanceTmpl
*:tianji:services:[sname]: serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanc eProductTagInService

Resource	Action	Description
*: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	OperateSRMachineInSe rvice
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	read	GetMachineSRInfoInSe rvice
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	delete	DeleteSRMachineActio nInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	read	GetMachinesSRInfoInS ervice
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	delete	DeleteSRMachinesActi onInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	write	OperateSRMachinesInS ervice
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]: apps:[app]:resources	read	GetAppResourceInService

Resource	Action	Description
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	TianjiLogsInService
*:tianji:services:[sname]: serviceinstances:[siname]:serverroles	read	GetServiceInstanceSe rverrolesInService
*:tianji:services:[sname]: serviceinstances:[siname]:schema	write	SetServiceInstanceSc hema
*:tianji:services:[sname]: serviceinstances:[siname]:schema	delete	DeleteServiceInstanc eSchema
*:tianji:services:[sname]: serviceinstances:[siname]:rollings:[version]	write	OperateRollingJobInS ervice
*:tianji:services:[sname]: serviceinstances:[siname]:rollings	read	ListRollingJobInService
*:tianji:services:[sname]: serviceinstances:[siname]:resources	read	GetInstanceResourceI nService
*:tianji:services:[sname]: serviceinstances:[siname]:machines:[machine]	read	GetMachineAllSRInfoI nService
*:tianji:services:[sname]: serviceinstances:[siname]	write	DeployServiceInstanc eInService
*:tianji:services:[sname]: serviceinstances:[siname]	read	GetServiceInstanceConf
*: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
Resource	Action	Description
--	--------	---------------------------------------
*: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	GetMachineFileListIn Service
*: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	GetMachineAppListInS ervice
*:tianji:services:[sname]: serverroles:[serverrole]: apps:[app]:dockerinspect	read	DockerInspect
*:tianji:services:[sname]: schemas:[schemaname]	write	PutServiceSchema
*:tianji:services:[sname]: schemas:[schemaname]	delete	DeleteServiceSchema
*:tianji:services:[sname]: resources	read	GetResourceInService
*:tianji:services:[sname]	delete	DeleteService
*:tianji:services:[sname]	write	CreateService
*:tianji:projects:[pname]: machinebuckets:[bname]: machines:[machine]	read	GetMachineBucketMach ineInfo
*:tianji:projects:[pname]: machinebuckets:[bname]: machines	read	GetMachineBucketMach ines

Resource	Action	Description
*:tianji:projects:[pname]: machinebuckets:[bname]	write	CreateMachineBucket
*:tianji:projects:[pname]: machinebuckets:[bname]	write	OperateMachineBucket Machines
*:tianji:projects:[pname]: machinebuckets:[bname]	delete	DeleteMachineBucket
*:tianji:projects:[pname]: machinebuckets:[bname]	read	GetMachineBucketMach inesLegacy
*:tianji:projects:[pname]: machinebuckets	read	GetMachineBucketList
*:tianji:projects:[pname]: projects:[pname]:clusters :[cname]:tssessions:[tssessionname]:tsses	terminal	UpdateTsSessionTssBy Cluster
*:tianji:projects:[pname]: projects:[pname]:clusters: [cname]:tssessions	terminal	CreateTsSessionByCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTj monTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:template	delete	DeleteServiceInstanc eTmplInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:template	write	SetServiceInstanceTm plInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:template	read	GetServiceInstanceTm plInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:tags:[tag]	write	AddServiceInstancePr oductTagInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:tags:[tag]	delete	DeleteServiceInstanc eProductTagInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: resources	read	GetServerroleResourc eInCluster
*: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	GetMachineAppFileDow nload
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]:apps: [app]:files:content	read	GetMachineAppFileCon tent
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]:apps :[app]:filelist	read	GetMachineFileList
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]:apps: [app]:dockerlogs	read	DockerLogsInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]:apps: [app]:debuglog	read	GetMachineDebugLog
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]:apps	read	GetMachineAppList
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	read	GetMachineSRInfoInCl uster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	write	OperateSRMachineInCl uster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines:[machine]	delete	DeleteSRMachineActio nInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	write	OperateSRMachinesInC luster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	delete	DeleteSRMachinesActi onInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: machines	read	GetAllMachineSRInfoI nCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: apps:[app]:resources	read	GetAppResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	TianjiLogsInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles:[serverrole]: apps:[app]:dockerinspect	read	DockerInspectInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:serverroles	read	GetServiceInstanceSe rverrolesInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:schema	delete	DeleteServiceInstanc eSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:schema	write	SetServiceInstanceSc hemaInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:resources	read	GetInstanceResourceI nCluster

Resource	Action	Description
*: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	GetServiceInstanceCo nfInCluster
*:tianji:projects:[pname]: clusters:[cname]:rollings: [version]	write	OperateRollingJob
*:tianji:projects:[pname]: clusters:[cname]:rollings	read	ListRollingJob
*:tianji:projects:[pname]:clusters:[cname]: resources	read	GetResourceInCluster
*:tianji:projects:[pname]: clusters:[cname]:quota	write	SetClusterQuotas
*:tianji:projects:[pname]:clusters:[cname]: machinesinfo	read	GetClusterMachineInfo
*:tianji:projects:[pname]:clusters:[cname]: machines:[machine]	read	GetMachineAllSRInfo
*:tianji:projects:[pname]:clusters:[cname]: machines:[machine]	write	SetMachineAction
*:tianji:projects:[pname]:clusters:[cname]: machines:[machine]	delete	DeleteMachineAction
*:tianji:projects:[pname]:clusters:[cname]: machines	write	OperateClusterMachines
*:tianji:projects:[pname]: clusters:[cname]:difflist	read	GetVersionDiffList

Resource	Action	Description
*:tianji:projects:[pname]: clusters:[cname]:diff	read	GetVersionDiff
*:tianji:projects:[pname]:clusters:[cname]: deploylogs:[version]	read	GetDeployLogInCluster
*:tianji:projects:[pname]:clusters:[cname]: deploylogs	read	GetDeployLogListInCl uster
*:tianji:projects:[pname]: clusters:[cname]:builds:[version]	read	GetBuildJob
*:tianji:projects:[pname]: clusters:[cname]:builds	read	ListBuildJob
*:tianji:projects:[pname]: clusters:[cname]	write	OperateCluster
*:tianji:projects:[pname]: clusters:[cname]	delete	DeleteCluster
*:tianji:projects:[pname]: clusters:[cname]	read	GetClusterConf
*:tianji:projects:[pname]: clusters:[cname]	write	DeployCluster
*:tianji:projects:[pname]	write	CreateProject
*:tianji:projects:[pname]	delete	DeleteProject
*:tianji:idcs:[idc]:rooms :[room]:racks:[rack]: rackunits:[rackunit]	write	CreateRackunit
*:tianji:idcs:[idc]:rooms :[room]:racks:[rack]: rackunits:[rackunit]	write	SetRackunitAttr
*:tianji:idcs:[idc]:rooms :[room]:racks:[rack]: rackunits:[rackunit]	delete	DeleteRackunit
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	SetRackAttr
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	write	CreateRack

Resource	Action	Description
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	delete	DeleteRack
*:tianji:idcs:[idc]:rooms:[room]	write	CreateRoom
*:tianji:idcs:[idc]:rooms:[room]	delete	DeleteRoom
*:tianji:idcs:[idc]:rooms:[room]	write	SetRoomAttr
*:tianji:idcs:[idc]	delete	DeleteIdc
*:tianji:idcs:[idc]	write	SetIdcAttr
*:tianji:idcs:[idc]	write	CreateIdc

1.3.8.2.2 Permission list of Webapp-rule

This topic describes the permissions of Webapp-rule.

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

1.3.8.2.3 Permission list of the workflow console

This topic describes the permissions of the workflow console.

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

1.3.8.2.4 Permission list of Tianjimon This topic describes the permission of Tianjimon.

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

1.4 Apsara Infrastructure Management Framework

1.4.1 Old version

1.4.1.1 What is Apsara Infrastructure Management Framework?

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

1.4.1.1.1 Overview

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

Apsara Infrastructure Management Framework also supports monitoring data and analyzing reports, which facilitates users to perform one-stop Operation & Maintenance (O&M) control. Based on the Apsara Infrastructure Management Framework services, automated O&M is implemented in the large-scale distribute d 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 TianjiClie nt as the agent on machines it manages. Then, TianjiMaster accepts and issues the upper-layer instructions to TianjiClient for execution. In the upper layer, Apsara Infrastructure Management Framework is divided into different components based on different functions, and then provides API server and portal for external use.

Core functions

• Network initialization in data centers

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

1.4.1.1.2 Basic concepts

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

project

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

cluster

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

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

service

A set of software, which provides relatively independent functions. A service is composed of one or more server roles and can be deployed on multiple clusters to form multiple sets of services and provide the corresponding service capabiliti es. 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 applicatio ns. If a service is deployed on a cluster, all the server roles of the service must be deployed to machines of this cluster. Multiple server roles, such as PanguMaster and TianjiClient, can be deployed on the same server. server role instance

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

application

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

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

rolling

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

service configuration template

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

associated service template

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

final status

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

dependency

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

upgrade

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

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

1.4.1.2 Log on to Apsara Infrastructure Management Framework

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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

Enter a user name
Enter the password
Log On



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 log on to ASO.
- 5. In the left-side navigation pane, select Products.

6. In the product list, select Apsara Infrastructure Management Framework.

1.4.1.3 Web page introduction

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

1.4.1.3.1 Introduction on the home page

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

Log on to Apsara Infrastructure Management Framework. The home page appears, as shown in Figure 1-7: Home page of Apsara Infrastructure Management Framework.

🐟 Tian Ji @ 15:21 () aliy Upgrade Task Summary Most-used Reports · Registration Vars of S Resource Apply Rep Machine Info Report 0.009 IP LIST Virtual Mach MM Monitor - Cluste Clusters 83 TOP: 15 Strewers Error Alarms OS Errors H/W Err Error Summary Rate of Abnormal Servers: 0.00% Server A... 0.00% OS Errors: 0.00% H/W Err... 0.00% Rate of Abnormal ServerRole Instances: 1.06% Alarms: 1.06% Abnormal: 46

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

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

see Table 1-9: Descriptions of functional areas.

Area		Description
Area 1	Top navigation bar	 Description 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 consists of the following sections: Cluster Operations: performs O&M operations on and manages clusters with the project permissions, such as viewing the cluster status. Service Operations: manages services with the service permissions, such as viewing the service list information. Machine Operations: maintains and manages all the machines in Apsara Infrastructure Management Framework, such as viewing the machine status.
		 configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects. Reports: displays the monitoring data in tables and provides the function of searching for different reports. Monitoring: effectively monitors metrics in the process of system operation and sends alert notifications for abnormal conditions. This menu includes the functions of displaying alert status, modifying alert rules, and searching for the alert history.

Table 1-9: Descriptions	of functional areas
-------------------------	---------------------

Area		Description
2	Function buttons in the upper -right corner	 O: TJDB Synchronization Time: the generated time of the data that is displayed on the current page. Final Status Computing Time: the computing time of the final-status data that is displayed on the current page. After data is generated, the system processes the data at maximum speed. As an asynchronous system, Apsara Infrastructure Management Framework has some latency. The time helps explain why the current data results are generated and determine whether the current system has a problem. English(US) : In the English environment, move the pointer over this and select another language. aliyuntest : The logon account information. Move the pointer over this and select Logout to log out of Apsara Infrastructure Management Framework.
3	Left-side navigation pane	In the left-side navigation pane, you can directly view the logical structure of the Apsara Infrastructure Management Framework model. You can view the corresponding detailed data analysis and operations by selecting different levels of nodes in the left- side navigation pane. For more information, see <i>Introduction</i> <i>on the left-side navigation pane</i> .

Area		Description
4	Home page	 Displays the summary of related tasks or information as follows: Upgrade Task Summary: the numbers and proportions of running, rolling back, and paused upgrade tasks. Cluster Summary: the numbers of machines, error alerts, operating system errors, and hardware errors for different clusters. Error Summary: the metrics for the rate of abnormal machines and the rate of abnormal server role instances. Most-used Reports: links of the most commonly used statistics reports, which facilitates you to view the report information.
5	Button used to collapse /expand the left -side navigation pane	If you are not required to use the left-side navigation pane when performing O&M operations, click the left-side navigation pane and increase the space of the content area.

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

Cluster

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

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

- Enter the cluster name in the search box to search for the cluster quickly. Fuzzy search is supported.
- Select a project from the Project drop-down list to display all the clusters in the project.
- Move the pointer over **a** at the right of a cluster and then perform operations on

the cluster as instructed.

 Click a cluster and all the machines and services in this cluster are displayed in the lower-left corner. Move the pointer over at the right of a machine or

service and then perform operations on the machine or service as instructed.

- Click the Machine tab in the lower-left corner. Double-click a machine to view all the server roles in the machine. Double-click a server role to view the applications and then double-click an application to view the log files.
- Click the Service tab in the lower-left corner. Double-click a service to view all the server roles in the service. Double-click a server role to view the machines, double-click a machine to view the applications, and double-click an application to view the log files.
- Double-click a log file. Move the pointer over at the right of the log file and then select Download to download the log file.

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

Service

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

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

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

then perform operations on the service instance as instructed.

Report

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

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

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

1.4.1.4 Cluster operations

This topic describes the actions about cluster operations.

1.4.1.4.1 View cluster configurations

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

Procedure

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

The Cluster Operations page displays the following information:

• Cluster

The cluster name. Click the cluster name to go to the *Cluster Dashboard* page.

· Scale-Out/Scale-In

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

• Abnormal Machine Count

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

• Final Status of Normal Machines

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

· Rolling

Displays whether the cluster has a running rolling task. Select Rolling Tasks to display clusters that have rolling tasks. Click the link to go to the *Rolling Task* page.

- 3. Select a project from the Project drop-down list and/or enter the cluster name in the Cluster field to search for clusters.
- 4. Find the cluster whose configurations you are about to view and then click Cluster Configuration in the Actions column. The Cluster Configuration page appears.

For more information about the Cluster Configuration page, see *Table 1-10: Cluster configurations*.

Category	Item	Description
Basic	Cluster	The cluster name.
Information	Project	The project to which the cluster belongs.
	Clone Switch	 Mock Clone: The system is not cloned when a machine is added to the cluster. Real Clone: The system is cloned when a machine is added to the cluster.
	Machines	The number of machines in the cluster. Click View Clustering Machines to view the machine list.
	Security Verification	The access control among processes . Generally, the non-production environment uses the default configurations and does not perform the verification. In other cases, customize the configurations based on actual requirements to enable or disable the verification.
	Cluster Type	 RDS NETFRAME T4: a special type that is required by the mixed deployment of e- commerce. Default: other conditions.
Deployment Plan	Service	The service deployed in the cluster.

Table 1-10: Cluster configurations

Category	Item	Description
	Dependency Service	The service that the current service depends on.
Service Information	Service Information	Select a service from the Service Information drop-down list and then the configurations of this service are displayed.
	Service Template	The template used by the service.
	Monitoring Template	The monitoring template used by the service.
	Machine Mappings	The machines included in the server role of the service.
	Software Version	The software version of the server role in the service.
	Availability Configuration	The availability configuration percentage of the server role in the service.
	Deployment Plan	The deployment plan of the server role in the service.
	Configuration Information	The configuration file used in the service.
	Role Attribute	Server roles and the corresponding parameters.

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

1.4.1.4.2 View the cluster dashboard

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. You have two ways to go to the Cluster Dashboard page:
 - In the left-side navigation pane, click the C tab. Move the pointer over 📑 at

the right of a cluster and then select Dashboard.

- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, click the cluster name.
- 3. On the Cluster Dashboard page, you can view the cluster information, including the basic information, final status information, rolling job information, dependencies, resource information, virtual machines, and monitoring information. For more information about the descriptions, see the following table.

Item	Description
Basic Cluster Information	 Displays the basic information of the cluster as follows: Project Name: the project name. Cluster Name: the cluster name. IDC: the data center to which the cluster belongs. Final Status Version: the latest version of the cluster. Cluster in Final Status: whether the cluster reaches the final status. Machines Not In Final Status: the number of machines that do not reach the final status in the cluster when the cluster does not reach the final status. Real/Pseudo Clone: whether to clone the system when a machine is added to the cluster. Expected Machines: the number of expected machines in the cluster. Actual Machines: the number of machines in the current environment. Machines Not Good: the number of machines whose status is not Good in the cluster. Actual Services: the number of services that are actually deployed in the cluster. Cluster Status: whether the cluster is starting or shutting down machines.

Item	Description
Machine Status Overview	The statistical chart of the machine status in the cluster
Machines in Final Status	The numbers of machines that reach the final status and those that do not reach the final status in each service of the cluster.
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
Disk_usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-system	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
Disk_IO-System	The statistical table of the disk input and output.
Service Instances	Displays the service instances deployed in the cluster and the related final status information.
	 Service Instance: the service instance deployed in the cluster. Final Status: whether the service instance reaches the final status. Expected Server Roles: the number of server roles that the service instance expects to deploy. Server Roles In Final Status: the number of server roles that reach the final status in the service
	 instance. Server Roles Going Offline: the number of server roles that are going offline in the service instance. Actions: Click Details to go to the Service Instance Information Dashboard page. For more information about the service instance dashboard, see <i>View the service instance dashboard</i>.

Item	Description
Upgrade Tasks	Displays the upgrade tasks related to the cluster.
	 Cluster Name: the name of the upgrade cluster. Type: the type of the upgrade task. The options include app (version upgrade) and config (configuration change). Git Version: the change version to which the upgrade task belongs. Description: the description about the change. Rolling Result: the result of the upgrade task. Submitted By: the person who submits the change. Submitted At: the time when the change is submitted. Start Time: the time to start the rolling. End Time: the time to finish the upgrade. Time Used: the time used for the upgrade. Actions: Click Details to go to the Rolling Task page. For more information about the rolling tasks, see View rolling tasks.
Cluster Resource Request Status	 Version: the resource request version. Msg: the exception message. Begintime: the start time of the resource request analysis. Endtime: the end time of the resource request analysis. Build Status: the build status of resources. Resource Process Status: the resource request status in the version.

Item	Description
Cluster Resource	 Service: the service name. Server Role: the server role name. App: the application of the server role. Name: the resource name. Type: the resource type. Status: the resource request status. Error Msg: the exception message. Parameters: the resource parameters. Result: the resource request result. Res: the resource ID. Reprocess Status: the status of interaction with Business Foundation System during the VIP resource request. Reprocess Result: the result 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. Reprocess Result: the result 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.
VM Mappings	 The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster. VM: the hostname of the virtual machine. Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed. Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.

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

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. You have three ways to go to the Cluster Operation and Maintenance Center page:
 - In the left-side navigation pane, click the C tab. Move the pointer over

at the right of a cluster and then select Cluster Operation and Maintenance Center.

- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, choose Monitoring > Cluster Operation and Maintenance Center in the Actions column at the right of a cluster.
- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, click a cluster name. On the Cluster Dashboard page, choose Operations Menu > Cluster Operation and Maintenance Center.

3. View the information on the Cluster Operation and Maintenance Center page.

Item	Description
SR not in Final Status	Displays all the server roles that do not reach the final status in the cluster. Click the number to expand a server role list, and click a server role in the list to display the information of machines included in the server role
Punning Tasks	machines included in the server role.
Kunning Tasks	Displays whether the cluster has running rolling tasks. Click Rolling to go to the Rolling Task page. For more information about the rolling task, see <i>View rolling tasks</i> .
Head Version Submitted At	The time when the head version is submitted. 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 detailed change contents. The head version analysis has the following statuses: Preparing: No new version is available now. Waiting: The latest version is found. The analysis module has not started up yet. Doing: The module is analyzing the application that requires change. done: The head version analysis is successfully completed. Failed: The head version analysis failed. The change contents cannot be parsed. If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version. Click the status to view the relevant information.

Item	Description
Service	Select a service deployed in the cluster from the drop- down list.
Server Role	Select a server role of a service in the cluster from the drop-down list.
	Note: After you select the service and server role, the information of machines related to the service or server role is displayed in the list.
Total Machines	The total number of machines in the cluster, or the total number of machines included in a specific server role of a specific service.
Scale-in/Scale-out	The number of machines or server roles that are scaled in or out.
Abnormal Machines	The number of abnormal machines that encounter each type of the following faults.
	 Ping Failed: A ping_monitor error is reported, and TianjiMaster cannot successfully ping the machine. No Heartbeat: TianjiClient on the machine does not regularly report data to indicate the status of this machine, which may be caused by the TianjiClient problem or network problem. Status Error: The machine has an error reported by the monitor or a fault of the critical or fatal level. Check the alert information and accordingly solve the issue.

Item	Description
Abnormal Services	 Description The number of machines with abnormal services. To determine if a service reaches the final status, see the following rules: The server role on the machine is in the GOOD status. Each application of the server role on the machine must keep the actual version the same as the head version. Before the Image Builder builds an application of the head version, Apsara Infrastructure Management Framework cannot determine the value of the head version and the service final status is unknown. This process is called the change preparation process. The
	service final status cannot be determined during the preparation process or upon a preparation failure.

Item	Description
Machines	Displays all the machines in the cluster or the machines included in a specific server role of a specific service.
	 Machine search: Click the search box to enter the machine in the displayed dialog box. Fuzzy or batch search is supported. Click the machine name to view the physical information of the machine in the displayed Machine Information dialog box. Click DashBoard to go to the Machine Details page. For more information about the machine details, see <i>View the machine dashboard</i>. Move the pointer over the blank area in the Final Status column or the Final SR Status column and then click Details to view the machine status, system service information, server role status on the machine, and exception message. If no service or server role is selected from the drop-down list, move the pointer over the blank area in the Running Status column and then click Details to view the machine over the blank area in the Running status information or exception message of the machine.
	 If a service and a server role are selected from the corresponding drop-down lists, move the pointer over the blank area in the SR Running Status column and then click Details to view the running status information or exception message of the server role on the machine. Click Error, Warning, or Good in the Monitoring Statistics column to view the monitored items of machines and monitored items of server roles. Click Terminal in the Actions column to log on to the machine and perform related operations. Click Machine Operation in the Actions column to restart, out-of-band restart, or clone the machine again.

1.4.1.4.4 View the service final status

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. You have two ways to go to the Service Final Status Query page:
 - In the left-side navigation pane, click the C tab. Move the pointer over \mathbf{I} at

the right of a cluster and then choose Monitoring > Service Final Status Query.

 In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, choose Monitoring > Service Final Status Query in the Actions column at the right of a cluster.

Item	Description	
Project Name	The name of the project to which the cluster belongs.	
Cluster Name	The cluster name.	
Head Version Submitted At	The time when the head version is submitted.	
Head Version Analysis	 The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses: Preparing: No new version is available now. Waiting: The latest version is found. The analysis module has not started up yet. Doing: The module is analyzing the application that requires change. done: The head version analysis is successfully completed. Failed: The head version analysis failed. The change contents cannot be parsed. If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.	
Cluster Rolling Status	Displays the information of the current rolling task in the cluster, if any. The rolling task may not be of the head version.	

3. View the information on the Service Final Status Query page.

Item	Description
Cluster Machine Final Status Statistics	The status of all machines in the cluster. Click View Details to go to the Cluster Operation and Maintenance Center page and view the detailed information of all machines. For more information about the cluster operation and maintenance center, see <i>View the cluster</i> <i>operation and maintenance center</i> .
Final Status of Cluster SR Version	The final status of cluster service version. Note: Take statistics of services that do not reach the final status, which is caused by version inconsistency or status exceptions. If services do not reach the final status because of machine problems, go to Cluster Machine Final Status Statistics to view the statistics.
Final Status of SR Version	The number of machines that do not reach the final status when a server role has tasks.

1.4.1.4.5 View operation logs

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

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. You have two ways to go to the Cluster Operation Logs page:
 - In the left-side navigation pane, click the C tab. Move the pointer over 👔 at

the right of a cluster and then choose Monitoring > Operation Logs.

- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, choose Monitoring > Operation Logs in the Actions column at the right of a cluster.
- 3. On the Cluster Operation Logs page, click Refresh. View the Git version, description, submitter, submitted time, and task status.

- 4. Optional: Complete the following steps to view the differences between versions on the Cluster Operation Logs page.
 - a) Find the log in the operation log list and then click View Release Changes in the Actions column.
 - b) On the Version Difference page, complete the following configurations:
 - Select Base Version: Select a base version.
 - Configuration Type: Select Extended Configuration or Cluster
 Configuration. Extended Configuration displays the configuration
 differences after the configuration on the cluster is combined with
 the configuration in the template. Cluster Configuration displays the
 configuration differences on the cluster.
 - c) Click Obtain Difference.

The differential file list is displayed.

d) Click each differential file to view the detailed differences.

1.4.1.5 Service operations

This topic describes the actions about service operations.

1.4.1.5.1 View the service list

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

Procedure

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

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

Item	Description
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.

1.4.1.5.2 View the service instance dashboard

The service instance dashboard allows you to view the basic information and statistics of a service instance.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, click the S tab.
- 3. Enter the service name in the search box. Services that meet the search condition are displayed.
- 4. Click a service name and then service instances in the service are displayed in the lower-left corner.
- 5. Move the pointer over **at the right of a service instance and then select**

Dashboard.

o. view the information on the service instance information dashboard pa	6.	tion on the Service Instance Informati	on Dashboar	d page
--	----	--	-------------	--------

Item	Description	
Service Instance Summary	Displays the basic information of the service instance as follows:	
	 Cluster Name: the name of the cluster to which the service instance belongs. Service Name: the name of the service to which the service instance belongs. Actual Machines: the number of machines in the current environment. Expected Machines: the number of machines that the service instance expects. Target Total Server Roles: the number of server roles that the service instance expects. Actual Server Roles: the number of server roles in the current environment. Template Name: the name of the service template used by the service instance. Schema: the name of the service schema used by the service instance. Monitoring System Template: the name of the service instance. 	
Server Role Statuses	The statistical chart of the current status of server roles in the service instance.	
Machine Statuses for Server Roles	The status statistics of machines where server roles are located.	
Service Monitoring Information	 Monitored Item: the name of the monitored item. Level: the level of the monitored item. Description: the description of the monitored contents. Updated At: the time when the data is updated. 	

Item	Description
Service Alert Status	 Alert Name Instance Information Alert Start Alert End Alert Duration Severity Level Occurrences: the number of times the alert is triggered.
Server Role List	 Server Role Current Status Expected Machines Machines In Final Status Machines Going Offline Rolling Task Status Time Used: the time used for running the rolling task. Actions: Click Details to go to the Server Role Dashboard page.
Service Alert History	 Alert Name Alert Time Instance Information Severity Level Contact Group
Service Dependencies	 The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role. Server Role: the server role name. Dependent Service: the service on which the server role depends. Dependent Server Role: the server role on which the server role depends. Dependent Cluster: the cluster to which the dependent server role belongs. Dependency in Final Status: whether the dependent server role reaches the final status.
1.4.1.5.3 View the server role dashboard

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

Procedure

- **1.** Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, click the S tab.
- 3. Enter the service name in the search box. Services that meet the search condition are displayed.
- 4. Click a service name and then service instances in the service are displayed in the lower-left corner.
- 5. Move the pointer over **a** at the right of a service instance and then select

Dashboard.

6. In the Server Role List section of the Service Instance Information Dashboard page, click Details in the Actions column.

7. View the information on the Server Role Dashboard page	ge.
---	-----

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

Item	Description
Machine Information	 Machine Name: the hostname of the machine. IP: the IP address of the machine. Machine Status: the machine status. Machine Action: the action that the machine is performing. Server Role Status: the status of the server role. Server Role Action: the action that the server role is performing. Current Version: the current version of the server role on the machine. Target Version: the expected version of the server role on the machine. Error Message: the exception message. Actions: Click Terminal to log on to the machine and perform operations. Click Restart to restart the server roles on the machine. Click Details to go to the Machine Details page. For more information about the machine details, see <i>View the machine dashboard</i>. Click Machine System View to go to the Machine Info Report page. For more information to restart, out of band restart. or clone the machine again
Server Role Monitoring Information of Machines	 Updated At: the time when the data is updated. Machine Name: the machine name. Monitored Item: the name of the monitored item. Level: the level of the monitored item. Description: the description of the monitored item.

Item	Description
VM Mappings	The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.
	$\cdot $ VM: the hostname of the virtual machine.
	• Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed.
	• Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.
Service Dependencies	The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role.
	 Dependent Service: the service on which the server role depends.
	• Dependent Server Role: the server role on which the server role depends.
	Dependent Cluster: the cluster to which the dependent server role belongs
	 Dependency in Final Status: whether the dependent server role reaches the final status.

1.4.1.6 Machine operations

This topic describes the actions about machine operations.

