Alibaba Cloud Apsara Stack Enterprise Operations and Maintenance Guide

Product Version: 2001, Internal: V3.11.0 Document Version: 20200918

C-J Alibaba Cloud

Legal disclaimer

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

- 1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud", "Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Anger Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.
O Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
၂) Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.
⑦ Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid Instance_ID
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

Table of Contents

1.0&M overview	38
2.Preparations before operations	40
2.1. Prepare an operations account	40
2.2. Log on to Apsara Stack Operations	40
2.3. Web page introduction	41
3.System settings	43
3.1. Default operations roles	43
3.2. ITIL Management	43
3.2.1. Overview	43
3.2.2. Dashboard	44
3.2.3. Services	44
3.2.3.1. Basic functions	44
3.2.3.1.1. Overview	44
3.2.3.1.2. Manage requests	45
3.2.3.1.3. Manage tasks	46
3.2.3.2. Manage incidents	46
3.2.3.2.1. Create an incident request	46
3.2.3.2.2. Manage incident requests	48
3.2.3.2.3. Manage incident tasks	49
3.2.3.3. Manage problems	50
3.2.3.3.1. Create a problem request	50
3.2.3.3.2. Manage problem requests	52
3.2.3.3.3. Manage problem tasks	53
3.2.4. Version control	54
3.2.5. Configure process templates	55
3.2.6. Configure CAB or ECAB	57

3.3. Configurations	58
3.3.1. Overview	58
3.3.2. Modify a configuration item of a product	58
3.3.3. Restore the configuration value of a modified config	59
3.3.4. Manage kernel configurations	59
3.3.5. Scan configurations	59
3.4. System Management	60
3.4.1. Overview	60
3.4.2. Department management	60
3.4.3. Role management	61
3.4.4. Logon policy management	62
3.4.5. User management	62
3.4.6. Two factor authentication	64
3.4.7. Application whitelist	67
3.4.8. Server password management	67
3.4.9. Operation logs	69
3.4.10. View the authorization information	69
3.4.11. Multi-cloud management	72
3.4.11.1. Add multi-cloud configurations	72
3.4.11.2. Modify the multi-cloud name	73
3.4.12. Menu settings	73
3.4.12.1. Add a level-1 menu	73
3.4.12.2. Add a submenu	74
3.4.12.3. Hide a menu	77
3.4.12.4. Modify a menu	77
3.4.12.5. Delete a menu	77
4.Monitoring	79
4.1. Daily monitoring	79

4.1.1. Operations and maintenance dashboard	79
4.1.2. Alert Monitoring	79
4.1.2.1. Dashboard	79
4.1.2.2. Alert events	82
4.1.2.3. Alert history	84
4.1.2.4. Alert configuration	84
4.1.2.4.1. Alert contacts	85
4.1.2.4.2. Alert contact groups	85
4.1.2.4.3. Static parameter settings	86
4.1.2.5. Alert overview	87
4.1.2.6. Alert subscription and push	87
4.1.2.7. Alert masking	89
4.1.2.7.1. Add a masking rule	89
4.1.2.7.2. Remove the masking	92
4.1.3. Physical servers	92
4.1.3.1. View the physical server information	92
4.1.3.2. Add a physical server	96
4.1.3.3. Modify a physical server	97
4.1.3.4. Export the physical server information	98
4.1.3.5. Delete a physical server	99
4.1.4. Inventory Management	100
4.1.4.1. View the ECS inventory	100
4.1.4.2. View the SLB inventory	105
4.1.4.3. View the RDS inventory	
	106
4.1.4.4. View the OSS inventory	
	106
4.1.4.4. View the OSS inventory	106 107

4.1.4.8. View the NAS inventory	109
4.1.4.9. View the HDFS inventory	110
4.1.5. Full Stack Monitor	111
4.1.5.1. SLA	111
4.1.5.1.1. View the current state of a cloud product	111
4.1.5.1.2. View the history data of a cloud product	111
4.1.5.1.3. View the availability of an instance	112
4.1.5.1.4. View the availability of a product	112
4.1.5.2. Operations full link logs	113
4.1.5.3. Correlation diagnosis and alarm	113
4.1.5.3.1. Full stack correlation alert	113
4.1.5.3.2. Server	114
4.1.5.3.3. Network device	115
4.1.5.3.4. ECS	116
4.1.5.3.5. RDS	117
4.1.5.3.6. SLB	118
4.1.5.3.7. VPC	122
4.1.5.4. Use case	123
4.1.5.4.1. Monitor and diagnose the VPC XGW host reso	123
4.1.6. Storage Operation Center	123
4.1.6.1. Pangu	123
4.1.6.1.1. Pangu grail	124
4.1.6.1.2. Cluster information	125
4.1.6.1.3. Node information	127
4.1.6.1.4. Pangu operation	128
4.1.6.1.5. Product settings	129
4.1.6.2. EBS	131
4.1.6.2.1. IO HANG fault analysis	131

4.1.6.2.2. Slow IO analysis	132
4.1.6.2.3. Inventory settings	134
5.Operations tools	
5.1. Offline Backup	136
5.1.1. Service configuration	136
5.1.1.1. Configure the backup server	136
5.1.1.2. Add a backup product	137
5.1.2. Backup service	137
5.1.2.1. Backup configuration	137
5.1.2.2. View the backup details	138
5.1.2.3. View the backup server status	139
5.1.3. View the backup status	139
5.1.4. Use cases	140
5.1.4.1. Offline backup of metadata	140
5.1.4.1.1. Preparations before the backup	140
5.1.4.1.2. Collect pangu information of each product	141
5.1.4.1.3. Configure the backup server	142
5.1.4.1.4. Add a backup product	143
5.1.4.1.5. Configure the backup	144
5.1.4.1.6. View the backup details	145
5.2. NOC	145
5.2.1. Overview	145
5.2.2. Dashboard	145
5.2.2.1. View the dashboard	145
5.2.2.2. View the network topology	147
5.2.2.3. Manage custom views	148
5.2.3. Network Service Provider	152
5.2.3.1. View access gateway instances	152

5.2.3.2. View operation logs	153
5.2.3.3. View network information of bare metals in VPC	154
5.2.3.4. O&M configurations	155
5.2.3.4.1. Apply for a bare metal in VPC	155
5.2.3.4.2. Release a bare metal in VPC	157
5.2.3.4.3. Delete a VPC route table entry	158
5.2.3.4.4. Delete a VBR route table entry	160
5.2.3.4.5. Delete a VPC router interface	161
5.2.3.4.6. Delete a VBR router interface	162
5.2.3.4.7. Delete a VBR	163
5.2.3.4.8. Delete a physical connection	164
5.2.3.4.9. Delete all resources with one click	166
5.2.3.4.10. View physical connection bandwidth	169
5.2.3.4.11. Modify the physical connection bandwidth	169
5.2.4. Resource management	171
5.2.4.1. Network elements	171
5.2.4.1.1. Device management	171
5.2.4.1.1.1. View the network monitoring information	171
5.2.4.1.1.2. View logs	173
5.2.4.1.1.3. Collection settings	173
5.2.4.1.2. Modify the device password	173
5.2.4.1.3. Configuration comparison	174
5.2.4.2. Fire wall	174
5.2.4.3. Service Load Balancers	176
5.2.4.3.1. View the cluster monitoring information	176
5.2.4.3.2. View the instance monitoring information	177
5.2.4.4. Collect IP addresses	177
5.2.4.5. View Anytunnel information	178

5.2.4.6. XGW management	178
5.2.4.7. IP address ranges	179
5.2.4.7.1. Import the planning file	179
5.2.4.7.2. Manually add the IP address pool information	
5.2.4.7.3. Modify the IP address pool information	180
5.2.4.7.4. Export the IP address pool information	180
5.2.4.7.5. Delete the IP address pool information	181
5.2.5. Alert management	181
5.2.5.1. View and process current alerts	181
5.2.5.2. View history alerts	181
5.2.5.3. Add a trap	182
5.2.5.4. View a trap	184
5.2.6. Network reconfiguration	184
5.2.6.1. Physical network integration	184
5.2.6.2. ASW scale-up	186
5.2.6.3. Push IPv6 configurations	187
5.2.7. Fault check	189
5.2.7.1. IP address conflict check	189
5.2.7.2. Leased line discovery	189
5.2.7.3. Network inspection	191
5.2.7.4. Configuration baseline audit	192
5.2.8. Use case	192
5.2.8.1. Troubleshoot network failures	192
5.3. Task Management	196
5.3.1. Overview	196
5.3.2. View the task overview	196
5.3.3. Create a task	197
5.3.4. View the execution status of a task	200

5.3.5. Start a task	201
5.3.6. Delete a task	202
5.3.7. Process tasks to be intervened	202
5.4. Log Management	203
5.4.1. Log configurations	203
5.4.1.1. Clear	203
5.4.1.1.1. Configure parameters for automatic log cleara	203
5.4.1.1.2. Configure the manual log clearance time	204
5.4.1.2. Project	204
5.4.1.2.1. Add a project	204
5.4.1.2.2. Delete a project	205
5.4.1.3. Agent	205
5.4.1.3.1. Add an agent	205
5.4.1.3.2. Modify an agent	207
5.4.1.3.3. Delete an agent	207
5.4.1.4. Bucket management	208
5.4.1.4.1. OSS configurations	208
5.4.1.4.2. NAS configurations	208
5.4.1.4.3. FTP configurations	209
5.4.2. Display logs	209
5.4.3. Log export	210
5.4.3.1. Export logs	210
5.4.3.2. View tasks	211
5.4.4. Log clearance	211
5.4.4.1. Containers	211
5.4.4.1.1. Obtain the watermark information of one or	211
5.4.4.1.2. Add a log clearance rule	212
5.4.4.1.3. Modify a log clearance rule	213

5.4.4.1.4. Delete a log clearance rule	213
5.4.4.1.5. Clear container logs	213
5.4.4.1.6. View clear records	214
5.4.4.2. Servers	214
5.4.4.2.1. Obtain the watermark information of one or	214
5.4.4.2.2. Add a log clearance rule	215
5.4.4.2.3. Modify a log clearance rule	215
5.4.4.2.4. Delete a log clearance rule	216
5.4.4.2.5. Clear server logs	216
5.4.4.2.6. View clear records	216
5.4.4.3. Import clearance rules of containers or servers	217
5.4.4.4. Export clearance rules of containers or servers	217
5.5. Apsara Stack Doctor (ASD)	217
5.5.1. Apsara Stack Doctor introduction	217
5.5.2. Log on to Apsara Stack Doctor	219
5.5.3. ASA	220
5.5.3.1. RPM Check	220
5.5.3.2. Virtual IP Check	221
5.5.3.3. Volume Check	222
5.5.3.4. NTP Check	223
5.5.3.5. IP Conflict Check	224
5.5.3.6. DNS Check	225
5.5.3.7. IP Details	225
5.5.3.8. Quota Check	226
5.5.3.9. Error Diagnostics	227
5.5.3.10. Versions	227
5.5.4. Support tools	227
5.5.4.1. Diagnose with the OS tool	227

5.5.4.2. Use Support Tools	228
5.5.4.3. Update Support Tools	230
5.5.4.4. Diagnose with inspection tools	230
5.5.4.5. Upload script files for EDAS diagnostics	231
5.5.4.6. EDAS diagnostics	232
5.5.5. Service Availability	233
5.5.5.1. View Service Availability	233
5.5.5.2. View Control Service Availability	233
5.5.6. Monitoring	235
5.5.6.1. View alert templates	235
5.5.6.2. View alert information	235
5.5.6.3. View the alert status	236
5.6. Apsara Infrastructure Management Framework	236
5.6.1. Old version	237
5.6.1.1. What is Apsara Infrastructure Management Frame	237
5.6.1.1.1. Overview	237
5.6.1.1.2. Basic concepts	237
5.6.1.2. Log on to Apsara Infrastructure Management Fra	239
5.6.1.3. Web page introduction	240
5.6.1.3.1. Introduction on the home page	240
5.6.1.3.2. Introduction on the left-side navigation pane	242
5.6.1.4. Cluster operations	244
5.6.1.4.1. View cluster configurations	244
5.6.1.4.2. View the cluster dashboard	246
5.6.1.4.3. View the cluster operation and maintenance	250
5.6.1.4.4. View the service final status	252
5.6.1.4.5. View operation logs	254
5.6.1.5. Service operations	254

5.6.1.5.1. View the service list	255
5.6.1.5.2. View the service instance dashboard	255
5.6.1.5.3. View the server role dashboard	257
5.6.1.6. Machine operations	<mark>260</mark>
5.6.1.6.1. View the machine dashboard	26 0
5.6.1.7. Monitoring center	262
5.6.1.7.1. Modify an alert rule	262
5.6.1.7.2. View the status of a monitoring instance	263
5.6.1.7.3. View the alert status	263
5.6.1.7.4. View alert rules	264
5.6.1.7.5. View the alert history	265
5.6.1.8. Tasks and deployment summary	266
5.6.1.8.1. View rolling tasks	266
5.6.1.8.2. View running tasks	267
5.6.1.8.3. View history tasks	268
5.6.1.8.4. View the deployment summary	268
5.6.1.9. Reports	270
5.6.1.9.1. View reports	270
5.6.1.9.2. Add a report to favorites	271
5.6.1.10. Appendix	272
5.6.1.10.1. Project component info report	272
5.6.1.10.2. IP list	272
5.6.1.10.3. Machine info report	273
5.6.1.10.4. Rolling info report	275
5.6.1.10.5. Machine RMA approval pending list	276
5.6.1.10.6. Registration vars of services	278
5.6.1.10.7. Virtual machine mappings	278
5.6.1.10.8. Service inspector report	278

5.6.1.10.9. Resource application report	279
5.6.1.10.10. Statuses of project components	280
5.6.1.10.11. Relationship of service dependency	282
5.6.1.10.12. Check report of network topology	283
5.6.1.10.13. Clone report of machines	283
5.6.1.10.14. Auto healing/install approval pending report	284
5.6.1.10.15. Machine power on or off statuses of clusters	284
5.6.2. New version	286
5.6.2.1. What is Apsara Infrastructure Management Fram	286
5.6.2.1.1. Introduction	286
5.6.2.1.2. Basic concepts	287
5.6.2.2. Log on to Apsara Infrastructure Management Fra	288
5.6.2.3. Homepage introduction	289
5.6.2.4. Project operations	292
5.6.2.5. Cluster operations	292
5.6.2.5.1. View the cluster list	292
5.6.2.5.2. View the cluster details	294
5.6.2.5.3. View operation logs	297
5.6.2.6. Service operations	298
5.6.2.6.1. View the service list	298
5.6.2.6.2. View the server role details	299
5.6.2.7. Machine operations	300
5.6.2.8. Monitoring center	302
5.6.2.8.1. View the monitoring instance status	302
5.6.2.8.2. View the alert status	302
5.6.2.8.3. View alert rules	303
5.6.2.8.4. View the alert history	304
5.6.2.9. View tasks	305

5.6.2.10. Reports	306
5.6.2.10.1. View reports	306
5.6.2.10.2. Add a report to favorites	307
5.6.2.11. Tools	308
5.6.2.11.1. Machine tools	308
5.6.2.11.2. IDC shutdown	309
5.6.2.12. Appendix	309
5.6.2.12.1. Project component info report	309
5.6.2.12.2. IP list	310
5.6.2.12.3. Machine info report	311
5.6.2.12.4. Rolling info report	312
5.6.2.12.5. Machine RMA approval pending list	314
5.6.2.12.6. Registration vars of services	315
5.6.2.12.7. Virtual machine mappings	316
5.6.2.12.8. Service inspector report	316
5.6.2.12.9. Resource application report	316
5.6.2.12.10. Statuses of project components	318
5.6.2.12.11. Relationship of service dependency	320
5.6.2.12.12. Check report of network topology	320
5.6.2.12.13. Clone report of machines	321
5.6.2.12.14. Auto healing/install approval pending repo	322
5.6.2.12.15. Machine power on or off statuses of cluste	322
6.Products	324
6.1. Product list	324
6.2. ISV access configurations	324
6.2.1. Configure the ISV access information	324
6.2.2. Modify the ISV access information	325
6.2.3. Delete the ISV access information	326

7.Network operations	327
7.1. Apsara Network Intelligence	327
7.1.1. What is Apsara Network Intelligence?	327
7.1.2. Log on to the Apsara Network Intelligence console	327
7.1.3. Query information	328
7.1.4. Manage cloud service instances	329
7.1.5. Tunnel VIP	330
7.1.5.1. Create a Layer-4 listener VIP	330
7.1.5.2. Query the tunnel VIP of a cloud service	330
7.1.6. Create a Direct Any Tunnel VIP	331
7.1.7. Leased line connection	331
7.1.7.1. Overview	331
7.1.7.2. Manage an access point	332
7.1.7.3. Manage an access device	333
7.1.7.4. Establish a leased line connection	334
7.1.7.5. Create a VBR	336
7.1.7.6. Create router interfaces	338
7.1.7.7. Create a routing table	339
7.1.8. Manage Business Foundation System flows in a VPC	341
7.1.9. Configure reverse access to cloud services	341
8.Operations of basic cloud products	343
8.1. Elastic Compute Service (ECS)	343
8.1.1. ECS overview	343
8.1.2. Log on to the Apsara Stack Operations console	343
8.1.3. ECS operations and maintenance	344
8.1.3.1. Overview	344
8.1.3.2. VM	344
8.1.3.2.1. Overview	344

8.1.3.2.2. Search for VMs	345
8.1.3.2.3. Start a VM	345
8.1.3.2.4. Stop a VM	346
8.1.3.2.5. Restart a VM	346
8.1.3.2.6. Cold migration	347
8.1.3.2.7. Reset a disk	348
8.1.3.3. Disks	348
8.1.3.3.1. Overview	348
8.1.3.3.2. Search for disks	349
8.1.3.3.3. View snapshots	349
8.1.3.3.4. Mount a disk	349
8.1.3.3.5. Detach a disk	350
8.1.3.3.6. Create a snapshot	350
8.1.3.4. Snapshots	351
8.1.3.4.1. Overview	351
8.1.3.4.2. Search for snapshots	351
8.1.3.4.3. Delete a snapshot	352
8.1.3.4.4. Create an image	352
8.1.3.5. Images	353
8.1.3.5.1. Overview	353
8.1.3.5.2. Search for images	353
8.1.3.6. Security groups	353
8.1.3.6.1. Overview	353
8.1.3.6.2. Search for security groups	354
8.1.3.6.3. Add security group rules	354
8.1.4. VM hot migration	355
8.1.4.1. Overview	355
8.1.4.2. Limits on hot migration	356

8.1.4.3. Complete hot migration on AG	356
8.1.4.4. Modify the position of the NC where the VM is lo	358
8.1.4.5. FAQ	358
8.1.5. Hot migration of disks	360
8.1.5.1. Overview	360
8.1.5.2. Limits	360
8.1.5.3. O&M after hot migration	361
8.1.6. Upgrade solution	
8.1.6.1. Overview	361
8.1.6.2. Limits on GPU clusters	361
8.1.6.3. Limits on FPGA clusters	362
8.1.7. Disk maintenance of an instance	362
8.1.7.1. Overview	362
8.1.7.2. Maintenance procedure	363
8.1.7.3. Additional instructions	370
8.1.8. Handle routine alarms	371
8.1.8.1. Overview	371
8.1.8.2. API proxy	372
8.1.8.3. API Server	373
8.1.8.4. RegionMaster	373
8.1.8.5. RMS	374
8.1.8.6. PYNC	375
8.1.8.7. Zookeeper	375
8.1.8.8. AG	376
8.1.8.9. Server groups	377
8.1.9. Inspection	377
8.1.9.1. Overview	377
8.1.9.2. Cluster basic health inspection	377

8.1.9.2.1. Overview	
8.1.9.2.2. Monitoring inspection	
8.1.9.2.3. Inspection of basic software package versions	
8.1.9.2.4. Basic public resources inspection	377
8.1.9.3. Cluster resource inspection	378
8.1.9.3.1. Overview	378
8.1.9.3.2. Cluster inventory inspection	378
8.1.9.3.3. VM inspection	380
8.2. Container Service	381
8.2.1. Components and features	381
8.2.1.1. Console	381
8.2.1.2. Troopers	381
8.2.1.3. Mirana	382
8.2.2. System restart	383
8.2.2.1. Restart a control node	383
8.3. Auto Scaling (ESS)	383
8.3.1. Log on to the Apsara Stack Operations console	383
8.3.2. Product resources and services	384
8.3.2.1. Application deployment	384
8.3.2.2. Troubleshooting	385
8.3.3. Inspection	386
8.3.3.1. Overview	386
8.3.3.2. Monitoring inspection	386
8.3.3.3. Basic software package version inspection	386
8.4. Resource Orchestration Service (ROS)	386
8.4.1. ROS component O&M	386
8.4.1.1. API Server	
8.4.1.2. Engine Server	387

8.4.1.3. RabbitMQ clusters	387
8.4.1.4. Notify Server	388
8.5. Object Storage Service (OSS)	389
8.5.1. Log on to the Apsara Stack Operations console	389
8.5.2. OSS operations and maintenance	390
8.5.2.1. User data	390
8.5.2.1.1. Basic bucket information	390
8.5.2.1.2. User data overview	390
8.5.2.1.3. Data monitoring	
8.5.2.2. Cluster data	393
8.5.2.2.1. Inventory monitoring	393
8.5.2.2.2. Bucket statistics	394
8.5.2.2.3. Object statistics	394
8.5.2.2.4. Data monitoring	395
8.5.2.2.5. Resource usage rankings	397
8.5.3. Tools and commands	398
8.5.3.1. Typical commands supported by tsar	398
8.5.3.2. Configure tsar for statistic collection	399
8.6. Table Store	399
8.6.1. Table Store Operations and Maintenance System	399
8.6.1.1. Overview	399
8.6.1.2. User data	399
8.6.1.2.1. Instance management	399
8.6.1.3. Cluster management	402
8.6.1.3.1. Cluster information	402
8.6.1.4. Inspection center	405
8.6.1.4.1. Abnormal resource usage	405
8.6.1.5. Monitoring center	406

8.6.1.5.1. Cluster monitoring	406
8.6.1.5.2. Application monitoring	407
8.6.1.5.3. Top requests	408
8.6.1.5.4. Request log search	408
8.6.1.6. System management	409
8.6.1.6.1. Manage tasks	409
8.6.1.6.2. View tasks	410
8.6.1.7. Platform audit	411
8.6.1.7.1. Operation logs	411
8.6.2. Cluster environments	412
8.6.3. System roles	412
8.6.4. Pre-partition a table	413
8.6.4.1. Pre-partitioning	413
8.6.4.2. View partitions	414
8.7. ApsaraDB for RDS	414
8.7.1. Architecture	414
8.7.1.1. System architecture	415
8.7.1.1.1. Backup system	415
8.7.1.1.2. Data migration system	415
8.7.1.1.3. Monitoring system	415
8.7.1.1.4. Control system	416
8.7.1.1.5. Task scheduling system	416
8.7.2. Log on to the Apsara Stack Operations console	416
8.7.3. Instance management	417
8.7.4. Manage hosts	410
	419
8.7.5. Security maintenance	
8.7.5. Security maintenance	420

8.8. AnalyticDB for PostgreSQL	420
8.8.1. Overview	420
8.8.2. Architecture	421
8.8.3. Routine maintenance	423
8.8.3.1. Check for data skew on a regular basis	423
8.8.3.2. Execute VACUUM and ANALYZE statements	424
8.8.4. Security maintenance	424
8.8.4.1. Network security maintenance	424
8.8.4.2. Account password maintenance	425
8.9. KVStore for Redis	425
8.9.1. O&M tools	425
8.9.2. Architecture diagram	425
8.9.3. Architecture	425
8.9.3.1. Architecture	425
8.9.3.1.1. Backup system	425
8.9.3.1.2. Data migration system	426
8.9.3.1.3. Monitoring system	426
8.9.3.1.4. Control system	426
8.9.3.1.5. Task scheduling system	426
8.9.4. Log on to the Apsara Stack Operations console	427
8.9.5. Instance management	428
8.9.6. Host management	428
8.9.7. Security maintenance	429
8.9.7.1. Network security maintenance	429
8.9.7.2. Password maintenance	429
8.10. ApsaraDB for MongoDB	429
8.10.1. Service architecture	429
8.10.1.1. System architecture	429