1.4.1.6.1 View the machine dashboard

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

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, click the C tab.
- 3. On the Machine tab in the lower-left corner, enter the machine name in the search box. Machines that meet the search condition are displayed.
- 4. Move the pointer over **a** at the right of a machine and then select Dashboard.

5. On the Machine Details page, view all the information of this machine. For more information, see the following table.

Item	Description
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
DISK Usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-System	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
DISK IO-System	The statistical table of the disk input and output.
Machine Summary	 Project Name: the name of the project to which the machine belongs. Cluster Name: the name of the cluster to which the machine belongs. Machine Name: the machine name. SN: the serial number of the machine. IP: the IP address of the machine. IDC: the data center of the machine. Room: the room in the data center where the machine is located. Rack: the rack where the machine is located. Unit in Rack: the location of the rack. Warranty: the warranty of the machine. Purchase Date: the date when the machine is purchased. Machine Status: the running status of the machine. Status: the hardware status of the machine. Disks: the disk size. Memory: the memory size. Manufacturer: the machine manufacturer. Model: the machine model. os: the operating system of the machine. part: the disk partition.
Server Role Status of	The distribution of the current status of all server roles
Machine	on the machine.

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

1.4.1.7 Monitoring center

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

1.4.1.7.1 Modify an alert rule

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. In the top navigation bar, choose Operations > Service Operations.
- 3. Enter the service name in the search box.
- 4. Find the service and then click Management in the Actions column.
- 5. Click the Monitoring Template tab.
- 6. Find the monitoring template that you are about to edit and then click Edit in the Actions column.
- 7. Configure the monitoring parameters based on actual conditions.
- 8. Click Save Change.

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

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. In the top navigation bar, choose Operations > Service Operations.
- 3. Enter the service name in the search box.
- 4. Find the service and then click Management in the Actions column.
- 5. Click the Monitoring Instance tab.

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

1.4.1.7.3 View the alert status

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

Procedure

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

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

1.4.1.7.4 View alert rules

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. In the top navigation bar, choose Monitoring > Alert Rules.
- 3. You can configure the service name, cluster name, and/or alert name to search for alert rules.
- 4. View the detailed alert rules on the Alert Rules page. See the following table for the alert rule descriptions.

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

1.4.1.7.5 View the alert history

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

Procedure

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

Item	Description
Service	The name of the service to which the alert belongs.
Cluster	The name of the cluster where the service is located.
Alert Instance	The name of the resource where the alert is triggered.
Status	Alerts have two statuses: Restored and Alerting.

Item	Description
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services.
	· P1
	• P2
	· P3
	• P4
Alert Name	The name of the generated alert.
	Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members that are notified when an alert is triggered.
Actions	Click Show to show the data before and after the alert time.

1.4.1.8 Tasks and deployment summary

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

1.4.1.8.1 View rolling tasks

You can view running rolling tasks and the corresponding status.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the top navigation bar, choose Operations > Cluster Operations.
- 3. Select Rolling Tasks to display clusters with rolling tasks.
- 4. In the search results, click rolling in the Rolling column.
- 5. On the displayed Rolling Task page, view the information in the Change Task list and Change Details list.

Item	Description
Change Version	The version that triggers the change of the rolling task.
Description	The description about the change.

Table 1-11: Change Task list

Item	Description
Head Version	The head version analysis is the process that Apsara
Third, old	Infrastructure Management Framework detects the latest
	cluster version and parses the version to detailed change
	contents. The head version analysis has the following statuses:
	\cdot Preparing: No new version is available now.
	\cdot Waiting: The latest version is found. The analysis module has
	not started up yet.
	\cdot Doing: The module is analyzing the application that requires
	change.
	\cdot done: The head version analysis is successfully completed.
	\cdot Failed: The head version analysis failed. The change contents
	cannot be parsed.
	If the status is not done, Apsara Infrastructure Management
	Framework cannot detect the change contents of server roles in
	the latest version.
Blocked Server Role	Server roles blocked in the rolling task. Generally, server roles are blocked because of dependencies.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.
Actions	Click View Difference to go to the Version Difference page. For
	more information, see View operation logs.
	Click Stop to stop the rolling task.
	Click Pause to pause the rolling task.

Table 1-12: Change Details list

Item	Description
Service Name	The name of the service where a change occurs.

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

1.4.1.8.2 View running tasks

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

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

1.4.1.8.3 View history tasks

You can view the historical running conditions of completed tasks.

Procedure

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

1.4.1.8.4 View the deployment summary

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

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. In the top navigation bar, choose Tasks > Deployment Summary.
 - View the deployment status and the duration of a certain status for each project.
 - Gray: wait to be deployed. It indicates that some services of the project depend on server roles or service instances that are being deployed, and

other service instances or server roles in the project have already been deployed.

- Blue: being deployed. It indicates that the project has not reached the final status for one time yet.
- Green: has reached the final status. It indicates that all clusters in the project have reached the final status.
- Orange: not reaches the final status. It indicates that a server role does not reach the final status for some reason after the project reaches the final status for the first time.
- Configure the global clone switch.
 - normal: Clone is allowed.
 - block: Clone is forbidden.
- Configure the global dependency switch.
 - normal: All configured dependencies are checked.
 - ignore: The dependency is not checked.
 - ignore_service: None of the service-level dependencies, including the server role dependencies across services, are checked, and only the server role-level dependencies are checked.

3. Click the Deployment Details tab to view the deployment details.

Item	Description
Status Statistics	 The general statistics of deployment conditions, including the total number of projects that are currently available. Click each status to display the projects in the corresponding status in the list. The projects have five deployment statuses: Final: All the clusters in the project have reached the final status. Deploying: The project has not reached the final status for one time yet. Waiting: Some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed. Non-final: A server role does not reach the final status for some reason after the project reaches the final status for the first time. Inspector Warning: An error is detected on service instances in the project during the inspection.
Start Time	The time when Apsara Infrastructure Management Framework starts the deployment.
Progress	The proportion of server roles that reach the final status to all the server roles in the current environment.
Deployment Status	The time indicates the deployment duration for the following statuses: Final, Deploying, Waiting, and Inspector Warning. The time indicates the duration before the final status is reached for the Non-final status. Click the time to view the details.

For more information, see the following table.

Г

Item	Description
Deployment Progress	The proportion of clusters, services, and server roles that reach the final status to the total clusters, services, and server roles in the project.
	Move the pointer over the blank area at the right of the
	data of roles and then click Details to view the deployment
	statuses of clusters, services, and server roles. The
	deployment statuses are indicated by icons, which are the
	same as those used for status statistics.
Resource	Total indicates the total number of resources related to the
Application	project.
Progress	• Done: the number of resources that have been successfully applied for.
	• Doing: the number of resources that are being applied for and retried. The number of retries (if any) is displayed
	next to the number of resources.
	• Block: the number of resources whose applications are blocked by other resources.
	• Failed: the number of resources whose applications failed.
Inspector Error	The number of inspection alerts for the current project.
Monitoring	The number of alerts generated for the machine monitor and
Information	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.

1.4.1.9 Reports

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

1.4.1.9.1 View reports

The Reports menu allows you to view the statistical data.

Context

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

• System reports: default and common reports in the system.

• All reports: includes the system reports and custom reports.

Procedure

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

the right of All Reports and then select View.

See the following table for the report descriptions.

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

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

For more information about the reports, see *Appendix*.

1.4.1.9.2 Add a report to favorites

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

Procedure

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

the right of All Reports and then select View.

- 3. Enter the name of the report that you are about to add to favorites in the search box.
- 4. At the right of the report, click Add to Favorites in the Actions column.
- 5. In the displayed Add to Favorites dialog box, enter tags for the report.
- 6. Click Add to Favorites.

1.4.1.10 Appendix

1.4.1.10.1 Project component info report

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

Item	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.

Item	Description
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

1.4.1.10.2 IP list

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

IP List of Physical Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.

IP List of Docker Applications

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.
Docker IP	The Docker IP address.

1.4.1.10.3 Machine info report

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

Machine Status

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

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

Expected Server Role List

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

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

Abnormal Monitoring Status

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

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

Server Role Version and Status on Machine

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

Item	Description
Machine Name	The machine name.

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

Monitoring Status

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

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

1.4.1.10.4 Rolling info report

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

Choose a rolling action

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

Item	Description
Cluster	The cluster name.
Git Version	The version of change that triggers the rolling task.
Description	The description about the change entered by a user when the user submits the change.

Item	Description
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Submitted By	The ID of the user who submits the change.
Rolling Task Status	The current status of the rolling task.
Submitted At	The time when the change is submitted.

Server Role in Job

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

Item	Description
Server Role	The server role name.
Server Role Status	The rolling status of the server role.
Error Message	The exception message of the rolling task.
Git Version	The version of change to which the rolling task belongs.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Approve Rate	The proportion of machines that have the rolling task approved by the decider.
Failure Rate	The proportion of machines that have the rolling task failed.
Success Rate	The proportion of machines that have the rolling task succeeded.

Server Role Rolling Build Information

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

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

Server Role Statuses on Machines

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

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

1.4.1.10.5 Machine RMA approval pending list

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

Machine

Displays the basic information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.

Item	Description
IP	The IP address of the machine.
State	The running status of the machine.
Action Name	The action on the machine.
Action Status	The status of the action on the machine.
Actions	The approval button.

Machine Serverrole

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

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

Machine Component

Displays the hard disk information of pending approval machines.

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

1.4.1.10.6 Registration vars of services This report displays values of all service registration variables.

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

1.4.1.10.7 Virtual machine mappings

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

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

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

1.4.1.10.8 Service inspector report

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

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

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

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

Change Mappings

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

Changed Resource List

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

Resource Status

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

Item	Description
Name	The resource name.
Туре	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 message.

1.4.1.10.10 Statuses of project components

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

Error State Component Table

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

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

Item	Description
Server Role Status	The current status of the server role.
Machine Status	The current status of the machine.

Server Role Alert Information

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

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

Machine Alert Information

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

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

Service Inspector Information

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

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

1.4.1.10.11 Relationship of service dependency

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

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

1.4.1.10.12 Check report of network topology

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

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Network Instance	The name of the network device.

Item	Description
Level	The alert level.
Description	The description about the alert information.

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

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

1.4.1.10.13 Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

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

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.

Item	Description
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

1.4.1.10.14 Auto healing/install approval pending report

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

1.4.1.10.15 Machine power on or off statuses of clusters After a cluster starts or shuts down machines, you can view the related information in this report.

Cluster Running Statuses

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

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

Server Role Power On or Off Statuses

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

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

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
Action Name	The startup or shutdown action that is being performed by the server role.
Action Status	The status of the action.

Statuses on Machines

Displays the running status of the selected server role on machines.

Select a row in the Server Role Power On or Off Statuses section to display the information of the corresponding server role in the list.

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

Machine Statuses

Displays the running statuses of machines in the selected cluster.

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

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

1.4.2 New version

1.4.2.1 What is Apsara Infrastructure Management Framework?

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

1.4.2.1.1 Introduction

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

Overview

Apsara Infrastructure Management Framework also supports monitoring data and analyzing reports, which facilitates users to perform one-stop Operation & Maintenance (O&M) control. Based on the Apsara Infrastructure Management Framework services, automated O&M is implemented in the large-scale distribute d 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 TianjiClie nt as the agent on machines it manages. Then, TianjiMaster accepts and issues the upper-layer instructions to TianjiClient for execution. In the upper layer, Apsara Infrastructure Management Framework is divided into different components based on different functions, and then provides API server and portal for external use.

Core functions

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

1.4.2.1.2 Basic concepts

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

project

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

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

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

service

A set of software, which provides relatively independent functions. A service is composed of one or more server roles and can be deployed on multiple clusters to form multiple sets of services and provide the corresponding service capabiliti es. 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 applicatio ns. If a service is deployed on a cluster, all the server roles of the service must be deployed to machines of this cluster. Multiple server roles, such as PanguMaster and TianjiClient, can be deployed on the same server.

server role instance

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

application

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

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

rolling

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

service configuration template

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

associated service template

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

desired state

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

dependency

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

upgrade

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

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

1.4.2.2 Log on to Apsara Infrastructure Management Framework

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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

Enter a user name
Enter the password
Log On



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 log on to ASO.
- 5. In the left-side navigation pane, select Products.

6. In the Product List, select Apsara Infrastructure Management Framework.

1.4.2.3 Homepage introduction

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

Log on to Apsara Infrastructure Management Framework. The homepage appears, as shown in the following figure.

Figure 1-9: Homepage of Apsara Infrastructure Management Framework



For more information about the descriptions of functional areas on the homepage, see the following table.

Area		Description
1	Left-side navigation pane	 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 consists of the following sections:
		 Project Operations: manages projects with the project permissions.
		 Cluster Operations: performs O&M operations on and manages clusters with the project permissions, such as viewing the cluster status.
		 Service Operations: manages services with the service permissions, such as viewing the service list information.
		- Machine Operations: maintains and manages all the machines in Apsara Infrastructure Management Framework, such as viewing the machine status.
		• Tasks: A rolling task is generated after you modify the configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects.
		 Reports: displays the monitoring data in tables and provides the function of searching for different reports.
		• Monitoring: effectively monitors metrics in the process of system operation and sends alert notifications for abnormal conditions. This menu includes the functions of displaying alert status, modifying alert rules, and searching for the alert history.
		• Tools: provides the machine tools and the IDC shutdown

function.

Table 1-13: Descriptions of functional area	as
---	----

Area		Description
2	Function buttons in the upper -right corner	 Search box: Supports global search. Enter a keyword in the search box to search for clusters, services, and machines. Move the pointer over the time and then you can view: TJDB Sync Time: the generated time of the data that is displayed on the current page. Desired State Calc Time: the calculation time of the desired-state data that is displayed on the current page. After data is generated, the system processes the data at maximum speed. As an asynchronous system, Apsara Infrastructure Management Framework has some latency. The time helps explain why the current data results are generated and determine whether the current system has a problem. English (US) : In the English environment, click this drop-down list to switch to another language. Click the avatar of the logon user and then select Exit to log out of Apsara Infrastructure Management Framework.
3	Status section of global resources	 Displays the overview of global resources. Clusters: displays the total number of clusters, the percentage of clusters that reach the desired state, and the number of abnormal clusters. Instances: displays the total number of instances, the percentage of instances that reach the desired state, and the number of abnormal instances. Machines: displays the total number of machines, the percentage of machines with the Normal state, and the number of abnormal machines. Move the pointer over the section and then click Show Detail to go to the Cluster Operations page, Service Operations page, or Machine Operations page.

Area		Description
4	Task status section	Displays the information of tasks submitted in the last week. Click the number at the right of a task status to go to the My Tasks page and then view tasks of the corresponding status. The top 5 latest tasks are displayed at the bottom of this section and you can click Details to view the task details.
5	Quick actions	Displays links of common quick actions, which allows you to perform operations quickly.
6	Expand/ collapse button	If you are not required to use the left-side navigation pane when performing O&M operations, click this button to collapse the left-side navigation pane and increase the space of the content area.

1.4.2.4 Project operations

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

Procedure

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

Operations	/ Project Operations										
Pro	ject Status Statistics							Dep	loy Data IDC Topol	logical Graph amtest73	
	amtest73	Desired State 43 Projects Not Desired State 19 Projects		62 Total Projects	×	16 Alerting	Ē	4 In Progress		Cother Reasons	
Pro	ject Status									All	
	apigateway	Alerting 2	In Progress 0	Not Desired State		aso	Alerting 16		In Progress 0	Not Desired State	
I	astc	Alerting 3	In Progress 1	Not Desired State		blink	Alerting 2		In Progress 0	Not Desired State	
	drds	Alerting 2	In Progress 0	Not Desired State		ecs	Alerting 20	8	In Progress 1	Not Desired State	

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

Click the drop-down list in the upper-right corner of the Project Status section. Enter a project name in the search box, and then select the name to

search for the project. You can view the numbers of alerts and running tasks for the project and whether the project reaches the desired state.

- View the details of a project
 - Find the project whose details you are about to view. Click the number at the right of Alerting. In the displayed Alert Information dialog box, view the specific monitoring metrics, monitoring types, and alert sources. Click the value in the Alert Source column to view the service details.
 - Find the project whose details you are about to view. Click the number at the right of In Progress. In the displayed Tasks dialog box, view the details of Upgrade Service and Machine Change.

1.4.2.5 Cluster operations

This topic describes the actions about cluster operations.

1.4.2.5.1 View the cluster list

The cluster list allows you to view all of the clusters and the corresponding information.

Procedure

1. Log on to Apsara Infrastructure Management Framework.

- 2. To view the cluster list, you can:
 - On the Homepage, move the pointer over the Clusters section and then click Show Detail in the upper-right corner.
 - In the left-side navigation pane, choose Operations > Cluster Operations.

Operations / Cluster Operations						
Clusters						
IDC amtest73	 ✓ Project All 	V	Clusters Enter a cluster name	Q		
Clusters	Region	Status 🍸	Machine Status	Server Role Status	Task Status 🍸	Actions
AcsNodeCluster-A-20191030-2881 acs	cn-qingdao-env3b-d02	Desired State	7 in Total Normal	14 in Total Normal	Successful	Operations
AliguardCluster-A-20191030-2895 yundun-advance	cn-qingdao-env3b-d02	Not Desired State	3 in Total Normal	8 in Total Abnormal: 1	Failed	Operations
BasicCluster-A-20191030-284c dauthProduct	cn-qingdao-env3b-d02	Desired State	2 in Total Normal	7 in Total Normal	Successful	Operations

Item	Description
Clusters	The cluster name. Click the cluster name to view the cluster details.
Region	The name of the region where the cluster is located.

On this page, you can view the following information.

Item	Description	1			
Status	Indicates whether the cluster reaches the desired state. Use to filter the clusters.				
	• Desired State: All the clusters of a project reach the desired state.				
	• Not Desired State: After a project reaches the desired state for the first time, a server role does not reach the desired state because of undefined reasons.				
Machine Status	The number of machines and the corresponding status in the cluster. Click the status to go to the Machines tab of the Cluster Details page.				
Server Role Status	The number in the clust the Cluster Status colu cluster in the upper-right tab of the C	er of server roles a er. Click the statu Details page. Clic mn to view all the he displayed dialo t corner of the dia luster Details pag	and the correspond is to go to the Servic ek Abnormal in the e abnormal server r og box. Click View E alog box to go to the ge.	ing status ces tab of Server Role oles in the Details in the Services	
	Server Role Sta	tus Ta	ask Status 🍸]	
	7 in Total Nor	mal S	uccessful		
	38 in Total Ab	Abnormal: 20 F	ailed View Details		
	33 in Total	Samuer Pala			
	38 in Total	tianji.TianjiClient#	Machine Error		
	56 in Total	tianji-sshtunnel-client.SSł	Machine Error		
	56 in Total	nuwa.NuwaConfig#	Machine Error		
	56 in Total	nuwa.NuwaProxy#	The version is inconsistent.		
	11 in Total	EcsTdc.Tdc#	Machine Error		
		EcsNbd.Nbd#	Machine Error		
	4 in Total N	ecs-NcMananer NcDown	M Machine Error Top 20		

Item	Description
Task Status	The status of the task submitted to the cluster. Use $\overline{\mathbb{T}}$ to
	filter the clusters. Click the status to view the task details.

1.4.2.5.2 View the cluster details

You can view the cluster statistics by viewing the cluster details.

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, choose Operations > Cluster Operations.
- 3. Select a project from the Project drop-down list or enter the cluster name in the Clusters field to search for the corresponding cluster.
- 4. Find the cluster whose configurations you are about to view. Click the cluster name or Operations in the Actions column at the right of the cluster to go to the Cluster Details page.

Clusters AcsNodeCluster-A-20191030-	2881		Edit AG Sher	nnong View Cluster Start/Shutdown	
Status: Desired State	Project: ac	s	Region: cn-qingdao-env3b-d02		
Included Server Roles: 14	Included M	achines: 7	Task Status: Successful View		
Clone Mode: Real Clone	System Con	ifiguration: default	Git Version: 6671ee6277039e6f995a13842d6e8eaeeb303		
Security Authentication: Disable	Type: Ordi	nary Cluster			
Services Machines Cluster Configuration Operation Log Cluster Resource Service Inspection All: 6 Normal (6) Reset 2 Deploy Service Services Enter a service name Q					
· ·					
Services	Status	Server Role	Service Template	Actions	
OS Services	Status	Server Role 1 in Total Normal	Service Template	Actions Details Upgrade	
os tianji	Normal Normal	Server Role 1 in Total Normal 1 in Total Normal	Service Template default default	Actions Details Upgrade Details Upgrade	
services os tianji hids-olient	Normal Normal Normal	Server Role 1 in Total Normal 1 in Total Normal 1 in Total Normal 1 in Total Normal	Service Template default default	Actions Details Upgrade Details Upgrade Details Upgrade Details Upgrade	
services os canj hids-client acs-acs_control	Status Normal Normal Normal Normal	Server Role 1 in Total Normal 1 in Total Normal 1 in Total Normal 9 in Total Normal	Service Template default default	Actions Details Upgrade Details Upgrade Details Upgrade Details Upgrade Unpublish Details Upgrade Unpublish	
Services os os banji Nids-client acs-aca_control sanj-dockerdaemon	Normal Normal Normal Normal Normal	Server Role 1 in Total Normal 1 in Total Normal 1 in Total Normal 0 in Total Normal 1 in Total Normal 1 in Total Normal	Service Template default default default	Actions Details Upgrade Details Upgrade Details Upgrade Unpublish Details Upgrade Unpublish Details Upgrade Unpublish Details Upgrade Unpublish	

Area	Item	Description
1	Status	 Desired State: All the clusters of this project reach the desired state. Not Desired State: After the project reaches the desired state for the first time, a server role does not reach the desired state because of undefined reasons.
	Project	The project to which the cluster belongs.
	Region	The region to which the cluster belongs.

Area	Item	Description
	Included Server Roles	The number of server roles included in the cluster.
	Included Machines	The number of machines included in the cluster.
	Task Status	The status of the current task. Click View to view the task details.
		 Successful: indicates the task is successful. Preparing: indicates data is being synchronized and the task is not started yet. In Progress: indicates the cluster has a changing task. Paused: indicates the task is paused Failed: indicates the task failed. Terminated: indicates the task is manually terminated.
	Clone Mode	 Mock Clone: The system is not cloned when a machine is added to the cluster. Real Clone: The system is cloned when a machine is added to the cluster.
	System Configuration	The name of the system service template used by the cluster.
	Git Version	The change version to which the cluster belongs.
	Security Authentication	The access control among processes. 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.

Area	Item	Description
	Туре	 Ordinary Cluster: an operations unit facing to machine groups, where multiple services can be deployed. Virtual Cluster: an operations unit facing to services, which can centrally manage software versions of machines of multiple physical clusters. RDS: a type of cluster that renders special cgroup configurations according to a certain rule. NETFRAME: a type of cluster that renders special configurations for the special scenario of Server Load Balancer (SLB). T4: a type of cluster that renders special configurations for the mixed deployment of e-commerce. Currently, Alibaba Cloud Apsara Stack only has ordinary clusters.
2	Services	 View the statuses of all the services in this cluster. You can also upgrade or unpublish a service. Normal: The service works properly. Not Deployed: No machine is deployed on the service. Changing: Some server roles in the service are changing. Operating: No server role is changing, but the machine where server roles are installed is performing the Operation and Maintenance (O&M) operations. Abnormal: No server role is changing or the machine where server roles are installed is not performing the O&M operations, but the service role status is not GOOD or the version that the service runs on the machine and the version configured in the configurations are different.

Area	Item	Description	
	Machines	View the running statuses and monitoring statuses of all the machines in this cluster. You can also view the details of server roles to which the machine belongs.	
	Cluster Configuration	The configuration file used in the cluster.	
	Operation Log	View the version differences.	
	Cluster Resource	Filter the resource whose details you are about to view according to certain conditions.	
	Service Inspection	View the inspection information of each service in the cluster.	

1.4.2.5.3 View operation logs

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

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. To view the operation logs of a cluster, you can:
 - Enter a cluster name in the search box in the upper-right corner of the page. Click Operations at the right of the cluster to go to the Cluster Details page. Click the Operation Log tab.
 - In the left-side navigation pane, choose Operations > Cluster Operations. On the Cluster Operations page, click Operations in the Actions column at the right of a cluster to go to the Cluster Details page. Click the Operation Log tab.

Services Machines	Cluster Configuration	Operation Log	Cluster Resource	Service Inspection		
Submission Time 12/04/19	12/11/19	Submitter Please input	Q	Services All \vee		Refresh
Description		Operation Type	Status	Git Version	Submitter	Actions
commit by tianji importer			Successful	4f19df6c535c0c718784815e2c49380D1e887fac	aliyuntest Dec 05, 2019, 23:39:18	View Version Differences Details
					total 1 items < 1 >	10/Page v Go to 1 Page

- 3. On the Operation Log tab, view the version differences.
 - a) Click View Version Differences in the Actions column at the right of a log.
 - b) On the Version Differences page, select a basic version from the Versus dropdown list. Then, the contents of the different file are automatically displayed.
 - c) Select each different file from the Different File drop-down list to view the detailed differences.

1.4.2.6 Service operations

1.4.2.6.1 View the service list

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

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. To view the service list, you can:
 - On the Homepage, move the pointer over the Instances section and then click Show Detail in the upper-right corner.
 - In the left-side navigation pane, choose Operations > Service Operations.

Operations / Service Operations				
Services Enter a service name				
Services	Description	Clusters	Included Service Templates	Actions
Ali-tianji-machine-decider		1 in Total Desired State: 1	0	Operations
EcsBssTools		3 in Total Desired State: 3	1	Operations
EcsNbd		5 in Total Desired State: 4 Not Desired State: 1	1	Operations
EcsRiver		3 in Total Desired State: 3	2	Operations
EcsRiverDBInit		1 in Total Desired State: 1	1	Operations
EcsRiverMaster		1 in Total Desired State: 1	1	Operations
EcsStorageMonitor		5 in Total Desired State: 4 Not Desired State: 1	1	Operations
EcsTdo		5 in Total Desired State: 4 Not Desired State: 1	3	Operations
RenderTestService1		0 in Total	0	Operations Delete
RenderTestService2		0 in Total	0	Operations Delete
			total 412 items < 1 2 3 4 42 > 10	Page 🗸 Go to 1 Page

	On this p	oage, you cai	ı view the	following	information.
--	-----------	---------------	------------	-----------	--------------

Item	Description
Services	The service name. Click the service name to view the service details.
Clusters	The number of clusters where the service is located and the corresponding cluster status.
Included Service Templates	The number of service templates this service includes.
Actions	Click Operations to go to the Service Details page.

3. Enter a service name in the search box and then the service that meets the condition is displayed in the list.

1.4.2.6.2 View the server role details

You can view the server role statistics by viewing the server role details.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, choose Operations > Service Operations.
- 3. Enter a service name in the search box and then the service that meets the condition is displayed in the list.
- 4. Click the service name or click Operations in the Actions column.

Servi	ices EcsNbd d Clusters: 5			Included Server Ro	oles: 2		Included Service T	emplates: 1		
Clus	ters Service Tem	plate								
P	roject All	~			Clusters Enter a cluster name	Q		Template	Please sele	et y
Templa	te Please select	~			Tag Please select	~				Batch Add Tag
	Clusters	Region	Status T	Server Role Status	Machine Status		Task Status 🍸	Template		Actions
	ECS-GPU-A-289b ecs	cn-qingdao-env3b-d02	Not Desired State	2 in Total Abnormal:	2 1 in Total Abno	ormal Server Roles: 0 ormal Machines: 1	Successful	TMPL_ECS_V707_TIANJ	LV4 Details	Operations Task Details
	ECS-IO11-A-ac1c ecs	cn-qingdao-env3b-d02	Desired State	2 in Total Normal	5 in Total Norm	ıal	Successful	TMPL_ECS_V707_TIANJ	_V4 Details	Operations Task Details
	ECS-IO7-A-80db ecs	cn-qingdao-env3b-d02	Desired State	2 in Total Normal	6 in Total Norm	sal	Successful	TMPL_ECS_V707_TIANJ	_V4 Details	Operations Task Details
	ECS-IO8-A-288c ecs	cn-qingdao-env3b-d02	Desired State	2 in Total Normal	4 in Total Norm	ıal	Successful	TMPL_ECS_V707_TIANJ	_V4 Details	Operations Task Details
	ECS-IO8-A-28a7 ecs	cn-qingdao-env3b-d02	Desired State	2 in Total Normal	13 in Total Nor	mal	Successful	TMPL_ECS_V707_TIANJ	_V4 Details	Operations Task Details
								total 5 items < 1 >	10/Page	✓ Go to 1 P.