8.10.1.1.1. Backup system	429
8.10.1.1.2. Data migration system	430
8.10.1.1.3. Monitoring system	430
8.10.1.1.4. Control system	431
8.10.1.1.5. Task scheduling system	431
8.10.2. ApsaraDB for MongoDB O&M overview	431
8.10.3. Log on to the Apsara Stack Operations console	431
8.10.4. Manage ApsaraDB for MongoDB instances	432
8.10.5. Host management	434
8.10.6. Security maintenance	434
8.10.6.1. Network security maintenance	434
8.10.6.2. Account password maintenance	435
8.11. Log Service	435
8.11.1. O&M methods	435
8.11.2. O&M	438
8.11.2.1. View logs on machines	438
8.11.2.2. Use Log Service Portal to view logs	444
8.12. Apsara Stack Security	446
8.12.1. Log on to the Apsara Infrastructure Management F	446
8.12.2. Routine operations and maintenance of Server Gua	447
8.12.2.1. Check the service status	447
8.12.2.1.1. Check the client status	447
8.12.2.1.2. Check the status of Aegiserver	447
8.12.2.1.3. Check the Server Guard Update Service status	449
8.12.2.1.4. Check the Defender module status	449
8.12.2.2. Restart Server Guard	450
8.12.3. Routine operations and maintenance of Network Tr	451
8.12.3.1. Check the service status	451

8.12.3.1.1. Basic inspection	451
8.12.3.1.2. Advanced inspection	452
8.12.3.2. Common operations and maintenance	453
8.12.3.2.1. Restart the Network Traffic Monitoring System	453
8.12.3.2.2. Uninstall Network Traffic Monitoring System	453
8.12.3.2.3. Disable TCP blocking	453
8.12.3.2.4. Enable TCPDump	454
8.12.4. Routine operations and maintenance of Anti-DDoS	454
8.12.4.1. Check the service status	454
8.12.4.1.1. Basic inspection	454
8.12.4.1.2. Advanced inspection	454
8.12.4.2. Common operations and maintenance	456
8.12.4.2.1. Restart Anti-DDoS Service	456
8.12.4.2.2. Troubleshoot common faults	457
8.12.5. Routine operations and maintenance of Threat Det	460
8.12.5.1. Check the service status	460
8.12.5.1.1. Basic inspection	460
8.12.5.1.2. Advanced inspection	460
8.12.5.2. Restart TDS	461
8.12.5.2. Restart TDS	
8.12.5.2. Restart TDS	
	462
8.12.6. Routine operations and maintenance of WAF	462 462
8.12.6. Routine operations and maintenance of WAF	462 462 462
8.12.6. Routine operations and maintenance of WAF 8.12.6.1. Check the service status 8.12.6.1.1. Basic inspection	462 462 462 462
 8.12.6. Routine operations and maintenance of WAF 8.12.6.1. Check the service status 8.12.6.1.1. Basic inspection 8.12.6.1.2. Advanced inspection 	462 462 462 462 464
 8.12.6. Routine operations and maintenance of WAF 8.12.6.1. Check the service status 8.12.6.1.1. Basic inspection 8.12.6.1.2. Advanced inspection 8.12.7. Routine operations and maintenance of Sensitive Da 	462 462 462 462 464 464
 8.12.6. Routine operations and maintenance of WAF 8.12.6.1. Check the service status 8.12.6.1.1. Basic inspection 8.12.6.1.2. Advanced inspection 8.12.7. Routine operations and maintenance of Sensitive Da 8.12.7.1. Check the service status 	462 462 462 462 464 464

8.12.7.1.4. Advanced inspection: Check the status of the	467
8.12.7.1.5. Advanced inspection: Check the status of the	468
8.12.7.2. Restart SDDP	468
8.12.8. Routine operations and maintenance of Apsara Sta	470
8.12.8.1. Check service status	470
8.12.8.1.1. Basic inspection	470
8.12.8.1.2. Advanced inspection	470
8.12.8.2. Restart the secure-console service	471
8.12.9. Routine operations and maintenance of secure-serv	471
8.12.9.1. Check the service status	471
8.12.9.1.1. Basic inspection	471
8.12.9.1.2. Advanced inspection: Check the secure-servic	472
8.12.9.1.3. Check the Dolphin service status	472
8.12.9.1.4. Check the data-sync service status	473
8.12.9.2. Restart secure-service	474
8.13. Apsara Stack DNS	
	475
8.13. Apsara Stack DNS	475 475
8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS	475 475 476
8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance	475 475 476 476
 8.13. Apsara Stack DNS	475 475 476 476 476
 8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance 8.13.2.1. View operational logs 8.13.2.2. Enable and disable a service 	475 475 476 476 476 476
 8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance 8.13.2.1. View operational logs 8.13.2.2. Enable and disable a service 8.13.2.3. Data backup 	475 475 476 476 476 476 476
 8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance 8.13.2.1. View operational logs 8.13.2.2. Enable and disable a service 8.13.2.3. Data backup 8.13.3. DNS API 	475 475 476 476 476 476 476
 8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance 8.13.2.1. View operational logs 8.13.2.2. Enable and disable a service 8.13.2.3. Data backup 8.13.3. DNS API 8.13.3.1. Manage the API system 	475 475 476 476 476 476 476 476
 8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance 8.13.2.1. View operational logs 8.13.2.2. Enable and disable a service 8.13.2.3. Data backup 8.13.3. DNS API 8.13.3.1. Manage the API system 8.13.3.2. Troubleshooting 	475 475 476 476 476 476 476 476 479
 8.13. Apsara Stack DNS 8.13.1. Introduction to Apsara Stack DNS 8.13.2. Maintenance 8.13.2.1. View operational logs 8.13.2.2. Enable and disable a service 8.13.2.3. Data backup 8.13.3. DNS API 8.13.3.1. Manage the API system 8.13.3.2. Troubleshooting 8.13.4. DNS system 	475 476 476 476 476 476 476 479 479

8.13.5. Log analysis	481
8.13.6. View and process data	481
9.Operations of middleware products	482
9.1. Enterprise Distributed Application Service (EDAS)	482
9.1.1. O&M overview	482
9.1.1.1. Architecture	482
9.1.1.2. O&M architecture	484
9.1.2. Overview of critical operations	485
9.1.3. Maintenance preparation	486
9.1.4. Routine maintenance	486
9.1.4.1. Inspection	487
9.1.4.1.1. Component inspection	488
9.1.4.1.1.1. Manual inspection	488
9.1.4.2. Monitoring	488
9.1.4.2.1. Monitoring logs	489
9.1.5. Troubleshooting	490
9.1.5.1. Alarm handling	490
9.1.5.1.1. CPU utilization alerts	490
9.1.5.1.2. Memory usage alarms	491
9.1.5.1.3. Disk usage alarms	491
9.1.5.2. Service continuity exceptions	492
9.1.5.2.1. EDAS monitoring exceptions	492
9.1.5.2.2. Excessive node logs	493
9.1.5.2.3. Console access failure	493
9.1.5.2.4. ECS instance import failure	494
9.1.5.2.5. TLog data collection errors	494
9.1.6. Log reference	496
9.1.6.1. EDAS console logs	497

9.1.6.2. EDAS admin logs	498
9.1.6.3. EDAS server logs	499
9.1.6.4. DiamondServer logs	500
9.1.6.5. Cai-fs logs	501
9.1.6.6. ConfigServer logs	501
9.1.6.7. Cai-address logs	502
9.1.6.8. EagleEye console logs	503
9.1.7. Configuration reference	503
9.1.7.1. Component configuration	503
9.1.7.2. JVM configuration	506
10.Operations of big data products	509
10.1. MaxCompute	
10.1.1. Concepts and architecture	509
10.1.2. O&M commands and tools	511
10.1.2.1. Before you start	512
10.1.2.2. odpscmd commands	512
10.1.2.3. Tunnel commands	514
10.1.2.4. LogView tool	520
10.1.2.4.1. Before you start	520
10.1.2.4.2. LogView introduction	520
10.1.2.4.3. Preliminary knowledge of LogView	521
10.1.2.4.4. Basic operations and examples	524
10.1.2.4.5. Best practices	526
10.1.2.5. Apsara Bigdata Manager	527
10.1.3. Routine O&M	527
10.1.3.1. Configurations	527
10.1.3.2. Routine inspections	528
10.1.3.3. Shut down a chunkserver, perform maintenance	532

10.1.3.4. Shut down a chunkserver for maintenance with	537
10.1.3.5. Adjust the virtual resources of the Apsara syste	539
10.1.3.6. Restart MaxCompute services	541
10.1.4. Common issues and solutions	542
10.1.4.1. View and allocate MaxCompute cluster resources	542
10.1.4.2. Common issues and data skew troubleshooting	551
10.2. DataWorks	559
10.2.1. Basic concepts and structure	559
10.2.1.1. What is DataWorks (base)?	559
10.2.1.2. Functions of base	560
10.2.1.3. Introduction to data analytics	560
10.2.1.4. DataWorks architecture in Apsara Stack V3	561
10.2.1.5. Directory of each service	563
10.2.2. Common administration tools and commands	563
10.2.2.1. Find the container that runs the service	563
10.2.2.2. Cluster resource list	564
10.2.2.3. Commands to restart services	564
10.2.2.4. View logs of a failed node	564
10.2.2.5. Rerun a task	564
10.2.2.6. Terminate a task	564
10.2.2.7. Filter tasks in the administration center	565
10.2.2.8. Commonly used Linux commands	565
10.2.2.9. View the slots usage of each resource group	566
10.2.3. Process daily administration operations	566
10.2.3.1. Daily check	566
10.2.3.1.1. Check the service status and the basic inform	566
10.2.3.1.2. Check the postgres database	567
10.2.3.1.3. Check the status of each gateway server	567

10.2.3.1.4. Check the case test report	567
10.2.3.2. View logs of the services	568
10.2.3.3. Scale out the node cluster that runs the base-b	568
10.2.3.4. Scale in the base-biz-gateway cluster	571
10.2.3.5. Restart the base-biz-alisa service	574
10.2.3.6. Restart the base-biz-phoenix service	574
10.2.3.7. Restart base-biz-tenant	575
10.2.3.8. Restart base-biz-gateway	575
10.2.3.9. Restart the base-biz-api service	576
10.2.3.10. Restart the base-redis service	576
10.2.3.11. Restart DataWorks Data Service	576
10.2.3.12. Restart DataWorks Data Management	577
10.2.4. Common issues and solutions	577
10.2.4.1. Nodes remain in the Pending (Resources) state	577
10.2.4.2. An out-of-memory (OOM) error occurs when syn	580
10.2.4.3. A task does not run at the specified time	581
10.2.4.4. The test service of base is not in the desired st	581
10.2.4.5. The Data Management page does not display t	582
10.2.4.6. Logs are not automatically cleaned up	582
10.2.4.7. The real-time analysis service is not in the desir	583
10.3. Realtime Compute	583
10.3.1. Job status	583
10.3.1.1. Overview	583
10.3.1.2. Task status	583
10.3.1.3. Health score	583
10.3.1.4. Job instantaneous values	583
10.3.1.5. Running topology	584
10.3.2. Curve charts	587

10.3.2.1. Overview	587
10.3.2.2. Overview	587
10.3.2.3. Advanced view	589
10.3.2.4. Processing delay	591
10.3.2.5. Throughput	591
10.3.2.6. Queue	591
10.3.2.7. Tracing	591
10.3.2.8. Process	592
10.3.2.9. JVM	592
10.3.3. FailOver	592
10.3.4. CheckPoints	593
10.3.5. JobManager	593
10.3.6. TaskExecutor	594
10.3.7. Data lineage	594
10.3.8. Properties and Parameters	594
10.3.9. Improve performance by automatic configuration	595
10.3.10. Improve performance by manual configuration	600
10.3.10.1. Overview	601
10.3.10.2. Optimize resource configuration	601
10.3.10.3. Improve performance based on job parameter	603
10.3.10.4. Optimize upstream and downstream data stora	603
10.3.10.5. Apply new configuration	604
10.3.10.6. Concepts	604
10.4. Quick BI	605
10.4.1. Introduction to O&M and tools	605
10.4.1.1. Introduction to operations and maintenance	605
10.4.1.2. Troubleshoot Quick BI issues by using the Apsar	605
10.4.2. Routine maintenance	608

10.4.2.1. Introduction to Quick BI components	608
10.4.2.2. Database initialization components	609
10.4.2.3. Cache components	609
10.4.2.4. Runtime components	610
10.4.2.5. Web service components	610
10.4.2.6. Automated testing components	611
10.5. Graph Analytics	611
10.5.1. Operations and maintenance tools and logon meth	611
10.5.1.1. Log on to Apsara Bigdata Manager	611
10.5.1.2. Log on to Apsara Infrastructure Management Fr	613
10.5.1.3. Log on to the Graph Analytics container	615
10.5.2. Operations and maintenance	616
10.5.2.1. Operations and maintenance based on BigData	616
10.5.2.1.1. View and handle cluster alerts	616
10.5.2.1.2. View cluster performance metrics	620
10.5.2.1.3. View server operation metrics	621
10.5.2.2. Operations and maintenance based on Apsara I	622
10.5.2.3. Operations and maintenance based on the Grap	624
10.5.2.3.1. View instances	624
10.5.2.3.2. Log files	625
10.5.2.3.3. Database logs	625
10.5.2.3.4. Stop the service	626
10.5.2.3.5. Restart the service	626
10.5.3. Security maintenance	627
10.5.3.1. Network security maintenance	627
10.5.3.2. Account password maintenance	627
10.5.4. Troubleshooting	627
10.5.4.1. Fault response mechanism	627

10.5.4.2. Troubleshooting methods	627
10.5.4.3. Common failure troubleshooting	627
10.5.4.4. Hardware troubleshooting	628
10.6. Machine Learning Platform for AI	628
10.6.1. Query server and application information	628
10.6.1.1. Apsara Stack Machine Learning Platform for AI	628
10.6.1.1.1. Query server information	628
10.6.1.1.2. Log on to a server	629
10.6.1.1.3. Query configurations	629
10.6.1.1.4. Restart an application service	630
10.6.1.2. Online model service	630
10.6.1.2.1. Query online model service information	630
10.6.1.2.2. Log on to the online model service container	631
10.6.1.2.3. Restart a pod	631
10.6.1.3. GPU cluster and task information	631
10.6.1.3.1. Query GPU cluster information	631
10.6.1.3.2. Query GPU task information	632
10.6.2. Maintenance and troubleshooting	632
10.6.2.1. Machine Learning Platform for AI maintenance	632
10.6.2.1.1. Run ServiceTest	632
10.6.2.1.2. Common faults and solutions	633
10.6.2.2. Online model service maintenance (must be act	633
10.6.2.3. GPU cluster maintenance (deep learning must b	634
10.7. E-MapReduce (EMR)	635
10.7.1. Methods for logging on to O&M platforms	635
10.7.1.1. Log on to Apsara Infrastructure Management Fra	635
10.7.2. Routine maintenance	637
10.7.2.1. O&M on Apsara Infrastructure Management Fram	637

10.7.3. Troubleshooting	638
10.7.3.1. Troubleshooting methods	638
11.Apsara Asapi Management system	639
11.1. Apsara Asapi Management system overview	639
11.2. Log on to the Apsara Asapi Management platform	640
11.3. Manage APIs	641
11.3.1. Register APIs	641
11.3.2. Modify information about APIs	642
11.3.3. Test APIs	643
11.3.4. Remove information about APIs	644
11.3.5. API design	645
11.3.5.1. Designers	645
11.3.5.2. Designer nodes	645
11.3.5.3. Design an API flow	646
11.4. Version management	646
11.4.1. Apsara Stack version management	646
11.4.1.1. Add information about versions	646
11.4.1.2. Select products for an Apsara Stack version	647
11.4.1.3. Compare versions	648
11.4.1.4. Remove information about Apsara Stack versions	650
11.4.2. Product baseline	650
11.4.3. Products	650
11.4.3.1. Add information about products	651
11.4.3.2. Add information about product versions	651
11.4.3.3. Import information about APIs	651
11.4.3.4. Set SDK versions	652
11.4.3.5. Modify product names and descriptions	652
11.4.3.6. View information about product versions	653

11.4.3.7. Modify information about product versions	653
11.4.3.8. Remove information about product versions	653
11.4.3.9. Delete products	654
11.4.3.10. Delete product APIs	654
11.4.4. SDKs	654
11.4.4.1. Customize SDKs	654
11.4.4.2. Modify SDKs	655
11.4.4.3. Delete SDKs	656
11.5. Test management	656
11.5.1. Test cases	656
11.5.1.1. Modify test cases	656
11.5.1.2. Run test cases	657
11.5.1.3. Delete test cases	658
11.5.2. Test sets	658
11.5.2.1. Create test sets	658
11.5.2.2. Associate test cases	658
11.5.2.3. Run test sets	659
11.5.2.4. Delete test sets	659
11.5.3. View execution history of test cases	659
11.6. System management	660
11.6.1. Metadatabase management	660
11.6.1.1. View information about added metadatabases	660
11.6.1.2. View connection information about metadatabas	661
11.6.1.3. Remove information about metadatabases	661
11.6.2. Server management	662
11.6.2.1. View information about added servers	662
11.6.2.2. Remove server information	663
11.6.3. Audit APIs	663

11.6.4. View logs	664
12.Appendix	665
12.1. Operation Access Manager (OAM)	665
12.1.1. OAM introduction	665
12.1.2. Instructions	665
12.1.3. Quick start	666
12.1.3.1. Log on to OAM	666
12.1.3.2. Create a group	668
12.1.3.3. Add group members	668
12.1.3.4. Add group roles	668
12.1.3.5. Create a role	669
12.1.3.6. Add inherited roles to a role	669
12.1.3.7. Add resources to a role	670
12.1.3.8. Add authorized users to a role	670
12.1.4. Manage groups	671
12.1.4.1. Modify the group information	671
12.1.4.2. View group role details	672
12.1.4.3. Delete a group	672
12.1.4.4. View authorized groups	672
12.1.5. Manage roles	673
12.1.5.1. Search for roles	673
12.1.5.2. Modify the role information	673
12.1.5.3. View the role inheritance tree	673
12.1.5.4. Transfer roles	674
12.1.5.5. Delete a role	674
12.1.5.6. View authorized roles	674
12.1.5.7. View all roles	675
12.1.6. Search for resources	675

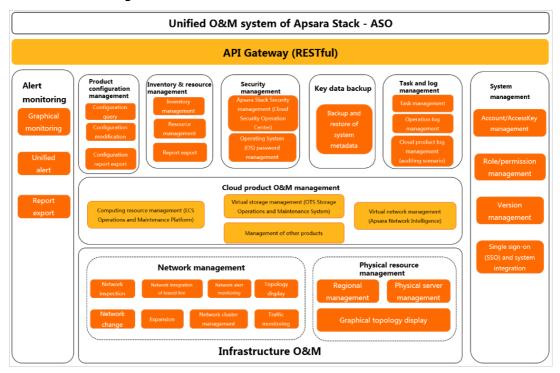
12.1.7. View the personal information	675
12.1.8. Appendix	676
12.1.8.1. Default roles and their functions	676
12.1.8.1.1. Default role of OAM	676
12.1.8.1.2. Default roles of Apsara Infrastructure Manag	676
12.1.8.1.3. Default roles of Webapp-rule	678
12.1.8.1.4. Default roles of the workflow console	679
12.1.8.1.5. Default role of Tianjimon	679
12.1.8.2. Operation permissions on O&M platforms	680
12.1.8.2.1. Permissions on Apsara Infrastructure Manage	680
12.1.8.2.2. Permission list of Webapp-rule	<mark>690</mark>
12.1.8.2.3. Permission list of the workflow console	69 0
12.1.8.2.4. Permissions on Monitoring System of Apsara	691

1.0&M overview

This topic describes the Alibaba Cloud Apsara Stack Operations & Maintenance (O&M) service management system.

According to the Information Technology Infrastructure Library (ITIL) and IT Service Management (ITSM) standards, the O&M process and requirements must be abstract and automation is implemented by using intelligent O&M tools for common high-frequency operations. For customized operations, interfaces and multi-level approval must be used to reduce the risks of O&M operations.

Alibaba Cloud Apsara Stack adopts the ISO20000 series standards and refers to the methods regulated by the Information Technology Service Standards (ITSS) and ITIL frameworks to form the Alibaba Cloud Apsara Stack O&M service management framework. The Alibaba Cloud Apsara Stack O&M service management framework is as shown in O&M service management framework.



O&M service management framework

Apsara Stack Operations (ASO) is a unified intelligent O&M platform of Alibaba Cloud Apsara Stack. Based on the understanding of cloud O&M from Alibaba Cloud Apsara Stack, cloud O&M is classified into three layers: infrastructure O&M, cloud product/service O&M, and business O&M. The O&M service management framework contains the full lifecycle management methods, management standards/regulations, management modes, management supporting tools, management objects, and process-based management methods of IT O&M services.

ASO provides you with a unified O&M portal. You can use the O&M portal to have a consistent O&M experience and a unified O&M entrance. ASO supports connecting with third-party platforms and provides unified API O&M capabilities to provide O&M system data in the form of API for third-party systems.

ASO performs unified 0&M management, such as automated deployment, upgrade, change, and configuration, on physical devices, operating systems, computing resources, network, storage, databases, middleware, and business applications in the cloud computing environment. ASO also provides the functions of alert monitoring and automatic analysis, diagnosis, and troubleshooting for faults, performance, and configurations. By analyzing, processing, and evaluating the running status and quality of cloud platforms, ASO guarantees the continuous and stable running of cloud computing business applications and provides services and support for 0&M service processes to build an improved 0&M service management platform.

Based on ITIL/ISO20000, with process-oriented, normalized, and standardized management as the method, adaption to various management modes as the aim, the O&M service management framework uses management supporting tools to implement the systematic management of the overall process of O&M services.

With the accumulated O&M experience and data collection of the three-layer system, Alibaba Cloud Apsara Stack aggregates data collected by O&M platforms to the Configuration Management Database (CMDB) of the platform. ASO, the intelligent O&M platform, consolidates, analyzes, and comprehensively processes the data and solidifies rich practical experience and O&M capabilities to the platform O&M tools. With the design concept of facing to the final status, ASO uses the unified O&M tools for the fault discovery, fault tracking, link display, ITIL process, and self-repaired faults of the platform to realize the ultimate goal of AlOps.

Process assurance and personnel management, except the tools, are essential to the O&M integrity. Alibaba Cloud Apsara Stack provides development on-site supporting services for major problems, on-site services, expert escort services, business consulting services, and business optimization services. By using the first-line, second-line, and third-line supporting systems to support platform problems of customers and providing the upgrade channel to support urgent problems of customers, Alibaba Cloud Apsara Stack uses ASO, the completely autonomous and controllable platform, to make sure that technical problems can be effectively solved in time.

The Alibaba Cloud Apsara Stack O&M service management system stipulates various entities involved with O&M activities and relationships between these entities. Related entities are organically organized and coordinated according to the O&M service management system, and required to provide different levels of O&M services according to the service agreements.

2. Preparations before operations

2.1. Prepare an operations account

Before performing Operations & Maintenance (O&M) operations in Apsara Stack Operations (ASO), make sure that you obtain an operations account with corresponding permissions from the administrator.

Create an operations account and grant permissions to the account as follows:

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

⑦ Note To divide permissions of the operations role at a finer granularity, the administrator can create a basic role according to Appendix > Operation Access Manager (OAM), grant permissions to the role, and then grant the role to the corresponding operations account.