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

Service Details ECS-IO11-A-ac1c / EcsNbd			
Server Role Enter a server role Q			Refresh
EcsNbd.Guestfsd# EcsNbd.Nbd#			
			Diagnostic Mode:
All: 5 Normal (5) Reset			
Machines Enter one or more hostnames/IP addresses Q			Batch Terminal
Machines	Server Role Status	Metric	Actions
a55g01009.cloud.g01.amtest73 10.3.1.90	Normal Details	View	Terminal Restart Server Role
a55g07004.cloud.g07.amtest73 10.3.3.115	Normal Details	View	Terminal Restart Server Role
a65g07112.cloud.g08.amtest73 10.3.3.116	Normal Details	View	Terminal Restart Server Role
a58g07211.cloud.g00.amtest73 10.3.3.117	Normal Details	View	Terminal Restart Server Role
a56g07215.cloud.g09.amtest73 10.3.3.118	Normal Details	View	Terminal Restart Server Role
		total 5 ite	ms < 1 > 10/Page v Go to 1 Page

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

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

1.4.2.7 Machine operations

You can view the machine statistics by viewing the machine list.

- **1.** Log on to Apsara Infrastructure Management Framework.
- 2. To view the machine list, you can:
 - On the Homepage, move the pointer over the Machines section and then click Show Detail in the upper-right corner.
 - In the left-side navigation pane, choose Operations > Machine Operations.

Operations)	Machine Operations						
Mac	hines						
Proje	ect All	V Clusters Enter a cli	uster name Q Ma	chines Enter one or more host	names/IP addresses Q		Batch Terminal
	Hostname	Clusters	Project	Region	Status 🕎	Machine Metrics	Actions
	a56g01001.cloud.g01.amtest73	tianji-A-2898	tianji	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management \geq
	a58g01002.cloud.g01.amtest73	AliguardCluster-A-20191030- 2895	yundun-advance	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~
	a56g01003.cloud.g01.amtest73	slbCluster-A-20191030-2865	slb	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~
	a56g01004.cloud.g01.amtest73	tianji-A-2898	tianji	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~
	a58g01005.cloud.g01.amtest73	BasicNcCluster-A-20191030- 288d	asto	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~
	a56g01006.cloud.g01.amtest73	BasicNcCluster-A-20191030- 288d	asto	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~
	a58g01007.cloud.g01.amtest73	ads-A-20191205-354e	ads	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~
	a58g01008.cloud.g01.amtest73	ads-A-20191205-354e	ads	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management >
	a56g01009.cloud.g01.amtest73	ECS-IO11-A-ac1c	ecs	cn-qingdao-env3b-d02	Normal Details	View	Operations Terminal Machine Management ~

3. Select a project or enter the cluster name or machine name to search for the corresponding machine.

Item	Description				
Hostname	Click the hostname to go to the Machine Details page.				
Status	The current status of the machine. Use 🕎 to filter the				
	machines. Click Details and then the Status Details of				
	Machine dialog box appears.				
Machine Metrics	Click View and then the Metrics dialog box appears.				
	Metrics X				
	Description (Markov, Markov,				
	Server Role Metric Machine Metrics				
	Server kole Metric Alert status Update Ime Lescription ads-service WorkerMonitor# postcheck monitor good Dec 11, 2019, 13.20.11				
	ads-service Workenfontion# smm_app_process good Dec 11, 2019, 14, 42, 45				
	ads-service Workenflorition# smm_process_oproup good Dec 11, 2019, 14:42:45				
	ads-service. WorkerMonitor# smm_actual_version good Dec 11, 2019, 14 42 45				
	ads-service.WorkerMontor# fianj_app_process good Dec 11, 2016, 13.12.23				
	ade-service WorkerMonton# sm_local_sr pood Dec 11, 2019, 15:13:22				
	ads-service.WorkerMonitor# sm_overwrite_app good Dec 11, 2019, 15:13:22				
	The Server Role Metric tab and Machine Metrics tab display the corresponding metrics, and you can view the specific alert status and updated time. Enter a keyword in the search box in the upper-right corner to search for a server role or metric. You can also select the status to filter metrics.				
Actions	 Click Operations to go to the Machine Details page. Click Terminal to log on to the machine and perform operations. You can select multiple machines and then click Batch Terminal in the upper-right corner to log on to multiple machines at a time. Click Machine Management to perform an out-of-band restart operation on the machine. 				

1.4.2.8 Monitoring center

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

1.4.2.8.1 View the monitoring instance status

You can view the status of a monitoring instance after it is deployed.

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, choose Operations > Service Operations.
- 3. Enter a service name in the search box to search for the corresponding service.
- 4. Click Operations in the Actions column at the right of the service.
- 5. On the Clusters tab, configure the conditions and then search for the cluster. Click Operations in the Actions column.
- 6. On the Cluster Details page, select the server role you are about to view and then click View in the Metric column. Then, view the statuses of server role metrics and machine metrics.

Services Machines Cluster Con	ifiguration Operation Log Cluster Resource Service Inspection	n					
Server Role Enter a server role C			Refresh				
EcsRiver.RiverCluster# EcsRiver.RiverClusterDBManager# EcsRiver.RiverServer#							
All: 3 hierard (2)			Diagnostic Mode:				
Machines Enter one or more hostnames/IP ac	ddresses Q		Batch Terminal				
Machines	Server Role Status	Metric	Actions				
	Normal Details	View	Terminal Restart Server Role				
	Normal Details	View	Terminal Restart Server Role				
	Normal Details	View	Terminal Restart Server Role				
		tota	il 3 items < 1 > 10/Page v Go to 1 Page				

1.4.2.8.2 View the alert status

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

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, choose Monitoring. On the Monitoring page, click Go to open the target page.
- 3. In the top navigation bar, choose Monitoring > Alert Status.

Alert Status							
Service All	•	Cluster All	Enter an alert name).	Time Range 12/10/19, 20:10:00 ~ 12/11/19, 20	0:10:00 Searc	h
Service	Cluster	Instance	Alert Status	Alert Level	Alert Name	Alert Time	Actions
tianji	slbCluster-A	cluster=sibCluster-A-20191030-2885,host=a	• Alerting	PI	memo_cluster_host	11/23/19, 13:03:00 Lasted for 18 Days 7 Hours 7 Minutes 35 Seco nds	Show
tianji	slbCluster-A	cluster=sibCluster-A-20191030-2885,host=a	Alerting	P1	memo_cluster_host	11/23/19, 13:03:00 Lasted for 18 Days 7 Hours 7 Minutes 35 Seco nds	Show
tianji	mongodb-A	cluster=mongodb-A-20191030-289a,host=a5	Alerting	PI	memo_cluster_host	11/23/19, 13:04:00 Lasted for 18 Days 7 Hours 6 Minutes 35 Seco nds	Show
tianji	mongodb-A	cluster=mongodb-A-20191030-289a,host=a5	Alerting	PI	memo_cluster_host	11/23/19, 13:04:00 Lasted for 18 Days 7 Hours 6 Minutes 35 Seco nds	Show

- 4. You can configure the service name, cluster name, alert name, or the time range when the alert is triggered to search for alerts.
- 5. On the Alert Status page, view the alert details. For more information about the alert status descriptions, see the following table.

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Instance	The name of the service instance being monitored. Click the instance to view the alert history of this instance.
Alert Status	Alerts have two statuses: Restored and Alerting.
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. • P1 • P2 • P3 • P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered and how long the alert has lasted.
Actions	Click Show to show the data before and after the alert time.

1.4.2.8.3 View alert rules

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

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

3. In the top navigation bar, choose Monitoring > Alert Rules.

Alert Rules							
Service All Clus	ster All	Enter an alert name.	Searc	h			
Service Cluster	Alert Name	Alert Conditions	Periods	Alert Contact	Status		
yundun-semawaf	semawaf_check_disk	\$Use>90	60	*	Running		
yundun-semawaf	semawaf_check_disk	\$Use>90	60	*	Running		
yundun-semawaf	app_vip_port_check_serverrole	Sstate!=0;Sstate!=0	60	*	Running		
yundun-semawaf	alert_ping_yundun-soc	\$rta_avg>500 \$loss_max>80;\$rta_avg>400 \$loss_max>60	60	*	Running		
yundun-consoleservice	check_auditLog_openapi	\$totalcount>9	300	*	Running		
yundun-consoleservice	check_sas_openapi	\$totalcount>9	300	*	Running		
yundun-consoleservice	check_aegis_openapi	\$totalcount>9	300	*	Running		
yundun-consoleservice	check_secureservice_openapi	\$totalcount>9	300	-	Running		
yundun-consoleservice	consoleservice_check_disk	long(\$size)>20971520	60	-	Running		
yundun-consoleservice	check_aegis_openapi	Stotalcount>9	300	*	Running		

- 4. You can configure the service name, cluster name, or alert name to search for alert rules.
- 5. On the Alert Rules page, view the detailed alert rules. For more information about the alert rule descriptions, see the following table.

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

1.4.2.8.4 View the alert history

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

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

3. In the top navigation bar, choose Monitoring > Alert History.

Alert His	story All	Notifica	tions S	uppressions								
Service	Ali		•	Cluster A	Al	-						
1 Hour	12 Hours	1 Day	1 Week	1 Month	3 Months	Custom	12/10/19, 20:16:00 ~ 1	2/11/19, 20:16:00				Search
Service		Cluster		Alert Instar	nce		Status	Alert Level	Alert Name	Alert Time	Alert Contact	Actions
drds-cons	sole			service=drds	s-console,serverro	e=drds-cons	sol ORestored	Restored	tianji_drds_prectrl_check_url	12/10/19, 20:38:13		Show
EcsTdc		ECS-IO8-A	-288c	cluster=ECS	-IO8-A-288c,serve	rrole=EcsTo	do OAlerting	P4	ecs_server_compute-cpu_usa ge	12/10/19, 20:39:49		Show
EcsTdc		ECS-IO8-A	-288c	cluster=ECS	-IO8-A-288c,serve	rrole=EcsTo	do ORestored	Restored	ecs_server_compute-cpu_usa ge	12/10/19, 20:41:49		Show
aso-syste	emMgr			service=aso	-systemMgr,server	role=aso-sy	ste OAlerting	P1	tianji_aso_auth_check_url	12/10/19, 21:48:28	*	Show
ecs-houyi		ECS-HOU'	YIRE	cluster=ECS	-HOUYIREGION-	A-28a2,serv	err OAlerting	P4	ecs-houyi_ecs_regionmaster- unknow_error	12/10/19, 21:57:39	*	Show
ecs-houyi		ECS-HOU'	YIRE	cluster=ECS	HOUYIREGION-	A-28a2,serv	err ORestored	Restored	ecs-houyi_ecs_regionmaster- unknow error	12/10/19, 22:08:39	*	Show

- 4. You can configure the service name, cluster name, time range, or period to search for alerts.
- 5. On the Alert History page, view the history alerts. For more information about the history alert descriptions, see the following table.

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

1.4.2.9 View tasks

The task list allows you to view the submitted tasks and the corresponding status.

Procedure

1. Log on to Apsara Infrastructure Management Framework

- 2. To view the task list, you can:
 - In the left-side navigation pane, choose Tasks > My Tasks.
 - In the left-side navigation pane, choose Tasks > Related Tasks.
- 3. You can use 🕎 to filter tasks in the Status column.
- 4. Find the task whose details you are about to view and then click the task name or click Details in the Actions column.
- 5. On the Task Details page, view the status and progress of each cluster and server role.

Tasks / My Tasks / Task Details							
Summary Task Status: Successful Duration: 12 minutes	Submission Time: Dec Task Description: Rem	11, 2019, 20:25:32 oveMachine: ['iZh5	i05w9770q3zmqilt	Submitter: aliyı xdZ', 'iZh5i066934zp0of5l54mzZ'	untest , 'iZh5i05w9770q3zmqilbxcZ', 'iZ	Ref Zh5i05w9770q3z	fresh
Server Role All v							
Clusters Q Region T	Status		Progress		Start Time	Actions	
hbase-A-20191210-ac17 cn-qingdao-env3b-d02	Successful		🕢 Build — (Change	Dec 11, 2019, 20:25:32	View Version Differences Operation Log	
					total 1 items < 1 > 10/	/Page v Go to 1	Page
Change Details Clusters hbase-A-20191210-ac17	Service	Upgrade (8)	Machine Change (4)			
Server Role 🔾	Services T	Upgrade Type	Status T	Progress		Actions	
rds-hbase.DbHBase# 🔗	rds-hbase	Configuration Change	Successful	🕑 Download — 🕑 Upgrade		Details	
rds-hbase.Dblnit# 🔗	rds-hbase	Configuration Change	Successful	🕑 Download — 🕑 Upgrade		Details	
rds-hbase.InitCluster# 🔗	rds-hbase	Configuration Change	Successful	🕑 Download — 🕑 Upgrade		Details	

1.4.2.10 Reports

1.4.2.10.1 View reports

The Reports module allows you to view the statistical data.

Context

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

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

Procedure

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

All Reports Favorites						
Fuzzy Search		Q				Permission Management C Refresh
Report 🖸	Group 🗹	Status	Public	Created By 🔽	Published At	Actions
XDB Instance Metric Info	Tianjimon	Published	Public	admin	Published at : 11/13/19, 23:46:28 Created at : 11/13/19, 23:46:28	Add to Favorites Request Group Permission
Alert Status Profile	Tianjimon	Published	Public	admin	Published at : 10/30/19, 13:14:46 Created at : 10/30/19, 13:14:46	Add to Favorites Request Group Permission
Server Role Action Statuses	Tianji	Published	Public	admin	Published at : 10/30/19, 13:14:46 Created at : 10/30/19, 13:14:46	Add to Favorites Request Group Permission
Machine and Server Role Statuses	Tianji	Published	Public	admin	Published at : 10/30/19, 13:14:48 Created at : 10/30/19, 13:14:48	Add to Favorites Request Group Permission

Item	Description
Report	The report name.
	Move the pointer over the down-arrow button next to Report to search for reports by report name.
Group	The group to which the report belongs.
	Move the pointer over the down-arrow button next to Group to filter reports by group name.
Status	Indicates whether the report is published. • Published
	• Not published
Public	Indicates whether the report is public.
	\cdot Public: All of the logon users can view the report.
	• Not public: Only the current logon user can view the report.
Created By	The person who creates the report.
Published At	The time when the report is published and created.
Actions	Click Add to Favorites to add this report to your favorites. Then, you can view the report by choosing Reports > Favorites in the top navigation bar.

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

3. Optional: Enter the name of the report that you are about to view in the search box.

4. Click the report name to go to the corresponding report details page.

For more information about the reports, see Appendix.

1.4.2.10.2 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. In the left-side navigation pane, choose Reports. On the Reports page, click Go to open the target page.
- 3. Enter the name of the report that you are about to add to favorites in the search box.
- 4. At the right of the report, click Add to Favorites in the Actions column.
- 5. In the displayed Add to Favorites dialog box, enter tags for the report.
- 6. Click Add to Favorites.

1.4.2.11 Tools

1.4.2.11.1 Machine tools

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

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

Operation scene	Description	Action
Scene 1: NC Scale-out (with existing machines)	Scales out an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 2: Host Scale-out (with existing machines)	Scales out the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 3: NC Scale-in	Scales in an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 4: Host Scale-in	Scales in the DockerHost #Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.

3. Select the operation scene according to actual situations.

Operation scene	Description	Action
Scene 5: VM Migration	Migrates virtual machines (VMs) from a host to another host.	Select a source host and a destination host. Select the VMs to be migrated in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 6: Host Switching	Switches from a standby host to a primary host.	Select a source host and a destination host. Click Submit and then click Confirm in the displayed dialog box.

1.4.2.11.2 IDC shutdown

If you are about to maintain the IDC or shut down all of the machines in the IDC, you must shut down the IDC.

Prerequisites



This is a high-risk operation, so proceed with caution.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, choose Tools > IDC Shutdown, and then click Go to open the target page.
- 3. On the Clusters Shutdown page, click Start Shutdown to shut down all of the machines in the IDC with one click.

1.4.2.12 Appendix

1.4.2.12.1 Project component info report

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

Item	Description
Project	The project name.

Item	Description
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

1.4.2.12.2 IP list

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

IP List of Physical Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.

IP List of Docker Applications

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.

Item	Description
Docker IP	The Docker IP address.

1.4.2.12.3 Machine info report

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

Machine Status

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

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

Expected Server Role List

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

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

Abnormal Monitoring Status

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

Item	Description
Machine Name	The machine name.

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

Server Role Version and Status on Machine

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

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

Monitoring Status

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

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

1.4.2.12.4 Rolling info report

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

Choose a rolling action

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

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

Server Role in Job

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

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

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

Server Role Rolling Build Information

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

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

Server Role Statuses on Machines

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

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

1.4.2.12.5 Machine RMA approval pending list

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

Machine

Displays the ba	sic information	n of pending app	roval machines.
-----------------	-----------------	------------------	-----------------

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

Machine Serverrole

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

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

Machine Component

Displays the hard disk information of pending approval machines.

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

1.4.2.12.6 Registration vars of services

This report displays values of all service registration variables.

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

1.4.2.12.7 Virtual machine mappings

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

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

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

1.4.2.12.8 Service inspector report

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

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

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

1.4.2.12.9 Resource application report

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

Change Mappings

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

Changed Resource List

Item	Description
Res	The resource ID.
Туре	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.
Parameters	The resource parameters.

Item	Description
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
АРР	The application of the server role.
Name	The resource name.
Туре	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 message.

1.4.2.12.10 Statuses of project components

This report displays the status of all server roles in an abnormal status on machines of the project, and the monitoring information (alert information reported by the

server role to Apsara Infrastructure Management Framework monitor) of server roles and machines.

Error State Component Table

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

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

Server Role Alert Information

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

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

Machine Alert Information

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

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

Service Inspector Information

Select a row in the Error State Component Table section to display the

corresponding information in the list.

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

1.4.2.12.11 Relationship of service dependency

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

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Dependent Service	The service on which the server role depends.
Dependent Server Role	The server role on which the server role depends.

Item	Description
Dependent Cluster	The cluster to which the dependent server role belongs.
Dependency in Final Status	Whether the dependent server role reaches the final status.

1.4.2.12.12 Check report of network topology

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

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

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

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

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

1.4.2.12.13 Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

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

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

1.4.2.12.14 Auto healing/install approval pending report

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

1.4.2.12.15 Machine power on or off statuses of clusters

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

Cluster Running Statuses

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

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

Server Role Power On or Off Statuses

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

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

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
Action Name	The startup or shutdown action that is being performed by the server role.
Action Status	The status of the action.

Statuses on Machines

Displays the running status of the selected server role on machines.

Select a row in the Server Role Power On or Off Statuses section to display the information of the corresponding server role in the list.

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

Machine Statuses

Displays the running statuses of machines in the selected cluster.

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

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

1.5 Network operations

1.5.1 What is Apsara Network Intelligence?

Apsara Network Intelligence is a system to analyze network traffic. 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.
- · Configure reverse access to cloud services.
- Configure leased lines through graphical interfaces and set up active and standby routers.
- Query the tunnel VIPs of cloud services.
- Create Layer 4 listeners.

1.5.2 Log on to the Apsara Network Intelligence console

This topic describes how to log on to the Apsara Network Intelligence console.

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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 1-10: Log on to ASO

Enter a user name
Enter the password
Log On

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 log on to ASO.
- 5. In the left-side navigation pane, click Products. On the right side of the page, click Apsara Network Intelligence.

1.5.3 Query information

You can enter an instance ID to query details of the instance.

Procedure

1. Log on to the Apsara Network Intelligence console.

- 2. Enter the ID of a VPC or an SLB instance to query details.
 - Enter the ID of a VPC to query VPC, VRouter, and VSwitch details.
 - VPC details

PC Resources / VPC Details									
Basic Information Subresource Information									
Configuration Information									
VPC ID	RegionNo	Status	Attached CENID	TunnellD					
vpc-q8c44na	cn-qingdao-env8d-d01	Created	None	24					
Created At	Modified At	Name	Description	Created by User					
2019-05-29 11:51:17	2019-05-29 11:51:21	muyan_vpc	None	Yes					
Enable ClassicLink	CIDR Block	User CIDR	Actions						
No	172.16.0.0/16	Details	Details						

- Information about VRouters, route tables, router interfaces, and VSwitches
- Enter the ID of an SLB instance to query instance details.
 - Information about SLB instance configurations, VIPs, specifications, and users

PR Resource / Stall Instance Details								
Instance Information Listence Information								
Configuration Information								
LB ID	Cluster	EIP Type		Gateway Type		SLB Mode	status	
Ib-q8ckib	cn-qingdao-env8d-d01	intranet		classic		fnat	active	
LVSs	Proxies	Created At		Modified At		After WAF/Anti-DDoS Protection	Actions	
				No data				
Cleaning Threshold				Black Hole Threshold	d			
None				None				
VIP(EIP)Information								
VIP(EIP)	Status	Tunnel ID		Service Unit Name		Primary IDC/LVS Name	Secondary IDC/LVS Name	
				No clata				
Specifications Information								
VIP MAX CONN LIMIT	VIP OUT bit/s	VIP IN bit/s	VIP QPS		VIP CPS	Specifications	Instance Type	
				No clata				
User Information								
User ID								
				No data				

- Listener information

Click Show in the Back-end Server/Health Check column to view details on backend servers.

VPC R	Checures / Ski Instance Details										
1	batanca Information Listenera Information										
	inter filter conditions.										
	Listener ID	Protocol	Frontend Port	Use Server Group	Use Primary/Secondary Server Group	Proxy Port	Port Redirection	Status	Back-end Server/Health Check	Created At	Modified At
	lb-q8ckibpdtql	tcp	80	No	No	None	None	running	Show	2019-05-16 03:14:45	2019-05-16 03:14:56
	Ib-q8ckibpdtql	top	22	No	No	None	None	running	Show	2019-05-16 03:14:36	2019-05-16 03:14:56
4											÷.

1.5.4 Manage cloud service instances

You can create a cloud service in a region or query the instance information of a region.

Procedure

- **1.** Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Virtual Private Cloud > VPC Instance Type Management.
- 3. Select the region from the Select Region drop-down list for which you want to create a cloud service instance. All cloud service instances in the specified region are displayed.
- 4. Click Add to add a cloud service type.

1.5.5 Tunnel VIP

1.5.5.1 Create a Layer-4 listener VIP

You can create Layer-4 listener VIPs to forward traffic for cloud services in your VPC.

Procedure

- 1. Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Server Load Balancer > VIP Management.
- 3. Click Create VIP.
- 4. On the Create VPC Instance tab, select Cloud Service, CIDR Type, and Tunnel Type.

The tunnel types are listed as follows:

- singleTunnel: specifies a single tunnel VIP that allows ECS instances in a single VPC to access external cloud services.
- anyTunnel: specifies a tunnel VIP that allows ECS instances in all VPCs to access a specified cloud service.
- 5. Click Create. On the Create SLB Instance tab, select a primary data center or use the default data center.

- 6. Click Create. On the Add Band-end Server to SLB Instance tab, configure the following parameters as needed:
 - VPC ID: specifies the ID of the VPC to which target ECS instances belong. This parameter must be configured if the network type of the ECS instances is VPC.
 - Back-end Servers: specifies the backend servers that you want to add. You
 can enter the information of only one backend server on each line. A backend
 server information entry contains the 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 100 is used.

SLB VIP Application				
Create VPC Instance		Create SLB Instance	Add Back-end Server to SLB Instance	Create Listener
	VPC ID : Back-end Servers:	Enter a VPC ID to add a VM in the VPC. 192,168.0.00 192,168.0.0		
				ОК

- 7. Click Create. On the Create SLB instance tab, select a primary data center or use the default data center.
- 8. Click OK. On the Create Listener tab, click Add to configure a UDP or TCP listener. Then, click Submit.
- 9. On the Publish Online tab, 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.

1.5.5.2 Query the tunnel VIP of a cloud service

You can query information such as creation time, connectivity, and VIP for cloud services that have Server Load Balancer (SLB) VIPs.

- **1.** Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Server Load Balancer > VIP Management.

3. On the Tunnel VIP Management page, select Region ID, Cloud Service, and Status. Click Search.

Tunnel VIP Management	unnel VIP Management										
* Region ID : on-qingdao-env8d	-801	~	V Cloud Service: gts V		\sim	Status: Running			\sim		
Enter filter conditions.										Search	
Region	Cloud Service	Cloud Instance I D	SLB Instance ID	LB VIP	Status	Created At	Modified At	Modified By	Port C onnec tivity	Actions	
cn-qingdao-erw8d-d01	gto	en-gingdao-env8	16h1h1d3436-cn-oinodao-env8d-	10.68	workin 9	2019-06-03 10:17:38	2019-06-03 10:18:06	aliyuntest	80	Actions \checkmark	
										< 1 >	

1.5.6 Create a Direct Any Tunnel VIP

You can create Direct Any Tunnel VIPs for cloud services in your VPC to allow traffic forwarding through XGW.

Procedure

- 1. Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Server Load Balancer > Direct Any Tunnel VIP Management.
- 3. On the Direct Any Tunnel VIP Management page, click Create Direct Any Tunnel VIP.
- 4. On the Create Direct Any Tunnel VIP page, configure the parameters for the Direct Any Tunnel VIP.



5. Click Create. Cloud service instances that have Direct Any Tunnel VIPs can forward traffic through XGW.

1.5.7 Leased line connection

1.5.7.1 Overview

You can connect a VPC to an IDC through a leased line.

Before connecting to a VPC through a leased line, you must confirm the initial CSW configurations meet the following conditions:

 \cdot You have uploaded the licenses required for VLAN functions onto the CSWs.

- You have set the management IP address on the loopback 100 interface of each CSW.
- You have configured the CSW uplink interfaces to ensure interoperability with the Layer 3 interfaces used by VPC APIs.
- You have deleted the default configuration of bridge-domain.
- You have enabled NETCONF and STelnet for CSWs. The configuration details are included in the CSW initial configuration template.
- You have configured the service type of CSW interfaces to tunnel.

You must also obtain the following account information:

- BID: specifies the ID of the account group. The BID for Mainland China users is 26842, and the BID for international users is 26888.
- UID: specifies the ID of the account to which the destination VPC belongs.

1.5.7.2 Manage an access point

Access points are Alibaba Cloud data centers located in different regions. Each region contains one or more access points. This topic describes how to query and modify information about access points of a region.

Query access point information

Perform the following steps to 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.

Access Points												
* Region: cn-qing	dao-env8d-d01 🗸	Access Point ID: ap-										
Search	Reset											
Access Point Id	Managing Region	Physical Region	Туре	Status	Name	Description	Physical Location	IDC Operator	Created At	Modified At	Actions	
ap-cn-qingdao-env8d-	cn-qingdao-env8d-d01	None	VPC	recommended	ap-cn-qingdao-	ap-cn-qingdao-env8d-	AMTEST61	Other	2019-04-30 06:50:00	2019-04-30 06:50:0	Modify	Show Details
•												÷
1-1/1												< 1 >

Modify access point information

Perform the following steps to modify the information about an access point:

- 1. Click Modify in the Actions column corresponding to an access point that you want to modify.
- 2. Modify access point information.
- 3. Click Modify.

The parameters are described as follows:

- Access Point Location: specifies the physical location of an access point. You can specify this parameter as needed.
- · Access Point IDC Operator: specifies the name of the data center operator.

Modify Access Point	×
* Access Point ID: ap-cn-qingdao-env8d-	
* Enter an access point name ap-cn-qingdao-env8	
* Description: ap-cn-qingdao-env	
* Access Point Status: 💿 Available i Busy 💿 Full i Unavailable	
* Access Point Location: AMTEST61	
* Access Point IDC Operator: Other	
Physical Region :	\sim
Modify Cancel	

1.5.7.3 Manage an access device

This topic describes how to query and modify information about access devices of a region.

Query access device information

Perform the following steps 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 Devices.

4. Enter the region and device ID of an access device that you want to query.



If Device ID is not set, the information about all devices in a region is queried.