2.2. 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-domain-id.com.
- Google Chrome browser (recommended).

Procedure

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

Log On	
<u>8</u>	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.

? Note Obtain the username and password used to log on to ASO from the deployment personnel or the administrator.

- 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. 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.



The description of each	area is as follows.
-------------------------	---------------------

Area		Description
1	Multi-cloud selection area	Click the drop-down list to switch the multi-cloud.

Operations and Maintenance Guide • Preparations before operations

Area		Description
2	Region selection area	Click the drop-down list to switch the region, which allows you to centrally manage each region.
3	Authorization information display area	Click to go to the Authorization page and then view the authorization conditions of services.
4	Help center	In the help center, you can view the alarm knowledge base and upload other documents in the .html format related to operations.
5	Language switching area	Click the drop-down list to switch the language.
6	Information area of the current logon user	Click the drop-down list to view the information of the current user, modify the password, and complete the logo settings and logon settings.
7	Expand button	Move the pointer over this button to expand the left-side navigation pane.
8	Left-side navigation pane	Click to select a specific Operations & Maintenance (O&M) operation.
9	Operating area	The information display area and operating area.

3.System settings

3.1. Default operations roles

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

For quick management, the following roles are preset in ASO: OAM super administrator, system administrator, security officer, auditor officer, and multi-cloud configuration administrator. For more information about these roles and their responsibilities, see the following table.

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

3.2. ITIL Management

3.2.1. 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:

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.

3.2.2. 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 the ASO console.
- 2. In the left-side navigation pane, choose ITIL Management > Dashboard.

3.2.3. Services

3.2.3.1. Basic functions

3.2.3.1.1. 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.

3.2.3.1.2. Manage requests

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

Procedure

- 1. Log on to the ASO console.
- 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.

- <u>2</u>: 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.

3.2.3.1.3. 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 **P**.

• 👩: Processed by me, indicating the history tasks that have been processed by you. After you

process the tasks under 👧 , they are displayed under 🧔.

Procedure

- 1. Log on to the ASO console.
- 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.

3.2.3.2. Manage incidents

3.2.3.2.1. 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 the ASO console.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the Request tab.
- 3. Click + New and then select Incident. Configure the incident request on the displayed page.

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.
	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
Priority	• Major
	• Minor
	• Remind
	• Cleared
	° System
Alarm Code	The alert ID.
Summary	The summary of this request.
Description	The detailed description about the request.

Configuration	Description
Suggestion	Optional. The suggestion about the request processing.

4. After completing the configurations, click Confirm.

3.2.3.2.2. 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 the ASO console.
- 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 (2) list. Click Recycle to cancel or logically delete the incident request. The incident request is in the recycle bin (
) list after being recycled.

• Restore the incident request

Perform this operation for incident requests in the recycle bin (
) list. Click Restore to restore the recycled incident request. After being restored, the incident request is in the in processing (
) list and restored to the status before the request is recycled.

• Delete the incident request

Perform this operation for incident requests in the recycle bin (m) list. Click Delete to

delete the incident request. After being deleted, the incident request is physically deleted and cannot be restored.

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

Note To check out the tasks in the task pool to the current username, click the (Task Pool) button and then click Detail at the right of the task. Click Check Out. In

the displayed dialog box, enter the Description and then click OK.

- 4. 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.

3.2.3.3. Manage problems

3.2.3.3.1. 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.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose ITIL Management > Services. Click the Request tab.
- 3. Click + New and then select **Problem**. Configure the problem request on the displayed page.

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 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.

Configuration

Description

4. After completing the configurations, click Confirm.

3.2.3.3.2. 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 the ASO console.
- 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.

⑦ 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 (2) list. Click **Recycle** to cancel or logically delete the problem request. The problem request is in the recycle bin (
1) list after being recycled.

• Restore the problem request

Perform this operation for problem requests in the recycle bin (m) list. Click Restore to restore the recycled problem request. After being restored, the problem request is in the in processing (2) list and restored to the status before the request is recycled.

• Delete the problem request

Perform this operation for problem requests in the recycle bin (m) list. Click Delete to

delete the problem request. After being deleted, the problem request is physically deleted and cannot be restored.

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

Note To check out the tasks in the task pool to the current username, click the
 (Task Pool) button and then click Detail at the right of the task. Click Check Out. In

the displayed dialog box, enter the Description and then click OK.

- 4. 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.

Configuration	Description
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.
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.

3.2.4. 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 the ASO console.
- 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.

ONOTE Before the search, click to synchronize the information to Apsara Stack

Operations (ASO).

3.2.5. 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 the ASO console. 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	Process Template-Incident
Catalogue Incident Problem Change Role Create Identity Reset Password Logout Identity Change Version Upgrade Hotfix Upgrade Configuration Upgrade	
	Regulation

Process

Currently, the following processes are supported:

- Incident
- Problem
- Change Role
- Create Identity
- Reset Password
- Logout Identity
- Change
- Version Upgrade
- Hotfix Upgrade
- Configuration Upgrade

Process template

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

- 🕟 is the start node of the process. A process usually starts with the request creation.
- 🐼 indicates the gateway. The gateway defines the process trend in different branches. In

the BPMN specification, gateways are classified into different types, such as inclusive gateway, exclusive gateway, parallel gateway, and hybrid gateway. Here it is the exclusive gateway, indicating that multiple routes have only one valid path.

- 💼 is the end node of the process. A process usually ends with archiving.
- Resolve indicates the phase. A phase is usually composed of roles with specific functions.
- is the route, indicating the process trend. A phase contains one or more egress routes and ingress routes.

The templates can be classified into the following three types:

• Incidents and problems

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

• Request fulfillment

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

• Change management

Change, Version Upgrade, Hotfix Upgrade, and Configuration Upgrade. The whole process has the following phases: Record, Preliminary Approval, Modify Information, 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.

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

• Assign by 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.
- $\circ~$ 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. Where, CAB is abbreviated from Change Advisory Board and ECAB is abbreviated from Emergency Change Advisory Board.

Click CAB/ECAB Configuration to go to the CAB/ECAB Configuration page. You can configure the CAB or ECAB based on business needs. For more information about how to configure the CAB or ECAB, see Configure CAB or ECAB.

3.2.6. 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 the ASO console.
- 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 to cancel the

configuration for the selected users.

3.3. Configurations

3.3.1. 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.

3.3.2. Modify a configuration item of a product

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

Procedure

- 1. Log on to the ASO console.
- 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.

3.3.3. Restore the configuration value of a modified configuration item

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

Procedure

- 1. Log on to the ASO console.
- 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.

3.3.4. Manage kernel configurations

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

Procedure

- 1. Log on to the ASO console.
- 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**.

3.3.5. 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 the ASO console.
- 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.

3.4. System Management

3.4.1. 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.

3.4.2. 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. Departments are displayed in a hierarchy and you can create sub-departments under each level of departments.

Procedure

- 1. Log on to the ASO console.
- 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.

3.4.3. 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.

Procedure

- 1. Log on to the ASO console.
- 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

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.

3.4.4. 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 the ASO console.
- 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 allowed 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.

3.4.5. 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:

Notice 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 the ASO console.
- 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

? Note 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

? Note 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

(?) Note To modify a user in ASO, you must have the ASO security officer role.

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

• Delete a user

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

(?) Note Deleted users are in the recycle bin. To restore a deleted user, click the Recycled tab. Find the user to be restored, click Cleared in the Actions column, and then click OK in the displayed dialog box.

• Bind a logon policy

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

• View personal information of the current user

In the upper-right corner, click 🔤 next to the logon username and then select Personal

Information. The appeared **Personal Information** dialog box displays the personal information of the current user.

• Add a custom logo

In the upper-right corner, click 🔤 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.

3.4.6. 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 configuring users in ASO, and then enter the key in the Google authenticator app of your mobile phone. The app dynamically generates a verification code based on the time and key for logon.

• USB key authentication

Install the drive and browser controls (currently, only Windows + IE 11 environment is supported) according to the third-party manufacturer instructions if you select this 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 third-party manufacturer, and enable the USB key authentication for the user when you configure the authentication method in ASO.

Upon logon, if the account enables the USB key authentication, the ASO frontend calls the browser controls, reads the certificate in 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 information or certificate provided upon logon. Therefore, add the authentication server configurations if you select these two authentication methods.

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

Procedure

- 1. Log on to the ASO console.
- 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. (Optional)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.

3.4.7. 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 application whitelist function allows you to access ASO in scenarios where no permissions are assigned. With the whitelist function enabled, the application can be accessed by all users who have successfully logged on. The application whitelist permissions 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 the ASO console.
- 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.

3.4.8. 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 the ASO console.
- 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.

• 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**.

3.4.9. 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 the ASO console.
- 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.

3.4.10. 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 the ASO console.
- 2. In the left-side navigation pane, choose **System Management** > **Authorization**. You are on the **Authorization Details** tab by default.

Authorization Details	Authorization Specificat	ion Details Auth	orization Specification Info	rmation			
Basic Information							
Authorization Version :				Authorization Type : 👘 Trial AuthorizationExpired/Quota Exceeded			
Customer ID :			ECS Instance ID :				
Customer User ID :				Cloud Platform Version :			
Customer Name :				Authorization Created At : Nov 2	21, 2019, 15:50:20		
Service Name	Service Content	Authorization Mode	Service Authorizations	Actual Authorizations	Software License Update and Tech Support Started At	Software License Update and Tech Support Expire At	Authorization Status
Virtual Private Cloud (VPC)	VPC Standard	Authorization Mode	1(SET)	1(SET)	Nov 21, 2019, 15:50:20	Jan 13, 2027, 15:50:20	
Container Service (CS)	Expansion Plan for Container Service Basic	Authorization Mode	2(SET)	2(SET)	Nov 21, 2019, 15:50:20	Jun 15, 2032, 15:50:20	
Graph Analytics	Graph Analytics Enterprise	Authorization Mode	1(SET)	1(SET)	Dec 21, 2019, 15:50:20	Mar 20, 2020, 15:50:20	
Enterprise Distributed Application Service (EDAS)	EDAS Pro	Authorization Mode	1(SET)	1(SET)	Apr 4, 2023, 15:50:20	Jul 3, 2023, 15:50:20	
Dataphin	Intelligence Edition	Authorization Mode	1(SET)	1(SET)	Nov 21, 2019, 15:50:20	May 9, 2022, 15:50:20	

3. Complete the following steps to view the authorization information.

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

• On the Authorization Details tab, view the basic information of authorization.

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 of each service, in the current Apsara Stack environment on this page.

For more information about the detailed authorization information and the corresponding description, see the following table.

Authorization information

Description

Authorization information	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 service authorization information 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.			
	Includes the service name, service content, authorization mode, service authorizations, actual authorizations, software license update and tech support start time, software license update and tech support expiration time, and real-time authorization status.			
	If the following information appears in the Authorization Status column of a service:			
Authorization information of	RENEW Service Expired			
a service	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.			

• Click the Authorization Specification Details tab to view the authorization specification information of a service.

For more information about the authorization specification information and the corresponding description, see the following table.

Column name	Description
Service Name	The name of an authorized service.
Specification Name	The specification name of an authorized service.
Specifications	The total number of current authorizations of a specification for a service.
Specification Quota	The authorization quota of a specification for a service.
Specification Status	The current authorization status of a specification for a service.

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

Select a service from the Service Name drop-down list and then select the time range. Click Search to view the current authorization specification information of a service, namely the maximum specifications, minimum specifications, and average specifications in the selected time range, and the time when the maximum specifications and minimum specifications occur.

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

3.4.11. Multi-cloud management

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

3.4.11.1. Add multi-cloud configurations

If a multi-cloud environment is used, the multi-cloud configuration administrator and the super administrator can add multi-cloud configurations. After adding the configurations, you can switch to different cloud data centers on an operations and maintenance platform, and then view or perform related Operations & Maintenance (O&M) operations.

Prerequisites

Before adding multi-cloud configurations, make sure that:

- Networks between clouds are interconnected. Each cloud shares an account, with the same username and password, with other clouds.
- You have the permissions of a multi-cloud configuration administrator or super administrator.

Procedure

1. Log on to the ASO console as a multi-cloud configuration administrator or super

administrator.

- 2. In the left-side navigation pane, choose System Management > Multi-cloud Management.
- 3. Click Add.

Multi-cloud Management		
Name Enter name Add		
Name	Console link	Actions
self	https:/	

4. Configure the console link of another cloud, and then click **OK**.

Configuration	Description
Name	The name of another cloud.
Console link	The console link of another cloud. Enter the correct information. Otherwise, an error message appears, indicating that the addition fails.

After the addition, you can log on to ASO with the account that each cloud shares and switch to different clouds for O&M operations.

3.4.11.2. Modify the multi-cloud name

After adding multi-cloud configurations, the multi-cloud configuration administrator or the super administrator can modify the multi-cloud name in the configurations.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose System Management > Multi-cloud Management.
- 3. (Optional)Enter the multi-cloud name to be modified in the search box and then click Search.
- 4. Find the multi-cloud configurations to be modified and then click **Modify** in the **Actions** column.
- 5. In the displayed dialog box, modify the multi-cloud name and then click OK.

3.4.12. Menu settings

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

3.4.12.1. Add a level-1 menu

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

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose System Management > Menu Settings.
- 3. Click Add.

4. On the displayed page, complete the configurations for the level-1 menu you are about to add.

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

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

Add Level-1 Menu		
Menu Icon		
• Menu Name		
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.

3.4.12.2. Add a submenu

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

Procedure

> Document Version:20200918

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

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

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

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

A	dd Submenu	×
	Menu Name	
	Menu Name(繁體)	
	Menu Name(English)	
	LJ	
	Menu Order	
	1	
	Show Deletable(Yes)	
	Link Address	
	Parent Menu	
	×	

iii. Click OK.

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

4. Click the button at the left of the level-1 menu to expand the level-2 menus. Add a level-3

menu. For more information, see the preceding step.

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

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

3.4.12.3. Hide a menu

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

Prerequisites

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

Procedure

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

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

• Hide a level-2 or level-3 menu

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

3.4.12.4. Modify a menu

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

Procedure

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

3.4.12.5. Delete a menu

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

Prerequisites

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

Procedure

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

4.Monitoring

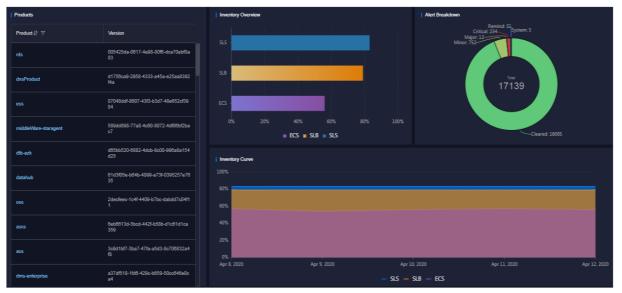
4.1. Daily monitoring

4.1.1. 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 the ASO console. 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.



4.1.2. Alert Monitoring

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

4.1.2.1. Dashboard

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

Context

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

Procedure

1. Log on to the ASO console.

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

The **Dashboard** page appears by default.

Enter the search content.	Search						
Basic							
Recovered	Total	Recovered	Total	Recovered	Total	Recovered	Total
0 ⊘	0 🗘	0⊘	0 🗘	5,450 🥑	5,454 🗘	1,449 📀	1,492 🗘
test	_ ů						
Recovered	Total						
0 🕢	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.

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

• Search for alerts

Enter a keyword, such as cluster, product, service, severity, status, 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 📲. On the displayed page, complete the configurations.

Enter the search content.	Search					Name:	Enter a shortcut name.	
						Conditions:	Service	
Recovered	Total	Recovered	Total	Recovered	Total		Product	
0 📀	0 0	0 0					Severity	
U	υų	U	0 🗘	5,450 ⊘	5,454 🗘		Status	
							Monitoring Metric Type	
							Enter the search content.	
							Start date ~ End date	
Recovered	Total							
0 ⊘	0 🗘							
		© 2009-2019 Alibaba	Cloud Computing Limited. All righ	ts reserved.			ок	Cancel

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

Configuration	Description
Name	The filter name to be displayed on the Dashboard page.
Conditions	 Configure the following filter conditions. Service: The service to which the alerts to be filtered belong. Product: The product to which the alerts to be filtered belong. Severity: The severity to which the alerts to be filtered belong. The alert severity has the following six 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. P3: indicates the minor 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 P3 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. Status: The current status of the alerts to be filtered. Monitoring Metric Type: The monitoring metric type to which the alerts to be filtered belong. Basic Critical Important Minor Enter the search content: Enter the information of the alerts to be filtered. Select the start date and end date of the alerts to be filtered.

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

• Modify a custom filter

After adding a custom filter, you can click 🗾 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 🛅 as required to delete it if it is no longer in

use.

4.1.2.2. 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 the ASO console.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Events.

Hardware & System Base Modules	Monitoring & Mana	gement Cloud Product	Timeout Alert						
By Monitoring Item V Monitoring Metric Type V Produc	t v S	Service \lor Severity \lor	Status v	Start Date	~ End (Date 🛱	Enter the sear	ch content.	Search
<u> </u>									
Monitoring Metric	Monitoring Type	Alert Details	Alerts				P4	P0	P5
postcheck_monitor_tianji_base-template	Event							746	0
testimage_monitortianji_base-template	Event								0
project_monitortianji_sub-template	Event							267	0
ping_monitortianji_base-template	Event								0
tianij_db_upgradetianij_base-template	Event								0

- 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 ____ at the top of the page to export the alert list.

4.1.2.3. 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 the ASO console.
- 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 **I** 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.

• View the original alert information

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

4.1.2.4. Alert configuration

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

4.1.2.4.1. Alert contacts

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

Procedure

- 1. Log on to the ASO console.
- 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.

4.1.2.4.2. Alert contact groups

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

Procedure

- 1. Log on to the ASO console.
- 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**.

4.1.2.4.3. 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 the ASO console.
- 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	×
Parameter Name	
Alarm Time Out	٦
Parameter Code	
ALARM_TIME_OUT	Γ
Parameter Value	
5	
Description	
Alarms that exceed a specified number of days are classified as overdue, Unit: day	

7. Then, click OK.

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

Unsolved Alerts - Last Week (i)					
120					
100					
80					
0					
Jan 5, 2020	Jan 6, 2020		8, 2020 Jan 9, 2020 2 • P3 • P4	Jan 10, 2020	Jan 11, 2020
Product Alerts (j)					
tianji 874	dts 13		rds 3	tlog 3	drds 2
P1 P2 P3 P4	P1 P2 P3 P4	P1 P2 P3 P4	P1 P2 P3 P4	P1 P2 P3 P4	P1 P2 P3 P4
837 13 0 24		3 0 124 0			

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

4.1.2.6. 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.

Procedure

- 1. Log on to the ASO console.
- 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 subscription is located.
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.

Configuration	Description
	 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.
Push Template	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.
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.

4.1.2.7. Alert masking

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

4.1.2.7.1. Add a masking rule

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

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Alert Monitoring > Alert Masking.
- 3. Click Add on the page.
- 4. On the Add page, complete the configurations to mask a certain type of alerts.

Description				
Optional. The product to which alerts to be masked belong.				
Optional. The cluster to which alerts to be masked belong.				
Optional. The name of the service to which alerts to be masked belong.				
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.				
Optional. The monitoring metric to which alerts to be masked belong.				
Optional. The alert details of the alerts to be masked. Example:				
{"serverrole":"ecs-yaochi.ServiceTest#","machine":"vm01001 2016074","level":"error"}				

Configuration	Description
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	\sim
Cluster	
Select	\sim
Service	
Select	
Alert Item	
Monitoring Metric	
Select	\sim
Alert Plan	
Enter data in JSON format.	
Severity	
Select	\sim
OK	cel

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.

4.1.2.7.2. Remove the masking

You can remove the masking for masked alerts.

Procedure

- 1. Log on to the ASO console.
- 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.

Alert Masking							
Product ~	Service	Alert Item		Search Add			
Product	Cluster	Service	Alert Item	Monitoring Metric	Alert Source	Severity	Actions
350							

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.

4.1.3. Physical servers

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

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

Resource Manag	Product	Server	Physical Vie	ew of Device					Physical Servers	: (23) Servers with Alerts: (1	39 Alerts: (512)
Physical Servers		٩	Proc	duct 🗸 Entera v	alue Q						
	+ cn-qingdao-env4	5-d01	Proc	duct	Hostname	Cluster	Group	IP Address	Host Server	Alerts	Operation
			mide	idleWare-staragent	vm010004029063	BasicCluster-A-2019102 8-eaea	StaragentInit		a56h11112.cloud.h12.a mtest72		
			mide	idleWare-staragent	vm010004024188	BasicCluster-A-2019102 8-eaea	Staragent2_1_Controlle r		a58g13108.cloud.g14.a mtest72		
			mide	idleWare-staragent	vm010004021248	BasicCluster-A-2019102 8-eaea	Staragent2Fs		a58g10104.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 B button to view the monitoring graph in full screen.
- Click the L button to download the monitoring graph to your local computer.
- Click the to manually refresh the monitoring data.
- Click the 👩 button 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 the ASO console.
- 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 data center > 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	ical View of Device				Physical Servers: (23) Servers v	ith Alerts: (139) Alerts: (512)
Enter a value	Q	Hostname V Enter a value	Q				+ 🗉
G10		Hostname	Device Function	IP Address	SN	Alerts	Operation
a58g11	0001.cloud.g10.amtest72	a58h11101.cloud.h12.amtest72	Worker				
-	0002.cloud.g10.amtest72 0003.cloud.g10.amtest72	a58h11010.cloud.h11.amtest72	Worker				
-	0004.cloud.g10.amtest72 0005.cloud.g10.amtest72	a56h11012.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 Mount 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 👩 button 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 the ASO console.
- 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 data center > 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 Se	erver Physica			Critical Alerts	Servers	1
Rack 🗸 Enter a	value Q	+ 🗉	- 100%	G11	Normal Servers 0	Moderate Alerts 0 Critical Alerts 1
amtest72					1	
G10					2	
					3	
					4	
			\sim		5	
					6	
					7	
_					8	
-					9	
					10 🔳	
	2				11	
	2				12	
					13	
					14	
					15	
					16	
					17	
					18	

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



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

You can switch the tab to view the monitoring 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 solution to view the monitoring graph in full screen.
- Click the L button to download the monitoring graph to your local computer.
- Click the 💽 button to manually refresh the monitoring data.
- Click the 👩 button 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.

4.1.3.2. 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 the ASO console.
- 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.

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

Configuration	Description
Zone	The name of the zone 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.

Configuration	Description
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.

6. Click OK.

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

- 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 w	ith Alerts: (139) Alerts: (512)
Q	Hostname V Enter a value	Q				+ 🗉
	Hostname	Device Function	IP Address	SN	Alerts	Operation
	a56h11101.cloud.h12.amtest72	Worker				Details Modify Delete
	a56h11010.cloud.h11.amtest72	Worker				Details Modify Delete
	a56h11012.cloud.h11.amtest72	Worker				Details Modify Delete

- 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 the ASO console.
- 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 data center > rack to find the physical server to be modified.

? Note 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 data center.

- 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.

cn-qingdao-env4b-d01 🗸			G10	1	-	Critical	Alerts	
Product Server Phys								
Rack V Enter a value Q	+ 0	- 100%	G1D	No		Details Modify Delete		Alerts 1
G10					•			

- 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.

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



3. In the upper-right corner, click the
button 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 the ASO console.
- 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.

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

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

deleted.

Note 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 data center.

- 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.

self	√ cr	01 🗸		H03		Servers	
Product	Server						
Rack 🗸	Enter a value	Q + B	- 100%	HO3	No	Details Modify Delete	Alerts 0
- amtest88							
			a a a a				

7. Click OK in the displayed dialog box.

4.1.4. Inventory Management

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

4.1.4.1. View the ECS inventory

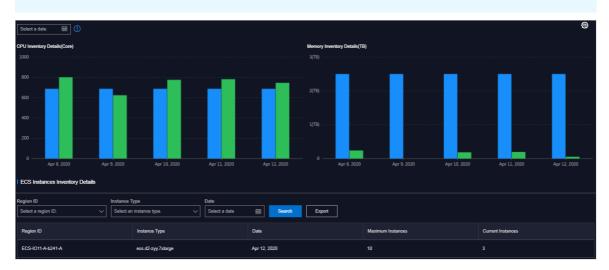
By viewing the Elastic Compute Service (ECS) inventory, you can know the current usage and surplus of ECS product resources to perform Operations & Maintenance (O&M) operations according to actual requirements.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Inventory Management > ECS.

Once 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 Region ID, 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	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	None	1	2.0	1
	None	2	4.0	1
N4 (Shared	None	4	8.0	2
General Type)	None	8	16.0	2
	None	16	32.0	2
	None	32	64.0	2
	None	1	4.0	1
	None	2	8.0	1
MN4 (Shared	None	4	16.0	2
Balanced Type)	None	8	32.0	3
	None	16	64.0	8
	None	32	128.0	8
	None	1	8.0	1
	None	2	16.0	1
E4 (Memory Optimized Type)	None	4	32.0	2
	None	8	64.0	3
	None	16	128.0	8
XN4 (Shared Intense Type)	None	1	1.0	1

Instance type family	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	440	4	30.0	2
	440	8	60.0	3
	880	8	60.0	3
gn5 (Compute Optimized Type	880	16	120.0	8
with GPU)	440	28	112.0	8
	1760	32	240.0	8
	880	56	224.0	8
	3520	56	480.0	8
	4 * 5500	8	32.0	3
	8 * 5500	16	64.0	8
	12 * 5500	24	96.0	8
d1 (Big Data Type)	12 * 5500	32	128.0	8
	16 * 5500	32	128.0	8
	12 * 5500	56	160.0	8
	28 * 5500	56	224.0	8
	None	4	30.0	2
	None	8	60.0	3
gn4 (Compute Optimized Type	None	32	48.0	8
with GPU)	None	8	60.0	3
	None	16	60.0	8
	None	56	96.0	8
	1*87	4	10.0	2
	1*175	8	20.0	3
ga1	1*350	16	40.0	8
(Visualization Compute	1*700	32	80.0	8

Operations and Maintenance Guide · Monitoring

Optimized Type with GPU) Instance type family	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	1*1400	56	160.0	8
	None	2	16.0	1
	None	4	32.0	2
se1ne (Memory Optimized Type with Enhanced	None	8	64.0	3
Network	None	16	128.0	8
Performance)	None	32	256.0	8
	None	56	480.0	8
	None	2	8.0	1
	None	4	16.0	2
sn2ne (General Purpose Type with Enhanced	None	8	32.0	3
Network Performance)	None	16	64.0	8
Performance)	None	32	128.0	8
	None	56	224.0	8
	None	2	4.0	1
sn1ne (Compute	None	4	8.0	2
Optimized Type with Enhanced Network	None	8	16.0	3
Performance)	None	16	32.0	8
	None	32	64.0	8
	None	2	8.0	1
anti (Comente	None	4	16.0	2
gn5i (Compute Optimized Type with GPU)	None	8	32.0	2
with GPU)	None	16	64.0	2
	None	56	224.0	2
	None	2	8.0	2
	None	4	16.0	3

Instance type family	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
g5 (General Purpose Type)	None	8	32.0	4
	None	16	64.0	8
	None	24	96.0	8
	None	32	128.0	8
	None	64	256.0	8
	None	88	352.0	15
	None	2	16.0	2
	None	4	32.0	3
se1 (Memory	None	8	64.0	4
Optimized Type)	None	16	128.0	8
	None	32	256.0	8
	None	56	480.0	8
f3 (Compute	None	16	64.0	8
Optimized Type with FPGA)	None	32	128.0	8
with FGA)	None	64	256.0	16
ebmg5 (General Purpose Type with ECS Bare Metal Instance)	None	96	384.0	32
	1 * 894	4	32.0	3
	1 * 1788	8	64.0	4
i2 (Local SSD Type)	2 * 1788	16	128.0	8
	4 * 1788	32	256.0	8
	8 * 1788	64	512.0	8
	None	60	990.0	8
	None	120	1980.0	15

re5 (High Memory Type) Instance type family	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	None	180	2970.0	15

4. (Optional)In the ECS Instances Inventory Details section, after searching for the corresponding data by Region ID, Instance Type, and Date, you can click Export to export the ECS inventory details to your local computer.

4.1.4.2. 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.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Inventory Management > SLB.

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

thresholds.



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.

4.1.4.3. 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 the ASO console.
- 2. In the left-side navigation pane, choose Inventory Management > RDS.

Note You can click may 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.

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

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

thresholds.

Inventory Availability History(TB) 500(TB)				Current Inventory Usage(TB)	0
400(TB)					
300(TB)					50
200(TB)					
100(TB)				40.	29TB
Jan 8, 2020	Jan 9, 2020	Jan 10, 2020	Jan 11, 2020	~	
OSS Bucket Inventory Details					
Date Select a date Select a date					
Date	Region ID	Total(TB)	Used(TB)	Available(TB)	Usage (%)

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.

4.1.4.5. 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.

Procedure

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

? Note You	u can click 👩 in	the upper-righ	t corner to conf	igure the globa	al quota.
Inventory Availability History(TB)				Current Inventory Usage(TB)	0
50(TB)					
40(TB)					
30(ТВ)					50
20(TB)					0%
10(TB)					50G
0 , Jan 8, 2020	Jan 9, 2020	Jan 10, 2020	Jan 11, 2020		
OTS Bucket Inventory Details					
Date					
Select a date 🗰 Search					
Date	Region ID	Total(TB)	Used(TB)	Available(TB)	Usage (%)

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.

4.1.4.6. 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 the ASO console.
- 2. In the left-side navigation pane, choose Inventory Management > SLS.

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

thresholds and global quota.

sls-inner sls-public	;				ø
History Inventory Records(TB)			rrrent Quota Details(G) 41.796875(T8)		
60(TB)			2.9296875(TB)		
50(TB) 40(TB)			44140625(TB) 1.953125(TB)		
30(TB)		1	46484375(TB)		
10(TB) 0					
Jan 8, 2020 Log Service Inventory Details		Jan 10, 2020			
Date Select a date 🗰 Select a date	earch				
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.

4.1.4.7. 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.

Context

? Note EBS is the Apsara Distributed File System storage provided by the base for ECS to use. The ECS IO cluster is the cluster of Apsara Distributed File System storage. Therefore, viewing the EBS inventory is viewing the EBS inventory in the ECS IO cluster.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Inventory Management > EBS.



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.

4.1.4.8. 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.

1. Log on to the ASO console.

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

Inventory Availability History(TB)				Current Inventory Usage(TB)	
70(ТВ) 60(ТВ)					
50(TB) 40(TB) 40(TB)					50
30(TB) 20(TB)					9% 78TB
10(TB) 0 Sep 10, 2020	Sep 11, 2020	Sep 12, 2020 Sep 13, 3	2020 Sep 14, 2020		
NAS Bucket Inventory Details					
Date Select a date 📾 Search					
Date	Region ID	Total(TB)	Used(TB)	Available(TB)	Usage (%)
Sep 14, 2020	cn-	64.53	5.78	58.75	8.96%

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.

4.1.4.9. 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 the ASO console.
- 2. In the left-side navigation pane, choose Inventory Management > DFS.

Inventory Avail	lability History(TB)					Current Inventory Usage(TB)		
500(TB)								
400(TB)								11.
300(TB)							F O(
200(TB)							5% 26.20TB	
100(TB) 0							20.2016	
		Jan 20, 2020	Jan 21, 2020	Jan 22, 2020	Jan 23, 2020	-		
DFS Bucke	et Inventory Details							
Date		_						
Select a date	e 📰 Search							

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.

4.1.5. 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.

4.1.5.1. 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.

4.1.5.1.1. 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 the ASO console.
- 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.

4.1.5.1.2. 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 the ASO console.
- 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.

4.1.5.1.3. 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 the ASO console.
- 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.

4.1.5.1.4. 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 the ASO console.
- 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.

4.1.5.2. 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 the ASO console.
- 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.

4.1.5.3. 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.

4.1.5.3.1. 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 the ASO console.

- 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.

(?) Note 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.

4.1.5.3.2. 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 the ASO console.
- 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.
- Indicates the detection does not obtain results.
- Indicates an exception at the warning level exists.
- Indicates a fatal exception exists.
- Indicates 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.

4.1.5.3.3. 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 the ASO console.
- 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.

4.1.5.3.4. 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.

Procedure

- 1. Log on to the ASO console.
- 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 Host NC 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.

4.1.5.3.5. 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.

Procedure

- 1. Log on to the ASO console.
- 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.

4.1.5.3.6. 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. You can also use the IP address of SLB LVS to search for the standalone monitoring information and bandwidth chart of the SLB LVS.

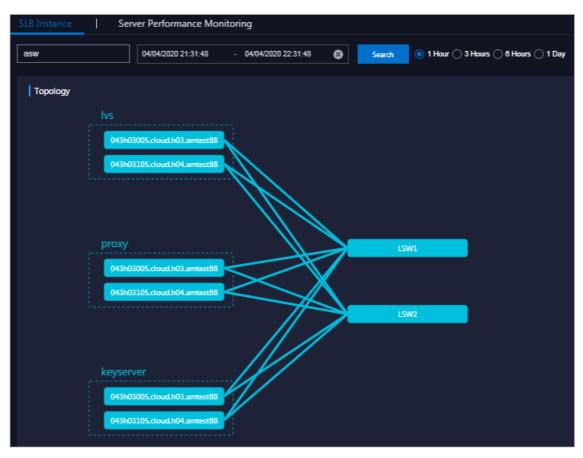
Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Full Stack Monitor > Correlation Diagnosis and Alarm.
- 3. Click the SLB tab. By default, you are on the SLB Instance sub-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. 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 section

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 section

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.

SLB Clusters	
Service Name	Instance ID
sib_keyserver	043h03005.cloud.h03.amtest88 043h03105.cloud.h04.amtest88
sib_lvs	043h03005.cloud.h03.amtest88 043h03105.cloud.h04.amtest88
slb_proxy	043h03005.cloud.h03.amtest88 043h03105.cloud.h04.amtest88

• Diagnostics section

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

Diagnostics			
Rapid Increase in Outgoing Bandwidth	Rapid Increase in Incoming Bandwidth	🤣 Rapid Increase in Outgoing Packets	Rapid Increase in Outgoing Packets
Rapid Decrease in Outgoing Bandwidth	Rapid Decrease in Incoming Bandwidth	Rapid Decrease in Outgoing Packets	Rapid Decrease in Incoming Packets
Active conn decrease list	🥪 Active conn increase list	Inactive conn decrease list	Inactive conn increase list

• SLB Bandwidth Chart section

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

• Full Stack Alert section

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

5. Click the Server Performance Monitoring sub-tab. Enter the IP address of the SLB LVS in the search box, select the time range, and then click Search to view the standalone performance data. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.

Onte LVS is the core forwarding component of SLB.

Where,

• LVS Standalone Traffic (per second) section

Move the pointer over the graph to view the traffic data, namely the inbound bps, outbound bps, number of new connections per second of this node, number of active connections of this node, and number of inactive connections of this node, of the LVS machine at a specific time.

For more information about the fields appeared when you move the pointer over the graph and the corresponding descriptions, see the following table.

Field	Description
lvs_bytin (bps)	The inbound bps.
lvs_byteout (bps)	The outbound bps.
lvs_conns	The number of new connections per second of this node.
lvs_lact	The number of active connections of this node.
lvs_linact	The number of inactive connections of this node.

• GWIP Address Configuration Status section

Move the pointer over the graph to view if the GWIP address of the LVS machine is configured at a specific time.

If the value of the GWIP address configuration status is 1, it indicates that the GWIP address is configured correctly and the SLB instance of the VPC type is available. If the value of the GWIP address configuration status is 2, it indicates that the GWIP address is not correctly configured and the SLB instance of the VPC type is unavailable.

• SLB Health Check Address Status section

Move the pointer over the graph to view the status of the SLB health check address at a specific time.

If the value of the SLB health check address status is 1, it indicates that the SLB health check address is configured correctly and SLB can check the health status of the backend RS. If the value of the SLB health check address status is 2, it indicates that the SLB health check address is not correctly configured and SLB cannot check the health status of the backend RS.

• SLB Forwarding Process CPU Utilization section

Move the pointer over the graph to view the utilization of the CPU with a specific number of the LVS machine.

• SLB Forwarding Statistics section

Move the pointer over the graph to view the SLB forwarding statistics of the LVS machine at a specific time.

For more information about the fields appeared when you move the pointer over the graph and the corresponding descriptions, see the following table.

Field	Description
acl_dropped-VALUE_DELTA	The incremental data of the forwarding number after being filtered out by the whitelist.
bad_icmp6_hdr- VALUE_DELTA	The incremental data of the number of icmp6 wrong headers.
bad_ipv4_hdr-VALUE_DELTA	The incremental data of the number of IPv4 wrong headers.
bad_ipv6_ext_hdr- VALUE_DELTA	The incremental data of the number of IPv6 wrong extension headers.
bad_ipv6_hdr-VALUE_DELTA	The incremental data of the number of IPv6 wrong headers.

• SLB Connections section

Move the pointer over the graph to view the number of SLB connections of the LVS machine at a specific time.

For more information about the fields appeared when you move the pointer over the graph and the corresponding descriptions, see the following table.

Field	Description
conns	The number of current connections.
max	The maximum number of connections.
template_conns	The number of session synchronized templates.
template_util	The utilization of session synchronized templates.
total	The total number of current connections and templates.

• SLB Health Check Type section

Move the pointer over the graph to view the SLB health check type and data.

For more information about the fields appeared when you move the pointer over the graph and the corresponding descriptions, see the following table.

Field	Description
checker-down	The number of times that the RS health check status is down.

Field	Description
checker-total	The total number of RS health check.
checker-total_down_wave	The wave that the RS health check status is down.
checker-total_up_rate	The percentage that the RS health check status is up.
checker-up	The number of times that the RS health check status is up.

4.1.5.3.7. VPC

You can use the global tunnel ID of Vlink to search for the Vlink traffic, or use the router interface ID to view the router interface information and the corresponding Vlink traffic. You can also use the IP address of the XGW host to search for the monitoring information and diagnostic information occupied by VPC XGW standalone resources.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Full Stack Monitor > Correlation Diagnosis and Alarm.
- 3. Click the VPC tab. By default, you are on the XGW Vlink sub-tab. The Topology section displays the topology of the XGW cluster.

? Note XGW is the core component used to implement the route forwarding function of VPC.

- 4. Complete the following steps to view the performance monitoring data of XGW Vlink.
 - Enter the global tunnel ID of Vlink in the search box, select the time range, and then click Search to view the Vlink 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 Vlink traffic. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
- 5. Click the XGW Host sub-tab. Enter the IP address of the XGW host in the search box, select the time range, and then click Search to view the XGW diagnostic information and bandwidth chart of the XGW host. You can select one hour, three hours, or six hours as the time range, or customize the time range.

Where,

• XGW Diagnostic Information section

View the diagnostic information of the XGW host. A green tick indicates the normal status, and a red cross indicates the abnormal status. Move the pointer over the abnormal diagnostic information to view the corresponding error message in the displayed dialog box.

• Bandwidth Chart section

Move the pointer over the **CPU** graph in the upper-left corner to view the usage of the CPU with a specific number of the XGW host. Move the pointer over the **MEM** graph in the upper-right corner to view the memory usage of the XGW host at a specific time.

Move the pointer over the graph at the bottom of the page to view the outbound pps, inbound pps, packet loss rate, outbound byte rate, and inbound byte rate of the XGW host at a specific time.

4.1.5.4. Use case

4.1.5.4.1. Monitor and diagnose the VPC XGW host

resources

This topic uses a common use case to tell people how to use the Correlation Diagnosis and Alarm module of VPC to monitor and diagnose the VPC XGW host resources.

Scenario

The network access latency of the current VPC is increased. You must make sure if the host resource occupancy of the core forwarding component of VPC is normal.

Procedure

Complete the following steps to troubleshoot the problem.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Full Stack Monitor > Correlation Diagnosis and Alarm.
- 3. Click the VPC tab and then click the XGW Host sub-tab.
- 4. Enter the IP address of the XGW host that you are about to view and then select the monitoring cycle. Click Search to search for each metric.
- 5. View the diagnostic information. If a red item exists, it indicates that the monitoring metric has an exception.
- 6. Move the pointer over the abnormal diagnostic information to view the exception details.
- 7. View the bandwidth chart and find that traffic sometimes increases or decreases sharply.

Result

By monitoring the XGW host resources, you can find that the problem may be caused by the sharp increase of traffic.

4.1.6. Storage Operation Center

The Storage Operation Center module consists of pangu and EBS.

4.1.6.1. Pangu

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

4.1.6.1.1. 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 the ASO console.
- 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 greater than zero.

✓ Overview									
Stor	Storage Server Health								
Clusters	14	Servers	257	Abnormal Disks	1	Log Warning Num	3		
Storage	2,199.24T	Masters	42	Abnormal Masters	0	Log Error Num	0		
Percentage	7.3500%	Chunk Servers	78	Abnormal Chunk Servers	O	Log Fatal Num	0		
Files	46,080,070			Abnormal Water Levels	0	Replica Error Num	0		

• Heatmap of Health

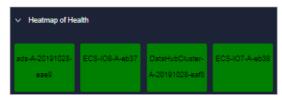
The Heatmap of Health section displays the health information of all the clusters in the selected product. Clusters in different health statuses are displayed in different colors.

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 corresponding product 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 6, 2020, I	00:00:00 ~ Jan 6, 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		sls			
	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				oss			

4.1.6.1.2. Cluster information

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

Procedure

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

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

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

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

• Overview

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

∨ Oveniew								
Sto	Storage Server Health							
Storage	34.86T	Servers		Abnormal Water Levels		Log Warning Num		
Percentage	17.5100%	Abnormal Masters/Masters	0/3	Abnormal Masters		Log Error Num		
Chunk Servers		Abnormal Chunk Servers/Chunk	0/5	Abnormal Chunk Servers		Log Fatal Num		
Files	214,849	Abnormal Disks/Disks	0/50	Abnormal Disks		Replica Error Num		

• Alarm Monitor

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

✓ Alarm Mor	itor		
			Alarm Log
Fuzzy Search:		Q	
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.

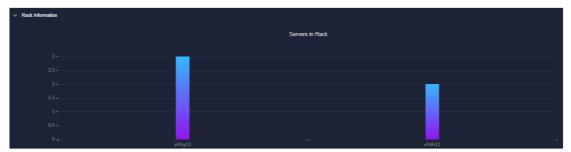
? Note Predicted Water Levels has values only if Historical Water Levels has a certain amount of data. Therefore, some clusters may only have historical water levels, without predicted water levels.



• Rack Information

Contains Servers in Rack and Storage. Where,

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



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



4.1.6.1.3. 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 the ASO console.
- 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.

? 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.

• 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.

Cluster Na	ame: ECS-IO7-A-eb38	~	
∨ Ma	ister Info		
Tota			
Serv	er		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	nk Server Info						
Total: Fuzzy	5 Normal:5 Search: Enter a keyword						Refresh
	Server	DiskBroken Disks/Disks	SSDCacheBroken Disks/Disks	Status	Backup	Storage (TB)	Usage(%)
+	a58g13210.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%
+	a58g13211.cloud.h14.amtest 72	0/10	0/10	NORMAL		13.8476	24.1900%
+	a58h11210.cloud.h13.amtest 72	0/10	0/10	NORMAL		13.8476	28.6300%
+	a58g13110.cloud.g14.amtest 72	0/10	0/10	NORMAL		13.8476	24.1000%

4.1.6.1.4. Pangu operation

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

- 1. Log on to the ASO console.
- 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-eb 33	ECS-IO8 -A-eb 37	ECS-IO7 -A-eb 38		

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

4.1.6.1.5. Product settings

By default, the system configures the alert threshold for each cluster. You can adjust the water level threshold, chunk server threshold, and disk threshold of each cluster based on business needs.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Storage Operation Center > Pangu > Product Settings.
- 3. Select the cluster you are about to configure from the **Cluster Name** drop-down list.
- 4. Click Modify at the bottom of the page to modify the threshold information of the cluster.

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

Configuration		Description
	Warn Threshold	If the storage capacity utilization of the cluster is greater than or equal to this value, the alert at the warning level is triggered and the heatmap of health of this cluster is displayed in yellow. You can enter a value greater than zero and less than or equal to 100. If not configured, the alert at the warning level is triggered when the water level of the cluster is greater than or equal to 65% by default.
	Error Threshold	If the storage capacity utilization of the cluster is greater than or equal to this value, the alert at the error level is triggered and the heatmap of health of this cluster is displayed in red. You can enter a value greater than zero and less than or equal to 100. If not configured, the alert at the error level is triggered when the water level of the cluster is greater than or equal to 85% by default.
Cluster Water Level		

Operations and Maintenance Guide • Monitoring

Configuration		Description
	Fatal Error Threshold	If the storage capacity utilization of the cluster is greater than or equal to this value, the alert at the fatal error level is triggered and the heatmap of health of this cluster is displayed in dark red. You can enter a value greater than zero and less than or equal to 100. If not configured, the alert at the fatal error level is triggered when the water level of the cluster is greater than or equal to 92% by default.
	Warn Threshold (Abnormal Chunk Server Quantity)	If the number of abnormal chunk servers is greater than or equal to this value, the alert at the warning level is triggered and the heatmap of health of this cluster is displayed in yellow. If not configured, the alert at the warning level is triggered when the number of abnormal chunk servers is greater than or equal to 1 by default.
Chunk Server	Error Threshold (Abnormal Chunk Server Ratio)	If the ratio of abnormal chunk servers to all the chunk servers is greater than this value, the alert at the error level is triggered and the heatmap of health of this cluster is displayed in red. If not configured, the alert at the error level is triggered when the ratio of abnormal chunk servers to all the chunk servers is greater than or equal to 10% by default.
	Warn Threshold (Abnormal Disk Quantity)	If the number of abnormal disks is greater than or equal to this value, the alert at the warning level is triggered and the heatmap of health of this cluster is displayed in yellow. If not configured, the alert at the warning level is triggered when the number of abnormal disks is greater than or equal to 1 by default.
Disk		

Configuration		Description
	Error Threshold (Abnormal Disk Ratio)	If the ratio of abnormal disks to all the disks is greater than this value, the alert at the error level is triggered and the heatmap of health of this cluster is displayed in red. If not configured, the alert at the error level is triggered when the ratio of abnormal disks to all the disks is greater than or equal to 10% by default.

Cluster Name: ECS-IO8-A-d10f V	
Threshold	
Cluster Water Level : (Warning value must be greater than zero, critical error value greater than	
Warn Threshold	
Error Threshold	
Fatal Error Threshold	
Chunk Server:	
Warn Threshold(Abnormal Chunk Server Quantity)	
Error Threshold(Abnormal Chunk Server Ratio)	
Disk:	
Warn Threshold(Abnormal Disk Quantity)	
Error Threshold(Abnormal Disk Ratio)	
Modify	

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

5. Then, click Save.

4.1.6.2. EBS

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

4.1.6.2.1. 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.

1. Log on to the ASO console.

- 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 user ID to which these VMs belong.

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

✓ Affected VM List				
Enter a keyword Q				
Cluster Name J]*	User ID↓}	Virtual Machine JN	Start Time ↓}	Recovery Time ↓
ECS-108-A-5679			2020-02-24 13:58:09	2020-02-25 13:48:13

• 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.

✓ VM Cluster Statistics	
Enter a keyword Q	
Cluster Name ↓	Number of Virtual Machines 🕸
ECS-108-A-5679	57

• 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.