5. Click Search.

Access Devices											
* Region: cn-qingdad	o-env8d-d01 🗸	Device ID: CSW-VM-V									
Search	Reset										
Device ID	Region	Access Point ID	Device Status	Physical Location	Access Method	Device Name	Description	Created At	Modified At	Actions	
CSW-VM-VPC-C	cn-qingdao-env8d-d01	ap-cn-qingdao-env8d-	available	AMTEST61	vlanToVxlanRo uting	CSW-VM-VPC-G1-	CSW-VM-VPC-G1-	2019-04-29 22:50:32	2019-04-29 22:50:32	Modify S	how Details
4											÷
1-1/1											< 1 >

6. Click Show Details in the Actions column to view the details of the access device.

Modify access device information

Perform the following steps to modify the information about an access device:

1. Click Modify in the Actions column corresponding to a device that you want to modify.

2.	Follow the on-screen	prompts to modify	the device information.
----	----------------------	-------------------	-------------------------

Modify Access Device	×
* Device ID:	CSW-VM-VPC-G1
* Region :	cn-qingdao-env8d-d01 v
* Device Status:	● Available 🔵 Full 🔵 Unavailable
* Access Device Location :	AMTEST61
* Specify whether to use XN	● Yes 🔘 No
* XNET Endpoint URL:	http://xnet.en
* XNET Device ID:	1
* Outer Source IP Encapsula	10.48
* Inner Source MAC Encapsu	00-00-5E-00-01-02
Device Management IP Addı	10.48.
Device Manufacturer:	Ruijie
Device Model:	RG-S6220
Device Name:	CSW-VM-VPC
Device Description:	CSW-VM-VPC-
	Modify Cancel

3. Click Modify.

1.5.7.4 Establish a leased line connection

A leased line can be obtained from a telecom operator to establish a physical connection between your on-premises data center and an Alibaba Cloud access point. This topic describes how to establish a leased line connection and query leased line information of a region.

- **1.** Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Express Connect > Network Environment Management.

3. Choose Network Environment Management > Leased Lines. On the page that appears, click Create Leased Line.

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 same as the CSW host name.
- Device Port: optional. If specified, the device port number must be the same as the CSW port number.
- UID: the ID of the account to which a destination VPC belongs.
- Access Point ID: the ID of the region where your data center is located.
- Redundant Leased Lines: a previously obtained leased line, to act as a redundancy for the leased line you are creating.

Name: The leased line name. It can be 2 to 128 characters in length and cannot description: The leased line description. It can be 2 to 128 characters in length and * BID: 26842 * UID: EnterUID * Region: cn-qingdao-env8d-d01 * Region: cn-qingdao-env8d-d01 * Region: cn-qingdao-env8d-d01 * Region: cn-qingdao-env8d-d01 * Access Point Type: VPC Access Point ease as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: VPC Access Point * VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start wit Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. alue range: (2-10000). * Port Type: Select V You can leave it empty if the value is unknown. Redundant Leased Lines Enteruid, bid,regionId • When establishing the second leased line, you can specify it as redundant one and upload its ID. If you do so, Alibaba Cloud al cates a separate access device for higher availability. <th>Create Leased Line</th> <th></th> <th>×</th>	Create Leased Line		×
Name: The leased line name. It can be 2 to 128 characters in length and cannot Description: The leased line description. It can be 2 to 128 characters in length and * BID: 26842 * UID: EnterUID * Region: cn-qingdao-env8d-d01 The region ID is used for managing access devices (which is not n cessarily the same as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: VPC Access Point • VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start wit Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s.' alue range: [2-10000]. * Port Type: Select V You can leave it empty if the value is unknown. Redundant Leased Lines Enteruid, bid,regionId • When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud al cates a separate access device for higher availability.			
Description: The leased line description. It can be 2 to 128 characters in length and * BID: 26842 * UID: EnterUID * Region: cn-qingdao-env8d-d01 * Region: cn-qingdao-env8d-d01 The region ID is used for managing access devices (which is not n cessarily the same as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: • VPC Access Point • VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. alue range: [2-10000]. * Port Type: Select v You can leave it empty if the value is unknown. Redundant Leased Lines; Enteruid, bid/regionId • When establishing the second leased line, you can specify it as redundant one and upload its ID. If you dos o, Alibaba Cloud al cates a separate access device for higher availability.	Name:	The leased line name. It can be 2 to 128 characters in length an	d cannc
 * BID: 26842 * UID: EnterUID * Region: cn-qingdao-env8d-d01 The region ID is used for managing access devices (which is not n cessarily the same as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: • VPC Access Point VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. alue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines Enteruid, bid,regionId When establishing the second leased line, you can specify it as redundant one and upload its ID. If you do so, Alibaba Cloud al cates a separate access device for higher availability. 	Description:	The leased line description. It can be 2 to 128 characters in leng	th and
 * UID: EnterUID * Region: cn-qingdao-env8d-d01 The region ID is used for managing access devices (which is not n cessarily the same as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: • VPC Access Point VPC -VPC access point, for leased lines that can access VPC netorks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. 'alue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines Enteruid, bid, regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 	* BID :	26842	
 * Region: cn-qingdao-env8d-d01 The region ID is used for managing access devices (which is not n cessarily the same as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: VPC Access Point VPC - VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start wit Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. 'alue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid, regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud al cates a separate access device for higher availability. 	* UID :	EnterUID	
The region ID is used for managing access devices (which is not n cessarily the same as the attached region ID of the access device, ut must be the same as the region ID of the access point). * Access Point Type: • VPC Access Point • VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID: Access Point ID: Device Name: Device names can be 2 to 256 characters in length and cannot start wi Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. 'alue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines Enteruid, bid,regionId • When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud al cates a separate access device for higher availability.	* Region:	cn-gingdao-env8d-d01	~
 * Access Point Type: VPC Access Point VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. 'alue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 	. negioni	The region ID is used for managing access devices (which is cessarily the same as the attached region ID of the access o ut must be the same as the region ID of the access point).	s not ne device, b
 VPC -VPC access point, for leased lines that can access VPC net orks * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. Value range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId When establishing the second leased line, you can specify it as redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 	* Access Point Type:	VPC Access Point	
 * Access Point ID: Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. Value range: [2-10000]. * Port Type: Select V You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 		 VPC -VPC access point, for leased lines that can access V orks 	PC netw
Access Point ID Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. 'alue range: [2-10000]. * Port Type: Select Vou can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId • When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability.	* Access Point ID :		
Device Name: Device names can be 2 to 256 characters in length and cannot start with Device Port: CSW Port Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. Talue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId • When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability.		Access Point ID	
Device Port: CSW Port Bandwidth: [2-10000] The inbound interface bandwidth of the leased line. Unit: Mbit/s. alue range: [2-10000]. * Port Type: Select Vou can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId • When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability.	Device Name:	Device names can be 2 to 256 characters in length and cannot	start wit
Bandwidth: [2-10000] Mbps The inbound interface bandwidth of the leased line. Unit: Mbit/s. 'alue range: [2-10000]. * Port Type: Select V You can leave it empty if the value is unknown. You can leave it empty if the value is unknown. Mbps Redundant Leased Lines Enteruid, bid, regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability.	Device Port:	CSW Port	
The inbound interface bandwidth of the leased line. Unit: Mbit/s. alue range: [2-10000]. * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId • When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability.	Bandwidth:	[2-10000]	Mbps
 * Port Type: Select You can leave it empty if the value is unknown. Redundant Leased Lines: Enteruid, bid,regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 		The inbound interface bandwidth of the leased line. Unit: Nalue range: [2-10000].	/bit/s. V
 You can leave it empty if the value is unknown. Redundant Leased Lines Enteruid, bid,regionId When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 	* Port Type:	Select	~
 Redundant Leased Lines: Enteruid, bid,regionId When establishing the second leased line, you can specify it as redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 		You can leave it empty if the value is unknown.	
 When establishing the second leased line, you can specify it as redundant one and upload its ID. If you do so, Alibaba Cloud all cates a separate access device for higher availability. 	Redundant Leased Lines	Enteruid, bid,regionId	
The leased line that you specify must exist and be in Allocated, onfirmed, or Enabled status. Create Cancel		 When establishing the second leased line, you can specify redundant one and upload its ID. If you do so, Alibaba Clicates a separate access device for higher availability. The leased line that you specify must exist and be in Allo onfirmed, or Enabled status. 	y it as a loud allo cated, C

When the leased line state is Confirmed, the line is created.

5. On the Leased Lines page, find the created leased line and choose Actions > Enable.

If the allocation process for a leased line persists for several minutes after you click Enable, choose Products > Network Controller > Business Foundation System Flow. On the page that appears, set Instance ID to the leased line ID, set Step Status to All, and click Search. Check the flow status in the search results. A flow in red indicates that the corresponding step has failed. Click Resend to restart the task, 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 availability group (AG). If an error is returned, you can check the xnet service logs to troubleshoot the issue based on the returned error.

1.5.7.5 Create a VBR

A virtual border router (VBR) is a router between customer-premises equipment (CPE) and a VPC, and functions as a data forwarding bridge from a VPC to an onpremises IDC. This topic describes how to create a VBR in a region and query VBR information of the region.

- 1. Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Express Connect > Network Environment Management.
- 3. Choose Network Environment Management > VBRs.

4. Click Create VBR.

reate VBR	>
* BID:	26842
* UID:	EnterUID
* Region:	cn-qingdao-env8d-d01 v
2	The ID of the region to which the instance belongs.
* Leased Line ID:	
* VLAN ID:	[1, 2999]
	The VLAN of the VBR leased line interface. • VLAN : [1, 2999]
	 Only the leased line owner can specify or modify VLAN.
Local Gateway IP Addre	Local Gateway IP Address
	 The local IP address of the leased line interface. It is required when the interface status is not waiting. Only the VBR owner can specify or modify the local IP address.
Peer Gateway IP Addre	Peer Gateway IP Address
	 The peer IP address of the leased line interface. It is required when the interface status is not waiting. Only the VBR owner can specify or modify the local IP address.
* Subnet Mask:	[(255.255.255.0) - (255.255.255.252)]
	 The subnet mask for the connection between the local IP address s and peer IP address. It is required when the interface status is not waiting.
	• Only the VBK owner can specify or modify the local IP address.
Name:	EnterName
	The leased line name. It can be 2 to 128 characters in length and c annot start with http:// or https://.
Description:	EnterDescription
	The leased line description. It can be 2 to 128 characters in length and cannot start with http:// or https://.
ownerBid:	EnterownerBid
ownerAliUid:	EnterownerAliUid

5. Follow the on-screen prompts to configure the VBR parameters.

The parameters are described as follows:

- Leased Line ID: specifies the ID of the leased line that the VBR connects to.
- VLAN ID: specifies the VLAN ID of the VBR. The value ranges from 0 to 2999.

When creating router interfaces, you can use VLAN IDs to identify subsidiaries or departments that use the leased line, thus implementing Layer 2 network isolation between them.

- Local Gateway IP: specifies the local IP address of the router interface for the leased line.
- Peer Gateway IP: specifies the peer IP address of the router interface for the leased line.
- Subnet Mask: specifies the subnet mask of the leased line between the local IP address and peer IP address.

Only two IP addresses are required. Therefore, you can enter a longer subnet mask.

6. Click Create.

When the VBR state is Active, the VBR is created.

VBRs	Create VBR
* Region: cn-qingdao-env5b-d01 v * BID: * UID: 115	
VBK IU: Search Reset	
VBR ID VLAN ID VLAN Interface ID Status Routing Table ID Local Gateway IP Peer Gateway IP Address Address Subn Actions	
vbr- 33 ri- extive vtb- 192.168. 192.168. 255.25 rf 1997 rg Release	Actions 🔨
1-1/1 Modify Terminate	•

You can click Release, Modify, Terminate, or Show Details in the Actions column to manage a VBR.

1.5.7.6 Create router interfaces

After you create a VBR, you must create a pair of router interfaces to connect the VBR and VPC. The connection initiator must be the VBR.

Procedure

1. Log on to the Apsara Network Intelligence console.

- 2. From the Products menu, choose Express Connect > Network Environment Management.
- 3. Choose Network Environment Management > Router Interfaces.
- 4. Click Create Router Interface.

5. Configure router interface parameters and click Submit.

Set Create Router Interface to Double. Configure the local router interface based on the created VBR information, and configure the peer router interface based on the destination VPC information.

Create Router Interface		×
1 Local End Informati	on 2 Peer Information	3 Results
Select Router Type:	Single ODouble	
Name:	name of router interface	
Description:	description of router interface	
* Bid :	EnterBid	
* Uid:	EnterUid	
* Region:	cn-qingdao-env8d-d01	\sim
* Router Type:	● VRouter ○ VBR	
Zone:	SelectZone	
* Router ID:		
* Role:	InitiatingSide	
* Specifications:		\sim
Health Check Source IP:	EnterHealth Check Source IP	
Health Check Destination	EnterHealth Check Destination IP	
Skip Inventory Check:	🔿 Yes 💿 No	
	Next Cancel	

When the router interface state is Active, the interface is created.

Router Int	terfaces	S							Cr	eate Router Interface
* Region :	cn-qingdao	-env5b-d01	× BID:	26		* UID :	11			
l	Search	Reset								
Local Router II	D	Local Rou ter Type	Local Router Interface ID	Router Int erface Sta tus	Local Access Point ID	Role	Peer Router ID	Peer Router Type	Actions	
vrt-fs 2		VRouter	ri-f9n	Active	None	Accepting Side	vbr-fi	VBR	Deactivate	Actions 🔨
vrt-f9r 2		VRouter	ri-f9ri	Inactive	None	Accepting Side	vbr-f9	VBR	Modify Attribute Modify Specification	
< 1-2/2									Show Details	age V Goto

1.5.7.7 Create a routing table

A routing table is a list of route entries on a VRouter. This topic describes how to create routing tables in a region and query the routing table information of a region.

- 1. Perform the following steps to add routes on a 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) Choose Function Modules > Routing Tables.
 - d) Set search conditions such as Region, BID, UID, Router Type, Routing Table ID, and Router ID, and click Search to query routing tables.
 - e) Click Add Route Entry in the Actions column corresponding to a routing table.
 - f) Specify a route entry destined for the CIDR block of a destination VPC, and click Create.

The parameters are described as follows:

- Destination CIDR Block: the destination CIDR block.
- Next Hop Type: the next hop type.

• Next Hop Instance ID: the ID of the next hop instance for the specified next hop type.

Add Routing Entry	×
* BID:	268
* UID:	119
* Routing Table ID:	vtb-f9r:
	Modify the routing table ID to which the routing entry belongs.
* Destination CIDR Block:	Enter a Destination CIDR Block
	The network mask, such as 255.255.255.0/24.
* ECMP:	🔿 Yes 💿 No
* Next Hop Type:	Instance \lor
	The next hot type. Valid values: Instance, Tunnel, HaVip, RouterInterface.Set the value to RouterInterface for ECMP.
* Next Hop ID:	
	The next hop interface ID for the route entry.
	Create Cancel

Figure 1-11: Add a route destined for a destination VPC

g) Repeat the preceding steps to add a route destined for a target IDC.

Note:

You can navigate to the VBRs page and locate the VLAN Interface ID area to obtain next hop router interface information.

- 2. Add a route destined for the router interface of a VBR in the VPC.
- 3. On the gateway of the on-premises IDC, configure a route destined for the VPC.

1.5.8 Manage Business Foundation System flows in a VPC

You can view the execution state of tasks in a VPC and restart the tasks as needed.

Procedure

1. Log on to the Apsara Network Intelligence console.

- 2. From the Products menu, choose Network Controller > Business Foundation System Flow.
- 3. Query the flow state of the task you want to view.

Enter a leased line ID in Instance ID and set Step Status to All to check the flow status. A flow in red indicates that the corresponding step has failed. Click Resend to restart the task, and then requery the flow status.

Figure 1-12: Flow Management page



1.5.9 Configure reverse access to cloud services

Cloud services cannot be accessed directly through external networks. You must configure reverse access to allow external networks to access cloud services through ECS instances.

Prerequisites

Log on to the Apsara Stack console. Navigate to the Personal Information page and obtain AccessKey ID and AccessKey Secret.

- **1.** Log on to the Apsara Network Intelligence console.
- 2. From the Products menu, choose Cloud Service Management > Cloud Service Reverse Access.
- 3. On the page that appears, enter AccessKey ID and AccessKey Secret and click OK. The Cloud Service Reverse Access page appears.
- 4. Click Create Cloud Service Reverse Access.
- 5. On the Allocate App ID tab, set Region, Name, and Description.
- 6. Click Continue. The following information is automatically created and displayed on the Create Address Pool tab: the application IDs of cloud services that allow reverse access and the address pools that are used for reverse access to the cloud services.

- 7. Click Continue. On the Add Server Address tab, configure an ECS instance to be used for reverse access.
 - VPC ID: specifies the ID of a VPC, an ECS instance, or a single-tunnel cloud service instance.
 - Server IP: specifies the IP address of the ECS instance to be used for reverse access.
- 8. Click Continue. On the Create Mapping IP tab, configure VSwitch ID and Mapping IP of the ECS instance in the destination VPC.
- 9. Click Continue. On the Complete Authorization tab, configure VPC ID, ECS Instance IP, and Instance Port for reverse access.

The value of Instance Port must be an integer value. You can specify multiple instance ports separated by commas (,). Example: 10,20,30. You can configure up to 10 instance ports.

2 Elastic Compute Service (ECS)

2.1 ECS overview

Elastic Compute Service (ECS) is a user-friendly computation service featuring elastic processing capabilities that can be managed more efficiently than physical servers. You can create instances, resize disks, and release any number of ECS instances at any time based on your business needs.

An ECS instance is a virtual computing environment that includes basic components such as the CPU, memory, and storage. Users perform operations on ECS instances. Instances are the core concept of ECS, and are operated from the ECS console. Other resources such as block storage, images, and snapshots can be used only after they are integrated with ECS instances. For more information, see *Figure*

2-1: ECS instance.





2.2 Log on to the Apsara Stack Operations console

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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

Log On	
<u>8</u>	Enter a user name
ß	Enter the password
	Log On



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 log on to ASO.

2.3 ECS operations and maintenance

2.3.1 Overview

The ECS Operations and Maintenance Platform is a platform for support engineers to operate and monitor ECS instances, help users troubleshoot problems with ECS instances, and ensure that ECS instances are properly operated and utilized.

2.3.2 VM

2.3.2.1 Overview

On the ECS Operations and Maintenance Platform page, the existing ECS VM information and available O&M functions are displayed. You can search for, start, and migrate a VM as needed.

2.3.2.2 Search for VMs

You can view the list of existing VMs and their information in the Apsara Stack Operations console.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose **Products**.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

4. On the VMs tab that appears, set the filter conditions and click View.

Region is a required filter condition.

ECS Operations and Maintenance Platform									
VMs Disks Snapshots Images Security Groups									
Region			mas () to separate	Public IP A	Address mmas () to separate	AliUid			
Security Group Status Select V								Vie	w Clear
Start	Stop Reboot	Stop and	I Migrate Mo	re	Disks 🎝	Security Groups 🖒			
	Host		VM IP Address	Public IP Address	Region	CPU (C) Memory (M)	Disk Information	Internet Bandwidth	Status
			10 mail 12 m 20 10				Andreas Angles Angles Angles Angles Angles Angles Angles Angles Angles		Running

5. In the VM list, click a VM ID. You can view the VM information in the VM Details message that appears.

2.3.2.3 Start a VM

You can start a VM in the same way as you start a real server.

Prerequisites

The VM to be started must be in the Stopped state.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose > Products.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

- 4. On the VMs tab that appears, set the filter conditions and click View.
- 5. In the VM list, select the VM to be started. Click Start above the list.
- 6. In the dialog box that appears, set Start.

You can select Normal or Repair.

Note:

If you want to reset the network settings for the VM, set Start to Repair. Otherwise, set Start to Normal.

7. Set Operation Reason. Click OK.

2.3.2.4 Stop a VM

You can stop a VM in the same way as you stop a real server.

Prerequisites

The VM to be stopped must be in the Running state.



This operation will interrupt the programs running on the VM. Perform this operation during off-peak hours to minimize the impact on services.

Procedure

1. Log on to the Apsara Stack Operations console.

2. In the left-side navigation pane, choose



- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.
- 4. On the VMs tab that appears, set the filter conditions and click View.
- 5. In the VM list, select the VM to be stopped. Click Stop above the list.
- 6. In the dialog box that appears, set Shutdown Policy.

You can select Non-force Shutdown or Force Shutdown.



When Force Shutdown is selected, the VM is shut down regardless of whether its processes have been stopped. We recommend that you do not select Force Shutdown unless Non-force Shutdown does not work.

7. Set Operation Reason. Click OK.

2.3.2.5 Restart a VM

You can restart a VM in the same way as you restart a real server.

Prerequisites

The VM to be restarted must be in the Running state.

Note:

This operation will interrupt the programs running on the VM. Perform this operation during off-peak hours to minimize the impact on services.

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose **Products**.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.
- 4. On the VMs tab that appears, set the filter conditions and click View.
- 5. In the VM list, select the VM to be restarted. Click Reboot above the list.

6. In the dialog box that appears, set Start and Shutdown Policy.

For the Start parameter, you can select Normal or Repair.

For the Shutdown Policy parameter, you can select Non-force Shutdown or Force Shutdown.

7. Set Operation Reason. Click OK.

2.3.2.6 Cold migration

You can perform cold migration on a VM while it is offline to implement failover in the Apsara Stack Operations console.

Prerequisites

Cold migration must be performed offline. Make sure that the VM is in the Stopped state before you migrate it.

Context

If a VM or an NC fails, you must fail over the VM by shutting the VM down and migrating it to a new NC. Failover can only be performed within the same zone. Cross-zone failover is not allowed.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose > Products.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

- 4. On the VMs tab that appears, set the filter conditions and click View.
- 5. In the VM list, select the VM to be migrated. Click Stop and Migrate above the list.
- 6. In the dialog box that appears, configure the parameters.

Parameter	Description
Switchable NC	The destination NC to which the VM is to be migrated.
Switchover Policy	The switchover policy. Valid values: • Force Migrate • Active Migrate

Parameter	Description
Start	The startup mode. Valid values: • Normal • Repair
Recover	 The recovery mode. Valid values: Start After Migration Stop After Migration Status Unchanged After Migration Status Unchanged After Migration takes effect only on VMs that are in the Pending state.

7. Set Operation Reason. Click OK.

2.3.2.7 Reset a disk

You can reset disks to restore them to their initial state as needed.

Prerequisites

- When you reset a disk, applications that are installed on the disk are lost. Before you perform a reset operation, make sure that you have backed up your data.
- To reset a disk, make sure that the VM to which it belongs is in the Stopped state.

Context

Resetting a disk only restores the disk to its initial state and does not reformat the disk. The image that is used to create the disk will still exist after the disk is reset.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose **Products**.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

- 4. On the VMs tab that appears, set the filter conditions and click View.
- 5. In the VM list, select the VM that contains the disk to be reset. Choose More > Reset Disk above the list.

6. In the dialog box that appears, select the disk to be reset and set Operation Reason. Click OK.

2.3.3 Disks

2.3.3.1 Overview

In an ECS instance, cloud disks can be considered as physical disks. You can mount, detach, and create snapshots for disks.

2.3.3.2 Search for disks

You can view the list of existing disks and their information in the Apsara Stack Operations console.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. Click the Disks tab.
- 3. Specify the filter conditions and click View.

Region is a required filter condition.

ECS Operations and Maintenance Platform											
VMs Disks Snapshots Images Security Groups											
• Region	Region Disk ID VM ID AllUid Use commas (.) to separate multiple Enter only one ID.										
Disk Type Disk Status Select V Select V							Clear				
	Disk ID	AliUid	Disk ID	Release Auto Snapshot	Independent Disk	Disk Size	Disk Type	VM ID	Mount Point	Region	Disk Status
+						-				12****	In use

2.3.3.3 View snapshots

You can view the list of existing snapshots and their information in the Apsara Stack Operations console.

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose



- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.
- 4. Click the Disks tab.
> Products.

> Products.

and Security / 2 Elastic Compute Service (ECS)

Operations and Maintenance Guide - Cloud Essentials

- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the disk whose snapshots you want to view, and choose > View

Snapshot.

The information of all snapshots on the disk is displayed.

2.3.3.4 Mount a disk

After a disk is created, you must mount the disk to a VM.

Context

You can mount only disks that are separately created to VMs.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

- 4. Click the Disks tab.
- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the disk to be mounted and choose **Sector** > Mount.

7. In the dialog box that appears, set VM ID and Operation Reason. Click OK.

2.3.3.5 Detach a disk

You can only detach data disks in the Apsara Stack Operations console. You cannot detach system disks or local disks.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose



4. Click the Disks tab.

- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the disk to be detached and choose **Detach** > Detach.
- 7. In the dialog box that appears, set Operation Reason. Click OK.

2.3.3.6 Create a snapshot

You can manually create a snapshot for a disk in the Apsara Stack Operations console.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose > Products.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

- 4. Click the Disks tab.
- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the disk for which you want to create a snapshot, and choose **Take** > Take

Snapshot.

You can create snapshots only for system disks.

	Disk ID	AliUid	Disk ID	Release Auto Snapshot	Independent Disk	Disk Size	Disk Type	VM ID	Mount Point	Region	Disk Status
+	d-bs005ma64vv otmptiffa	100000021	20211-19970	No	Yes	50G	DATA	i-bs005ma64vv otmpwfe3p	/dev/xvdc	cn-qingdao-env 12-d01	In use
_	d-bs005t7bs10 8gwzwc56c	100000021	20211-19969	No	Yes	200G	DATA	i-bs005ma64vv otmpwfe3p	/dev/xvdb	cn-qingdao-env 12-d01	In use
	View Snapshot Detach Take Snapshot										
-	d-bs005ma64vv otmptiff9	100000021	20211-19968	Yes	No	100G	SYSTEM	i-bs005ma64vv otmpwfe3p	/dev/xvda	cn-qingdao-env 12-d01	In use
	View Snapshot	🗘 🛛 Take S	Snapshot								

7. In the dialog box that appears, set Snapshot Name, Snapshot Description, and Operation Reason. Click OK.

2.3.4 Snapshots

2.3.4.1 Overview

A snapshot stores the data stored on a disk for a certain point in time. Snapshots can be used to back up data or create a custom image.

When using disks, note the following points:

- When writing or saving data to a disk, we recommend that you use the data on one disk as the basic data for another disk.
- Although the disk provides secure data storage, you must still ensure that stored data is complete. However, data can be stored incorrectly due to an application error or malicious usage of vulnerabilities in the application. For these cases, a mechanism is required to ensure that data can be recovered to the desired state.

Alibaba Cloud allows you to create snapshots to retain copies of data on a disk for specific points in time.

2.3.4.2 Search for snapshots

You can view the list of existing snapshots and their information in the Apsara Stack Operations console.

Procedure

- 1. Log on to the Apsara Stack Operations console.
- 2. Click the Snapshots tab.
- 3. Specify the filter conditions and click View.

Region and AliUid are required filter conditions.

ECS C	Operations and Maintena	nce Platform						
VMs	VMs Disks Snapshots Images Security Groups							
• Region		• AliUid		Disk ID Enter only one ID.		Snapshot ID		View Clear
	Snapshot ID	Snapshot Background ID	Region	Created At	Snapshot Type	Snapshot Size	Disk ID	Progress
	No data is available							

2.3.4.3 Delete a snapshot

You can delete snapshots that are no longer needed in the Apsara Stack Operations console.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose



- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.
- 4. Click the Snapshots tab.
- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the snapshot to be deleted and choose **Delete**.
- 7. In the dialog box that appears, set Operation Reason. Click OK.

2.3.4.4 Create an image

You can create a custom image from a snapshot. The image includes the operating system and environment variables of the snapshot.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose > Products.
- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.

- 4. Click the Snapshots tab.
- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the snapshot from which you want to create an image, and choose

Create Image.

7. In the dialog box that appears, set Image Name, Image Version, Image Description, and Operation Reason. Specify whether the system disk for which the snapshot was taken is based on a public image or a custom image. Click OK.

2.3.5 Images

2.3.5.1 Overview

An ECS image is a template that contains software configurations such as the ECS instance operating system and the programs and servers for applications. You must specify an ECS image to create an instance. The operating system and software provided by the image will be installed on the instance that you create.

2.3.5.2 Search for images

You can view the list of existing images and their information in the Apsara Stack Operations console.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. Click the Images tab.
- 3. Specify the filter conditions and click View.

Region is a required filter condition.

ECS	ECS Operations and Maintenance Platform							
VMs	Disks Sn	apshots Images	Security Groups					
• Regio	n ~	Image Type Select	AliUid	Image ID	mmas (,) to separate multiple	View Clear		
	Image ID	Image Name	Snapshot ID	Created At	Image Type	Operating System Type		
+	1000.0000.0000	and a second second		Sep 20, 2019, 18:26:22		Windows Server 2008		

2.3.6 Security groups

2.3.6.1 Overview

A security group is a virtual firewall that provides Stateful Packet Inspection (SPI). Security groups provide virtual firewall-like functionality and are used for network access control for one or more ECS instances. They are important means of network security isolation and are used to divide security domains on the cloud.

Security group rules can permit the inbound and outbound traffic of the ECS instances associated with the security group. You can authorize or cancel security group rules at any time. Changes to security group rules are automatically applied to ECS instances that are members of the security group.

When you configure security group rules, ensure that the rules are concise and easy to manage. If you associate an instance with multiple security groups, hundreds

of rules may apply to the instance, which may cause connection errors when you access the instance.

2.3.6.2 Search for security groups

You can view the list of current security groups and their information in the Apsara Stack Operations console.