✓ Device Cluster Statistics	
Enter a keyword Q	
Cluster Name.↓	Number of Device ↓
ECS-IO8-A-5879	57

4.1.6.2.2. 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.

> Document Version:20200918

- 1. Log on to the ASO console.
- 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 IP address, 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 IP, Virtual Machine, Device ID, Storage Type, Start Time, Recovery Time, Number of Slow IO, and Reason as needed.

✓ Slow IO List								
Enter a keyword	٩							
Cluster Name ↓	NC IP 11	Virtual Machine 1	Device ID ↓	Storage Type ↓	Start Time ↓	Recovery Time ↓	Number of Slow IO ↓	Reason↓

• 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 address, 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, you can 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.

✓ Top Ten NC				
Enter a keyword				
NC IP 1	Cluster Name √	Slow IO ↓	Percentage	Major Reason ↓

• 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, you can 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, you can 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.

`	V Top Five Cluster Statistics						
	Enter a keyword Q						
	Top Five Cluster ↓	Number of Device ↓	Number of Slow IO 1	Percentage	Major Problem J↑		

- 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, you can 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.

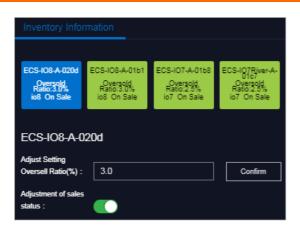
∨ Reason		
Enter a keyword		
Reason J	Number of Slow IO J	Percentage

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

4.1.6.2.3. 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.

- 1. Log on to the ASO console.
- 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.



- 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.

5.Operations tools 5.1. 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.

5.1.1. Service configuration

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

5.1.1.1. Configure the backup server

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

- 1. Log on to the ASO console.
- 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. 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. 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 Monitoring Path	The storage path of backup files on the backup 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.

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

Product Management				
				Add
Product	Backup Items	Backup Script	Retry Times	Actions
pangu	datahub_pangu_master	metadata_backup.py		Modify Delete

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_backup.py.
Retry Times	The number of retries. 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.

5.1.2. Backup service

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

5.1.2.1. 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 and ots pangu, according to different situations of Apsara Stack.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Offline Backup > 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.
- 3. 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 / <i>pr</i> <i>oduct name</i> _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 / <i>pr</i> <i>oduct name_</i> 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.				

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

5.1.2.2. 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 the ASO console.
- 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.

5.1.2.3. 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 the ASO console.
- 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.

5.1.3. 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 the ASO console.
- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Service Status > State.

		Backup Server Stat	Backup Server Status:				
In Proce	ess:0 Items Comp	olete:18 Items	e-out:6 Items	Failed:1 Items	Usage (%)	Memory	
Product	Backup Items	Completed Items	Failed Items	Latest			
angu	datahub_pangu_master	18					
rto_test_item_1 26	auto_test			Not Started			
					Usage (%)	Disk Status	
							Tin

- 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.
 - $\circ\;$ View the statuses of the latest backup items of the current product.

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

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

5.1.4. Use cases

5.1.4.1. Offline backup of metadata

To guarantee the availability of cloud platforms, you must back up the pangu data of each product.

5.1.4.1.1. Preparations before the backup

This topic describes the preparations before the backup.

Before the backup, prepare the following machines as required:

• Prepare an independent buffer machine as the backup server.

If no buffer machine exists in the environment, select the physical machine with large disk space and good network performance in the environment. Otherwise, the security of the backup data cannot be guaranteed.

Offline backup files cannot be stored on backed up objects. Therefore, if the on-site environment does not have extra physical machines or sufficient disk capacity, you must increase physical machines or disk space before the offline backup.

• A transfer machine is required for backup products to store one-time backup data and backup scripts of each product.

No other requirements are for the transfer machine.

• The network of the backup server must be connected with the network of the Docker container where the offline backup service is located to make sure that the backup container in the ASO cluster can log on to the transfer machine and backup server by using SSH, without providing the username and password.

5.1.4.1.2. Collect pangu information of each product

Collect the pangu information of products to be backed up to add the backup product information in Apsara Stack Operations (ASO).

Context

This topic uses the customized product names oss, ecs, ads, and ots and collects the information of these products. The products whose pangu information you are about to collect are subject to the on-site environment.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the top navigation bar, choose Operations > Service Operations.
- 3. Enter pangu in the search box to search for the pangu service.
- 4. At the right of the pangu service, click **Management** in the **Actions** column to go to the details page of the pangu service.
- 5. Click the Service Instance tab.
- 6. Click the instance name to go to the service instance information dashboard page.

Here use the service instance ECS-IO7-A-xx as an example.

- 7. In the Server Role List section, find PanguMaster#.
- 8. At the right of PanguMaster#, click Details in the Actions column.

Server Role List								
Server Role	Current Status	Expected Machines	Machines In Final Status	Machines Going Offline	Rolling Task Status	Time Used	Actions	
PanguChunkserver#	In Final Status	11	11	0	no rolling		Details	
PanguMaster#	In Final Status	3	3	0	no rolling		Details	
PanguMonitor#	In Final Status	20	20	0	no rolling		Details	
PanguSupervisor#	In Final Status	2	2	0	no rolling		Details	
PanguTools#	In Final Status	15	15	0	no rolling		Details	

9. In the Machine Information section, view and record the IP address of the pangu master.

Record any one IP address of the three PanguMaster#.

Machine Information 0									
Machine Name \equiv	IP	Machine Status	Machine Action	Server Role	Server Role	Current Versi	Target Version	Error Message	Actions
vm010101020075		good		good PROBATIO		e99e7f24d84c73b	e99e7f24d84c73b		Terminal Restart Details Machine System View Machine Operation
vm010101024066		good		good PROBATIO		e99e7f24d84c73b	e99e7f24d84c73b		Terminal Restart Details Machine System View Machine Operation
vm010101028082		good		good PROBATIO		e99e7f24d84c73b	e99e7f24d84c73b		Terminal Restart Details Machine System View Machine Operation

10. See steps 6 to 9 and click each instance name in sequence to view and record the pangu information of each product. The records are similar to the following table.

Cluster name	pangumaster IP address	Product name
AdvanceOssCluster-A-xx	10.10.10.1	055
ECS-I07-A-xx	10.10.10.2	ecs

Cluster name	pangumaster IP address	Product name
ads-A-xx	10.10.10.3	ads
otsv3_p-A-xx	10.10.10.4	ots

? Note You can customize the product name. Make sure the product name is a unique one and recognizable at a certain degree.

5.1.4.1.3. Configure the backup server

You must configure the backup server in Apsara Stack Operations (ASO).

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Service Configuration > Backup Service Configuration.
- 4. At the right of the backup server, click **Modify** in the **Actions** column and then configure the backup server information.

Configuration	Description
Backup Server IP Address	The IP address of the actual backup server. The backup server must be an independent physical server, managed by Apsara Infrastructure Management Framework, and has its network connected with other servers in Apsara Stack. 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 Monitoring Path	The storage path of backup files on the actual backup 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.

Backup Service Configuration			
Backup Server IP Address	Backup Server Monitoring Path	Backup Retention (day)	Actions
	/apsara/backup/		Save Cancel

5.1.4.1.4. Add a backup product

You can add the backup product information in Apsara Stack Operations (ASO).

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Service Configuration > Product Management.
- 4. On the Product Management page, 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 value based on the collected product information in Collect pangu information of each product in the format of <i>back</i> <i>up product name_</i> pangu. For example, oss_pangu.
Backup Script	For example, enter metadata_backup.py.
Retry Times	We recommend that you enter 3.

Add Product		×
* Product		
You must specify this field.		
# Backup Items		
You must specify this field.		
Backup Script		
You must specify this field.		
• Retry Times times		
	Cancel	ОК

6. Repeat steps 4 to 5 to add all the backup items in sequence.

Product Management				
				Add
Product	Backup Items	Backup Script	Retry Times	Actions
pangu	datahub_pangu_master	metadata_backup.py	2	Modify Delete

5.1.4.1.5. Configure the backup

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

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, select Offline Backup.
- 3. Choose Backup Service > Backup Configuration.
- 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 machine.
Backup File Folder	A folder on the transfer machine. 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 / <i>pr</i> <i>oduct name</i> _pangu/bin.
Script Execution Folder	A folder on the transfer machine. 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 / <i>pr</i> <i>oduct name_</i> pangu/bin.
Script Parameters	You must enter the value in the format ofip=xxx.xxx.xxx.xxx , in which the IP address is any IP address of pangu master recorded in Collect pangu information of each product.
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. Repeat steps 4 to 5 to configure all the backup items.

5.1.4.1.6. View the backup details

After configuring the backup items, you can view if each backup item works properly in Apsara Stack Operations (ASO).

Procedure

- 1. Log on to the ASO console.
- 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.

If the state of a backup item is **Complete**, it indicates that the backup item works properly.

If the backup is finished, view the MD5 value of the backup file to check if the MD5 value of the offline backup service is the same as that of the backup server. If yes, the backup is successful.

5.2. NOC

5.2.1. 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.

5.2.2. Dashboard

5.2.2.1. View the dashboard

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

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

	Item		
D	Device Overview	The model distribution of used network devices.	

ltem		Description
Device Management	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.
	Alerts	The total number of alerts generated by network devices.
Network Monitoring	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 utilization of SLB clusters.
	XGW Overview	The bandwidth utilization of XGW clusters.

Dashboard Networ	k Topology Custom View	N						
Device Management		Network Monitoring				Traffic Dashboard		
Device Overview						SLB Overview		
Total Devices by Model		Alerts		Alerting Devices		Cluster Bandwidth Utilization		
		0		0	/2	30M 11H 0H 12H S00Mbits/s 400Mbits/s 400Mbits/s 300Mbits/s		
	Нзс	Alarm Details				300Mbits/s 250Mbits/s		
🔵 H3C 🛑 Mellanox	Avocent	Time	Device Name	Alert Item	Details	200Mbits/s 150Mbits/s		
Ports Usage						100Mbits/s Jan 1, 1970, 08:00:00		
Ports Utilization							<u> </u>	
o DSW	56%					XGW Overview		
• ASW	31%					Cluster Bandwidth Utilization	30M 1H 6H 12H	
• LSW	0%							
Error Packets by PortTOP5								
	Time Range: 1 Month 🗸						H . A . A .	
Ranking Device	Port Total Error Packets					2Mbps 2Mbps	MMM	M/M

5.2.2.2. View the network topology

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

Procedure

- 1. Log on to the ASO console.
- 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 or Dynamic Topology as the Topology Type as needed.

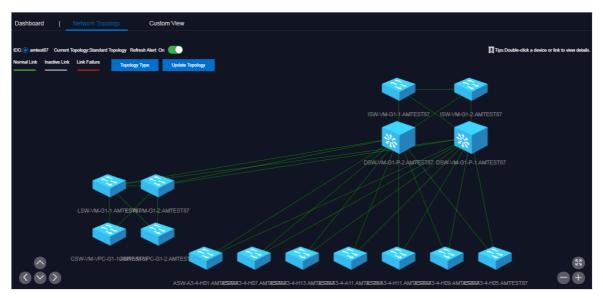
If the dynamic topology is inconsistent with the standard topology, a message appears in the upper-right corner when you are on the **Network Topology** tab and disappears after a few seconds. You can click **Update Topology** to update the standard topology.

? Note

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

- Green: The link works properly.
- Red: The link has a fault.
- 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 the device status or link status in the topology is 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.

5.2.2.3. Manage custom views

You can create a custom view based on business needs to configure how to display the independent monitoring data set as needed. By configuring the contents and rules to display in the view, you can summarize and display the monitoring data and graph information you are interested in.

Create a view

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Custom View tab.
- 4. Create a view.
 - i. Click Create View.

Dashboard	I	Netwo	rk Topology	Custom View				
Select a view.			02/18/2020 14:57:04	- 02/16/2020 20:57	7:04 😵	Search	Create View	Delete View

ii. 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.

5. Then, add a subview. By default, no subviews exist in a view after you create the view.

i. In the search box, select the view and then click Search.

Dashboard	Ne	etwork Topology	Custom View			
test01		V 02/16/2020 15:40:56	- 02/18/2020 21:40:58	8	earch Create Vi	ew Delete View
]	+			
		L	<u> </u>			

ii. Click the + button.

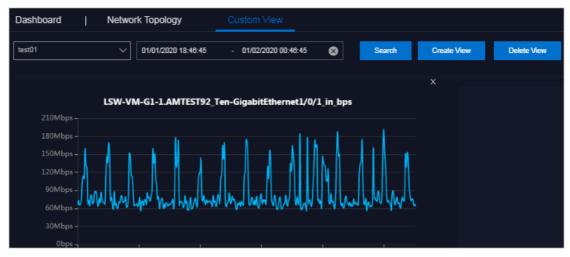
Configuration	Description
Device	Required. Select the device to be monitored from the drop- down list.
Monitoring Object	 Required. Select the monitoring object from the drop-down list. interface: the switch interface, including the water level, packet error, and packet loss of the interface. hardware: the switch hardware, including the memory usage and CPU usage. capacity: others, which is not supported currently.
Monitoring Metric	Required. Select the corresponding monitoring metric from the drop-down list according to the selected monitoring object.
Monitoring Submetric	Optional. Select the corresponding monitoring submetric from the drop-down list according to the selected monitoring metric.

iii. On the displayed page, select the device, monitoring object, monitoring metric, and monitoring submetric.

		×
Device	ASW-A3-4-H05.AMTEST87 V	
Monitoring Object	interface 🗸	
object		
Monitoring Metric	FortyGigE1/0/49 V	
incluso		
Monitoring Submetric	Select ^	
	out_bps	
	in_bps	
	in_pps	
	out_pps	
	in_out_bps	
	in_out_pps	

iv. Click OK.

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



v. You can add other subviews as needed.

Delete a subview

- 1. Log on to the ASO console.
- 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**.

Delete a view

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

1. Log on to the ASO console.

- 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 on top of the page.
- 6. In the displayed dialog box, click **OK**.

5.2.3. Network Service Provider

5.2.3.1. View access gateway instances

On the Instance Management tab, you can view the summary, namely the access gateway name, IBGP role, and created time, of access gateway instances in the current system.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > VPC.
- 3. Click the Instance Management tab.
- 4. Enter the region ID in the upper-left corner.

? Note To view the instances in other regions, click Reset in the upper-right corner and then enter the ID of another region.

5. Click Display Device List to view the access gateway device list in the current environment.

Note If new devices exist, click Scan for New Devices and then click Display Device List.

Column name	Description
Access Gateway Name	The access gateway name in the current system.
IBGP Role	 The role of the access gateway in the environment. Where, PR-Active: indicates the role of the current access gateway device is PR active device. Client: indicates the role of the current access gateway device is not PR active device.
Created At	Indicates the time when the current VSwitch acts as an access gateway instance.

Instance Management	I	Operation Logs	Bare-Met	al Networks	I	O&M			R	tegion:	NSP Version:
Region ID								Display Device List	Scan for New Devices		Reset
Access Gateway Name			IB	GP Role			Created At				

5.2.3.2. View operation logs

You can view API operation logs of bare metals based on Operations & Maintenance (O&M) needs.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > VPC.
- 3. Click the **Operation Logs** tab.
- 4. Configure the filter conditions, such as region ID, VSwitch ID, bare metal name, access gateway name, and time range, and then search for operation logs that meet the filter conditions.

For more information about the filter conditions and their descriptions, see the following table.

Filter condition	Description
Region ID	The name of the region in the current environment.
VSwitch ID	The ID of the VSwitch used when the bare metal is applied for or released in VPC.
Bare Metal Name	The name of the bare metal applied for or released in VPC. To identify the bare metal as a unique one in the region, here use the serial number of the bare metal.
VLAN ID	Currently not supported. We recommend that you do not configure this filter condition.
Access Gateway Name	The name of the access gateway to be searched for.
Request ID	Currently not supported. We recommend that you do not configure this filter condition.
Created At	The time range of the API operation to be searched for.

? Note To modify the filter conditions and search for operation logs again, click Clear in the upper-right corner and configure the filter conditions again.

Instance Manag	gement	Operation Logs	Bare-Metal Networks	O&M				Region: NSP Version:
Region ID		VSwitch ID		Bare Metal Name		VLAN ID		
Access Gateway Na	ime	Request ID		Created At				
				Start Date -	End Date	 		Search Clear
ID	Created At	API Operation	VSwitch ID	Access Gateway Name	Port	Bare Metal Name	Status	Actions

Column name	Description
ID	The index of the operation log.
Created At	The time when the current operation happens.
	The category of the API operation, such as Apply for Bare Metal in VPC and Release Bare Metal in VPC.
	Where,
	• add: indicates that a bare metal is applied for in VPC.
	• del: indicates that a bare metal is released in VPC.
API Operation	• del_pc: indicates that a physical connection is deleted.
	 del_vbr: indicates that a Virtual Border Router (VBR) is deleted.
	• del_router_intf: indicates that a router interface is deleted.
	 del_route_entry: indicates that a route table entry is deleted.
VSwitch ID	The ID of the VSwitch used when the bare metal is applied for or released in VPC.
Access Gateway Name	The name of the access gateway involved with the current operation.
Port	The port to which the bare metal belongs.
Bare Metal Name	The name of the bare metal applied for or released in VPC. To identify the bare metal as a unique one in the region, here is the serial number of the bare metal.
Status	The status of the API operation. success indicates the operation is successful. If the API operation is in progress, the value indicates the real-time status of the API operation. If the API operation is complete but the value is not success, you can view the failure information in this column.

For more information about the search results, see the following table.

5. In the search results, find an operation log and then click **View Details** in the **Actions** column. Then, view the details of the API operation.

5.2.3.3. View network information of bare metals in VPC

On the **Bare-Metal Networks** tab, you can view the information of bare metals added to VPC in the system.

Procedure

1. Log on to the ASO console.

- 2. In the left-side navigation pane, choose NOC > VPC.
- 3. Click the **Bare-Metal Networks** tab. By default, the system displays the network information of bare metals in the current system on different pages.
- 4. Configure the filter conditions, namely region ID, VPC ID, VSwitch ID, VBR ID, BD ID, access gateway name, and time range, and then search for the information of bare metals that meet the filter conditions.

Filter condition	Description
Region ID	The name of region in the current environment.
VPC ID	The ID of the VPC to which the bare metal to be searched for belongs.
VSwitch ID	The ID of the VSwitch to which the bare metal to be searched for belongs.
VBR ID	The VBR ID of the physical connection created on HSW by the VPC to which the bare metal to be searched for belongs.
BD ID	The value of the hardware bridge-domain (BD) to which the bare metal is added.
Access Gateway Name	The name of the access gateway to which the bare metal to be searched for belongs.
Created At	The time range when the current bare metal is allocated to the VPC.

? Note To modify the filter conditions and search for bare metals again, click Clear in the upper-right corner and configure the filter conditions again.

Instance Management	Operation	Logs Bare-Metal Network	ks	O&M					Region: NSP Version
Region ID		VPC ID		VSwitch ID		VBR ID			
BD ID		Access Gateway Name		Created At					
				Start Date	- End Date				Search Clear
Bare Metal Name	VPC ID	Created At	VSwitch II		Access Gateway Name	Port	VBR ID	BD ID	Actions

5. In the search results, find a bare metal and then click **View Details** in the **Actions** column. Then, view the details of the bare metal.

5.2.3.4. O&M configurations

5.2.3.4.1. Apply for a bare metal in VPC

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to add the physical port of the access gateway to which a bare metal belongs to VPC.

Prerequisites

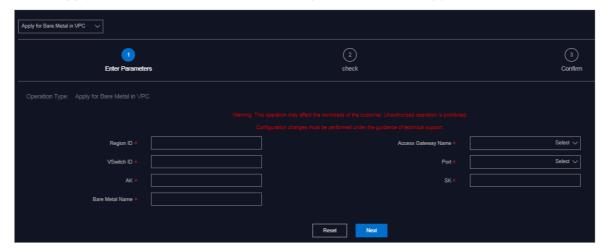
Notice You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Before using this function, note that:

- Generally, you cannot use this function to apply for a bare metal in VPC, but use the bare metal controller to call an API to open the bare metal network.
- By using this function, you can only connect the bare metal to the access gateway port, but do nothing to the bare metal. To configure the business network port IP address and routing information of the bare metal, ask the corresponding product team for guidance.

Procedure

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Apply for Bare Metal in VPC from the drop-down list in the upper-left corner.



4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Access Gateway Name	Select the name of the access gateway to which the bare metal is connected.
VSwitch ID	Enter the ID of the VS witch to which the bare metal is to be added. You can obtain this value from the VPC console.
Port	Select the port of the access gateway to which the bare metal is connected.
AK and SK	Obtain the AK and SK on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VPC belongs.
Bare Metal Name	The name of the bare metal to be applied for. Here, enter the serial number of the bare metal.

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can search for the bare metal based on the bare metal name and view that the bare metal is added to the VPC on the **Bare-Metal Networks** tab.

5.2.3.4.2. Release a bare metal in VPC

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to disconnect the physical port of the access gateway to which a bare metal belongs from VPC.

Prerequisites

Notice You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Before using this function, note that:

- Generally, you cannot use this function to release a bare metal in VPC, but use the bare metal controller to call an API to delete the bare metal network.
- By using this function, you can only connect the bare metal to the access gateway port, but do nothing to the bare metal. To configure the business network port IP address and routing information of the bare metal, ask the corresponding product team for guidance.

Procedure

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Release Bare Metal in VPC from the drop-down list in the upper-left corner.

Release Bare Metal in VPC		
1 Enter Parameters	2 check	3 Confirm
Operation Type: Release Bare Metal in VPC		
War		
Region ID		
Bare Metal Name *		
	Reset	

4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Bare Metal Name	The name of the bare metal to be released. Here, enter the serial number of the bare metal.

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can search for the bare metal based on the bare metal name and view that the bare metal does not exist in VPC on the **Bare-Metal Networks** tab.

5.2.3.4.3. Delete a VPC route table entry

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete a route table entry pointing to the bare metal subnet in VPC.

Prerequisites

Notice You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Before using this function, note that:

- Generally, you cannot use this function to delete a VPC route table entry. This operation is only for emergency situations.
- This operation can only delete a route table entry at a time. To delete multiple route table entries, you must perform this operation multiple times.

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Delete Route Table Entry from the drop-down list in the upper-left corner.

Operations and Maintenance Guide · Operations tools

Delete Route Table Entry				
1		2		3
Enter Paramete	ิส	check		Confirm
Operation Type: Delete Route Table Entry				
Region ID		Routing Table I		
Routing Interface ID *		Routing destination CID	R*	
AK				
		Reset		

4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Routing Table ID	The VPC route table ID, which can be obtained from the VPC console. For more information about how to obtain the route table ID, see the <i>VPC User Guide</i> .
Routing Interface ID	The VPC router interface ID, which can be obtained from the VPC console. For more information about how to obtain the router interface ID, see the <i>VPC User Guide</i> .
Routing destination CIDR	The destination CIDR block that VPC points to, which can be obtained from the VPC console. For more information about how to obtain the routing destination CIDR block, see the <i>VPC U ser Guide</i> .
AK and SK	The organization AK and SK of the VPC to which the bare metal belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VPC belongs.

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can log on to the VPC console and view that the route table entry of the entered destination CIDR block does not exist in VPC, indicating that the route table entry is deleted.

7. (Optional)In actual fault scenarios, if multiple route table entries exist in the VPC route table, repeat step 3 to step 6 to delete other route table entries.

5.2.3.4.4. Delete a VBR route table entry

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete the default route table entry of Virtual Border Router (VBR).