Context

You can modify security group rules to allow or deny inbound and outbound traffic between the security group and the public or internal network. You can add or delete security group rules in each security group at any time. Changes to security group rules automatically apply to the ECS instances in the security group.



- If two security group rules differ only in action, the deny rule takes precedence over the allow rule.
- No rule in a security group can allow outbound access from an ECS instance while denying inbound access to the ECS instance.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. Click the Security Groups tab.
- 3. Specify the filter conditions and click View.

Region is a required filter condition.

ECS	ECS Operations and Maintenance Platform						
VMs	Disks Snapsho	ots Images Se	ecurity Groups				
• Regio	n Secu	urity Group ID se commas (.) to separate multiple	MID Enter only one ID.		View Clear		
	Security Group ID	Security Group Name	Created At	VPC ID	Region		
+		With the state of the second					

2.3.6.3 Add security group rules

You can add rules to security groups based on your needs.

Procedure

1. Log on to the Apsara Stack Operations console.

2. In the left-side navigation pane, choose



- 3. In the left-side navigation pane, choose ECS > ECS Operations and Maintenance Platform.
- 4. Click the Security Groups tab.
- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the target security group and choose Add Rule.
- 7. In the dialog box that appears, configure the parameters.

For more information about the parameter configurations, see Table 2-1: Security

group rule parameters.

Table 2-1: Security group rule parameters

Parameter	Description
Protocol	 TCP UDP ICMP GRE ALL: All protocols are supported.
Rule Priority (1-100)	The smaller the value, the higher the priority.
Network Type	 Public: the public network Internal: the internal network
Authorization Policy	 Accept: accepts the packet. Drop: drops the packet. Reject: rejects the packet.
Port Number Range	Valid values: 1 to 65535. Example: 1/200, 80/80, or -1/-1.
Access Direction	 Ingress: allows inbound traffic Egress: allows outbound traffic
IP Address Range	Enter an IP address or a CIDR block, such as 10 .0.0.0, 0.0.0.0/0, or 192.168.0.0/24. Only IPv4 addresses and IPv4 CIDR blocks are supported.
Security Group ID	Enter the ID of the associated security group.

Parameter	Description
Operation Reason	Optional. Enter a reason for the operation.

8. Click OK.

2.4 VM hot migration

2.4.1 Overview

Hot migration is the process of migrating a running VM from one host to another. During migration, the VM runs normally and its services are not aware that any migration task is occurring. However, these services can detect a very short interruption between 100 and 1,000 ms.

Scenarios

During system operations and maintenance, hot migration is typically used for the following scenarios:

- Active O&M: The host is faulty and must be repaired, but the fault does not affect the operation of the system. You can use hot migration to migrate the VM to another host and repair the faulty host in offline mode.
- Server load balancing: When a host is experiencing a high load, you can migrate some of its VMs to other idle hosts to reduce resource consumption on the source host.
- Other scenarios where a VM must be migrated without affecting its business operations.

2.4.2 Limits on hot migration

Before performing hot migration, you must understand the limits.

The hot migration feature of Apsara Stack is subject to the following limits:

- Only the go2hyapi command can be used to implement hot migration in the KVM virtualization environment. ECS Operations and Maintenance Platform does not support hot migration.
- Only standard ECS instances support hot migration. ECS provides a list of migratable images. Alibaba Cloud does not take any responsibility for errors that occur when migrating a VM that is not included in the list of migratable images.

- If a VM is used as an RS to provide SLB or as a client to access SLB, the previous session will be closed after hot migration. New sessions created after migration are not affected.
- Migration can only be performed between hosts of the same type. Furthermore, each host must be running the same versions of software.
- Hot migration is not supported in DPDK avs scenarios.
- VMs using local storage solutions do not support hot migration. This is because after a VM is migrated to another host, it can no longer access the previous local storage space.
- VMs that use GPU, FPGA, or other (passthrough or SR-IOV) devices do not support hot migration.



VMs created in Apsara Stack versions earlier than V3.3 do not support hot migration. Hot migration becomes available after you restart the VMs.

2.4.3 Complete hot migration on AG

In Apsara Stack Operations, you can start and cancel hot migration operations as needed through the command line interface.

Trigger hot migration

After hot migration is triggered, you can run the go2which command or use ECS Operations and Maintenance Platform to check that the VM enters the migrating state. When hot migration is completed, the VM restores 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: I	Parameter	description
--------------	-----------	-------------

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A

Parameter	Function	Impact	Value
nc_id	Designates the destination NC to migrate the VM to.	If the NC does not support the specifications of the VM, the migration will fail.	N/A
rate	The amount of host bandwidth to be allocated for migration tasks.	The migration will use the bandwidth resources of the hosts.	 10 GB network: 80 MB 1 GB network: 40 MB
downtime	The maximum allowable downtime caused by migration. The default value is 300 ms.	The service downtime caused by migration is affected.	200 ms to 2,000 ms
no_check_image	Forcibly migrates the images that are not supported.	Performing this operation may violate the SLA.	false
no_check_load	Forcibly migrates images even when the load threshold requirements are not met.	Downtime cannot be controlled when this parameter is set to false.	false

Cancel hot migration

Run the following command to cancel a hot migration task:

```
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

Parameter	Function	Impact	Value
region_id	The ID of the region where the target VM is located.	N/A	N/A

2.4.4 Modify the position of the NC where the VM is located

When an exception occurs during hot migration and the migration cannot be rolled back through ECS Operations and Maintenance Platform, you can modify the VM state to trigger rollback.

Trigger rollback

If an exception occurs during hot migration, run the following command to trigger rollback:

```
go2hyapi call_api manually_change_migration_status == Functions
usage: == |- call_api manually_change_migration_status <vm_name> <
region_id> <where>
```

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
region_id	The ID of the region where the target VM is located.	N/A	N/A
where	The ID of the NC where the VM is located.	N/A	N/A

Table 2-4: Parameter description

2.4.5 FAQ

This topic lists common problems that you may encounter during hot migration and how to resolve them.

- Which parameters are required to call the Server Controller API to perform a hot migration?
 - Vm_name: VM name
 - nc_id

- · What preparations should I make before performing a hot migration operation?
 - Confirm that the VM is in the running state.
 - Confirm the destination of the VM migration.
- · Can hot migration be canceled? How can I cancel hot migration?

Yes. If the API request is successful and the migration has not completed, run the go2hyapi cancel_live_migrate_vm vm_name=[vm_name] region_id=[region_id] command to cancel the hot migration. If the VM has completed its migration to the destination NC, it is too late to cancel the hot migration.

You can get the value of region_id by running the go2which [vm_name] command to view region_info.

• The VM is still in the migrating state after the hot migration has completed, and the cancel_live_migrate_vm command is not working. What should I do?

You can run the virsh query-migrate [domid] command on the source NC of the VM to check whether the VM is still being migrated. If the VM is still being migrated, a piece of JSON information will be returned. If the VM has finished migration, run the following command on the AG to modify the state of the VM:

```
go2hyapi manually_change_migration_status vm_name=[vm_name] where=[
nc_id for the VM] region_id=[region_id]
```

domid is the name of the VM instance. You can run the virsh list|grep vm_name **command to view it.**

· How can I confirm whether the VM is migrated successfully?

On the destination NC of the VM, run the sudo virsh list|grep [vm_name] command. If the VM instance exists and is not in the running state, the migration is successful.

- · When an exception occurs during hot migration, which logs should I refer to?
 - View the Libvirt bottom layer migration log on the NC.

Run the /var/log/libvirt/libvirt.log command to view information about the migration process, such as vport offline, detach, delete, and relay route.

- Run the following command to view the API management log of Server Controller on the AG:

/var/log/houyi/pync/houyipync.log

- View the Qemu log.
- Run the following command to view the regionmaster log on the VM:

regionmaster/logs/regionmaster/error.log

• A VM fails to start after hot migration. Is the VM still in the pending state?

If error vport update nc conf by vpc master fails dest_nc_id:xxx is returned, it indicates that a VPC fault has occurred and the underlying task is interrupted.

• During hot migration, the API returns the following error message: distributed lock fail. What are the possible causes of this issue?

The API has been called too many times within a short period of time. Wait several minutes and then try again.

• What are some common scenarios where migration fails? How can I resolve these issues?

Scenario	Cause	Solution
The load is too high and the VM migration does not pass the pressure inspection.	Long service interruption	You can run no_check_l oad=true to skip this inspection.
The VM fails to pass image inspection.	It is not an Alibaba Cloud- specified image.	You can run no_check_i mage=true to skip this inspection. Be aware of the risks involved.

Table 2-5: Hot migration issues

2.5 Hot migration of disks

2.5.1 Overview

Hot migration seeks to facilitate operations and maintenance of online clusters and improve service operation. Hot migration provides online migration capabilities for virtual disks. This function can also quickly copy data to new locations, enhancing the flexibility of services.

2.5.2 Limits

Before performing hot migration on a disk, you need to understand the limits.

Limits

- Only disks of the river type support hot migration.
- The source and destination clusters for hot migration must belong to the same OSS domain.
- Disk sharing is not supported.
- Hot migration is not supported on disks whose capacity is greater than 2 TB.
- Format and capacity changes are not supported.
- Hot migration is only supported within the same zone.
- Due to how hot migration is implemented internally, the names of the source and destination clusters must be less than 15 bytes in length.

Note:

- The data of the original source disk will remain on the disk after hot migration has completed. You can use the pu tool to delete the remaining data. Job recycling is unavailable.
- During migration, an I/O latency of less than 1 second is considered normal.
- Migration cannot be rolled back.
- Migration will consume network bandwidth, so you must take measures to limit concurrent traffic during migration.

Migration operation

For more information about the APIs related to disk hot migration, see "Disk hot migration" in *ECS Developer Guide*.

2.5.3 O&M after hot migration

The original source disk data remains on the source disk after hot migration and data backup operations are completed. To release disk space, delete the data from the source disk. After the data is deleted from the source disk, the space will be released at a later time.

Procedure

- 1. On the compute cluster AG, run the go2houyiregiondbrnd -e 'select task_id from device_migrate_log where status="complete"' command to obtain task: allTaskIds.
- 2. On the compute cluster AG, run the go2riverdbrnd -e 'select task_id, src_pangu_path,dst_pangu_path from migration_log where task_id in (\$ allTaskIds) and status=2 and src_recycled=0 and DATE(gmt_finish) < DATE_ADD(CURDATE(), INTERVAL -1 DAY)' command.
- 3. Perform the following operations for each set of <task_id,src_pangu_path,dst_pangu_path>:
 - a) Run the /apsara/deploy/bsutil rlm --dir=\$dst_pangu_path|grep 'not-loaded'|wc -l command on the host that runs the bstools role in the storage cluster. If the command output is not 0, proceed to the next step.
 - b) Run the /apsara/deploy/bsutil delete-image --dir=\$src_pangu_path command on the host that runs the bstools role in the storage cluster.
 - c) Run the /apsara/river/river_admin migrate recycle \$task_id command on the host that runs the river role in the storage cluster.

2.6 Upgrade solution

2.6.1 Overview

For both hot and cold migration of GPU and FPGA clusters, you must understand the limitations that apply to cluster upgrades.

2.6.2 Limits on GPU clusters

Before upgrading a GPU cluster, you must understand the limits.

The upgrade of GPU clusters in Apsara Stack are subject to the following limits:

• GPU clusters are only supported in Apsara Stack 3.3 or later versions.

- To upgrade a GPU cluster, you must restart the NC server.
- VMs that use GPU, FPGA, or other passthrough or SR-IOV devices do not support hot migration.
- The GN5I, GN5E, and GN4 type GPU clusters do not have the specifications of local disk instances and only support offline cold migration.
- When you perform a forced cold migration on GN5 and GA1 type GPU clusters that have specifications of local disk instances, the local disk will be reformatte d, resulting in data loss. These disks must be backed up before they can be migrated.

2.6.3 Limits on FPGA clusters

Before upgrading an FPGA cluster, you must understand the limits.

The upgrade of FPGA clusters in Apsara Stack are subject to the following limits:

- FPGA clusters are only supported in Apsara Stack 3.5 or later versions.
- VMs in an FPGA cluster must be shut down before the cluster can be upgraded.
- The FPGA service relies on Redis to a great extent. If the Redis service is interrupted during the hot upgrade of Apsara Stack, the FPGA service will be interrupted. The FPGA service will recover after the Redis service is restored . However, if a Redis instance fails to be created, you must restart the FPGA service after the Redis service is restored.

2.7 Disk maintenance of an instance

2.7.1 Overview

This topic describes the limits on, procedure of, and related information about disk maintenance for an instance.

Application scope

- Applicable only to D1 disks.
- · Applicable only to disks whose mount point is /apsarapangu/disk*.
- The mount point of a physical disk on an NC does not change during the course of maintenance.
- Applicable to Apsara Stack 3.1 to 3.6.
- Currently applicable only to the N41S1-6T servers.

Background information

A disk is damaged, and you want to repair the physical disk and recreate the data disk without migrating data.

Impact

To restore the physical disk without migrating data, you must shut down the VM associated with the damaged disk.

Potential risks

- The data on the replaced physical disk is all lost.
- A problem occurs during the next startup if the disk UUID is written to the fstab file in the VM. This problem occurs in any scenario where the disk-mounting relationship changes.
- Strictly follow the procedure.

Environment inspection

Use a tool to inspect the entire cluster environment.

2.7.2 Maintenance procedure

This topic describes the maintenance procedure to repair a disk attached to an instance.

Procedure

1. Log on to the AG with the admin account to search for NC-related information.

Run the following command to obtain the NC ID based on the NC IP address:

go2ncinfo {nc_ip}

{nc_ip} is the IP address of the host where the disk to be repaired is located.

Example:

- Host IP address: 10.10.3.5
- · Host name: c43b07003.cloud.b07.amtest1221
- File name and mount point of the host with a damaged disk: /dev/sdb1 / apsarapangu/disk1
- · AG: vm010010016025
- Run the go2ncinfo 10.10.3.5 command to obtain the NC ID.
- NC ID: 21765-26

[:/hor	me/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
\$ go2ncinfo 10.10.3.5	
nc_id: 2170	55-26
ip: 10.1	10.3.5
hostname: c43	p07003.cloud.b07.amtest1221
biz_status: free	
priority: 8	
health: 5	

- 2. Use the AG through Server Controller to check which VMs are affected by this physical disk.
 - We recommend that you run the following command on the API to identify 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 obtain the region. {nc_id} is the NC ID of the

host obtained in the previous step, and {mount_point} is the mount point of the disk on the host.

You can also run the following command in /etc/houyi/script/

local_disk_ops.py to identify the affected VMs. 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_de
vice_id={mount_point}
```

{nc_id} is the NC ID of the host obtained in the previous step, and {mount_poin
t} 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
```



If an error is reported when the API is used, you must run the following command instead. The local_disk_ops.py script is in the /home/admin directory in this environment.

```
/home/tops/bin/python local_disk_ops.py --action=query_vms_
by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26 storage_de
vice_id=/apsarapangu/disk1
```



You can see that only the i-5wf05ykw7mic5aq65dv2 instance runs on this disk and is in the running state.

3. Shut down the VMs on the AG by using Server Controller.

a) If the VMs are in the running state, you need to shut them down first.

Run the following command:

go2hyapi stop_vm vm_name={vm_name}

{vm_name} is the ID of the running VM obtained in the preceding step.

Example:

go2hyapi stop_vm vm_name=i-5wf05ykw7mic5aq65dv2



Wait until the VM status changes to Stopped.

[admin@	:/home/admin] [:cn-neimeng-env10-d01:io11:vpc:21765]
<pre>\$ /home/tops/bin/ provid=21765-26 ct</pre>	python local_disk_ops.pyac	tion=query_vms_by_physical_disklogfile=/tmp/
[{'vm_name': 'i-5	wf05 2', 'status	': 'stopped'}]

- b) If the VM is in the pending or stopped state, you do not need to shut it down.
- c) If the VM is in another state, you must wait until its status changes to running, pending, or stopped. Alternatively, you can carry out an inspection.

4. Use Server Controller to check the local data disk associated with the physical disk.

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 obtained NC ID of the host, 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.

Example:

```
/home/tops/bin/python local_disk_ops.py --action=query_loca
l_disks_by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26
storage_device_id=/apsarapangu/disk1
```



Only the local data disk with the ID 1000-3388 is associated.

- 5. Replace the damaged physical disk on the NC.
 - a) Check the device file name of the damaged 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 serial numbers (SN) of the NC and the hard disk.

A. In the Apsara Infrastructure Management Framework console, check the SN of the NC in the corresponding cluster operation and maintenance 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 obtained earlier.

Example:

smartctl -a /dev/sdb1 | grep 'Serial Number'



The SN of /dev/sdb1: K1K3EPKD

c) Remove the original disk.

The on-site engineer will 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 configurations. Server model of the existing disk: N41S1-6T and V53. The N41S1-6T mode is a hard disk drive (HDD) and supports hot swapping. The V53 model is a solid state drive (SSD), and requires the machine to be shut down before it can be swapped.

The following operations are only applicable to the N41S1-6T model.

Example:

			<u>,</u>	<u></u>
C4-3.NT12	B07	06	CVXKB7CD00[,	N41S1-6T.22

The N41S1-6T model supports hot swapping and 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 need to remove the /dev/sdb1 hard disk from slot 1. The SN of the hard disk should be consistent with the K1K3EPKD SN obtained earlier.



- d) Insert a new disk.
- e) Partition and mount the disk, and modify the label and the 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 ID of the hard disk.

Example:

Disk /dev, Units = se Sector siz I/O size (Disk labe Disk ident	/sdb: 60 ectors o ze (logi (minimum l type: tifier:	001.2 GB, (of 1 * 512 ical/physic n/optimal): dos 0x00000000	5001175126010 = 512 bytes cal): 512 byt : 4096 bytes	5 bytes, 11 tes / 4096 / 4096 byt	7210451 bytes es	l68 sectors
Device /dev/sdb1 Partition	Boot 1 does	Start 1 not start	End 4294967295 on physical	Blocks 2147483647 sector bou	Id + ee ndary.	System GPT

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 parted command is used to partition the 5.5 TB hard disk.

mklabel gpt

Use the GPT to form a 5.5 TB partition.

(cparted) 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? (ves/No2 ves

Run the mkpart primary 1049k -1 command to configure a 5.5 TB primary partition that starts at 1,049 KB and ends at the capacity limit of the hard disk.

print is used to display the capacity of the configured partition. quit is used to exit the parted program.



C. Format the partition.

```
mkfs -t {filesystem_type} {device_name}
```

{filesystem_type} is the type of the file system to be formatted. {device_nam
e} is the name of the partition to be formatted.

Example:

mkfs -t ext4 /dev/sdb1

<pre>[root@(:/root] #mkfs.ext4 /dev/sdb1 mke2fs 1.42.9 (28-Dec-2013) Filesystem label= OS type: Linux Block size=4096 (log=2) Eracment_size_4096 (log=2)</pre>					
<pre>#mkfs.ext4 /dev/sdb1 mke2fs 1.42.9 (28-Dec-2013) Filesystem label= OS type: Linux Block size=4096 (log=2) Fracement size=4096 (log=2)</pre>					
<pre>mke2fs 1.42.9 (28-Dec-2013) Filesystem label= OS type: Linux Block size=4096 (log=2) Fracmont size=4096 (log=2)</pre>					
Filesystem label= OS type: Linux Block size=4096 (log=2)					
OS type: Linux Block size=4096 (log=2)					
Block size=4096 (log=2)					
Fragment size=4096 (log=2)					
Fragment Size=4090 (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					
uniting journal (J2700 Diocks), done					
Writing superblocks and filesystem accounting information: done					
Writing superblocks and filesystem accounting information: done					
Writing superblocks and filesystem accounting information: done [root@:/root]					
Writing superblocks and filesystem accounting information: done [root@:/root] #lsblk -T					
Writing superblocks and filesystem accounting information: done [root@:/root] #lsblk -T NAME FSTYPE LABEL UUID MOUNTPOINT					
Writing superblocks and filesystem accounting information: done [root@:/root] #lsblk -T NAME FSTYPE LABEL UUID MOUNTPOINT sda					
<pre>writing superblocks and filesystem accounting information: done [root@:/root] #lsblk -T NAME FSTYPE LABEL UUID MOUNTPOINT sda -sdal ext4 /boot 1fd12aa3-8f54-4bb0-ald3-a29595f391b8 /boot</pre>					
<pre>writing superblocks and filesystem accounting information: done [root@:/root] #lsblk -T NAME FSTYPE LABEL UUID MOUNTPOINT sdasda1 ext4 /boot 1fd12aa3-8f54-4bb0-a1d3-a29595f391b8 /bootsda2 ext4 / 3ac491f4-c2a4-4372-a4c3-3b3605b8a6da /</pre>					
<pre>writing superblocks and filesystem accounting information: done [root@:/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</pre>					
<pre>writing superblocks and filesystem accounting information: done [root@:/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 </pre>					
<pre>writing superblocks and filesystem accounting information: done [root@</pre>					
writing superblocks and filesystem accounting information: done [root@					
<pre>writing superblocks and filesystem accounting information: done [root@</pre>					
writing superblocks and filesystem accounting information: done [root@					
<pre>writing superblocks and filesystem accounting information: done [root@</pre>					
<pre>writing superblocks and filesystem accounting information: done [root@</pre>					
<pre>writing superblocks and filesystem accounting information: done [root@</pre>					
<pre>writing superblocks and filesystem accounting information: done [root@</pre>					

D. Mount the hard disk to the original directory.

The server supports hot swapping. If you remove and insert the same hard disk, it will be automatically mounted to the original directory. If a new

disk is inserted, it must be mounted manually. In this example, you must manually mount the disk.

```
mount {device_name} {mount_point}
```

{device_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.

Device files in the /etc/fstab directory are identified by their labels, so you must change the label of the new disk.

e2label {device_name} {label_name}

{device _name} is the device file name, and {label_name} is the label name.

Example:

The label of the removed disk is disk1, so you must change the label of the new disk to disk1.



e2label /dev/sdb1 disk1



F. Mount the disk based on the definitions in the fstab file.

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.



6. Use Server Controller to reset the data disk obtained earlier.

```
$/home/tops/bin/python local_disk_ops.py --action=reset_loca
l_disk_after_change_physical_disk --logfile=/tmp/tmp.log disk_id={
disk_id}
```

Note:

Exercise caution when performing the operation. The {disk_id} parameter must be the data disk obtained earlier based on the damaged disk.

Example:

```
/home/tops/bin/python local_disk_ops.py --action=reset_loca
l_disk_after_change_physical_disk --logfile=/tmp/tmp.log disk_id=
1000-3388
```



OK indicates that the disk is reset successfully.

7. Start the VM by using Server Controller.

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



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.

2.7.3 Additional instructions

This topic describes the scripts used for specific solutions during local disk maintenance.

Instructions for local_disk_ops

• Run the following command to view the script:

/home/tops/bin/python local_disk_ops.py -h

• Log description:

When a script is executed, a detailed log is recorded in a 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. The default log file has the same base name as the script and has the extension of .log.

For example, if you run the /home/tops/bin/python local_disk_ops.py --action= xxx arg1=value1 command, script execution is recorded in the local_disk_ops.log file.

• Error description:

If an error occurs when you execute a script, an error log is output to the current shell. Perform inspections based on the specific error information. Format of error message:

Error time Error (erroneous script 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 the value of the storage_device_id parameter is not specified.

Example 2:

\$ /home/tops/bin/python local_disk_ops.py --action=query_vms_by_physica
l_disk nc_id=1-1 storage_device_id=/apsarapangu/disk20

2018-03-13 21:23:42,764 ERROR (local_disk_ops.py: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 value is incorrect.

• For more information about this error, see *Maintenance procedure*.

2.8 Handle routine alarms

2.8.1 Overview

This topic describes the definition of each key metric and how to handle alerts.

The metrics monitored in ECS can be categorized into three types:

- Basic metrics: These metrics are used to monitor the CPU, memory, and correlated service processes of hosts.
- Connectivity metrics: These metrics are used to monitor the connectivity between different components and the connectivity between different networks.
- Service metrics: These metrics are used for service monitoring, such as the state of various types of API requests.

Metric type	Function	Solution
Basic metric/ service availabiliMonitors the basic performanc 		When CPU utilization is too high: identify which process consumes a large amount of CPU resources. If it is a key process, evaluate whether it can be restarted.
		When the memory usage is too high (for key services): dump the memory data, request the back-end R&D team to analyze the data, and restart the application.
Connectivi ty metric	Checks the connectivity between each module and its related modules.	 First, check the health status of the corresponding modules. For example, check whether the host works normally and whether services, ports, and domain names are normal. If two modules that are connected to each other are healthy, check the network connectivity between them.
Service metric	Monitors aspects of key request calls such as the latency, total number, failures of API requests, and database SQL exceptions.	 In case of an API request failure, you must view the corresponding logs to identify the cause of the failure. In case of a database SQL exception, check whether the exception was caused by a database exception (system breakdown or high connection count) or a problem with the application. If it is an applicatio n problem, forward the error information to the back-end R&D team for troubleshooting.

Table 2-6: Description of metric types

2.8.2 API proxy

This topic describes the metrics of API proxy.

Metric	Alert item	Description
check_apip roxy_dns	Database HA switchover occurs or not	Checks whether Server Controller database switchover occurs. If so, nginx will be reloaded automatically.
check_apip roxy_conn_new	check_apip roxy_conn_new	Checks the connectivity to the Server Controller database.
		Checks the connectivity to the API Server:
		 Checks whether the API Server is down. Checks the network connectivity.
check_apip roxy_proc_new	check_apip roxy_proc_new	Checks the memory usage and CPU utilization for nginx and memcache processes.

2.8.3 API Server

The topic describes the metrics of the API Server.

Table 2-8: Metric description

Metric	Alert Item	Solution
check_API Server_proc_new	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage

Metric	Alert Item	Solution
check_API Server_con n_new	Checks the connectivity between the API Server and Server Controller database.	Checks whether the corresponding component is down. If the corresponding component is down, fix the issue by taking necessary O&M
	Checks the connectivity between the API Server and TAIR.	measures. If the database is down, contact DBA to fix the issue. Checks whether the VIP is connected
	Checks the connectivity between the API Server and RegionMaster.	to the corresponding component. If not, contact the network engineer to fix it.
	Checks the connectivity between the API Server and the RMS.	пх п.
check_API Server_perf	Monitors metrics for API requests, such as the latency, total number of API requests , and number of failed API requests.	It is primarily used to identify faults.
check_API Server_errorlog	Checks database exceptions and instance creation failures.	 If an exception occurs to the database, contact DBA to check whether the database is normal. If the creation of an instance fails, locate the cause of the failure.

2.8.4 RegionMaster

This topic describes the metrics of RegionMaster.

Table 2-9: Metric	description
-------------------	-------------

Metric	Alert item	Description
check_regionmaster_p roc	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage.
check_regionmaster_w ork	rms_connectivity	Checks the connectivity to RMS.

Metric	Alert item	Description
	regiondb_connectivity	Checks the connectivity to the houyiregiondb database.
	houyi_connectivity	Checks the connectivity to the Server Controller database.
	tair_connectivity	Checks the connectivity to TAIR.
check_zookeeper_work	status	Checks the operating state of the Zookeeper process on the Server Controller.
check_regionmaster_e	errorlog_for_db	Checks whether the SQL
rrorlog	check_regionmaster_errorlog	statements are properly executed.
check_workflow_master	Checks the operating state of the master in the workflow process.	-
check_workflow_worke r	Checks the operating state of the worker in the workflow process.	-

2.8.5 RMS

This topic describes the metrics of RMS.

Table 2-10: Metric description

Metric	Alert item	Description
check_rms_proc	Checks the process status, CPU utilizatio n, and memory usage of RMS.	-
check_rabbitmq_proc	Checks the process status, CPU utilizatio n, and memory usage of the rabbitmq cluster.	-

Metric	Alert item	Description
check_rabb itmq_status	Checks the number of queues, exchanges , and bindings in the rabbitmq cluster.	Follow the maintenance guide for the rabbitmq cluster.
check_rabb itmq_queues	Checks whether messages are accumulated.	ether If messages are accumulated, it are will also check for the cause. ed.
	Check whether there are consumers.	If there are no consumers, check whether Regionmaster and APIserver are operating normally . If they are operating normally, check whether there is a problem with the rabbitmq cluster.

2.8.6 PYNC

This topic describes the metrics that are monitored for PYNC.

Table 2-11: Metric description

Metric	Alert item	Description
check_vm_s tart_failed	Checks the causes of a VM startup fault.	You do not need to handle it immediately. It is typically caused by custom images.
check_pync	Checks the CPU utilization and memory usage of PYNC.	-
	PYNC has too many open file handles.	-
	PYNC process count.	PYNC must have four processes.

Metric	Alert item	Description
	It has been long since pyncVmMoni tor.LOG was last updated at \${ pync_monitor_log_last_updated}.	 Checks for reasons why a log has not updated for a long period of time, such as: Whether a PYNC process has encountered a problem. Whether the NC is running a key process called Uninterruptible Sleep.

2.8.7 Zookeeper

This topic describes the metrics of Zookeeper.

Table 2-12: Metric description

Metric	Alert item	Description
check_zookeeper_proc	proc	The process does not exist.
		The memory usage or CPU utilization is too high.

2.8.8 AG

This topic describes the metrics of AGs.

Table 2-13: Metric description

Metric	Alert item	Description
disk_usage	apsara_90	/apsara disk usage.
	homeadmin_90	Usage of /home/admin.
check_system_ag	mem_85	Memory usage.
	cpu_98	CPU utilization.
	df_98	Disk usage of the root directory.

Metric	Alert item	Description
check_ag_d isk_usage	check_ag_disk_usage	Disk usage.
check_nc_d own_new	check_recover_failed	 Checks the causes of a VM migration fault. Possible causes include: No resources are available in the cluster. A VM does not belong to any cluster.
	check_repeat_recovered	Continuous VM migration.
	check_continuous_nc_down	Checks continuous NC downtime.
	check_nc_down_with_vm	 The state of the NC in the database is nc_down, but there are still VMs operating normally on the NC. Checks the NC for hardware faults: If a hardware fault occurs , you must perform operations and maintenance to resolve the fault. If no hardware fault is detected, restore the NC and change its state to locked.
check_ag_f htd_new	Checks whether the FHT downtime migration tool, mostly used by local disks, is operating normally.	If the tool does not exist, download the FHT downtime migration tool.

2.8.9 Server groups

This topic describes the metrics that are monitored for server groups.

Table 2-14: Metric description

Metric	Alert item	Description
check_pync	pync_mem	Monitors the memory usage of PYNC.
	pync_cpu	Monitors the CPU utilization of PYNC.
Metric	Alert item	Description
--------	----------------------------------	--
	pync_nofile	Monitors the number of PYNC handles.
	pync_nproc	Monitors the number of PYNC processes.
	pync_monitor_log_not _updated	Monitors the status of PYNC scheduled tasks.

2.9 Inspection

2.9.1 Overview

ECS inspection includes cluster basic health inspection and cluster resources inspection.

2.9.2 Cluster basic health inspection

2.9.2.1 Overview

Cluster basic health inspection includes monitoring inspection, inspection of basic software package versions, and basic public resources inspection.

2.9.2.2 Monitoring inspection

This topic describes basic monitoring inspections and connectivity monitoring inspections.

2.9.2.3 Inspection of basic software package versions

This topic describes the version inspections of Server Controller components, Apsara system, virtualization packages, and basic service packages.

2.9.2.4 Basic public resources inspection

This topic describes ISO inspections and basic image inspections.

ISO inspection

ECS Operations and Maintenance System provides two basic ISO files for each region:

- · linux-virt-release-xxxx.iso
- windows-virt-release-xxxx.iso

You can run the following command to search the database for relevant informatio

n:

```
$ houyiregiondb
mysql>select name,os_type,version,path,oss_info from iso_resource
where os_type! =''\G
```

Parameters in the command are as follows:

- name: the name of the ISO file, such as xxxx.iso.
- os_type: the operating system (OS) type of an image.
- path: the path on the Apsara Distributed File System cloud disk where the ISO file is stored. You can run the /apsara/deploy/pu meta \$path command to check whether the ISO exists in the files of Apsara Distributed File System.
- oss_info: the path on the local OSS disk where the ISO file is stored. To search for this path, you must provide relevant information to OSS support engineers for inspection.

Basic image inspection

• Run the following command to check the state of a basic image in the database:

```
houyiregiondb
mysql>select image_no,status,visibility,platform,
region_no from image;
```

• Check whether the basic image is usable. You can call the create_instance API to use relevant images to create a VM and manually check whether the VM can operate normally.

2.9.3 Cluster resource inspection

2.9.3.1 Overview

Cluster resource inspection includes cluster inventory inspection and VM inspection.

2.9.3.2 Cluster inventory inspection

This topic describes the inspections of cluster inventory resources. Cluster inventory resources are specified by the number of VMs that can be created by

using the remaining resources in the cluster. You can use the database to obtain the cluster inventory resources.

Suppose you need to inspect the inventory resources of a cluster based on 16-core 64 GB VMs. Run the following command to obtain the inventory resources of the cluster:

```
$ houyiregiondb
mysql> select sum( least ( floor(available_cpu/16),floor(available_
memory/64/1024))) from nc_resource,nc where nc.cluster_id=$id and nc.
biz_status='free' and nc.id=nc_resource.id;
```

If the current cluster contains a relatively large VM, ensure that the cluster has enough free resources to handle the VM, as well as an available host with sufficient resources for backup. This host will be the migration destination of the large VM in case the current host goes down. Otherwise, the large VM cannot be migrated when its host goes down, and you will have to either use hot migration to transfer resources or release redundant VMs in the cluster.

NC state inspection

NC state inspection mainly checks whether the state of a host is normal in the database and Apsara Infrastructure Management Framework.

- A host can be in one of the following states in Apsara Infrastructure Management Framework:
 - Good: indicates that the host is in a normal working state.
 - Error: indicates that the host has an active monitoring alert.
 - Probation: indicates that the host is in the probationary period and may fail.
 - OS _error: indicates that the host has failed and is being cloned.
 - Hw_error: indicates that the hardware of a host has failed and is being repaired.
 - OS _probation: indicates the host is recovering from a fault or hardware failure and is in a probationary period. If the host recovers within the probationary period, the state will change to probation. If the host fails to recover within the probationary period (an error is reported), the state will change to OS _error.



The Good state is considered to be the stable state, and all other states are considered to be unstable states.

- · Cluster definitions for Apsara Infrastructure Management Framework:
 - Default cluster: the cluster where NCs are placed when they go offline.
 - Non-default cluster: the cluster for online NCs.

An NC that is operating normally is placed in a non-default cluster, and is in the Good state.

The mappings of host states between the ECS database and Apsara Infrastructure Management Framework are described in *Table 2-15: Mappings of host states between the ECS database and Apsara Infrastructure Management Framework*.

Table 2-15: Mappings of host states between the ECS database and Apsara Infrastructure Management Framework

Host states in ECS database	Cluster	Host state	Scenario
mlock	Non-default cluster	Unstable	A host that goes online is immediately and proactivel y locked.
locked	Non-default cluster	Unstable	An NC needs to be unlocked •
free	Non-default cluster	Stable	A host operates normally.
nc_down	Non-default cluster	Unstable	A host operates normally or is in downtime.
offline	Default cluster	Unstable	A host goes offline from business attributes.

2.9.3.3 VM inspection

This topic describes pending VM inspections, VM state inspections, and VM resource inspections.

Pending VM inspection

This type of inspection focuses on VMs that have been in the pending state for a long period of time. When a VM has been in the pending state for a long period of time, it is considered a redundant resource. Contact the user to handle it. VM state inspection

This type of inspection focuses on the VM state consistency. For example, a VM is displayed as stopped in the database, but is displayed as running in NC. During the inspection, the VM states recorded in the database and on the host are checked. If the VM states are inconsistent, corresponding operations are performed.

• Run the following command to obtain the VM state in a database:

houyiregiondb -Ne "select status from vm where name='\$name'"

• Run the following command to obtain the VM state on a host:

sudo virsh list | grep \$name

VM resource inspection

After the configuration of a VM is changed, the system checks whether the configuration of the VM recorded in the database is consistent with that used on the host.

• Run the following command to obtain the VM configuration in a database:

houyiregiondb -Ne "select vcpu, memory from vm where name='\$name'"

• Run the following command to obtain the VM configuration on a host:

sudo virsh list | grep \$name

Obtain information about CPU and memory by viewing the corresponding fields.

3 Auto Scaling (ESS)

3.1 Log on to the Apsara Stack Operations console

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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 3-1: Log on to ASO





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 log on to ASO.

3.2 Product resources and services

3.2.1 Application deployment

All the applications in the ESS Business Foundation System are stateless. You must restart the applications by running the docker restart command.

 \cdot 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.
 - Triggers scheduled tasks and monitoring tasks.
- Coordinator

Coordinator is the open API layer that provides public-facing services. It maintains persistent requests and issues tasks.

• Worker

- Worker executes all scaling-related tasks, such as creating ECS instances , adding instances to SLB backend server groups and RDS whitelists, and synchronizing CloudMonitor 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.

3.2.2 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 3-1: Failed services and their impacts* describes the details.

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.

Table 3-1: Failed services and their impacts

Operations and Maintenance Guide - Cloud Essentials and Security / 3 Auto Scaling (ESS)

Service	Impact				
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.				
tianjimon (Monitoring System of Apsara Infrastructure Management Framework)	Some services are unavailable.				

3.3 Inspection

3.3.1 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

3.3.2 Monitoring inspection

This topic describes basic monitoring and connectivity monitoring inspection.

3.3.3 Basic software package version inspection

Version inspection for trigger, coordinator, worker, and base services.

4 Object Storage Service (OSS)

4.1 Log on to the Apsara Stack Operations console

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

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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 4-1: Log on to ASO





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 log on to ASO.

4.2 OSS operations and maintenance

4.2.1 User data

4.2.1.1 Basic bucket information

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 the Apsara Stack Operations console.

- 2. In the left-side navigation pane, choose Products > OSS > User Data.
- 3. On the Bucket Basic Information tab, select the bucket you want to view.

4. Click View, as shown in the following figure.

Products	Alibaba Cloud Accou	nt 🗸 allyuntest View
Product List	Bucket Basic Information User Data Overview	Data Monitoring
ECS Operations and	ECS Operations and Maintenance Platform	
Image Upload RDS	Bucket Name:	alka-ors
✓ OSS	User Account:	aliyuntest(99999999)
User Data	Enterprise/Individual Name:	
Cluster Data	Application:	fie -
✓ MPS	BID:	26842
Batch Retranscoding	Current Capacity:	Standard 199462673468, IA: 06, AR: 06
Apsara Distributed Fil	Storage Type:	standard
ISV Access Configur	Objects:	Standard: 2504947, IA: 0, AR: 0
	Log Service:	Inactivated
	Location:	oss-cn-qingdao-em/4b-d01-a(OssHybridCluster-A-20191028-eac5)
	Permissions:	private
	Anti-Leech Settings:	Configureditules: [Empty Refer: Allowed] [Refer List 0 Entries]
	CORS Settings:	0 Rules Configured

4.2.1.2 User data overview

You can query data statistics and trends, including resource usage and basic attributes of resources by UID, Alibaba Cloud Account, Bucket Name, or Bucket MD5.

Context

The User Data Overview tab is displayed only when you search by UID or Alibaba Cloud account. On the User Data Overview tab, you can specify a date to view total usage of various resources in all buckets owned by the user account.

You can collect resource statistics by total storage capacity, total inbound or outbound traffic through the public network, internal network, or CDN, or total charged requests.

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > User Data.
- 3. On the User Data Overview tab, you can view resource usage such as total storage capacity, total inbound and outbound traffic, and total charged requests by Alibaba Cloud Account or UID.

4. Set Date. Click OK. Click View, as shown in the following figure.

Products		Alibaba Cloud Account 🗸					/iew	
Product List	Bucket Basic Information	User Data Overview Data	User Data Overview Data Monitoring					
✓ ECS ECS Operations and	Select Date ECS Operations and Maintenance Platform							
Image Upload			- 6	Tatal Outban	nd Traffe		atal Observat Deservation	
BMS				Total Outbound Traffic				
RDS	0B	()B	0B			0K	
- OSS	STD IA	AR Public II Network N	nternal CDN letwork	Public Network	Internal CDN Network			PS 2 Requests
Cluster Data		08	0B	08	08			
- MPS	Show Storage Data Show Traffic D	Data						
User Configurations	Bucket Name	Region	Total Storage Capacity	Standard	IA (occupied size)	IA (charged)	Archive (occupied size)	Archive (charged)
Batch Retranscoding								

4.2.1.3 Data monitoring

This topic describes how to monitor OSS data in the Apsara Stack Operations console.

Context

You can query resource running statuses and usage such as the storage capacity, traffic, SLA, HTTP status, latency, QPS, and image processing capacity by UID, Alibaba Cloud Account, Bucket Name, or Bucket MD5. You can also query the resource usage and trends based on a specified time range.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > User Data.
- 3. On the Data Monitoring tab, set Bucket Name, Specify Time Range, and Monitoring Items.

Note:

Metric descriptions:

- SLA: indicates the service level availability metric for OSS. Formula: SLA = Non-5xx request count per 10s or hour/Total valid request count × 100%.
- HTTP Status: collects statistics for the percentages of the numbers of 5xx, 403 , 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
- Latency: collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.

- Storage Capacity: collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.
- Image Processing Capacity: collects statistics for the number of processed images.



By default, this metric is not displayed. You can select this metric from the Monitoring Items drop-down list.

- Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
- QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.

4. Click View.

The following example describes typical operations on the data monitoring trend chart:

• If you query data monitoring information by user, you can click the bucket name in the trend chart to show or hide the curve.

Figure 4-2: Data monitoring 1



• Move the pointer over the trend chart to display data at a specific point in time.



Figure 4-3: Data monitoring 2

4.2.2 Cluster data

4.2.2.1 Inventory monitoring

Metrics of inventory monitoring include the total capacity, available capacity, used capacity, backup ratio, and inventory usage.

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > Cluster data.

3. On the Inventory Monitoring tab, you can view statistics by Apsara Distributed File System, metric data, or KV data usage.

Products	Cluster Data											
Product List	Inventory Monitoring	nembary Monitoring Bucket Statistics Object Statistics Data Monitoring Resource Usage Ranking										
✓ ECS	Report Type: Storage Inventory Data Dimension: Apsara Distribute	sport Type:										
ECS Operations and							Data Increment	N(TB)				
Image Upload	Region T	Cluster ↓	Total Capacity(TB) ↓	Used Capacity(TB) √	Unused Capacity(TB) ↓	Utilization J ¹				Actions		
v oss	cn-qingdao-env4b-d01	osshybridcluster-a-20191028-e ac5	505.39	40.25	465.14	7.96%		0.42	8.06			
User Data Cluster Data	cn-qingdao-env4b-d01	osshybridcluster-a-20191028-a b52	519.83	26.84	492.99	5.16%	-0.02					
▼ MPS												
User Configurations	(i) Remarks:											
Apsara Distributed Fil	2. The data is green w space.	when the Apsara Distributed File Syst	em utilization is 70%-85%, yellow w	hen the utilization is over 85%, and r	ed when Apsara Distributed File Syst	tem expires in 30 days or the physica	I space of Apsara I	Distributed File Sys	tem is two times la	ger than the OSS logical		
ISV Access Configur												

Aside from basic cluster information such as the cluster name and region, you can also view metrics based on the following dimensions:

- Apsara Distributed File System Data: includes the actual total capacity for storage (including the total capacity for multiple data backups), used capacity , remaining capacity (available), usage, and backup ratio.
- Metric Data: includes the bucket storage used by users who use ECS instances and other instances.
- KV Data: includes the logic KV data, KV data in the recycle bin, and data increment (by day, week, or month).

4.2.2.2 Bucket statistics

This topic describes how to collect statistics for the number of buckets by cluster.

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > Cluster Data.

3. On the Bucket Statistics tab, select Report, Current Overall Statistics, or Growth Trend to view bucket statistics.

Products Products	Cluster Data				
Product List	Inventory Monitoring Bucket Statistics	Object Statistics Data Monitoring R	esource Usage Ranking		
▼ ECS	Display Method: Report Statistics V Specify Time Range: 01/15	2020 - 01/22/2020 📾 View			
ECS Operations and					
Image Upload	Region	Cluster	Active Users	Active Buckets	
RDS	cn-qingdao-env4b-d01	o: 8-eb52			
▼ 088	cn-qingdao-env4b-d01	o: ac5	207	1981	
User Data	Total Number v	ithout Duplicates		1988	
Cluster Data				Prev 1 Next >	
User Configurations					
Batch Retranscoding					
Apsara Distributed Fil					
ISV Access Configur					

- If you select Report, specify the time range.
- You can select Current Overall Statistics to query statistics of last hour.
- If you select Growth Trend, you can specify a time range of seven days, 30 days, three months, six months, or one year.
- 4. Click View.

4.2.2.3 Object statistics

This topic describes how to view the statistics for the number and trend of objects by cluster.

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > Cluster Data.

3. On the Object Statistics tab, select Current Overall Statistics or Growth Trend to view object statistics.

Products	Cluster Data				
Product List	Inventory Monitoring Bucket Statistics Object Statistics	Data Monitoring Resource Usage Ranking			
ECS ECS Operations and Image Upload	Display Method: Current Overall Statistics Current Overall Statistics				
RDS	Region	Cluster	Objects		
- OSS	cn-qingdao-env3b-d02	o∈ ==	113991990mull		
User Data	cn-qingdao-env3b-d02	e: 967	118/484/mul 40107104/ul 12554185/ul		
Cluster Data	cn-qingdao-env4b-d01	eac5			
User Configurations	cn-qingdao-env4b-d01	o:eb52			
Batch Retranscoding	Ta	lal	178458119muli		
Apsara Distributed Fil			<pre></pre>		
ISV Access Configur	III Active User Comparison				
	Ten thousand	Ten thousand			
	Ten thousand	Ten thousand			
	Ten thousand	Ten thousand			
	Ten thousand	Ten thousand			
	Ten thousand	/Ten thousand			
	/Ten thousand	Ten thousand	······		
	Ten thousand ossinger or ossingero	Nybridcluster	env3b-d02 cn-qingdao-env4b-d01		

- You can select Current Overall Statistics to query statistics of last hour.
- If you select Growth Trend, you can specify a time range of seven days, 30

days, three months, six months, or one year.

4. Click View.

4.2.2.4 Data monitoring

This topic describes how to collect statistics for each metric by cluster.

Context

Cluster data metrics are similar to user data metrics except that the object of cluster data metrics is the data collected by cluster.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > Cluster Data.
- 3. On the Data Monitoring tab, set Monitoring Items and Specify Time Range. Click View.



Metric descriptions:

 SLA: indicates the service level availability metric for OSS. Formula: SLA = Non-5xx request count per 10s or hour/Total valid request count × 100%.

- Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
- QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.
- Latency: collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.
- HTTP Status: collects statistics for the percentages of the numbers of 5xx, 403 , 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
- Storage Capacity: collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.

4. Move the pointer over the trend chart to display data at a specific point in time.



Figure 4-4: Data monitoring 1

Metric descriptions:

- SLA: indicates the service level availability metric for OSS. Formula: SLA = Non
 -5xx request count per 10s or hour/Total valid request count × 100%.
- HTTP Status: collects statistics for the percentages of the numbers of 5xx, 403, 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
- Latency: collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.
- Storage Capacity: collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.
- Image Processing Capacity: collects statistics for the number of processed images.



By default, this metric is not displayed. You can select this metric from the Monitoring Items drop-down list.

- Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
- · QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.

The following example describes typical operations on the data monitoring trend chart:

· If you query data monitoring information by user, you can click the bucket name in the trend chart to show or hide the curve.



- · Move the pointer over the trend chart to display data at a specific point in time.



Figure 4-6: Data monitoring 2

Figure 4-5: Data monitoring 2

4.2.2.5 Resource usage rankings

This topic describes how to collect usage of resources by cluster. This way, administrators can monitor users that consume more resources.

Context

Data resources can be ranked based on the following metrics:

- Total Requests
- Request Errors
- Public Inbound Traffic and Public Outbound Traffic
- Internal Inbound Traffic and Internal Outbound Traffic
- CDN Uplink Traffic and CDN Downlink Traffic
- Storage Capacity, Storage Increment, and Storage Decrement

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > OSS > Cluster Data.
- 3. On the Resource Usage Ranking tab, select Report or Trend from the Display Mode drop-down list. Select a number from the Top drop-down list. Set Specify Time Range and Monitoring Items to view resource usage.

Products	Clu	ister Data											
Product List		wentory Monitoring Bucket Statistics Object Statistics Data Monitoring Resource Usage Ranking											
▼ ECS	Displ	play Method Top Specify Time Range											
ECS Operations and	Rep	Angent 🗸 🔢 🗸 🔰 0122/2020 12:18:00 — 0122/2020 14:100 🚍 junit ins mage them in unit for inter time that a day blink ins specific time ingen with a daybyet											
Image Upload	Moni												
RDS	10	Rela Request Errors × Public Indound Traffic × Public Outbound Traffic × Internal Indound Traffic × ICON Uplink Traffic × CON Uplink Traffic × Storage Capacity × Storage Increment × Storage Decrement × V											
▼ 0\$\$		Tela Doniedt TrD11											
User Data		Total Requests-TOP10 Request Errors-TOP10											
Cluster Data		Bucket		Total Requests	Bucket		Request Errors						
▼ MPS				3.46Ten thousand			6289						
User Configurations				3.36Ten thousand									
Batch Retranscoding				1.57Ten thousand									
Apsara Distributed Fil				1.12Ten thousand									
of Access contiguin.				1.02Ten thousand									
				4456			0						
				4132									
				2340									
				2064									
				© 2009-2019 Alibaba Cloud Computing Limited. All rights rese	erved.								

- In report mode, you can view the top 10, 30, or 50 buckets by resource usage.
- In trend mode, you can view the top 10 buckets by resource usage.
- 4. Click View.

4.3 Tools and commands

4.3.1 Typical commands supported by tsar

You can use tsar to perform operations and maintenance on OSS. This topic describes typical commands supported by tsar.

tsar allows you to run the following commands:

- · View help details of tsar
 - Command: tsar -help
- · View the NGINX operation data of each minute from the past two days

Command: tsar -n 2 -i 1 -nginx

In this command, -n 2 indicates the data generated in the past two days. -i 1 indicates one result record generated each minute.

 View the tsar load status and operation data of each minute from the past two days

Command: tsar --load -n 2 -i 1

4.3.2 Configure tsar for statistic collection

You can configure tsar to collect data generated when NGINX runs.

Run the following command to configure tsar for statistic collection:

cat /etc/tsar/tsar.conf |grep nginx

The following figure shows that the status of mod_nginx is on.

admin /home/admin
\$cat /etc/tsar/tsar.conf |grep nginx

Ensure that this item is in the on state.

mod_nginx on
 Ensure that this item is in the on state.
 output_stdio_mod_mod_swap,mod_partition,mod_cpu,mod_mem,mod_lvs,mod_haproxy,mod_traffic,mod_squid,mod_load,mod_tcp,mod_udp,
 output_stdio_mod_mod_swap,mod_partition,mod_cpu,mod_mem,mod_lvs,mod_haproxy,mod_traffic,mod_squid,mod_load,mod_tcp,mod_udp,
 output_stdio_mod_mod_swap,mod_partition,mod_cpu,mod_mem,mod_lvs,mod_haproxy,mod_traffic,mod_squid,mod_load,mod_tcp,mod_udp,
 output_stdio_mod_mod_swap,mod_swap,mod_partition,mod_cpu,mod_mem,mod_lvs,mod_haproxy,mod_traffic,mod_squid,mod_load,mod_tcp,mod_udp,
 output_stdio_mod_swap,mod mod_tcpx,mod_apache,mod_pcsw,mod_io,mod_percpu,mod_nginx,mod_tcprt

5 Table Store

5.1 Table Store Operations and Maintenance System

5.1.1 Overview

Table Store Operations and Maintenance System helps locate problems during O&M and notifies users of the current running status of their services. Appropriate use of Table Store Operations and Maintenance System can significantly improve O&M efficiency.

The endpoint of Table Store Operations and Maintenance System is in the format of "chiji.ots.\${global:intranet-domain}."

Table Store Operations and Maintenance System consists of the following modules: User Data, Cluster Management, Inspection Center, Monitoring Center, System Management, and Platform Audit. These modules provide comprehensive O&M functions to meet different requirements.

5.1.2 User data

5.1.2.1 Instance management

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 instances.

You can specify a region and a cluster to view the instances, and the basic information of each instance in the specified cluster.

In	Instance Management											
Filter	Filter: Instance ID 🔻											
Regi	tegion: tn-qmgdan-arm8d-d01(APSA/UA_STACR) * Cluster: ots-ssd-a-20181031-25ce * Search											
Ва	atch Update									Add Instance	Refresh	
							Click to v	view the Modified	instance info	rmation pa	ge.	
	Region	Cluster	User ID	Instance II)	Instance Name	Туре	At	Description	Ope	ration	
•	cn- qingdao- enelid- dD1	ots-ssd-a- 20181031- 25ce	999999999	_gDSYO6	uQza/AUSiqgilit	AsToo1938	PUBLIC	2018- 12-27 10:21:37	AsToolBoxOts	Update	Delete	

On the Instance Management page, you can:

- View the instances in the cluster.
- View instance descriptions.
- View the links to details of instances by clicking instance names.
- 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 filtering conditions.

In	Instance Management												
Filter	Filter: Instance ID VQFtaLlaYyICdK7aga2A. Search												
Regio	legion: The Cluster: The Search												
Ba	itch Update								Add Instance	Refresh			
	Region	Cluster	User ID	Instance ID	Instance Name	Туре	Modified At	Description	Opera	ation			
	m-qingdao- erv6d-d01	tionji-a- 25ee	1265544520690241	_qRuUiYylCdKTuqu2A	asootsins	INTERNAL	2018-12- 11 14:32:30		Update	Delete			

The available filtering conditions include:

- Instance ID
- Instance name
- User ID
- Apsara Stack account

• View instance details.

- Instance overview

Click the otssmoke96 instance to go to the Details tab. This tab provides detailed information about the instance, such as the instance monitoring link, intranet and Internet URLs, and statistics on tables in the instance.

₿ ¶ Instan	nce asootsi	ns
Details	Tables	
Monitoring asootsins M	g onitoring	Click to view the instance monitoring page.
Endpoint		
Public Netw	ork:	
Private Netv	vori	and the constraint of the second second second second
Table Stati	istics	
Total Tables	/Total Data:	0/0B

- Table information

Click the Tables tab to view table information such as the max version, TTL, read CU, write CU, and timestamp.

1	≣ Instance odps								
	Details Tables								
,	Table Name	Max Version	TTL(s)	Read CU -	Write CU 👻	Partitions	Data Size ▼	Pangu Data Size ▲	Timestamp
	ODPS_META_X_META_HISTORY	1	-1		0	1	0B	59.4MB	2019-02-02 11:00:13
	ODPS_META_X_CHANGE_LOGS	1	-1	0	0	1	0B	2720.8KB	2019-02-02 11:00:13

• View table details.

- Details

On the Tables tab, click the test_base_monitor table. On the Details tab, 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								

- Partitions

You can obtain the basic information of a partition, such as the partition ID and worker information. You can also specify filtering conditions to filter the partitions that meet your requirements.

+-0 0-→	Table ODPS_META_X_META_HISTORY													
De	Details Partitions													
Search: Worker V Search														
ID	Partition ID	Start Key	End Key	Worker	Pangu File Size	Data Size	Youchao Files ▼	Timestamp						
1	1 1891d981-771c-45af-b239- 84312b750ba9		\xfd\xfd\xfd\xfd\xfd\xf a36f01001.cloud.f01.amtes		59.4MB	0B	9	2019-02-02 11:00:13						

The available filtering conditions include:

- Worker (For more information, see the value in the Worker column.)
- Partition ID

5.1.3 Cluster management

5.1.3.1 Cluster information

You can obtain cluster information through cluster searches, cluster usage, and top requests.

Function description

• Clusters

Clus	Cluster Information										
Region:	All	•	ОСМС	OCM Cluster Synchronization							
Status		Cluster	Region	Storage Type	Operation						
		Click to vie	ew the cluster informatic	on page.							
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 from which to obtain 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 Table Store Operations and Maintenance System to obtain all clusters in the regions.
- Cluster deletion: You can use this function to remove a cluster from Table
 Store Operations and Maintenance System after you confirm that the cluster is offline.

• Cluster details

Clust	Cluster Information											
Region:	All	•	OCM C	luster Synchronization	Refresh							
Status		Cluster	Region	Storage Type	Operation							
using		ots-hy-a-20181217-2e46	ew the cluster informatic cn-qingdao-env8	n page. 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:

- Overview: provides the basic information of a cluster.

Cluster ots-hy-a-201	81217-2e46
Overview Top Res	source Usage
*Region: cn-gingdac-anvöd-	d01jAPSAJA_STACK) Cluster: da-hy-e-20181217-2e46 Switch Cluster
Region Description	APSARA_STACK
Region	on-gingdao-envild-d01
Cluster	ots-hy-e-20101217-2e46
Armory App	mock_armory
Gateway	mock_ag
Cluster Type	public

- Top: provides top request information by partition and table.

Overview	Тор	Resource U	lsage Clic	k-to view the information (page of top request	
Top Parti	Click to vi tions by I	iew the table infor Pangu File Size	mation page.	Top Partitions by	Youchao Files	
			lick to view the partition	e formation page.		1
Table N	lame	Partition ID	Pangu File Size 🔻	• Table Name	Partition ID	Youchao Files 🔻
Top Table	es bv Par	nau File Size		Top Tables by Yo	uchao Files	
	Click	to view the instand	ce information page <mark>. Mor</mark>	e		[
						-

Resource Usage: provides cluster usage details. Typically, the usage statistics collection task is automatically triggered in the back-end at specific intervals. In special cases, you can click Collect Data to manually trigger the usage statistics collection task. After the usage statistics collection task is 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 information.)

🛱 Clust	er ots-h	y-a-20181	217-2	e46									
Overview	v Тор	Resou	urce Usa	ige		Click to r	nan	ually collect r	esou	rce usage ir	nfo	rmation	
Collected At: ~													
Check Result :													
Storage Resource Usage													
Total Di	isk Size	Tota	I File Siz	e	Recy	cle Bin Size	Table Size Fr		Free	Free Space D		isk Usage Ratio (%)	
											%		
Gap Siz	ze H	osts Total/Ma	aster/OT	SServer/S	qlWork	er	Hybrid Deployment		С	Cluster Type		Scale-out Requirement	
	11												
OTSSer	ver Reso	ource Usa	ge										
Hosts	Failed Hosts	Avg/Max (Usage (%	CPU)	Increase CPU Cor	d 'es	Avg/Max NetIn (MB/s)	In Ex	creased Hosts Due ccessive NetIn	to	Avg/Max NetOut (MB/s))	Increased Hosts Due to Excessive NetOut	
		1				1				/			

5.1.4 Inspection center

5.1.4.1 Abnormal resource usage

You can click Abnormal Resource Usage in the left-side navigation pane to locate all cluster abnormalities and their causes.

Function description

Abnorm	al Resourc	e Usage							
									Collect Data
Cluster Name						Abnorm	al Resource	e Usage	
	Date	Total Disk Size	Total File Size	Gap	Recycle Bin Size	Table Size	Free Space	Disk Usage Ratio (%)	Scale-out Requirement
tarý-a- 25ee	2019- 02-02	64.46TB	6.21TB	3.25TB	1.64TB	1.32TB	48.80TB	24.31%	Reach Safe Level in -1Days, Growth Rate:-35.27GB/Days Reach Safe Level in -1Days, Growth Rate:-35.28GB/Days

You can click Abnormal Resource 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 back-end at specific intervals. In special cases (such as a failure in back-end task execution), you can click Collect Data to manually trigger usage statistics collection. The collection action is performed asynchronously. After the usage statistics collection task is complete, refresh the page to display the latest usage statistics.

5.1.5 Monitoring center

5.1.5.1 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 within a specified time range, and determine whether a cluster service is healthy based on the metrics in the following dimensions.



5.1.5.2 Application monitoring

You can check the instance-level and table-level metrics to determine whether a service that belongs to a user is abnormal.

Function description

You can check the following metrics to determine whether a service for a specified user is in the healthy state.



The Instance field is required. Table and Operation fields are optional.



5.1.5.3 Top requests

You can view the top request distribution of clusters by monitoring level and dimension.

Function description

Four monitoring levels are supported for top requests: Instance, Instance-Operation, Instance-Table, and Instance-Table-Operation. You can view the top request details of a cluster based on 13 different metrics, such as the total number of requests and the total number of rows.

Top Requests	Top Requests												
*Region: m-ging	Region: cn-qingdo-envild-d01(APSARA_STACK) • Cluster: 5arji-a-25ce				*Time: 2	019-02-02 12:40:19	•	2019-02-0	2 13:40:19	1 Hour	٣		
*Monitoring Level	Instance	•	*SortBy: Tota	I Requests	¥	*TopN: 100)		Search				
Top Requests													
Торіс	Total Reque sts ↓	Total Ro ws ▼	Total Failed R ows -	Public Upli nk 👻	Public Downl	Internal Upli nk 👻	Internal Downl	Read C U -	Write C U 👻	Total Lat ency Max Avg	SQLWorker La tency Max Avg	HTTP Status	SQL Status
(instanceName=n etric	1,643,542	73,033,4 06	0	OB	OB	19.3GB	1308.2MB	245,919	73,070, 441	614,911 us 13,686 u s	613,801 us 12,844 us	{"200":1643542}	{"0":73175642}
(instanceName=o dps,	186,686	185,768	0	08	OB	45.4MB	100.7MB	180,059	11,366	203,426 us 885 us	203,268 us 748 us	{"200":186686}	{"0":186686}

5.1.5.4 Request log search

You can search for a log entry based on a request ID to streamline problem investigation.

Function description

Query all log information about a request based on the request ID.

Request Log Search												
*Region:	m-gingdas-enr@d-d01(APSARA_STACK) * *Cluster:	tanji-a-25os 🔹	*Request ID:		Search							
Log Searc	Log Search Result											
Host	Timestamp		File	Content								

5.1.6 System management

5.1.6.1 Manage tasks

You can maintain the back-end tasks in Table Store Operations and Maintenance System.

Function description

After Table Store Operations and Maintenance System is deployed in the Apsara Stack environment, the back-end tasks that collect usage statistics are automatically integrated. You can perform the following operations on the backend tasks:

 \cdot View task details such as the specific parameters and running time of each task.

• Enable or disable a task.



Disabled tasks no longer run automatically.

• Run a task immediately.

The following figure shows the monitoring task details page. Based on the monitoring rules, the task collects usage statistics at 2:00 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				

5.1.6.2 View tasks

You can view the execution status of back-end tasks and locate the causes of task exceptions.

The following figure shows the execution status of back-end tasks in Table Store Operations and Maintenance System. You can view the tasks, which have either succeeded or failed.

View Tasks						
All Tasks	Host VIP/Net A	pplication Resource	Usage Remote HTTP			
Time Range:	2019-02-02	То 2019-02-02	Check		All	
Status	Name	Туре	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	

Click View All or View Abnormal in the Operation column corresponding to the abnormal task to view the specific cause of a task failure, as shown in the following figure.



5.1.7 Platform audit

5.1.7.1 Operation logs

You can view the management and control operation logs of Table Store Operations and Maintenance System.
Function description

The Operation Log page provides the operation logs of Table Store Operations and Maintenance System. You can query audit records generated within a specified time range and filter the records as required. This helps management personnel obtain information about the platform status.

Operation Log												
Time Range: 2018-12-31 00:00:00 To 2019-02-02 01:	05:00 Add Condition Operator	Y	Check Chiji Log									
Operation Log	Operation Name	IP	Operator	Time								
/ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.128.219	aliyuntest	est 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	2019-01-18 13:39:37									
/ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.64.187	2019-01-18 13:36:08									

5.2 Cluster environments

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 back-end of the Table Store O&M platform
- · chiji: the Table Store O&M platform frequently used for fault location
- · TableStoreSqlInner/TableStoreSql: the Table Store back-end tool

5.3 System roles

- TableStoreOCM
 - OCMInit: the OCM initialization tool used to create tables and bind POP APIs
 - OCM: the service node of OCM
 - ServiceTest: the service test image of OCM
- TableStoreInner/TableStore
 - InitCluster: the process of adding cluster information to OCM, including the domain name, the cluster type, and the pre-configured Table Store account information
 - LogSearchAgent: the log collection service node of Table Store
 - 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 alert service module of Table Store
 - OTSFrontServer: the frontend server of Table Store, which can be NGINX, OTS Server, or Replication Server
 - OTSServer: the OTS server
 - OTSTEngine: the NGINX service for Table Store frontend servers
 - PortalAgServer: the backend service for Table Store 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 alerts but no longer provides services
 - TableStoreAdmin: all O&M tools of Table Store, including the splitting and merging tools
- TableStorePortal
 - PortalApiServer: the backend service for Table Store Operations and Maintenance System
- TableStoreSqlInner/TableStoreSql
 - Tools: the backend tools for Table Store, such as sqlonline_console
 - UpgradeSql: the backend hot upgrade tool for Table Store

5.4 Pre-partition a table

5.4.1 Pre-partitioning

When you create a table, Table Store automatically creates a partition for the table . This partition can be configured to automatically split based on the data size or data access load as your business develops. A table with only one partition may be unable to provide sufficient service capabilities during a stress test or data import. In this scenario, you must pre-partition the table.

Pre-partitioning rules

You can estimate the number of partitions required based on the standard size of 10 GB per partition. However, considering other factors such as the number of hosts and concurrent write operations by developers, we recommend that the total number of partitions do not exceed 256. If data can be written into the table evenly, you can partition the table equally based on the number of partitions required.

Note:

When data is written into the table, the system automatically splits the table to ensure sufficient partitions are available as the data increases.

Pre-partitioning methods

You can use split_merge.py to pre-partition a data table. You can obtain split_merge.py from /apsara/TableStoreAdmin/split on the host of TableStoreAdmin in TableStoreInner. You can use any of the following methods to partition a data table:

· Specify a split point

python2.7 split_merge.py split_table -p point1 point2 ... table name

- · Specify the number of partitions and the partition key format
 - The partition key is of the int type.

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_digit table name
```

- The partition key starts with an MD5 hash in lowercase. The MD5 hash can contain digits and lowercase letters a to f.

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_hex_lower table name
```

- The partition key starts with an MD5 hash in uppercase. The MD5 hash can contain digits and uppercase letters A to F.

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_hex_upper table name
```

- The partition key is encoded in Base64. The key can contain digits, letters, plus signs (+), and forward slashes (/).

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_base64 table name
```

- -- only_plan: generates split points but does not split the table. -- force:

directly splits the table without manual confirmation.

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_digit --only_plan table name
```

• Split a partition based on the existing data

```
python2.7 split_merge.py split_partition -n PART_COUNT (number of
partitions) partition_id
```

Note:

You can also use the preceding methods to partition a table that already has data.

5.4.2 View partitions

You can view the partitions of a data table in Table Store Operations and Maintenance System. On the homepage of Table Store Operations and Maintenance System, choose User Data > Instance Management from the left-side navigation pane. On the Instance Management page that appears, set Region and Cluster. Click Search. Locate an instance and click the instance name. On the Details tab that appears, click the Tables tab. On the displayed tab, click a table. On the Details tab that appears, click the Partition tab. You can view the information of all partitions in the Table. The information contains the partition ID, range, worker, Apsara Distributed File System file size, and data size. The partition size displayed may not be the current partition size because the data is updated only after the files are merged in the backend of the system. The Apsara Distributed File System file size is the compressed data size. The actual storage space is three times the file size because the data is stored in three copies.

6 ApsaraDB for RDS

6.1 Architecture

6.1.1 System architecture

6.1.1.1 Backup system

ApsaraDB for RDS can back up databases at any time and restore them to any point in time based on the backup policy, which makes the data more traceable.

Automatic backup

ApsaraDB RDS for MySQL supports both physical and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can create temporary backup files when necessary. Temporary backup files are retained for seven days.

Log management

ApsaraDB RDS for MySQL automatically generates binlogs and allows you to download them for local incremental backup.

Instance cloning

A cloned instance is a new instance with the same content as the primary instance , including data and settings. This feature allows you to restore data of the primary instance or create multiple instances that are the same as the primary instance.

6.1.1.2 Data migration system

ApsaraDB for RDS provides Data Transmission Service (DTS) to help you migrate databases.

Replicate databases between instances

ApsaraDB for RDS allows you to migrate databases from one instance to another.

Migrate data to or from RDS instances

ApsaraDB for RDS provides professional tools and migration wizards to help you migrate data to or from RDS instances.

Download backup files

ApsaraDB for RDS retains backup files for seven days. During this period, you can log on to the RDS console to download the files.

6.1.1.3 Monitoring system

RDS provides multi-dimensional monitoring services across the physical, network, and application layers to ensure business availability.

Performance monitoring

RDS provides nearly 20 metrics for system performance monitoring, such as disk capacity, IOPS, connections, CPU utilization, network traffic, TPS, QPS, and cache hit rate. You can obtain the running status information for any instances within the past year.

SQL auditing

The system records the SQL statements and related information sent to RDS instances, such as the connection IP address, database name, access account, execution time, and number of records returned. You can use SQL auditing to check instance security and locate problems.

Threshold alerts

RDS provides alert SMS notifications if status or performance exceptions occur in the instance.

These exceptions can be involved in instance locking, disk capacity, IOPS, connections, and CPU. You can configure alert thresholds and up to 50 alert recipients (of which five are effective at a time). When an instance exceeds the threshold, an SMS notification is sent to the alert recipients.

Web operation logs

The system logs all modification operations in the RDS console for administrators to check. These logs are retained for a maximum of 30 days.

6.1.1.4 Control system

If a host or instance does not respond, the RDS high-availability (HA) component checks for exceptions and fails over services within 30 seconds to guarantee that applications run normally.

6.1.1.5 Task scheduling system

You can use the RDS console or API operations to create and delete instances, or switch instances between the internal network and Internet. All instance operations are scheduled, traced, and displayed as tasks.

6.2 Log on to the Apsara Stack Operations console

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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 6-1: Log on to ASO





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 log on to ASO.

6.3 Instance management

You can view instance details, logs, and user information.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > RDS.

3. On the Instance Management tab of RDS, you can perform the following operations:

• View instances

View instances that belong to the account on the Instance Management tab, as shown in *Figure 6-2: Instances*.

Figure 6-2: Instances

Instance Name 🗸 Enter		(٦								
Instance Name	Availa	CPU Perfor	QPS Perfor	IOPS Perfor	Conne	Disk Usage	Instance Status	Datab Type	Actions		
	Yes	Yes				0		mysql	User Information Create Backup		
	Yes	Yes		2 %				redis	User Information Create Backup		

• View instance details

Click the ID of an instance to view details, as shown in *Figure 6-3: Instance details*. You can switch your service between primary and secondary instances and query history operations on this page.



If data is not synchronized between the primary and secondary instances, a forced switchover may result in data loss. Proceed with caution.

	Backup ID:
Instance Name:	Database: mysql 5.6
Backup Switch: On	No Persistent Backup: No Persistent Data
Retention Days: 30	Estimated Time:
Database List: All Databases	Backup Time: 18:00
Backup Status: Not Starled	Next Backup: Sep 27, 2019, 18:00:00
Backup Method: Physical Backup	Backup Type: Full Backup
Secondary Server IP:	IDC:
Backup Start At: -	Backup Uploading Start At -
Backup Source: Secondary Database Only	Log Uploading Start At: Sep 5, 2019, 17:36:12
Backup Compression: Table Compression	
Backup Period: V Monday V Tuesday V Wednesday V Thursday V Friday V Saturda	y 🗾 Sunday
Create Single Database Backup	

• View user information

On the Instance Management tab, click User Information in the Actions column corresponding to an instance, as shown in *Figure 6-4: User information*.

User Information 5													
	User Information:												
Instance Name	Instance Status	Database Typ e	Instance Usa ge Type			IOPS Utilization							
	CREATING	Redis			- %		%		%	s			
	CREATING	Redis	-		- %		s		- %	×			
	CREATING	Redis	-		- %		- %		- %	s			
	CREATING	Redis			- %		*		*	s			
	CREATING	Redis			- %		*		*	×			
	CREATING	Redis	-		- %				*	×			
Contraction (1997)	CREATING	Redis	1000		*		- %		%	s			

Figure 6-4: User information

Create backups

For ApsaraDB RDS for MySQL instances, click Create Backup in the Actions column to view the backup information, as shown in *Figure 6-5: Backup information*.

You can also click Create Single Database Backup on the Backup Information page to back up a single database.

Figure 6-5: Backup information

	Backup ID:
Instance Name:	Database: mysql 5.6
Backup Switch: On	No Persistent Backup: No Persistent Data
Retention Days: 30	Estimated Time:
Database List: All Databases	Backup Time: 18:00
Backup Status: Not Started	Next Backup: Sep 27, 2019, 18:00:00
Backup Method: Physical Backup	Backup Type: Full Backup
Secondary Server IP:	IDC:
Backup Start At: -	Backup Uploading Start At: -
Backup Source: Secondary Database Only	Log Uploading Start At: Sep 5, 2019, 17:36:12
Backup Compression: Table Compression	
Backup Period: 🖉 Monday 🖉 Tuesday 😴 Wednesday 📝 Thursday 📝 Friday 😴 Satur Note:	tay 🔽 Sunday
Create Single Database Backup	

6.4 Manage hosts

You can view and manage hosts.

Procedure

- 1. Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > RDS.
- 3. On the Host Management tab of RDS, you can view all host information.

RDS	RDS														
Instance Manager	Instance Management Host Management														
Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Ver sion	Database Engine								
	Normal offline	cn-qingdao-env8d-d0 1				5.6	MySQL								
100000	Normal offline	cn-qingdao-env8d-d0 1				5.6	MySQL								
	Normal offline	cn-qingdao-env8d-d0 1				5.6	MySQL								
10000	Normal offline	cn-qingdao-env8d-d0 1				5.6	MySQL								

4. Click a hostname to go to the RDS Instance page. You can view all instances on this host.

RDS In	RDS Instance 5																
Instanc e Lock Mode	O&M E nd Tim e	Instanc e Type	RDS In stance ID	Instanc e ID	Instanc e Spec ificatio n Code	Tempo rary In stance	Host I D	Instanc e Link Type	Databa se Eng ine	Instanc e Nam e	Instanc e Disk Storag e	RDS In stance Port	O&M S tart Ti me	Associ ated UI D	Instanc e Role	Databa se Eng ine Ver sion	Instanc e Statu s
															🗸 Prev	1 2	Next >

6.5 Security maintenance

6.5.1 Network security maintenance

Network security maintenance consists of device and network security maintenance.

Device security

Check network devices and enable their security management protocols and configurations of devices.

Check for timely updates to secure versions of network device software.

For more information about the security maintenance method, see the device documentation.

Network security

Based on your network considerations, select the intrusion detection system (IDS) or intrusion prevention system (IPS) to detect abnormal Internet and Intranet traffic and protect against attacks.

6.5.2 Account password maintenance

Account passwords include RDS system passwords and device passwords.

To ensure account security, you must periodically change the system and device passwords, and use passwords with high complexity.

7 AnalyticDB for PostgreSQL

7.1 Overview

Purpose

This guide summarizes possible problems that you may encounter during O&M operations and provides solutions for you.

If you encounter system problems not covered in this guide, you can submit a ticket to Alibaba Cloud for technical support.

Requirements

You must possess IT skills including computer network knowledge, computer operation knowledge, problem analysis, and troubleshooting.

Additionally, you must pass the pre-job training of the Alibaba Cloud system to learn necessary Alibaba Cloud system knowledge, including but not limited to system principles, networking, features, and the use of maintenance tools.

Note that during maintenance operations, you must comply with operating procedures to ensure personal and system security. User data must be kept strictly confidential and must not be copied or disseminated without the written consent of the users.

Precautions

To ensure a stable system and avoid unexpected events, you must follow the following guidelines.

· Hierarchical permission management

Permissions on networks, devices, systems, and data are granted based on the services and roles of the O&M personnel to prevent system faults caused by unauthorized operations.

• System security

Before performing any system operations, you must be aware of their impacts.

You must record all problems encountered during operations for problem analysis and troubleshooting.

- Personal and data security
 - You must take safety measures in accordance with the device manuals when operating electrical equipment.
 - You must use secure devices to access the business network.
 - Unauthorized data replication and dissemination are prohibited.

Support

You can contact Alibaba Cloud technical support for help.

7.2 Architecture

Physical cluster architecture

The following figure shows the physical cluster architecture of AnalyticDB for PostgreSQL.

Figure 7-1: Physical cluster architecture



You can create multiple instances within a physical cluster of AnalyticDB for PostgreSQL. Each cluster includes two components: the coordinator node and the compute node.

• The coordinator node is used for access from applications. It receives connection requests and SQL query requests from clients and dispatches computing tasks to

compute nodes. The cluster deploys a secondary node of the coordinator node on an independent physical server and replicates data from the primary node to the secondary node for failover. The secondary node does not accept external connections.

 Compute nodes are independent instances in AnalyticDB for PostgreSQL. Data is evenly distributed across compute nodes by hash value or RANDOM function , and is analyzed and computed in parallel. Each compute node consists of a primary node and a secondary node for automatic failover.

Logical architecture of an instance

You can create multiple instances within a cluster of AnalyticDB for PostgreSQL. The following figure shows the logical architecture of an instance.



Figure 7-2: Logical architecture of an instance

Data is distributed across compute nodes by hash value or RANDOM function of a specified distributed column. Each compute node consists of a primary node and a secondary node to ensure dual-copy storage. High-performance network communication is supported across nodes. When the coordinator node receives a request from the application, the coordinator node parses and optimizes SQL statements to generate a distributed execution plan. After the coordinator node sends the execution plan to the compute nodes, the compute nodes will perform an MPP execution of the plan.

7.3 Routine maintenance

7.3.1 Check for data skew on a regular basis

You must check for data skew on a regular basis during maintenance to prevent the instance from being read-only due to excessive data in some compute nodes.

You can use the following methods to locate data skew. The procedure is as follows.

- 1. For a single table or database, you can view the space occupied within each compute node to determine whether data has been skewed.
 - a. Execute the following statement to determine whether the data in a database has been skewed:

```
SELECT pg_size_pretty(pg_database_size('postgres')) FROM
gp_dist_random('gp_id');
```

You can view the space occupied by the dbname database in each compute node after the statement is executed. If the space occupied in one or more compute nodes is significantly greater than that of other compute nodes, it indicates the data in this database is skewed.

b. Execute the following statement to determine whether the data in a table has been skewed:

```
SELECT pg_size_pretty(pg_relation_size('tblname')) FROM gp_dist_ra
ndom('gp_id');
```

Using the preceding statement, you can view the space occupied by the tblname table within each compute node after the statement is executed. If the space occupied within one or more compute nodes is significantly greater than that of other compute nodes, it indicates the data in this table is skewed. You must modify the partition key to redistribute the data.

- 2. You can use the system views to determine whether data has been skewed.
 - a. Execute the following statement to check whether the storage space is skewed. The principle of this method is similar to that of the preceding space-viewing method:

SELECT * FROM gp_toolkit.gp_skew_coefficients

You can use the view to check the data volume of rows in a table. The larger the table, the more time it will take for the check to complete.

b. Use the gp_toolkit.gp_skew_idle_fractions view to calculate the percentage of idle system resources during a table scan to check whether the data is skewed:

SELECT * FROM gp_toolkit.gp_skew_idle_fractions

For more information, see Checking for Uneven Data Distribution.

7.3.2 Execute VACUUM and ANALYZE statements

You can execute VACUUM and ANALYZE statements on a regular basis for frequently updated tables and databases. You can also execute VACUUM and ANALYZE statements after you have performed a large number of update or write operations to prevent the operations from consuming excessive resources and storage space.

7.4 Security maintenance

7.4.1 Network security maintenance

Regular maintenance will help ensure the security of networks and devices.

Device security

Check network devices and enable the security management protocols and configurations for the devices you want to secure. Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner. For more information about security maintenance methods, see the product documentation of each device.

Network security

You can select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check public and internal traffic, and defend the network against abnormal behaviors and attacks.

7.4.2 Account password maintenance

Account passwords include the superuser password of AnalyticDB for PostgreSQL and the password of the host operating system.

To ensure account security, use complex passwords and periodically change the passwords of systems and devices.

8 KVStore for Redis

8.1 O&M tool

The Apsara Stack Operation console provides the following operations and maintenance (O&M) features for KVStore for Redis:

- Instance management: allows you to view instance details, instance logs, and user information.
- · Host management: allows you to view and manage hosts.



8.2 Architecture diagram

8.3 Architecture

8.3.1 Architecture

8.3.1.1 Backup system

Automatic backup

KVStore for Redis supports full backup. You can flexibly configure backup start time based on off-peak hours of your business. The system retains backup files for seven days or fewer.

Temporary backup

You can create temporary backups as needed. The system retains backup files for seven days or fewer.

8.3.1.2 Data migration system

Migrate data to and from KVStore for Redis

KVStore for Redis provides professional tools and migration wizards to help you migrate data to or out of KVStore for Redis.

Download backup files

KVStore for Redis retains backup files for seven days or fewer. During this period, you can log on to the KVStore for Redis console to download the files.

8.3.1.3 Monitoring system

Performance monitoring

KVStore for Redis provides a variety of system performance metrics, including disk capacity, memory usage, connections, CPU usage, network traffic, QPS, and request command operations. You can check the running status information within a period of one year for an instance.

Threshold alerts

KVStore for Redis can notify you of alerts by means of SMS messages in the case of exceptions in instance status or performance.

These exceptions involve instance locked status, disk capacity, input/output operations per second (IOPS), connections, and CPU usage. You can customize alert thresholds and configure 50 alert contacts or fewer. Five of these alert contacts can take effect at the same time. When an instance exceeds the threshold, the system sends SMS messages to the corresponding alert contacts.

Web operation logs

The system keeps logs for all changes in the KVStore for Redis console. Therefore, the administrator can check these logs. The system retains logs for 30 days or fewer

8.3.1.4 Control system

After a host or instance crashes, the KVStore for Redis high-availability (HA) component checks for the exception and performs the failover operation within 30 seconds. This guarantees that applications run normally and the KVStore for Redis service is highly available.

8.3.1.5 Task scheduling system

You can use the KVStore for Redis console or KVStore for Redis API operations to create and delete instances or switch instances between the internal and public networks. The backend schedules, traces, and displays all instance operations as tasks.

8.4 Log on to the Apsara Stack Operations console

Prerequisites

- ASO access address in the format of http://region-id.aso.intranet-domainid.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 8-1: Log on to ASO

Enter a user name
Enter the password
Log On



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 log on to ASO.

8.5 Instance management

You can view instance details, logs, and user information.

Procedure

- **1.** Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > RDS to go to the RDS page. Click the Instance Management tab. On the Instance Management tab, you can perform these operations:
 - View the list of instances.
 - On the Instance Management tab, you can view the instances under your account.
 - View the details of an instance.
 - Click the ID of a target instance to view the details of the instance.
 - View user information.

Click User Information in the Actions column.

8.6 Host management

Host management allows you to view and manage hosts.

Procedure

1. Log on to the Apsara Stack Operations console.

2. In the left-side navigation pane, choose Products > RDS to go to the RDS page. Click the Host Management tab to view the information about all hosts.

RDS	RDS													
Instance Managen	nent Host M	anagement												
Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine							
all and the second second second			$\frac{1}{2} = \frac{1}{2} = \frac{1}$											
			(1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,				ayora.							
CETT														
and the state of the second second			***		-		-							
				10.00	-									
estru			001											
		© 2009-2018 Alibaba (Cloud Computing Limited. All rig	hts reserved.										

3. Click a host name to go to the RDS Instance page. You can view all instances on this host.

RDS Insta	RDS Instance 5															
Instance Lock Mode	O&M End Time	Instance Type	RDS Instance ID	Instance ID	Instance Specifi Code	Tempo Instance	Host ID	Instance Link Type	Datab Engine	Instance Name	Instance Disk Storage	RDS Instance Port	O8M Start Time	Instance Role	Datab Engine Version	Instance Status
0									14.64							iling:
0				-			-					-		-		-
0					***		-		inin.			-		-		iniş.
0			-	-					-	202 2 ¹				-		
0									-	202 201				-		-
				62	009-2018 Aliba	aba Cloud Con	nputing Limite	d. All rights res	erved.							

8.7 Security maintenance

8.7.1 Network security maintenance

Network security maintenance involves device security and network security.

Device security

Check network devices, and enable security management protocols and configurat ions for these devices.

Check software versions of network devices and update them to more secure versions in time.

For more information about security maintenance methods, see documents of related devices.

Network security

Based on your network conditions, select the intrusion detection system (IDS) or intrusion prevention system (IPS) to detect abnormal Internet and intranet traffic and protect against abnormal behavior and attacks in real time.

8.7.2 Password maintenance

Passwords include system passwords and device passwords in KVStore for Redis.

To secure your account, you must periodically change the system and device passwords, and use complex passwords.

9 Apsara Stack Security

9.1 Log on to the Apsara Infrastructure Management Framework console

This section describes how to log on to the Apsara Infrastructure Management Framework console.

Prerequisites

You have obtained the URL of the Apsara Stack Operations console and the username and password to log on to the console from your system administrator.

Procedure

- 1. In the browser address bar, enter https://Apsara Stack Operations URL, and press Enter.
- 2. On the logon page, enter the username and password, and click Log On.
- 3. In the left-side navigation pane, choose Products.
- 4. In the product list, click Apsara Infrastructure Management Framework to go to the Apsara Infrastructure Management Framework console.