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Delete Route Table Entry from the drop-down list in the upper-left corner.
- 4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Routing Table ID	The VBR route table ID. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the Bare-Metal Networks tab based on the bare metal name and then click View Details. The VBR Route Table ID in the details is the value of this field. If the bare metal involved with the VBR is not added to VPC, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the VBR Route Table ID in the details is the value of this field.
Routing Interface ID	The VBR router interface ID. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the Bare-Metal Networks tab based on the bare metal name and then click View Details. The VBR RI in the details is the value of this field. If the bare metal involved with the VBR is not added to VPC, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the VBR RI in the details is the value of this field.
Routing destination CIDR	Fixed value: 0.0.0.0/0

Configuration	Description
AK and SK	The AK and SK of the organization to which the VBR belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VBR belongs.

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can click **Apsara Network Intelligence** of the **Products** module. On the home page of Apsara Network Intelligence, enter the VBR route table ID and search for the VBR route table. You can view that the route table entry 0.0.0.0/0 does not exist in the VBR route table, indicating the route table entry is deleted.

5.2.3.4.5. Delete a VPC router interface

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete a VPC router interface.

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select **Delete Router Interface** from the drop-down list in the upper-left corner.

Delete Router Interface		
enter Parameters	(2) check	3 Confirm
Operation Type: Delete Router Interface		
Warn		
Region ID		
Router Interface ID I		
AK		
SK		
	Reset	

Configuration	Description
Region ID	The name of the region in the current environment.
Router Interface ID	The VPC router interface ID. On the Operation Logs tab, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add . Click View Details and the VPC RI in the details is the value of this field.
AK and SK	The organization AK and SK of the VPC to which the bare metal belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VPC belongs.

4. Complete the configurations. For more information, see the following table.

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can click **Apsara Network Intelligence** of the **Products** module. On the home page of Apsara Network Intelligence, enter the VPC router interface ID and search for the router interface. No search result appears, indicating the router interface is deleted.

5.2.3.4.6. Delete a VBR router interface

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete a VBR router interface.

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Delete Router Interface from the drop-down list in the upper-left corner.
- 4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Router Interface ID	The VBR router interface ID. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the Bare-Metal Networks tab based on the bare metal name and then click View Details. The VBR RI in the details is the value of this field. If the bare metal involved with the VBR is not added to VPC, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the VBR RI in the details is the value of this field.
AK and SK	The organization AK and SK of the VPC to which the bare metal belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VPC belongs.

Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can click **Apsara Network Intelligence** of the **Products** module. On the home page of Apsara Network Intelligence, enter the VBR router interface ID and search for the router interface. No search result appears, indicating the router interface is deleted.

5.2.3.4.7. Delete a VBR

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete a Virtual Border Router (VBR).

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Procedure

1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.

2. Click the O&M tab.

3. Select **Delete VBR** from the drop-down list in the upper-left corner.

Delete VBR V		
1 Enter Parameters	2 check	3 Confirm
Operation Type: Delete VBR		
Wa		
Region IC		
VBRID		
AF		
SK		
	Reset	

4. Complete the configurations. For more information, see the following table.

Configuration	Description	
Region ID	The name of the region in the current environment.	
VBR ID	The ID of the VBR to be deleted. On the Operation Logs tab, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add. Click View Details and the VBR ID in the details is the value of this field.	
AK and SK	The AK and SK of the organization to which the VBR belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VBR belongs.	

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can click **Apsara Network Intelligence** of the **Products** module. On the home page of Apsara Network Intelligence, enter the VBR ID and search for the VBR. No search result appears, indicating the VBR is deleted.

5.2.3.4.8. Delete a physical connection

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete a physical connection.

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Procedure

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select **Delete Express Connect Circuit** from the drop-down list in the upper-left corner.

Delete Express Connect Cir V			
1 Enter Parameters	2 check	3 Confirm	
Region ID -			
Express Connect Circuit ID			
AF			
SK			
	Reset		

4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Express Connect Circuit ID	The ID of the physical connection to be deleted. On the Operation Logs tab, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add. Click View Details and the Express Connect Circuit ID in the details is the value of this field.
AK and SK	The AK and SK of the organization to which the VBR belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VBR belongs.

Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, you can click **Apsara Network Intelligence** of the **Products** module. On the home page of Apsara Network Intelligence, enter the physical connection ID and search for the physical connection. No search result appears, indicating the physical connection is deleted.

5.2.3.4.9. Delete all resources with one click

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to delete all resources, namely the Virtual Private Cloud (VPC) route table entries, Virtual Border Router (VBR) route table entries, VPC router interfaces, VBR router interfaces, VBRs, and physical connections, with one click.

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Procedure

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Delete ALL Resources from the drop-down list in the upper-left corner.

Deleta ALL Resources V			
1 Enter Parameter	s	2 check	3 Confirm
Operation Type: Delete ALL Resources			
Region ID *		Access Gateway Name	Select V
VPC Routing Interface ID		VPC Routing Table ID	
VBR Route Table ID		VPC CIDR 1	
VPC CIDR 2		VBR Routing Interface ID	
VBR ID		Express Connect Circuit ID	
VLAN ID] Trunk ID	
		Reset	

4. Complete the configurations. For more information, see the following table.

Configuration	Description
Region ID	The name of the region in the current environment.
Access Gateway Name	Select the name of the access gateway to which the bare metal is connected.

Configuration	Description
AK and SK	The AK and SK of the organization to which the VBR belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VBR belongs.
VPC Routing Interface ID	The router interface ID of VPC, which can be obtained from the VPC console. For more information about how to obtain the router interface ID, see the <i>VPC User Guide</i> .
VPC Routing Table ID	The route table ID of VPC, which can be obtained from the VPC console. For more information about how to obtain the route table ID, see the <i>VPC User Guide</i> .
VBR Route Table ID	The route table ID of VBR. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the Bare-Metal Networks tab based on the bare metal name and then click View Details. The VBR Route Table ID in the details is the value of this field. If the bare metal involved with the VBR is not added to VPC, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the VBR Route Table ID in the details is the value of this field.
VPC CIDR 1	The destination CIDR block 1 that VPC points to, which can be obtained from the VPC console. For more information about how to obtain the VPC CIDR block 1, see the VPC User Guide.
VPC CIDR 2	The destination CIDR block 2 that VPC points to, which can be obtained from the VPC console. For more information about how to obtain the VPC CIDR block 2, see the VPC User Guide.

Configuration	Description
	The router interface ID of VBR. If the bare metal involved with the VBR is added to VPC, you can search for the bare metal on the Bare-Metal Networks tab based on the bare metal name and then click View Details. The VBR RI in the details is the value of this field.
VBR Routing Interface ID	If the bare metal involved with the VBR is not added to VPC, configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the VBR RI in the details is the value of this field.
VBR ID	The ID of the VBR to be deleted. Configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the VBR ID in the details is the value of this field.
Express Connect Circuit ID	The ID of the physical connection to be deleted. Configure the bare metal name and created time to search for operation logs and find an operation log whose API Operation is add on the Operation Logs tab. Click View Details and the Express Connect Circuit ID in the details is the value of this field.
VLAN ID	Fixed value: 10
Trunk ID	You are not required to enter this value.

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, see the method to check the result in Delete a VPC route table entry, Delete a VBR route table entry, Delete a VPC router interface, Delete a VBR, and Delete a physical connection to check if the VPC route table entries, VBR route table entries, VPC router interfaces, VBR router interfaces, VBRs, and physical connections are deleted.

5.2.3.4.10. View physical connection bandwidth

You can view the physical connection bandwidth when the access gateway instance is connected to VPC in the system based on Operations & Maintenance (O&M) needs.

Procedure

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select View Express Connect Bandwidth from the drop-down list in the upper-left corner.

View Express Connect Band	. ~		
Specifications	Bandwidth (bit/s)	Region ID	
Large.1	1G	Router Interface ID *	
Large.2	2G	AK *	
Large.5	5G		
Xlarge.1	10G		
Xlarge.2	20G		Search Clear
Xlarge.4	40G		
Xlarge.5	50G		
Xlarge.8	80G		Express Connect Bandwidth:
Xlarge.10	100G		

4. Configure the filter conditions and then click Search.

Filter condition	Description	
Region ID	The name of the region in the current environment.	
Router Interface ID	Namely the VBR RI. To obtain this value, search for the bare metal on the Bare-Metal Networks tab based on the specific VPC ID and access gateway name, and then click View Details.	
AK and SK	You can obtain the AK and SK on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VBR belongs.	

The system displays the physical connection bandwidth information that meets the filter conditions.

The obtained bandwidth information is the specification of the physical connection bandwidth on HSW of the current VPC. View the table on the left and obtain the bandwidth (bit/s) based on the specification.

5.2.3.4.11. Modify the physical connection bandwidth

In Operations & Maintenance (O&M) emergency scenarios, you can use this function to modify the physical connection bandwidth.

Prerequisites

Notice This operation is only for emergency situations. You must use this function under the guidance of developers. Otherwise, the normal operation of the business will be affected.

Procedure

- 1. In the left-side navigation pane of Apsara Stack Operations (ASO), choose NOC > VPC.
- 2. Click the O&M tab.
- 3. Select Change Express Connect from the drop-down list in the upper-left corner.

Char	Change Express Connect V				
		1 Enter Parameters	2 check		3 Confirm
0					
	Specifications	Bandwidth (bit/s)	Region ID		
	Large.1	1G	Router Interface ID +		
	Large.2	2G	Router Interface Specifications *	Select V	
	Large.5	5G			
	Xlarge.1	10G	AK *		
	Xlarge.2	20G	SK *		
	Xlarge.4	40G			
	Xlarge.5	50G			
	Xlarge.8	80G			
	Xlarge.10	100G			
			Reset		

4. Complete the configurations. For more information, see the following table.

Configuration	Description	
Region ID	The name of the region in the current environment.	
Router Interface ID	The router interface ID that the physical connection bandwidth to be modified corresponds. Configure the VPC ID and access gateway name to search for the bare metal on the Bare-Metal Networks tab. Click View Details and the VBR RI in the details is the value of this field.	
Router Interface Specifications		
AK and SK	The AK and SK of the organization to which the VPC belongs, which can be obtained on the Organizations page of the Apsara Stack Cloud Management (ASCM) console according to the organization to which the VPC belongs.	

? Note If the configured information is incorrect, click Reset at the bottom of the page and configure the information again.

- 5. Then, click Next.
- 6. Check the information. If the information is correct, click **Confirm**. The system starts to push the configurations. After the configurations are pushed, the message Result: Successful appears.

After the configurations are pushed, see View physical connection bandwidth to check if the physical connection bandwidth is modified.

5.2.4. Resource management

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

5.2.4.1. Network elements

Network elements are network devices, including switches and routers. 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.

5.2.4.1.1. 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.

5.2.4.1.1.1. 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.

Procedure

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

? Note 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.
- View the port information, CPU utilization, memory usage, aggregation 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. On the **Port** tab, view the port list, port working status, and other link information of the device.
 - c. On the CPU Utilization tab, view the CPU utilization of the device.
 - d. On the Memory Usage tab, view the memory usage of the device.
 - e. On the Aggregation Port Management tab, view the aggregation port information of the device. At the right of a port, click View in the Operation column to view the water level of the aggregation port.
 - f. On the Alert Info tab, 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 on the **Alert Info** tab, 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.

Network Monitoring			
Device Name : ASW-A3-4-G07.AMTEST15 ち			
Device Name : ASW+3-4-607 ANTEST15 Manufacturer : HSC Online Duration : 181 Days Last Updated Time : Mar 1, 2020, 18:17:54	IP Model : 56800-540F Ping Status : Accessible Release Version : 2432P01	SN Role : ASW SNMP Status : Accessible Software Version : 7.1045	
Port CPU Utilization Memory Usage Aggregation Port Management	Alert Info		
Name		Operation	
Bridge-Aggregation1			
Bridge-Aggregation2			
Bridge-Aggregation3		View	

- View the traffic 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. On the Port tab, find the port that you are about to view, 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.

Network Monitoring		
Port Name : Ter-GigabitEthernet2/0/15 🖕		
Port Name : Tes-GigaldEthenet20015 Admin Status : Up End Port : eth3	Port Speed - 1 0000 Operation Status : Up End Port Alas :	Port Alias : Liai, SERVER-15 End Device : #346077015.doud g07.amlest82 Last Updated Time : Mar 1, 2020, 18:17-07
		03/01/2020 17:25:01 - 03/01/2020 18:25:01 😵 Search
Water Level		
Quick Query : 5MIN 30MIN 1H 6H		

5.2.4.1.1.2. 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 the ASO console.
- 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.

5.2.4.1.1.3. Collection settings

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

5.2.4.1.2. Modify the device password

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

Procedure

1. Log on to the ASO console.

- 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 upper-right 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.

• 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.

5.2.4.1.3. Configuration comparison

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

Procedure

- 1. Log on to the ASO console.
- 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.

5.2.4.2. Fire wall

If the cloud firewall is deployed in your environment, you can use the Fire Wall function to isolate or restore the firewall.

Prerequisites

Notice Confirm with the administrator that the cloud firewall is deployed in your environment. Otherwise, you cannot use the Fire Wall function to isolate or restore the firewall.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Fire Wall.
- 3. Select the operation type, firewall type, and data center from the drop-down lists as required and then click **Confirm**.

You can select one of the following operation types:

- **Isolate Firewall:** physically isolates the firewall from the network structure. If the cloud firewall service has an exception, the system removes the firewall device from the network forwarding path, making sure that the normal business traffic forwarding is not affected by faults.
- **Restore Firewall:** restores the firewall from the network isolated status to the normal status. After the exception of the cloud firewall is recovered, the system adds the firewall device back to the network forwarding path, making sure that the firewall is restored to the initial online status.

One-Click Isolation	amtest88	Confirm		
Select I	Device		Configuration Check	Result Check
_ ichu-1 _ ichu-2				
			Clear Selection Next	

- 4. On the Select Device tab, select the devices and then click Next.
- 5. On the **Configuration Check** tab, check the selected devices and template information. If the information is correct, click **Confirm**.



6. Click **OK** in the displayed dialog box.

Then, the system automatically isolates or restores the firewall in the selected devices based on the configuration template.

The results are automatically displayed on the Result Check tab.

7. On the **Result Check** tab, click **Details** in the **Details** column at the right of each device to view the corresponding result.

One-Click Isolation		
Restore Firewall V ICFW V amtest88 V C	anfirm	
Select Device	Configuration Check	Result Check
Product: idw Operation: Restore Firewall		
Device	Result	Details
		Dura
ISW-VM-G1-1.AMTEST88		Details
ISW-VM-G1-1.AMTEST88		Utdas

8. Click **Complete** to complete this operation.

5.2.4.3. 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.

5.2.4.3.1. 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.

- 1. Log on to the ASO console.
- 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	Instance Monitoring					
cn-qingdao-env4b-d01	Search					
						٩
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		
IP	Status : online	Local IP Address Range :
NIC : dummy0	LVS GROUP ID : 1	Proxy Check Type : http
SITE ID : 1	Active Connection Limit : 10000	Inactive Connection Limit : 0
Inbound Limit (Bit/s): 1048576	Outbound Limit (Bit/s) : No Limit	Inbound Limit (PPS) : 10000
Outbound Limit (PPS) : No Limit		

5.2.4.3.2. 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 the ASO console.
- 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 dropdown 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.

5.2.4.4. 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.

- 1. Log on to the ASO console.
- 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.

5.2.4.5. View Anytunnel information

You can view the Anytunnel information to know the Anytunnel resources registered by projects in the current environment or if a project registers Anytunnel. The system allows you to search for the registration information of Anytunnel resources based on project, cluster, service instance, and server role. You can use the global query function to search for the usage of all the Anytunnel resources in the current environment.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Management Function.
- 3. Then, you can:
 - Click Query Details to view all the Anytunnel information in the environment.



 Select the project, cluster, service instance, or server role, and then click Query AnyTunnel Information to view the Anytunnel information that meets the search conditions.

? Note To modify the search conditions, click Clear Conditions and then select the search conditions again.

5.2.4.6. XGW management

The XGW Management module allows you to view the basic information, running status, aggregated traffic, and water level of each device node of XGW network products.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Resource Management > XGW Management.
- 3. Select the cluster to be viewed from the drop-down list and then click Search. The system displays the basic information, aggregated traffic, and water level of all the device nodes in the selected cluster. By default, the water level information in the last one hour is displayed. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range to search for the water level information.

ode Information				
vpcAzoneCluster-A-2019111 V				
Hostname	IDC		State	Details
043			good	
04			good	
9 1 Hour 🔿 3 Hours 🔿 6 Hours 🔿 1 Day			04/04/2020 20:41:33 - 04/04/2020 21:41:3	3 🛞 Sear
		1bps 0.8bps		
2Mbps	$\sim \wedge \wedge \wedge \wedge$	0.6bps		
2Mbps 2Mbps		0.4bps		
2Mbps	r 4, 2020, 21:06:00 Apr 4, 2020, 21:18:00 Apr 4, 20	0.2bps 0bps 120, 21:30:00 Apr 4, 2020, 20:42:00 Apr 4, 20	20, 20:54:00 Apr 4, 2020, 21:06:00 Apr 4, 2020, 21:18	:00 Apr 4, 2020, 21
	Rate 🔵 outByteRate		packetLossRateIn packetLossRateOut	

- 4. Find a device node and click View in the Details column.
- 5. On the displayed page, view the water level traffic of the device node.

Traffic	
700Kbps 650Kbps 500Kbps 500Kbps 400Kbps 350Kbps 400Kbps 350Kbps 400Kbp	lbps 0.8bps 0.4bps 0.4bps 0.2bps 0.2bps Apr 4, 2020, 204200 Apr 4, 2020, 205400 Apr 4, 2020, 21:06:00 Apr 4, 2020, 21:18:00 Apr 4, 2020, 21:30:00
💿 inByteRate 🕒 cutByteRate	🔵 packetLossRateIn 🌘 packetLossRateOut
1.15kpp 1.05kpp 1.05kpp 900bps 900	
🔵 packetRateln 😑 packetRateOut	

5.2.4.7. 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.

5.2.4.7.1. 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 the ASO console.
- 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.

5.2.4.7.2. 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 the ASO console.
- 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.

5.2.4.7.3. 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 the ASO console.
- 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.

5.2.4.7.4. Export the IP address pool information

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

- 1. Log on to the ASO console.
- 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.

5.2.4.7.5. 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 the ASO console.
- 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.

5.2.5. Alert management

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

5.2.5.1. View and process current alerts

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

Procedure

- 1. Log on to the ASO console.
- 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.

5.2.5.2. View history alerts

You can view history alerts on the History Alerts tab.

- 1. Log on to the ASO console.
- 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.

5.2.5.3. Add a trap

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

Context

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

Procedure

- 1. Log on to the ASO console.
- 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.	-

Configuration	Description	Example
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.
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 information of the alert item, such as a system message or a message indicating the location of the state machine or the current status. Configure the value 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.	-
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:	C Event Alert None	
Тгар Туре:	Select V				
Trap Index:		+ Submit	Trap Msg:		e

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**.

5.2.5.4. View a trap

You can view a trap configured in the current system.

Procedure

- 1. Log on to the ASO console.
- 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 upper-right 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.

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

5.2.6. Network reconfiguration

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

5.2.6.1. 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.

- 1. Log on to the ASO console.
- 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.
 - i. 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.
 - ii. 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.
 - iii. 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.
 - i. In the Configure Interfaces step, click Edit.
 - ii. 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.
 - iii. Click Save Project in the upper-right corner to save the information.
- 9. Click Next.
- 10. Configure the route parameters.
 - i. In the Configure Routes step, click Edit.
 - ii. 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.
 - iii. Click Save Project in the upper-right corner to save the information.
- 11. Click Next.
- 12. Configure the route policies.
 - i. In the Configure Route Policies step, click Edit.
 - ii. 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.
 - iii. Click Save Project in the upper-right corner to save the information.
- 13. Click Next.

14. In 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.

5.2.6.2. 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 scale-up parameters, Network Operation Center (NOC) automatically generates the configuration and pushes the configuration to a specific device for automatic scale-up.

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 the ASO console.
- 2. In the left-side navigation pane, choose NOC > Network Reconfiguration > ASW Scale-up.
- 3. Select devices to be implemented.
 - i. 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.
 - ii. 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.
 - i. In the **Disable DSW Port** step, click **Port Settings** at the right of the device to be implemented.
 - ii. Disable the corresponding port and then click Implement.
 - iii. In the displayed dialog box, click OK to run the script commands.
- 6. Click Next.
- 7. Configure the DSW ports.
 - i. In the **Configure DSW Port** step, click **Edit** at the right of the device to be implemented. The **Interface Parameter Configuration** list is displayed.

- ii. 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.
- iii. After adding the interface parameter, click Implement at the right of the device.
- iv. 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.
 - i. In the **Configure BGP** step, click **Edit** at the right of the device to be implemented. The **Interface Parameter Configuration** list is displayed.
 - ii. 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.
 - iii. After adding the interface parameter, click Implement at the right of the device.
 - iv. 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.
- 13. Enable the DSW ports.
 - i. In the Enable DSW Port step, click Port Settings at the right of the device to be implemented.
 - ii. Enable the corresponding port and then click Implement.
 - iii. In the displayed dialog box, click **OK** to run the script commands.
- 14. Click Next.
- 15. Perform the scale-up test.
 - i. In the **Test Scale-up** step, click **Select** at the right of the device to be implemented. The route table is displayed.
 - ii. 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.
 - iii. Click Test and then the system returns the test results.

5.2.6.3. Push IPv6 configurations

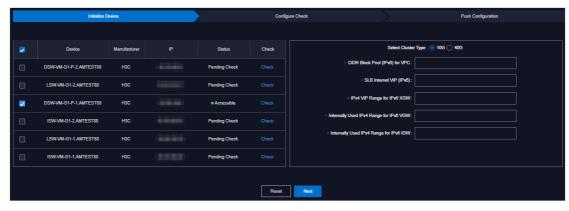
The system can automatically push IPv6 configurations. After network operations engineers configure the IPv6 parameters in the IPV6 Configuration module, the system automatically generates the IPv6 configurations and pushes the configurations to specific devices.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Network Reconfiguration > IPV6

Configuration.

- 3. On the **Initialize Device** step, select the devices to be configured and complete the configurations to complete the initialization.
 - i. Find the device to be configured in the device list and click **Check** in the **Check** column to check if the device is accessible. You can check multiple devices to be configured.
 - ii. Select one or more devices whose Status is Accessible after you click Check.
 - iii. Complete the configurations on the right.

Configuration	Description
Select Cluster Type	10G or 40G . Select the option according to the planned cluster type.
CIDR Block Pool (IPv6) for VPC	The VPC CIDR block pool in the format of IPv6.
SLB Internet VIP (IPv6)	The SLB Internet VIP address in the format of IPv6.
IPv4 VIP Range for IPv6 XGW	The XGW VIP address range in the format of IPv4.
Internally Used IPv4 Range for IPv6 VGW	The internally used IP address range for VGW in the format of IPv4.
Internally Used IPv4 Range for IPv6 IGW	The internally used IP address range for IGW in the format of IPv4.



4. Click Next.

5. On the **Configure Check** step, check the configurations.

During the configuration check, the system automatically checks the current configurations of the selected devices and generates the IPv6 configuration script based on the check results. Click **View** at the right of the script file to view the generated configuration script, or click **Download** to download the configuration script to your local computer.

Note If you select multiple devices on the Initialize Device step, click Batch
 Download to download multiple configuration scripts to your local computer at a time.

Generally, the following results appear during the configuration check.

- The configuration is generated. Pending Pushing
- $\circ~$ Failed to check the configuration. No BGP processes have been found.
- $\circ\;$ Failed to check the configuration. Failed to generate the configuration.
- $\circ~$ Failed to check the configuration. The IPv6 configuration already exists.
- 6. Click Next.

The system checks if the configuration pushing function is enabled. If not, the message Contact the onsite manager to enable the function before you continue appears. If yes, check if the pushing condition is met based on the configuration check results and generation conditions of IPv6 configuration scripts.