9.2 Routine operations and maintenance of Server Guard

9.2.1 Check the service status

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.



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.

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.alibaba-inc.com/aqs/aegiserverlite@sha256:f9d292f54c
58646b672a8533a0d78fba534d26d376a194034e8840c70d9aa0b3 "/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 ? Sl 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_MESS AGE.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.

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.

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.

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/aegiserever/bin/jbossctl restart

d. Run the following command to check whether the process has been successfully restarted:

curl 127.0.0.1:7001/checkpreload.htm

- Restart Server Guard Update Service:
 - a. Run the following command to view the Java process ID:

ps aux |grep aegisupdate

b. Run the following command to stop the current process:

kill -9 process

c. Run the following command to restart the process:

sudo -u admin /home/admin/aegisupdate/bin/jbossctl restart

d. Run the following command to check whether the process has been successfully restarted:

curl 127.0.0.1:7001/checkpreload.htm

- Restart the Defender service of Server Guard.
 - a. Run the following command to view the Java process ID:

ps aux |grep secure-service

b. Run the following command to stop the current process:

kill -9 process

c. Run the following command to restart the process:

sudo -u admin /home/admin/secure-service/bin/jbossctl restart

d. Run the following command to check whether the process has been successfully restarted:

curl 127.0.0.1:7001/checkpreload.htm

9.3 Routine operations and maintenance of Network Traffic Monitoring System

9.3.1 Check the service status

9.3.1.1 Basic inspection

When you inspect Network Traffic Monitoring System, check whether the service has reached its final state.

Procedure

- 1. Log on to the Apsara Infrastructure Management Framework console.
- 2. Choose Operations > Project Operations. On the page that appears, enter yundunadvance in the search bar. Click Details in the Actions column corresponding to the yundun-advance cluster.
- 3. On the Cluster Operations page that appears, select BeaverCluster.
- 4. On the Cluster Dashboard page that appears, scroll down to the Service Instances section and check whether yundun-beaver-advance has reached its final state.

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

- **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-tolast 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.

9.3.2 Common operations and maintenance

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

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

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

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

9.4 Routine operations and maintenance of Anti-DDoS Service

9.4.1 Check the service status

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.

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_d efender_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*

9-1: Check the status of Anti-DDoS Service.

Figure 9-1: Check the status of Anti-DDoS Service

1	[root@1]]].cloud. /home/admin]
2	<pre>#aliguard_defender_check</pre>
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 9-2: Reinjection route error message* appears, the reinjection route is faulty.

Figure 9-2: 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 9-3: BGP routing error message* appears, the BGP protocol (for traffic routing) is faulty.

Figure 9-3: 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.



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.

[rootdcloud.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		

9.4.2 Common operations and maintenance

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.

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:

aliguard_host_ip	port	aliguard_port_ip	csr_port_ip
10.1.4.12	TO	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	TO	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

Figure 9-4: Anti-DDoS Service configuration example

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 9-5: Unbind NICs*.

Figure 9-5: 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 9-6: Bind NICs*.

Figure 9-6: 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 9-4: 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.

9.5 Routine operations and maintenance of Threat Detection Service

9.5.1 Check the service status

9.5.1.1 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.

9.5.1.2 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.

9.5.2 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

9.6 Routine operations and maintenance of WAF

9.6.1 Check the service status

9.6.1.1 Basic inspection

The basic inspection feature of Web Application Firewall (WAF) focuses on whether the service has reached the final status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.

2. Choose Operations > Project Operations.

- 3. In the Fuzzy Search search box, enter yundun-semawaf. The search results are displayed.
- 4. Click Details in the Actions column. The Cluster Operations page is displayed.
- 5. In the cluster list, click the cluster name that starts with SemaWafCluster.
- 6. In the Service Instances area on the Cluster Dashboard page, check whether the yundun-semawaf service instance is in final status.

Note:

If the Final Status column for an instance is True, the instance has reached final status.

9.6.1.2 Advanced inspection

The advanced inspection feature of Web Application Firewall (WAF) focuses on the system status and service status.

Procedure

- **1.** Log on to the Apsara Infrastructure Management Framework console.
- 2. Log on to two WAF physical machines respectively.
 - a) In Apsara Infrastructure Management Framework, choose Operations > Project Operations.
 - b) In the Fuzzy Search search box, enter yundun-semawaf. Click Details in the Actions column, and the Cluster Operations page is displayed.
 - c) Click the SemaWafCluster cluster.
 - d) In Service Instances, select yundun-semawaf, and click Details. The Service Instance Information Dashboard page is displayed.
 - e) In Server Role List, select YundunSemawafApp#, and click Details. The Server Role Dashboard page is displayed.
 - f) In Machine Information, click Terminal to log on to two WAF physical machines respectively.

- 3. Check the system status.
 - a) Check the system logs.

Run the dmesg -T |tail -30 command to check for exception logs.

- b) Check the system load.
 - Run the free -h command to check whether the memory usage is normal.
 - Run the df -h command to check whether the disk usage is normal.
 - Run the uptime command to check whether the system load average is normal.
 - Run the top command to check whether the CPU usage is normal.
- 4. Check the service status.



The following check is based on the WAF installation directory, which is /home/ safeline by default.

- a) Run the cd /home/safeline command to open the installation directory.
- b) Check the minion service.
 - A. Run the systemctl status minion command to check the execution time and status of the minion service.
 - B. Run the tail -100 logs/minion/minion.log command to check for exception logs.
- c) Check the mgt-api service.
 - A. Run the docker logs --tail 50 mgt-api command to check for exception logs.
 - B. Run the docker exec -it mgt-api supervisorctl status command to check whether the service runs normally and whether uptime is normal.
 - C. Run the tail -50 logs/management/gunicorn.log command to check for exception logs.
 - D. Run the tail -50 logs/management/daphne.log command to check for exception logs.
 - E. Run the tail -50 logs/management/scheduler.log command to check for exception logs.
 - F. Run the tail -50 logs/management/dramatiq.log command to check for exception logs.
- d) Check the Redis service.

Run the docker logs --tail 50 mgt-redis command to check for exception logs.

- e) Check the detector service.
 - A. Run the docker logs --tail 50 detector-srv command to check for exception logs.
 - B. Run the tail -50 logs/detector/snserver.log command to check for exception logs.
 - C. Run the curl 127.0.0.1:8001/stat | grep num command to check whether the service responds normally and whether the real-time request processing data is normal. For example, check the req_num_to

tal parameter, which indicates the number of requests that have been processed within the last five seconds.

- f) Check the tengine service.
 - A. Run the docker logs --tail 50 tengine command to check for exception logs.
 - B. Run the tail -50 logs/nginx/error.log command to check for exception logs.
- g) Check the mario service.
 - A. Run the docker logs --tail 50 mario command to check for exception logs.
 - B. Run the tail -50 logs/mario/mario.log command to check for exception logs.
 - C. Run the curl 127.0.0.1:3335/api/v1/state command to check whether the service responds normally and whether the real-time request processing data is normal. For example, check whether the num_pendin g parameter remains at a high value of nearly 10,000, or whether the num_processed_last_10s parameter, which indicates the number of requests that have been processed within the last 10 seconds, is normal.

9.7 Routine operations and maintenance of Sensitive Data Discovery and Protection

9.7.1 Check the service status

9.7.1.1 Basic inspection

During the basic inspection of Sensitive Data Discovery and Protection (SDDP), 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.
- 3. In the Fuzzy Search field, enter yundun-sddp.
- 4. Click Details in the Actions column of the yundun-sddp project to go to the Cluster Operations page.

- 5. In the cluster list, click the cluster name that starts with SddpCluster.
- 6. In the Service Instances section of the Cluster Dashboard page, check whether the yundun-sddp service instance is in the final status.

9.7.1.2 Advanced inspection: Check the status of the SddpService service

This topic describes how to check the running status of the SddpService service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.

- 2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
 - a) Choose Operations > Project Operations.
 - b) In the Fuzzy Search field, enter yundun-sddp. Click Details in the Actions column of the yundun-sddp project to go to the Cluster Operations page.
 - c) In the cluster list, click the cluster name that starts with SddpCluster.
 - d) In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.

Service Instances C 2							
Service Instance	Final Status	Expected Server Roles	Server Roles In Final	Server Roles Going O	Actions		
hids-client	True	1	1	0	Actions - Details		
os	True				Actions - Details		
tianji	True	1	1	0	Actions - Details		
tianji-dockerdaemon	True	1	1	0	Actions - Details		
yundun-sddp	True	9	9	0	Actions Details		

e) In the Server Role List section, find SddpService# and click Details in the Actions column to go to the Server Role Dashboard page.

Server Role List	Server Role List								
Server Role	Current Status	Expected Machi	Machines In Fin	Machines Goin	Rolling Task St	Time Used	Actions		
SddpAlgorithm#	In Final Status	1	1	0	no rolling		Details		
SddpData#	In Final Status	2	2	0	no rolling		Details		
SddpDatamask#	In Final Status	2	2	0	no rolling		Details		
SddpDbInit#	In Final Status	1	1	0	no rolling		Details		
SddpLog#	In Final Status	2	2	0	no rolling		Details		
SddpPrivilege#	In Final Status	2	2	0	no rolling		Details		
SddpRuleEngine#	In Final Status	2	2	0	no rolling		Details		
SddpService#	In Final Status	2	2	0	no rolling		Details		
ServiceTest#	In Final Status	1	1	0	no rolling		Details		

f) In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.

Machine	Machine Information								
Machi	IP	Mac	Mac	Serv	Serv	Curr	Targ	Error	Actions
a56g101	10	good		good P		2fb869ef	2fb869ef		Terminal Restart Details Machine System View Machine Operation
a56h1116	10	good		good P		2fb869ef	2fb869ef		Terminal Restart Details Machine System View Machine Operation

3. Log on to two Docker containers of the SddpService service, respectively.

Run the sudo docker exec -it \$(sudo docker ps | grep SddpService | awk '{print \$1}') bash command.

4. Check the process status of the SddpService service.

Run the ps aux | grep java | grep yundun-sddp-service **command. If any** record is returned, the service is normal.

#ps aux grep java grep yundun-sddp-service
root 162 0.1 30.7 7224188 2579604 ? Sl May31 26:35 /opt/taobao/java/bin/java -Dspring.profiles.ag
ve=cloud -server -Xms4g -Xmx4g -Xmn2g -XX:MetaspaceSize=256m -XX:MaxMetaspaceSize=512m -XX:MaxDirectMemorySize=
-XX:SurvivorRatio=10 -XX:+UseConcMarkSweepGC -XX:CMSMaxAbortablePrecleanTime=5000 -XX:+CMSClassUnloadingEnabled
X:CMSInitiatingOccupancyFraction=80 -XX:+UseCMSInitiatingOccupancyOnly -XX:+ExplicitGCInvokesConcurrent -Dsun.r
dgc.server.gcInterval=2592000000 -Dsun.rmi.dgc.client.gcInterval=2592000000 -XX:ParallelGCThreads=4 -Xloggc:/roc
logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/root/lo
/java.hprof -Djava.awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun.net.client.defaultReadT
out=30000 -DJM.LOG.PATH=/root/logs -DJM.SNAPSHOT.PATH=/root/snapshots -Dfile.encoding=UTF-8 -Dhsf.publish.delay6
true -Dproject.name=yundun-sddp-service -Dpandora.boot.wait=true -Dlog4j.defaultInitOverride=true -Dserver.port
01 -Dmanagement.port=7002 -Dmanagement.server.port=7002 -Dpandora.location=/home/admin/yundun-sddp-service/targ
taobao-hsf.sar -classpath /home/admin/yundun-sddp-service/target/yundun-sddp-service -Dapp.location=/home/admin
ndun-sddp-service/target/yundun-sddp-service -Djava.endorsed.dirs= -Djava.io.tmpdir=/home/admin/yundun-sddp-ser
e/.default/temp com.taobao.pandora.boot.loader.SarLauncher

5. Check the health status.

Run the curl 127.0.0.1:7001/checkpreload.htm command. If the response is success, the service is normal.



- 6. View related logs.
 - View common logs in the /home/admin/yundun-sddp-service/logs/common-log.log file.
 - View application logs in the /home/admin/yundun-sddp-service/logs/ application.log file.
 - View front-end request logs in the /home/admin/yundun-sddp-service/logs/ common-request.log file.
 - View system logs in the /home/admin/yundun-sddp-service/logs/servicestdout.log file.

9.7.1.3 Advanced inspection: Check the status of the SddpData service

This topic describes how to check the running status of the SddpData service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.

- 2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
 - a) Choose Operations > Project Operations.
 - b) In the Fuzzy Search field, enter yundun-sddp. Click Details in the Actions column of the yundun-sddp project to go to the Cluster Operations page.
 - c) In the cluster list, click the cluster name that starts with SddpCluster.
 - d) In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.
 - e) In the Server Role List section, find SddpData# and click Details in the Actions column to go to the Server Role Dashboard page.
 - f) In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.
- 3. Log on to two Docker containers of the SddpData service, respectively.

Run the sudo docker exec -it \$(sudo docker ps | grep SddpData | awk '{ print \$1}') bash command.

4. Check the process status of the SddpData service.

Run the ps aux | grep yundun-sddp-data command. If any record is returned, the service is normal.

5. View related logs.

View logs in the /home/admin/yundun-sddp-data/logs/sddp.log file.

9.7.1.4 Advanced inspection: Check the status of the SddpPrivilege service

This topic describes how to check the running status of the SddpPrivilege service.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.

- 2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
 - a) Choose Operations > Project Operations.
 - b) In the Fuzzy Search field, enter yundun-sddp. Click Details in the Actions column of the yundun-sddp project to go to the Cluster Operations page.
 - c) In the cluster list, click the cluster name that starts with SddpCluster.
 - d) In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.
 - e) In the Server Role List section, find SddpPrivilege# and click Details in the Actions column to go to the Server Role Dashboard page.
 - f) In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.
- 3. Log on to two Docker containers of the SddpPrivilege service, respectively. Run the sudo docker exec -it \$(sudo docker ps | grep SddpPrivilege | awk '{print \$1}') bash command.
- 4. Check the process status of the SddpPrivilege service.

Run the ps aux | grep java | grep yundun-sddp-privilege command. If any record is returned, the service is normal.

5. Check the health status.

Run the curl 127.0.0.1:7001/checkpreload.htm command. If the response is success, the service is normal.

- 6. View related logs.
 - View exception logs in the /home/admin/yundun-sddp-privilege/logs/ exception.log file.
 - View application logs in the /home/admin/yundun-sddp-privilege/logs/ application.log file.
 - View task logs in the /home/admin/yundun-sddp-privilege/logs/task.log
 file.
 - View system logs in the /home/admin/yundun-sddp-privilege/logs/servicestdout.log file.

9.7.1.5 Advanced inspection: Check the status of the SddpLog service

This topic describes how to check the running status of the SddpLog service.

Procedure

- **1.** Log on to the Apsara Infrastructure Management Framework console.
- 2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
 - a) Choose Operations > Project Operations.
 - b) In the Fuzzy Search field, enter yundun-sddp. Click Details in the Actions column of the yundun-sddp project to go to the Cluster Operations page.
 - c) In the cluster list, click the cluster name that starts with SddpCluster.
 - d) In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.
 - e) In the Server Role List section, find SddpLog# and click Details in the Actions column to go to the Server Role Dashboard page.
 - f) In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.
- 3. Log on to two Docker containers of the SddpLog service, respectively.

Run the sudo docker exec -it \$(sudo docker ps | grep SddpLog | awk '{ print \$1}') bash **command.**

4. Check the process status of the SddpLog service.

Run the ps aux | grep java | grep yundun-sddp-log. If any record is returned, the service is normal.

5. Check the health status.

Run the curl 127.0.0.1:7001/checkpreload.htm command. If the response is success, the service is normal.

- 6. View related logs.
 - View exception logs in the /home/admin/yundun-sddp-log/logs/exception.
 log file.
 - View application logs in the /home/admin/yundun-sddp-log/logs/applicatio
 n.log file.
 - View debug logs in the /home/admin/yundun-sddp-log/logs/debug.log file.
 - View system logs in the /home/admin/yundun-sddp-log/logs/service-stdout
 . log file.

9.7.2 Restart SDDP

This topic describes how to restart Sensitive Data Discovery and Protection (SDDP) when a fault occurs.

Procedure

- 1. Run the ssh Server IP address command to log on to the server that hosts SDDP.
- 2. Run the following command to find the image ID of the service:

docker ps -a |grep service name

3. Run the following command to log on to the Docker container:

docker exec -it [imageId] /bin/bash

4. Restart related services.

- Restart the yundun-sddp-service service.
 - a. Run the following command to stop the current process:

kill -9 \$(ps -ef | grep java | grep yundun-sddp-service | grep -v
grep | awk '{print\$2}')

b. Run the following command to restart the process:

/bin/bash /home/admin/start.sh

c. Run the following command to check whether the process is restarted:

curl 127.0.0.1:7001/check.htm

If the response is success, the service is normal.

- Restart the yundun-sddp-log service.
 - a. Run the following command to stop the current process:

kill -9 \$(ps -ef | grep java | grep yundun-sddp-log | grep -v grep | awk '{print \$2}')

b. Run the following command to restart the process:

/bin/bash /home/admin/start.sh

c. Run the following command to check whether the process is restarted:

curl 127.0.0.1:7001/check.htm

- If the response is success, the service is normal.
- Restart the yundun-sddp-privilege service.
 - a. Run the following command to stop the current process:

kill -9 \$(ps -ef | grep java | grep yundun-sddp-privilege | grep -v grep | awk '{print \$2}')

b. Run the following command to restart the process:

/bin/bash /home/admin/start.sh

c. Run the following command to check whether the process is restarted:

curl 127.0.0.1:7001/check.htm

If the response is success, the service is normal.

· Restart the yundun-sddp-data service.

a. Run the following command to stop the current process:

```
kill -9 $(ps -ef | grep yundun-sddp-data | grep -v grep | awk '{
print $2}')
```

b. Run the following command to restart the process:

/bin/bash /home/admin/yundun-sddp-data/start.sh

c. Check whether the process is restarted.

Run the ps aux | grep yundun-sddp-data command. If any record is returned, the service is normal.

9.8 Routine operations and maintenance of Apsara Stack Security Center

9.8.1 Check service status

9.8.1.1 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.

9.8.1.2 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.

9.8.2 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

9.9 Routine operations and maintenance of secure-service

9.9.1 Check the service status

9.9.1.1 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 yundunadvance, 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.

9.9.1.2 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/secureservice/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

9.9.1.3 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/commonerror.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.

9.9.1.4 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.

9.9.2 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

10 Apsara Stack DNS

10.1 Introduction to Apsara Stack DNS

This topic describes Apsara Stack DNS and the features of its modules.

Database management system

The database management system compares the versions in the baseline configurat ion with those in the database to better manage databases. This allows you to validate the database version in each update.

API system

The API system determines the business logic of all calls and manages all data and tasks. This system is written in Java.

DNS

The DNS system consists of BIND and Agent. Agent receives and processes task information passed from the API system. Agent parses the tasks into commands, and then delivers the commands to the BIND system.

10.2 Maintenance

10.2.1 View operational logs

During operations and maintenance, you can query and view logs that are stored at specific locations in different systems to troubleshoot errors.

The operational logs of the API service are stored in the /home/admin/gdns/logs/ directory. You can query logs as needed.

The operational logs of the Agent service are stored in the /var/log/dns/ directory of the DNS server. Each log contains log entries of a specific day.

The operational logs of the BIND service are stored in the /var/named/chroot/var/ log/ **directory of the DNS server.**

10.2.2 Enable and disable a service

You can log on to the API server as an administrator and run the /home/admin/gdns /bin/appctl.sh restart command to restart the API service. We recommend that you run the command on one server at a time to ensure that another server can provide services. You can specify the start, stop, and restart parameters in the preceding command.

Apsara Stack DNS provides services by using anycast IP addresses. You must run the service ospfd stop command to disable the OSPF service before you run the service named stop command to disable the DNS service.

You must run the service named start command to enable the DNS service before you run the service ospfd start command to enable the OSPF service.

You can run the /usr/local/AgentService/agent -s start command to enable the Agent service. If you receive a message that indicates the PID file already exists, delete the /var/dns/dns.pid file and run the command again.

You can run the /usr/local/AgentService/agent -s stop command to disable the Agent service.

10.2.3 Data backup

If you need to back up data before updating the service, copy the /var/named/ and /etc/named/ directories to a backup location. When you need to restore your data, copy the backup data to the original directories. Do not trigger automatic update during a data restoration process. Otherwise, data inconsistency may occur.

10.3 DNS API

10.3.1 Manage the API system

You can manage the API system in the Apsara Infrastructure Management Framework console. To log on to the server in which the API system resides, choose Operations > Server Operations in the Apsara Infrastructure Management Framework console.

Context

To determine whether a service role is running as expected, follow these steps:

Procedure

- 1. In the Apsara Infrastructure Management Framework console, check whether the API is at desired state.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Tasks > Deployment Summary to open the Deployment Summary page.
 - c) Click Deployment Details.
 - d) On the Deployment Details page, find the dnsProduct project.
 - e) Find the dnsServerRole# service role, and click Details in the Deployment Progress column to check whether the service role is at desired state.

If a green check mark is displayed after dnsServerRole#, then dnsServerRole# is at desired state.

Figure 10-1: View API status

dnsProduct	Final 4 Days 19 Hours	Cluster: 2 / 2 Service: 9 / 9 Role: 12 / 12 Details
drds	Final 4 Days 7 Hours	C 🚓 dnsCluster-A-20 ⊘ 🗠 dnsService 💿 🛧 ServiceTes# 📀
dts	Final 3 Days 23 Hours	A standardCluster⊘ ≪ hids-client ⊘ + bindServerRole# ⊘
ecs	Final 1 Hour 24 Minutes	C «tianji (C * discontinuation)
edas	Final 4 Days 21 Hours	c tianji-dockerdae ⊘
elasticsearch	Final 11 Hours 57 Minutes	c
emr	Final 4 Days 21 Hours	c
ess	Final 3 Days 22 Hours	q

- 2. Obtain the IP addresses of servers where the API services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Operations > Cluster Operations.
 - c) Click a cluster URL to open the Cluster Dashboard page.
 - d) On the Cluster Dashboard page, choose Operations Menu > Cluster Operation and Maintenance Center.

Figure 10-2: Cluster Operation and Maintenance Center

Cluster Dashboard	Operations Menu 👻	
	Change Machine	
Basic Cluster Information	Deploy Service	2.2
Title	Upgrade Service	
Project Name	Upgrade Service (Simple Mode)	
	Service Authorization	
Cluster Name	Offline Service	
IDC	Configuration Files	
Final Status Version		25d5
Cluster in Final Status	Cluster Operation and Maintenance Center	
	Service Final Status Query	
Machines Not In Final Status	Cluster Configuration	
Real/Pseudo Clone	Operation Logs	
Expected Machines		

e) On the Cluster Operation and Maintenance Center page, view and obtain the IP addresses of servers that are deployed with the API service.

Cluster Operations > Cluster Operation and Maintenance Center Cluster Operation and Maintenance Center (Cluster: <u>dnsCluster-A-20190827-4eb3</u>)								
SR not in Final Status: N/A Running Tasks: No rolling is available. Head Version Submitted At: 09/26/19, 10:29:12 Head Version Analysis 0 : done								
Service dnsSe	ervice	Server Role dnsServerRole#	out: 0 Abnom	nal Pir	ng Failed: 0 Abnormat	TJ-Client: 0		
Machines	2 Expected Mac: 2	Scale-out SR: in: 0	out: 0 Machin	es 0 No Sta	atus Error: 0	Other SRs: 0		
🖵 Machines								
Machine Search	Supports multiple machine search.	Q, Res	et					
Machine	0 Final SR Status 🗹	OSR Running Statu	IS 🖬 Action 🖻	Action Status 🗹	Monitoring Statistics	Actions		
wm0100120120	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 7	Terminal Approval Action Restart Server Role		
m0100120160	048 Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 7	Terminal Approval Action Restart Server Role		
Batch Termina	al				Items per F	Page 10 V « < 1 > »		

Figure 10-3: View the IP addresses of servers

- 3. Log on to the DNS API server. Run the curl http://localhost/checkpreload. htm command, and check whether the command output is "success".
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Operations > Server Operations.
 - c) Click Terminal in the Actions column of a server to log on to the server.
 - d) Run the curl http://localhost/checkpreload.htm command on the server where the API service is deployed and check whether the command output is "success".

Figure 10-4: Verify the server



10.3.2 Troubleshooting

Procedure

1. View logs stored in /home/admin/gdns/logs/.
- 2. Check whether the API service is running. If an error occurs when you call an API operation, check the log to troubleshoot the error.
- 3. If the API service is running, but its features do not function as expected, check the application.log file.

10.4 DNS system

10.4.1 Check whether a server role is normal

Procedure

- 1. In the Apsara Infrastructure Management Framework console, check whether the Apsara Stack DNS system is in its final state.
 - 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, click Deployment Details.
 - d) On the Deployment Details page, find dnsProduct.
 - e) Click Details in the Deployment Progress column to check whether the bindServerRole# role is in its final state.

Figure 10-5: Checking whether the bindServerRole# server role is in its final state

dnsProduct	Final 4 Days 19 Hours	Cluster: 2/2 Service: 9	/ 9 Role: 12 / 12 Det	ails
drds	Final 4 Days 7 Hours	C 🚠 dnsCluster-A-20 ⊘	≪ dnsService ⊘	♣ ServiceTest#
dts	Final 3 Days 23 Hours	at standardCluster ⊘ C	∝ hids-client ⊘	♣ bindServerRole#
			~ os 📀	♣ dnsServerRole#
ecs	Final 1 Hour 24 Minutes	c	🕫 tianji 📀	-✿ dnsServiceDbInit# ⊙
edas	Final 4 Days 21 Hours	с	vs tianji-dockerdae ⊘	-∿ monitorSrDemo# ⊘
elasticsearch	Final 11 Hours 57 Minutes	с		
emr	Final 4 Days 21 Hours	с		
ess	Final 3 Days 22 Hours	c		

- 2. Obtain the IP addresses of the servers where DNS services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Operations > Cluster Operations.
 - c) Click a cluster URL to go to the Cluster Dashboard page.
 - d) On the Cluster Dashboard page, choose Operations Menu > Cluster Operation and Maintenance Center.

Figure 10-6: Cluster Operation and Maintenance Center

Cluster Dashboard	Operations Menu 👻	
	Change Machine	
Basic Cluster Information	Deploy Service	2.2
Title	Upgrade Service	
Draiget Name	Upgrade Service (Simple Mode)	
	Service Authorization	_
Cluster Name	Offline Service	
IDC	Configuration Files	
Final Status Version		25d5
Cluster in Final Status	Cluster Operation and Maintenance Center	- 1
	Service Final Status Query	
Machines Not In Final Status	Cluster Configuration	
Real/Pseudo Clone	Operation Logs	
Expected Machines		

e) On the Cluster Operation and Maintenance Center page, view and obtain IP addresses of all the servers that are assigned with the bindServerRole# role.

- 3. Log on to the DNS server, run the python /bind/hello/check_health.py|echo
 - **\$? command, and check whether the command output is 0.**
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - **b**) Choose Operations > Machine Operations.
 - c) Select a server and click Terminal to log on to the server.
 - d) Run the python /bind/hello/check_health.py|echo \$? command on each server that is assigned with the bindServerRole# role and check whether the command output is 0.

Figure 10-7: Verifying the server



10.4.2 Troubleshooting

Procedure

- **1.** Check the operational logs of the BIND service that are stored in the /var/named/ chroot/var/log/ directory, and determine whether errors have occurred.
- 2. Check the operational logs of the Agent service that are stored in the /var/log/ dns/directory, and determine whether errors have occurred.
- 3. Run the named-checkconf command to check whether errors have occurred in the configuration file.

10.4.3 Errors and exceptions

Error: exit code 1

Run the health check script to view the cause of this error.

Common causes include:

- The DNS service is not running.
- The Agent service is not running.
- The OSPF service is not running, or anycast and public IP addresses cannot be advertised because of a network information retrieval error.
- Failed to run the task.

10.5 Log analysis

Query log entries by request ID

After you send a request, you will receive a response that contains the request ID. The request ID can be used in the following scenarios:

- 1. Query the tasks that are associated with the current request from the database.
- 2. Retrieve the execution results and error messages of the current request from the API system log.
- 3. Retrieve the results of the current request from the log of bindServerRole#, and verify the results with information that is retrieved from multiple other systems.

10.6 View and process data

Context

You can view task records and execution results.

Procedure

- 1. Log on to the API server to view database connection details.
- 2. Run the use genesisdns command of MySQL to log on to the database and then run the select * from task command to retrieve the progress and status of each task.