- If the configuration check is successful and the IPv6 configuration scripts are generated in the previous step, a dialog box appears. Click **Continue** to automatically push the configuration scripts to the selected devices.
- If the configuration check result is Failed to check the configuration. No BGP processes have been found., Failed to check the configuration. Failed to generate the configuration., or Failed to check the configuration. The IPv6 configuration already exists. in the previous step, a dialog box appears, and the system does not automatically push the configurations.
- 7. After the configurations are pushed, view the pushing result on the **Push Configuration** step.

If the system indicates that it is pushing the configurations, click **Refresh** to refresh the pushing result.

After the configurations are pushed, click **View** to view the current running configurations of the selected devices to check if the IPv6 configurations are pushed.

5.2.7. Fault check

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

5.2.7.1. 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 the ASO console.
- 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.

5.2.7.2. 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 the ASO console.
- 2. In the left-side navigation pane, choose NOC > Fault Check > Leased Line Discovery.
- 3. Select a discovery source.
 - i. 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.
 - ii. 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.
 - i. In the Configure Parameters step, click Edit. The Configure Parameters list is displayed.
 - ii. 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.
- 10. In 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.
- 14. In 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.

5.2.7.3. 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 the ASO console.
- 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.
 - i. 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.
 - ii. 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.

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

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

7. Click Next.

- 8. Select check items.
 - i. 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.

- ii. Click Save Project in the upper-right corner to save the current information.
- 9. Click Next.
- 10. In 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.

5.2.7.4. 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 the ASO console.
- 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.

5.2.8. Use case

5.2.8.1. Troubleshoot network failures

This topic uses a common use case to tell people how to use the NOC module to troubleshoot network failures.

Scenario

> Document Version:20200918

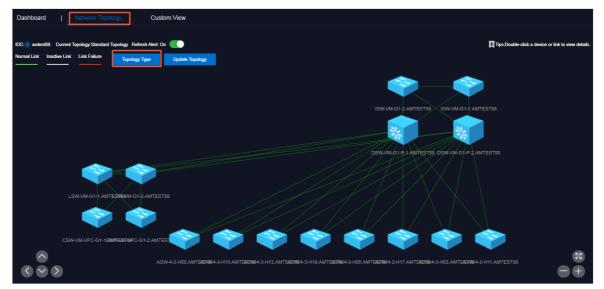
If a cloud product has latency of visits and number of retransmissions increased, you must make sure whether this is caused by network failures.

Procedure

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Dashboard.
- 3. Click the Network Topology tab.
- 4. Click Topology Type and select Standard Topology.

Wait five seconds. After the page finishes loading, the system displays the network-wide topology and device connections of the AZ in the current environment.

If device alerts are not triggered in the network, the device icon is blue, the link between devices is green, and the device name is white in the topology. If device alerts are triggered in the network, the topology updates the alert information in the current network every five seconds and displays the updated alert information.



5. If the device name or the link becomes red in the topology, it indicates that alerts are detected in the network device or link port. Double-click the device icon whose device name becomes red, and then you can view the basic information of this device and the network alert information related to the device on the right.

Network De	evice Informa	ation	
Device Name	DSW-VM-G1-P-1.	AMTEST88	
IP			
Role Node Alerts	DSW		
Node Alerts			
Alert Time	Alert Name	Alert Item	Alert Details
72	linkDown	FortyGigE1 0/20	/ Details
100	bgpBackwa rdTransNoti fication		Details
72	linkDown	Ten-Gigabi Ethernet0/ 0/2:2	

In this example, the port connected to this DSW has a linkDown alert and a bgp peer down alert. An ASW is identified based on the bgp peer IP address. Therefore, you can determine that a link between DSW and ASW has a problem, which causes the port down and triggers the alerts.

6. Click the red link in the topology, and then you can view one or more actual physical links contained in the logical link and the alert information of the link between devices on the right.

Link Status	;		×
Links			
Source Device	Source Port	Destination Device	Destination Port
	Ten-Gigabit		Ten-Gigabit
Alerts			
Alert Time	Alert Name	Alert Item	Alert Details
	linkDown	Ten-Gigabit Ethernet0/C/ 2:2	Details

In this example, the logical link connected to the two devices contains four actual end-toend links. The port 0/0/2:2 has a port **linkDown** alert. Then, you can continue to log on to the device and check if this is caused by the low optical power or damaged module.

7. After the problem involved with the preceding alerts is solved, the system automatically updates the alert information. If the fault is repaired, the alert automatically disappears and the topology is restored to the normal status, namely no device name or link is red.

Use the Alert Management module to troubleshoot the problem as a supplement

If the device name or link in the topology becomes red, namely a problem exists in the network device or link, you can choose NOC > Alert Management > Alert Dashboard and view the current alerts that are not recovered in the network on the Current Alerts tab.

The Current Alerts tab displays more detailed alert information.

If an alert is for test or generated within the plan because of cutover, you can click **Ignore** or **Delete** in the **Actions** column at the right of the alert to ignore or completely delete the alert.

Use the syslog log query tool to troubleshoot the problem as a supplement

If the device name or link in the topology becomes red, namely a problem exists in the network device or link, and you confirm that the device alert is not an expected one or within the plan because of cutover by using the Alert Management module, you must view the detailed exception logs. Here, use the syslog log query tool of the switch to search for logs.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.
- 3. On the Device Management tab, click the Syslogs tab.
- 4. In the upper-right corner, select the device and time range you are about to view and then click **Search** to search for logs within the selected time range. By default, you can search for

at most 1000 logs.

- 5. In the upper-left corner, you can enter the keyword in the search box and then click the search icon to search for specific logs in the search results.
- 6. After the search, to export logs to open a ticket or submit logs to the device manufacturer for location, click **Export to CSV** in the upper-right corner to save logs as a .csv file to your local computer.

5.3. Task Management

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

5.3.1. 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.

5.3.2. 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.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Task Management > Task Overview.

Pending for Intervention			Completed	Task Name	Task Name Task Description		Ac	Actions	
1 ≖	0 🖬	2 👩	6⊘	test01	test	Dec 30, 20	19, 10:59:45 Dei		
Create Task			Create Task						
Running Status in Last 7 Days									
1									
0.8				Running Tasks(Running tir	me more than 1 day)				
0.6				Task Name	Task Description	Target Group	Start Time	Running Duration	
0.4									
0.2									

- 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.

 $\circ~$ 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.

5.3.3. Create a task

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

- 1. Log on to the ASO console.
- 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.

Configuration	Description
	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.
Target Group	 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 Z 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 .
	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
Execution Batch	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
	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.
Execution Method	• 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.
	Click Add Script. Select one or more .tar packages to upload the script file. After the upload, you can delete and reupload the script.
Add 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 🚄					
aso 🗸	vm010004024196 ×	vm010004028255 ×	n010004021104 ×	vm010004028003 \times	~
	vm010004020034 \times	vm010004029054 \times	n010004021096 ×	vm010004024238 \times	
Execution Batch @) Single-Machine	Order 🐵			
Manual Execution					
+Add Script Supported Extension	c .tar				
					Create

5. Then, click Create.

Result

The created task is displayed in the task list.

5.3.4. 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 the ASO console.
- 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.

? Note 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 Status 🗸 Start Date - 6	End 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, 2019, 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, 10:51:10 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			Modify Start Delete

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			vm01000402	0034	
		vm010004028003			vm01000402	9054	
					vm01000402	1096	
Batch5							
vm010004024236							
						Ci	ose

5.3.5. 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 the ASO console.
- 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 Sta	atus 🗸 Start Date - En	nd Date 📾 Query	Create		
Task Name	Task Description	Time	Task Status	Target Group	Actions
test03		Created At: Dec 30, 2019, 14:34:17 Start Time :Dec 30, 2019, 14:39:47 End Time :Dec 30, 2019, 14:40:08			Modify Start Delete
vest02		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, 10:59:45 Start Time : Dec 30, 2019, 14:29:30 End Time :	Pending for Intervention		

5.3.6. Delete a task

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

Procedure

- 1. Log on to the ASO console.
- 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.

5.3.7. 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.

Procedure

- 1. Log on to the ASO console.
- 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	

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

5.4. 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.

5.4.1. Log configurations

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

5.4.1.1. Clear

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

5.4.1.1.1. Configure parameters for automatic log

clearance

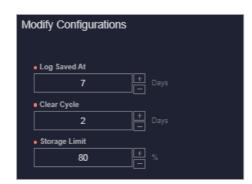
To avoid logs filling up the disk space, the system supports automatically clearing logs. You can configure the parameters for automatic log clearance based on business needs.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Log Management > Deploy. By default, you are on the Clear tab.

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

- 3. In the Actions column, click Modify.
- 4. On the displayed page, complete the following configurations.

Configuration	Description
Log Saved At	The system clears logs saved before the configured time.
Clear Cycle	The system automatically clears logs according to this cycle.
Storage Limit	If 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.



5. Click OK.

5.4.1.1.2. 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 the ASO console.
- 2. In the left-side navigation pane, choose Log Management > Deploy. By default, you are on the Clear tab.

Clear Project Agent Bucket Ma	nagement		
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.

5.4.1.2. Project

You can add or delete projects on the **Project** tab.

5.4.1.2.1. 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.

- 1. Log on to the ASO console.
- 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.

⑦ Note 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.

5.4.1.2.2. Delete a project

You can delete one or more projects.

Procedure

- 1. Log on to the ASO console.
- 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.

5.4.1.3. Agent

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

5.4.1.3.1. 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 the ASO console.
- 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.

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

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 (,).
Dath	
	 csv: Commas (,) are used to separate properties and the first property is instanceId. For example,
	uw905d8ny00drzx9****,2019-08-15 00:00:07,15.75.128.85,64 05,0,0,15.74.181.5

Product		
Select Product		
• Cluster		
Select Cluster		
• Service		
Select Service		
Service Role		
Select Service		~
 Path(Multiple paths are supported three paths are supported.) 	I. You can split the paths by c	ommas (,). A maximum of

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

5.4.1.3.2. 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 the ASO console.
- 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.



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

5.4.1.3.3. Delete an agent

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

- 1. Log on to the ASO console.
- 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.

5.4.1.4. Bucket management

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

5.4.1.4.1. 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 the ASO console.
- 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 <i>OSS Developer Guide</i> .
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 <i>OSS Developer Guide</i> .
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.

5.4.1.4.2. 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 the ASO console.
- 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 <i>NAS Developer Guide</i> .
Path	The path on the NAS server, which is used to store the log backup file.

7. Then, click OK.

5.4.1.4.3. FTP configurations

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

Procedure

- 1. Log on to the ASO console.
- 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.

5.4.2. 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 the ASO console.
- 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.

5.4.3. Log export

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

5.4.3.1. 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.

- 1. Log on to the ASO console.
- 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.

5.4.3.2. View tasks

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

Procedure

- 1. Log on to the ASO console.
- 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.

5.4.4. Log clearance

The Log Clearance module allows you to clear logs in a specific log file of containers or servers specified in the system.

5.4.4.1. 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.

5.4.4.1.1. 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.

- 1. Log on to the ASO console.
- 2. In the left-side navigation pane, choose Log Management > Log Clearance.

Containers	Servers							
Product Select Product	Cluster V Select C	Service uster V Select Service	Service Role Select Service Role	✓ Search	Add	Obtain Watermarks	Xear 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.

5.4.4.1.2. Add a log clearance rule

You can add a clearance rule for a specific log file in the container as needed.

Procedure

- 1. Log on to the ASO console.
- 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.

Configuration	Description
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.

5.4.4.1.3. Modify a log clearance rule

You can adjust a configured log clearance rule in time based on business needs.

Procedure

- 1. Log on to the ASO console.
- 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.

5.4.4.1.4. 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 the ASO console.
- 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**.

5.4.4.1.5. 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 the ASO console.
- 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.

5.4.4.1.6. View clear records

After clearing logs, you can view the detailed clear records.

Procedure

- 1. Log on to the ASO console.
- 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.

5.4.4.2. 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.

5.4.4.2.1. 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 the ASO console.
- 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.

5.4.4.2.2. 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 the ASO console.
- 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.

5.4.4.2.3. Modify a log clearance rule

You can adjust a configured log clearance rule in time based on business needs.

- 1. Log on to the ASO console.
- 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.

5.4.4.2.4. 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 the ASO console.
- 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**.

5.4.4.2.5. 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 the ASO console.
- 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.

5.4.4.2.6. View clear records

After clearing logs, you can view the detailed clear records.

- 1. Log on to the ASO console.
- 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.

5.4.4.3. 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 the ASO console.
- 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.

5.4.4.4. 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.

Procedure

- 1. Log on to the ASO console.
- 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.

5.5. Apsara Stack Doctor (ASD)

5.5.1. Apsara Stack Doctor introduction

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

Basic features

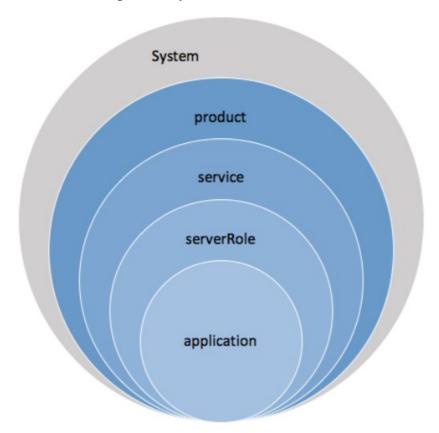
- Provides data filtering, analysis, and processing for O&M data consumers.
- Provides encapsulation, orchestration, and rights management of O&M operations.
- Provides O&M experience accumulation and archiving capabilities.
- Provides troubleshooting, pre-diagnosis, health check, and early warning capabilities.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

Benefits

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

Terms

Apsara Stack has five levels of release granularity, as shown in Levels of release granularity.



Levels of release granularity

• system

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

• product

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

• service

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

• server role (sr)

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

• application (app)

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

- docker: a Docker image that is built from source code.
- $\circ\;$ 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.

5.5.2. Log on to Apsara Stack Doctor

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

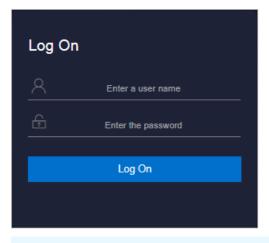
Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

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



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

3. Enter your username and password.

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

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

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

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

5.5.3. ASA

ASA is a tool provided to help you improve the efficiency in testing, operating, maintaining, and releasing cloud products in Apsara Stack while ensuring the stability of version qualities. ASA retains the features of Apsara Stack V2, including inspection, scanning, and version tracking. This continues and precipitates all the long-term experience of Apsara Stack.

5.5.3.1. 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.

- 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

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.

5.5.3.2. 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.

- 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
	9090	21069		ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal

Parameters on the Virtual IP Check page

Description
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. o normal: indicates that the virtual IP address is correctly bound to the IP address of the backend service. o abnormal: indicates that the virtual IP address is not bound to the backend IP address is not bound to the backend IP address properly.

5.5.3.3. 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.

Operations and Maintenance Guide · Operations tools

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

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.

5.5.3.4. 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

Parameters on the NTP Check page

Parameter	Description
Host	The name of a host.
Time Offset	The time offset. Unit: milliseconds.

5.5.3.5. IP Conflict Check

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

Procedure

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose ASA > IP Conflict Check.

IP	Physical Host	Server Role	Туре	Virtual Host

Parameters on the IP Conflict Check page

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

5.5.3.6. 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

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.

5.5.3.7. 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.

- 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

Parameters on the IP Details page

Parameter	Description
IP	The IP address of a resource.
Virtual Host	The 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.

5.5.3.8. 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.DrdsManager,drds-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.

5.5.3.9. 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.

Procedure

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

5.5.3.10. 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.
 - $\circ~$ Click the Version Tree tab to view information related to version trees.

5.5.4. Support tools

5.5.4.1. Diagnose with the OS tool

The OS tool allows you to perform OS diagnostics on physical machines in Apsara Stack.

Context

The OS tool allows you to diagnose the following metrics: disk file metadata usage, memory usage, process statuses, time synchronization, kernel errors, 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, unusual network packets, TCP connection status exceptions, TCP queue exceptions, network packet loss, bonding exception, NIC exception, SN retrieval exceptions, OOB IP 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. The section below the search bar displays the physical machines.
- 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 completed, click **View Report** in the **Actions** column to view the diagnostic result.
- 8. (Optional)For more information, click View Result or Download Report in the Actions column.

5.5.4.2. Use Support Tools

Support Tools allows you to diagnose some services and export diagnostic reports.

- 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. The following table lists the supported diagnostic items.

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.
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.

? Note 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 S	earch by hostname or IP add	Iress. 🛞 Se	arch 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.

5.5.4.3. Update Support Tools

When the Support Tools toolkit has updates, you can update it to the latest version by uploading files.

Procedure

- 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	
File: L Choose	File	
Verification Code:	Enter a verification code.	
	Upload File	

5.5.4.4. Diagnose with inspection tools

You can use inspection tools to diagnose and inspect services, such as Apsara File Storage NAS (NAS), Block Storage, and Apsara Name Service and Distributed Lock Synchronization System.

> Document Version:20200918

- 1. Log on to Apsara Stack Doctor.
- 2. In the left-side navigation pane, choose Support Tools > Inspection Tool.
- 3. Select the target service from the Product drop-down list and click Search. The search result appears in the section below. Apsara Stack Doctor (ASD) supports diagnostics for services, including NAS, Block Storage, and Apsara Name Service and Distributed Lock Synchronization System.
 - NAS diagnostics

It allows you to collect NAS information, including disk status, KV (key-value) status, KV server spacing, version, recycle bin, memory, and TCP.

• EBS diagnostics

It allows you to collect the utilization information about storage clusters.

o Diagnostics of Apsara Name Service and Distributed Lock Synchronization System

It allows you to check the following information about this service:

- The health status of the E2E service link.
- The disk space of the service.
- Whether the nuwazk log is properly stored.
- Whether the nuwaproxy log is properly stored.
- 4. 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.

Produc	t: nas	✓ Search	Get Gateway List	Run Diagno	stics	
-	Admin Gateway	IP	Diagnostics Execution	Status	Executed At	Actions
	vm010017051065					Run Diagnostics Download Inspection Log

5. After the diagnostics is complete, you can click **Download Report** in the **Actions** column corresponding to the machine to download the diagnostic results to your local machine.

5.5.4.5. 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	~		
* Service :	edas-hsf	~		
Server Role :	Hsflnit	~		
Existing Scripts:				
* Script File :	⊥ 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.

5.5.4.6. 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.

Procedure

- 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.

? Note 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.

roduct: e	edas V	Service: Select an iten	n. 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.

5.5.5. Service Availability

5.5.5.1. 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.

- 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. The following table describes the service statuses.

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.
Service Availability	
Product: All V Search	

Product: All V Search							
Product	Product Status	Exception Message	Checkpoint Time				

5.5.5.2. 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.

5.5.6. Monitoring

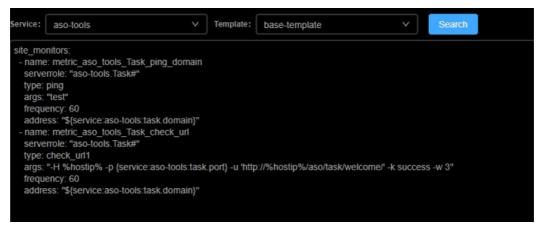
The Monitoring module allows you to view alert templates, alerts, and alert status in the system.

5.5.6.1. View alert templates

Alert templates are used to configure alert monitoring settings. You can filter alert template content by service and template.

Procedure

- 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.



5.5.6.2. 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	m	V Cluster:	Enter a cluster na	me	Alert: Ente	er an alert nam	ne	⊗ Search
Alert Name	Metric	Alert Rule	Monitoring Dimension	Subject	Data Source	Enabled	Monitor Interval	Alert Status
private.C_dnsCluster-A- 20191021-165c_dnsSer vice_timon-service_avail able_alarmdnsService _base-template	dnsService_tjm on-check_servic e_available_clu ster_serverrole	{"name":"dnsService_tim on-service_available_ala mr","template":"base-tem plate"}	["{ \"cluster \":\"\$\$CLUSTER \$\$\", \"serverrole \":\"dnsService.b indSererRole#\" }"]	private. service not available ala rm	sourceType:METRIC project.tjm_dnsService	true	60	INSUFFICIENT_DATA
private.C_dnsCluster-A- 20191021-165c_dnsSer vice_tjimon-max_open_fil e_alarmdnsService_b ase-template	dnsService_tim on-check_log_k eyword_max_o pen_file_serverr ole	("name":"dnsService_tjm on-max_open_file_alar m","template"."base-tem plate"}	["{ \"cluster \":\"\$\$CLUSTER \$\$\", \"serverrole \":\"dnsService.b indServerRole# \" }"]	Alarm-02.659.0 002.00006	sourceType:METRIC project.tjm_dnsService	true	60	INSUFFICIENT_DAT/

• 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.

5.5.6.3. View the alert status

After alerts are triggered, you can view the status of all alerts in the system.

Procedure

- 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 cluster	r name 🛞 Alert:	Enter an alert name	Alert Status: Select a st	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-template"}	serverrole=drds-console.ServiceTest#, machine=vm010148065201,level=erro r	
private.testimage_monitor_alarm_tia 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.

5.6. Apsara Infrastructure Management Framework

5.6.1. Old version

5.6.1.1. What is Apsara Infrastructure Management

Framework?

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

5.6.1.1.1. Overview

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

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

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

Core functions

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

5.6.1.1.2. Basic concepts

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

project

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

cluster

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

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

service

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

service instance

A service that is deployed on a cluster.

server role

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

server role instance

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

application

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

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

rolling

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

service configuration template

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

associated service template

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

final status

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

dependency

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

upgrade

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

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

5.6.1.2. Log on to Apsara Infrastructure Management

Framework

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

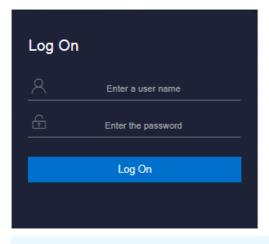
Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

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



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

3. Enter your username and password.

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

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

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

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

5.6.1.3. Web page introduction

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

5.6.1.3.1. Introduction on the home page

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

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

🐟 Tian Ji @ 15:21 @ a #C 3 Most-used Reports · Registration Vars of 4 · Resource Apply Rep Machine Info Report Running 100.00% Paused 0.00% 0.009 · IP LIST JVM Monitor - Clu ers 83 TOP 15 127 WIN Error Alarms OS Errors H/W Errors Error Summary Rate of Abnormal Servers: 0.00% 30 Server A... 0.00% OS Errors: 0.00% H/W Err... 0.00% Rate of Abnormal ServerRole Instances: 1.06% Alarms: 1.06% normal: 46

Home page of Apsara Infrastructure Management Framework

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

Descriptions of functional areas

	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.
Top navigation	 Machine Operations: maintains and manages all the machines in Apsara Infrastructure Management Framework, such as viewing the machine status.
bar	• Tasks: A rolling task is generated after you modify the configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects.
	 Reports: displays the monitoring data in tables and provides the function of searching for different reports.
	• Monitoring: effectively monitors metrics in the process of system operation and sends alert notifications for abnormal conditions. This menu includes the functions of displaying alert status, modifying alert rules, and searching for the alert history.
	•

Area		Description
2	Function buttons in the upper- right corner	 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. aliguntest - : 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 Introduction on the left-side navigation pane.
4	Home page	 Displays the summary of related tasks or information as follows: Upgrade Task Summary: the numbers and proportions of running, rolling back, and paused upgrade tasks. Cluster Summary: the numbers of machines, error alerts, operating system errors, and hardware errors for different clusters. Error Summary: the metrics for the rate of abnormal machines and the rate of abnormal server role instances. Most-used Reports: links of the most commonly used statistics reports, which facilitates you to view the report information.
5	Button used to collapse/e xpand the left-side navigation pane	If you are not required to use the left-side navigation pane when performing O&M operations, click = to collapse the left-side navigation pane and increase the space of the content area.

5.6.1.3.2. 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 **T** 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 **T** 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 **T** 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 **i** at the right of a service instance and then perform operations on the service instance as instructed.

Report

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

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

• Enter the report name in the search box to search for the report quickly. Fuzzy search is

supported.

• Click All Reports or Favorites to display groups of different categories in the lower-left corner. Double-click a group to view all the reports in this group. Double-click a report to view the report details on the right pane.

5.6.1.4. Cluster operations

This topic describes the actions about cluster operations.

5.6.1.4.1. View cluster configurations

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

Procedure

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

The Cluster Operations page displays the following information:

• Cluster

The cluster name. Click the cluster name to go to the Cluster Dashboard page.

• Scale-Out/Scale-In

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

• Abnormal Machine Count

The statistics of machines whose status is not Good in the cluster. Click the link to go to the Cluster Operation and Maintenance Center page.

• Final Status of Normal Machines

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

• Rolling

Displays whether the cluster has a running rolling task. Select **Rolling Tasks** to display clusters that have rolling tasks. Click the link to go to the **Rolling Task** page.

- 3. (Optional)Select a project from the **Project** drop-down list and/or enter the cluster name in the **Cluster** field to search for clusters.
- 4. Find the cluster whose configurations you are about to view and then click **Cluster Configuration** in the **Actions** column. The **Cluster Configuration** page appears.

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

Cluster configurations

Category	ltem	Description
	Cluster	The cluster name.

Category	Item	Description
	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.
Basic Information	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.
	Service	The service deployed in the cluster.
Deployment Plan	Dependency Service	The service that the current service depends on.
	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.
Service Information	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.

Category	ltem	Description
	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.

5.6.1.4.2. 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 **1** 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.

ltem

Description

ltem	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. Actual Services: the number of services that are actually deployed in the cluster. Actual Server Roles: the number of server roles that are actually deployed in the cluster. Cluster Status: whether the cluster is starting or shutting down machines. 		
Machine Status Overview	The statistical chart of the machine status in the cluster.		
Machines in Final Status	The numbers of machines that reach the final status and those that do not reach the final status in each service of the cluster.		
Load-System	The system load chart of the cluster.		
CPU-System	The CPU load chart.		
Mem-System	The memory load chart.		
Disk_usage-System	The statistical table of the disk usage.		
Traffic-System	The system traffic chart.		
TCP State-system	The TCP request status chart.		
	The chart of TCP retransmission amount.		
TCP Retrans-System	The chart of TCP retransmission amount.		

ltem	Description
Service Instances	 Displays the service instances deployed in the cluster and the related final status information. Service Instance: the service instance deployed in the cluster. Final Status: whether the service instance reaches the final status. Expected Server Roles: the number of server roles that the service instance expects to deploy. Server Roles In Final Status: the number of server roles that reach the final status in the service instance. Server Roles Going Offline: the number of server roles that are going offline in the service instance. Actions: Click Details to go to the Service Instance Information Dashboard page. For more information about the service instance dashboard.
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. Start Time: the time to start the rolling. End Time: the time used for the upgrade. Actions: Click Details to go to the Rolling Task page. For more information about the rolling task, see View rolling tasks.
Cluster Resource Request Status	 Version: the resource request version. Msg: the exception message. Begintime: the start time of the resource request analysis. Endtime: the end time of the resource request analysis. Build Status: the build status of resources. Resource Process Status: the resource request status in the version.

ltem	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.
	• Refer Version List : the version that uses the resource. The information of virtual machines in the cluster. Data is available
VM Mappings	 only when virtual machines are deployed in the cluster. VM: the hostname of the virtual machine. Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed. Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.
Service Dependencies	 The dependencies of service instances and server roles in the cluster, and the final status information of the dependent service or server role. Service: the service name. Server Role: the server role name. Dependent Service: the service on which the server role depends. Dependent Server Role: the server role on which the server role depends. Dependent Cluster: the cluster to which the dependent server role belongs. Dependency in Final Status: whether the dependent server role reaches the final status.

5.6.1.4.3. 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.

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

3. View the information on the **Cluster Operation and Maintenance Center** page.

ltem	Description
Head Version Analysis	 The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses: Preparing: No new version is available now. Waiting: The latest version is found. The analysis module has not started up yet. Doing: The module is analyzing the application that requires change. done: The head version analysis is successfully completed. Failed: The head version analysis failed. The change contents cannot be parsed. If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version. Click the status to view the relevant information.
Service	Select a service deployed in the cluster from the drop-down list.
Server Role	Select a server role of a service in the cluster from the drop-down list. O 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.

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

5.6.1.4.4. View the service final status

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

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. You have two ways to go to the Service Final Status Query page:
 - In the left-side navigation pane, click the C tab. Move the pointer over 🛐 at the right of a

cluster and then choose Monitoring > Service Final Status Query.

- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, choose Monitoring > Service Final Status Query in the Actions column at the right of a cluster.
- 3. View the information on the Service Final Status Query page.

ltem	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.
	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.
Head Version Analysis	• Waiting: The latest version is found. The analysis module has not started up yet.
	• Doing: The module is analyzing the application that requires change.
	• done: The head version analysis is successfully completed.
	• Failed: The head version analysis failed. The change contents cannot be parsed.
	If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.
Cluster Rolling Status	Displays the information of the current rolling task in the cluster, if any. The rolling task may not be of the head version.
Cluster Machine Final Status Statistics	The status of all machines in the cluster. Click View Details to go to the Cluster Operation and Maintenance Center page and view the detailed information of all machines. For more information about the cluster operation and maintenance center, see View the cluster operation and maintenance center.

ltem	Description
	The final status of cluster service version.
Final Status of Cluster SR Version	⑦ Note Take statistics of services that do not reach the final status, which is caused by version inconsistency or status exceptions. If services do not reach the final status because of machine problems, go to Cluster Machine Final Status Statistics to view the statistics.
Final Status of SR Version	The number of machines that do not reach the final status when a server role has tasks.

5.6.1.4.5. View operation logs

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

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. You have two ways to go to the Cluster Operation Logs page:
 - In the left-side navigation pane, click the C tab. Move the pointer over 🛐 at the right of a

cluster and then choose Monitoring > Operation Logs.

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

The differential file list is displayed.

iv. Click each differential file to view the detailed differences.

5.6.1.5. Service operations

This topic describes the actions about service operations.

5.6.1.5.1. 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.

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

5.6.1.5.2. View the service instance dashboard

The service instance dashboard allows you to view the basic information and statistics of a service instance.

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

ltem	Description
Item Service Instance Summary	 Description 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 monitoring system template used by the service instance.
Server Role Statuses	The statistical chart of the current status of server roles in the service instance.
Machine Statuses for Server Roles	The status statistics of machines where server roles are located.
Service Monitoring Information	 Monitored Item: the name of the monitored item. Level: the level of the monitored item. Description: the description of the monitored contents. Updated At: the time when the data is updated.
Service Alert Status	 Alert Name Instance Information Alert Start Alert End Alert Duration Severity Level Occurrences: the number of times the alert is triggered.

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

5.6.1.5.3. View the server role dashboard

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

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

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

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

ltem	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 Details to go to the Machine Details page. For more information about the machine details, see View the machine dashboard. Click Machine System View to go to the Machine Info Report page. For more information about the machine info report, see Machine info report. Click Machine Operation to restart, out of band restart, or clone the machine again.
Server Role Monitoring Information of Machines	 Updated At: the time when the data is updated. Machine Name: the machine name. Monitored Item: the name of the monitored item. Level: the level of the monitored item. Description: the description of the monitored item.
VM Mappings	 The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster. VM: the hostname of the virtual machine. Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed. Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.

ltem	Description
	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.
Service Dependencies	• Dependent Server Role: the server role on which the server role depends.
	• Dependent Cluster: the cluster to which the dependent server role belongs.
	• Dependency in Final Status: whether the dependent server role reaches the final status.

5.6.1.6. Machine operations

This topic describes the actions about machine operations.

5.6.1.6.1. View the machine dashboard

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

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side navigation pane, click the C tab.
- 3. (Optional)On the Machine tab in the lower-left corner, enter the machine name in the search box. Machines that meet the search condition are displayed.
- 4. Move the pointer over **T** 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.

ltem	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.

ltem	Description
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. CPUs: the number of CPUs for the machine. Disks: the disk size. Memory: the memory size. Manufacturer: the machine manufacturer. Model: the machine model. os: the operating system of the machine. part: the disk partition.
Server Role Status of Machine	The distribution of the current status of all server roles on the machine.
Machine Monitoring Information	 Monitored Item: the name of the monitored item. Level: the level of the monitored item. Description: the description of the monitored contents. Updated At: the time when the monitoring information is updated.

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

5.6.1.7. Monitoring center

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

5.6.1.7.1. Modify an alert rule

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

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

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

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

5.6.1.7.2. View the status of a monitoring instance

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

Procedure

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

5.6.1.7.3. View the alert status

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

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

ltem	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.

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

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

ltem	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.

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

5.6.1.7.5. 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.

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

ltem	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 P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members that are notified when an alert is triggered.
Actions	Click Show to show the data before and after the alert time.

5.6.1.8. Tasks and deployment summary

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

5.6.1.8.1. View rolling tasks

You can view running rolling tasks and the corresponding status.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the top navigation bar, choose **Operations** > **Cluster Operations**.
- 3. Select Rolling Tasks to display clusters with rolling tasks.
- 4. In the search results, click rolling in the Rolling column.
- 5. On the displayed **Rolling Task** page, view the information in the **Change Task** list and **Change Details** list.

Change Task list

ltem	Description
Change Version	The version that triggers the change of the rolling task.
Description	The description about the change.
	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.
Head Version	• Waiting: The latest version is found. The analysis module has not started up yet.
Analysis	• Doing: The module is analyzing the application that requires change.
	• done: The head version analysis is successfully completed.
	 Failed: The head version analysis failed. The change contents cannot be parsed.
	If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.
Blocked Server Role	Server roles blocked in the rolling task. Generally, server roles are blocked because of dependencies.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.

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

Change Details list

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

5.6.1.8.2. View running tasks

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

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

5.6.1.8.3. View history tasks

You can view the historical running conditions of completed tasks.

Procedure

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

5.6.1.8.4. 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.

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the top navigation bar, choose Tasks > Deployment Summary.
 - View the deployment status and the duration of a certain status for each project.
 - Gray: wait to be deployed. It indicates that some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed.
 - Blue: being deployed. It indicates that the project has not reached the final status for one time yet.
 - Green: has reached the final status. It indicates that all clusters in the project have reached the final status.
 - Orange: not reaches the final status. It indicates that a server role does not reach the final status for some reason after the project reaches the final status for the first time.
 - Configure the global clone switch.
 - normal: Clone is allowed.
 - block: Clone is forbidden.
 - Configure the global dependency switch.

- normal: All configured dependencies are checked.
- ignore: The dependency is not checked.
- ignore_service: None of the service-level dependencies, including the server role dependencies across services, are checked, and only the server role-level dependencies are checked.
- 3. Click the Deployment Details tab to view the deployment details.

For more information, see the following table.

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

ltem	Description
Resource Application Progress	 Total indicates the total number of resources related to the project. Done: the number of resources that have been successfully applied for. Doing: the number of resources that are being applied for and retried. The number of retries (if any) is displayed next to the number of resources. Block: the number of resources whose applications are blocked by other resources. Failed: the number of resources whose applications failed.
Inspector Error	The number of inspection alerts for the current project.
Monitoring Information	The number of alerts generated for the machine monitor and the machine server role monitor in the current project.
Dependency	Click the icon to view the project services that depend on other services, and the current deployment status of the services that are depended on.

5.6.1.9. Reports

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

5.6.1.9.1. 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 **T** at the right of All **Reports** and then select **View**.

See the following table for the report descriptions.

ltem

Description

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

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

5.6.1.9.2. Add a report to favorites

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

- 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 **T** at the right of All **Reports** and then select **View**.
- 3. (Optional)Enter the name of the report that you are about to add to favorites in the search box.
- 4. At the right of the report, click Add to Favorites in the Actions column.
- 5. In the displayed Add to Favorites dialog box, enter tags for the report.
- 6. Click Add to Favorites.

5.6.1.10. Appendix

5.6.1.10.1. Project component info report

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

ltem	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

5.6.1.10.2. IP list

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

IP List of Physical Machines

ltem	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

ltem	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.
Docker IP	The Docker IP address.

5.6.1.10.3. 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.

ltem	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.

ltem	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.

ltem	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.

ltem	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.

ltem	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.

5.6.1.10.4. 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.

ltem	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.

ltem	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.

ltem	Description
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.

ltem	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.

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

5.6.1.10.5. 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.

Operations and Maintenance Guide · Operations tools

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.

ltem	Description
Component	The hard disk on the machine.
State	The running status of the hard disk.
Action Name	The action on the hard disk.
Action Status	The status of the action on the hard disk.
Actions	The approval button.

5.6.1.10.6. Registration vars of services

This report displays values of all service registration variables.

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

5.6.1.10.7. 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.

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

5.6.1.10.8. Service inspector report

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

ltem	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.

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

5.6.1.10.9. 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

ltem	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

ltem	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.

ltem	Description
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

ltem	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.

5.6.1.10.10. 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.

ltem	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.

ltem	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.

ltem	Description
Cluster	The cluster name.
Machine Name	The machine name.

ltem	Description
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.

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

5.6.1.10.11. 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.

ltem	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.

ltem	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.

5.6.1.10.12. 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.

ltem	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.

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

5.6.1.10.13. Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

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

ltem	Description
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

ltem	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

5.6.1.10.14. 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.

5.6.1.10.15. 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.

Item	Description
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.

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

5.6.2. New version

5.6.2.1. What is Apsara Infrastructure Management

Framework?

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

5.6.2.1.1. Introduction

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

Overview

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

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

Core functions

- Network initialization in data centers
- Server installation and maintenance process management
- Deployment, expansion, and upgrade of cloud products

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

5.6.2.1.2. Basic concepts

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

project

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

cluster

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

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

service

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

service instance

A service that is deployed on a cluster.

server role

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

server role instance

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

application

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

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

rolling

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

service configuration template

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

associated service template

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

desired state

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

dependency

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

upgrade

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

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

5.6.2.2. Log on to Apsara Infrastructure Management

Framework

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

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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

Enter a user name
Enter the password
Log On

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

3. Enter your username and password.

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

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

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

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

5.6.2.3. 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.

spage 🥶	Clusters			🛞 Instances			Machines		
ations >	189 Clusters	95.24% Desired State	9 Abnormal	1082 Instances	94.92% Desired State	55 Abnormal	1023 Machines	97.85% Normal	22 Abnorma
rs 🚺 📛									
oring M	l y Tasks Tasks in Last Wee	k: 670					Quick Actions		
<u> </u>	• Failed: 61	•	In Progress: 20	• Preparing: 2	• Term	inated: 7	0	egg	
	commit by tianji importer Status: Preparing Continued: a few seconds	4	commit by tianji import Status: Preparing Continued: a few secon		commit by tianji importer Status: In Progress Continued: a minute		Project Operations	OAM Permission Management	

Homepage of Apsara Infrastructure Management Framework

For more information about the descriptions of functional areas on the homepage, see the following table.

Descriptions of functional areas

Area		Description
Area 1	Left-side navigation pane	 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: 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.

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.
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/col lapse 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.

5.6.2.4. 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									
Project Status Statis	stics						Deploy Data IDC Topo	logical Graph amtest73	
amtest	Desired State 43 P Not Desired State		62 Total Projects		16 Alerting	E 4 In Progre	55	Cther Reasons	
Project Status								All	
apigateway	Alerting 2	In Progress 0	Not Desired State	1	aso	Alerting 16	In Progress 0	Not Desired State	
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	1	ecs	Alerting 20 E 2	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.

5.6.2.5. Cluster operations

This topic describes the actions about cluster operations.

5.6.2.5.1. View the cluster list

The cluster list allows you to view all of the clusters and the corresponding information.

- 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	~	Clusters Enter a cluster name	Q		
Clusters	Region	Status 🕎	Machine Status	Server Role Status	Task Status 🐺	Actions
-2881 acs	cn	Desired State	7 in Total Normal	14 in Total Normal	Successful	Operations
-2895 yundun-advance	cn	Not Desired State	3 in Total Normal	8 in Total Abnormal: 1	Failed	Operations
-284c dauthProduct	cn	Desired State	2 in Total Normal	7 in Total Normal	Successful	Operations

On this page, you can view the following information.

ltem	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.			
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.			

ltem	Description
Server Role Status	The number of server roles and the corresponding status in the cluster Click the status to go to the Services tab of the Cluster Details page. Click Abnormal in the Server Role Status column to view all the abnormal server roles in the cluster in the displayed dialog box. Click View Details in the upper-right corner of the dialog box to go to the Services tab of the Cluster Details page. Server Role Status Task Status T 7 in Total Normal Successful 38 in Total Abnormal: 20 Failed 33 in Total Abnormal: 20 Failed 33 in Total Server Roles View Details 33 in Total Server Roles View Details
	56 in Total tianji-sshtunnel-client.SSH1 Machine Error 56 in Total nuwa.NuwaConfig# Machine Error
	56 in Total The version is inconsistent.
	EcsTdc.Tdc# Machine Error 11 in Total EcsNbd.Nbd# Machine Error
	4 in Total N ecs-NcManager NcDownM Machina Error Top 20
Task Status	The status of the task submitted to the cluster. Use $\overline{\mathbf{T}}$ to filter the clusters. Click the status to view the task details.

5.6.2.5.2. 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. (Optional)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.

Clust	ers 2881			Edit AG She	ennong View Cluster Start/Shutdow	
tatus:	Desired State	Project:	acs	Region: cn-	•	
cludec	luded Server Roles: 14 Included Machines: 7			Task Status: Successful View		
Clone Mode: Real Clone System Configuration: default Git Version: 6671ee6277039e					995a13842d6e8eaeeb303	
Security Authentication: Disable Type: Ordinary Cluster						
<i>.</i> .						
Servi	ices Machines Cluster Configura	ation Operation Log Clust	ter Resource Service Inspection			
	Machines Cluster Configura Normal (6) Reset		ter Resource Service Inspection			
.ll: 6		ation Operation Log Clust	ter Resource Service Inspection		Deploy Service Batch Upgra	
II: 6 ervices	Normal (6) Reset		ter Resource Service Inspection Server Role	Service Template	Deploy Service Batch Upgra	
II: 6 ervices	Normal (6) Reset	2		Service Template default		
II: 6 ervices	Normal (6) Reset Enter a service name	2 Status	server Role		Actions	
II: 6 ervices	Normal (6) Reset Enter a service name Services 03	2 Status Normal	Server Role 1 in Total Normal	default	Actions Details Upgrade Details Upgrade	
II: 6 ervices	Normal (6) Reset Enter a service name Q Services os tianj	2 Status Normal Normal	Server Role 1 in Total Normal 1 in Total Normal	default	Actions Details Upgrade Details Upgrade Details Upgrade Unpublis	
II: 6 ervices	Normal (6) Reset Enter a service name Q Services Sanji hids-client	2 Status Normal Normal	Server Role 1 in Total Normal 1 in Total Normal 1 in Total Normal 1 in Total Normal	default	Actions Details Upgrade	

Area	ltem	Description
	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.
	Included Server Roles	The number of server roles included in the cluster.
	Included Machines	The number of machines included in the cluster.
		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.
	Task Status	• 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.

Area	ltem	Description
	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.
1	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.
	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.

Area	ltem	Description
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 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 server role status is not GOOD or the version that the service runs on the machine and the version configured in the configurations are different.
	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.

5.6.2.5.3. 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		
ubmission Time 12/04/19	12/11/19	Submitter Please input	Q	Services All ~		Refresh
Description		Operation Type	Status	Git Version	Submitter	Actions
commit by tianji importer			Successful	4f19dt0c535cOc718784815e2c4938001e887fac	aliyuntest Dec 05, 2019, 23:39:18	View Version Differences Details

- 3. On the Operation Log tab, view the version differences.
 - i. Click View Version Differences in the Actions column at the right of a log.
 - ii. On the Version Differences page, select a basic version from the Versus drop-down list. Then, the contents of the different file are automatically displayed.
 - iii. Select each different file from the **Different File** drop-down list to view the detailed differences.

5.6.2.6. Service operations

5.6.2.6.1. 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**.

vices Enter a service name	Q			
ervices	Description	Clusters	Included Service Templates	Actions
li-tianji-machine-decider		1 in Total Desired State: 1	0	Operations
csBssTools		3 in Total Desired State: 3	1	Operations
csNbd		5 in Total Desired State: 4 Not Desired State: 1	1	Operations
csRiver		3 in Total Desired State: 3	2	Operations
csRiverDBInit		1 in Total Desired State: 1	1	Operations
csRiverMaster		1 in Total Desired State: 1	1	Operations
csStorageMonitor		5 in Total Desired State: 4 Not Desired State: 1	1	Operations
csTdc		5 in Total Desired State: 4 Not Desired State: 1	3	Operations
lenderTestService1		0 in Total	0	Operations Delete
enderTestService2		0 in Total	0	Operations Delete

On this page, you can view the following information.

ltem	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.

ltem	Description
Included Service Templates	The number of service templates this service includes.
Actions	Click Operations to go to the Service Details page.

3. (Optional)Enter a service name in the search box and then the service that meets the condition is displayed in the list.

5.6.2.6.2. 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. (Optional)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.

Services EcsNbd Included Clusters: 5			Included Server Roles:	2	Included Service	Templates: 1	
Clusters Service Ten Project All Template Please select	nplate v			Enter a cluster name Q Tag Please select >		Template Please set	ect v Batch Add Tags
Clusters	Region	Status T	Server Role Status	Machine Status	Task Status 🝸	Template	Actions
289b ecs	cn	Not Desired State	2 in Total Abnormal: 2	1 in Total Abnormal Server Roles: 0 Abnormal Machines: 1	Successful	TMPL_ECS_V707_TIANJI_V4 Details	Operations Task Details
ecs	cn	Desired State	2 in Total Normal	5 in Total Normal	Successful	TMPL_ECS_V707_TIANJI_V4 Details	Operations Task Details
ecs 60db	cn	Desired State	2 in Total Normal	6 in Total Normal	Successful	TMPL_ECS_V707_TIANJI_V4 Details	Operations Task Details
288c ecs	cn	Desired State	2 in Total Normal	4 in Total Normal	Successful	TMPL_ECS_V707_TIANJI_V4 Details	Operations Task Details
	cn-	Desired State	2 in Total Normal	13 in Total Normal	Successful	TMPL_ECS_V707_TIANJI_V4 Details	Operations Task Details
						total 5 items < 1 > 10/Page	✓ Go to 1 Pa

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
amtest73	Normal Details	View	Terminal Restart Server Role
amtest73	Normal Details	View	Terminal Restart Server Role
amtest73	Normal Details	View	Terminal Restart Server Role
amtest73	Normal Details	View	Terminal Restart Server Role
amtest73	Normal Details	View	Terminal Restart Server Role
		total C	items < 1 > 10/Page < 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.

ltem	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.

5.6.2.7. 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**.

ations / 1	Machine Operations						
Machi Projec		V Clusters Enter	r a cluster name Q	Machines Enter one or more hos	stnames/IP addresses Q		Batch Termin
	Hostname	Clusters	Project	Region	Status 🕎	Machine Metrics	Actions
	amte	st73 2898	tianji	cn-	Normal Details	View	Operations Terminal Machine Management ~
	amte	st73 2895	yundun-advance	cn	Normal Details	View	Operations Terminal Machine Management \sim
	amte	st73 28	65 sib	cn-	Normal Details	View	Operations Terminal Machine Management ~
	.amte	st73 2898	tianji	cn	Normal Details	View	Operations Terminal Machine Management \sim
	amte	st73 288d	asto	cn	Normal Details	View	Operations Terminal Machine Management \sim
	amte	st73 288d	astc	cn	Normal Details	View	Operations Terminal Machine Management ~
	amte	st73 354e	ads	cn	Normal Details	View	Operations Terminal Machine Management \sim
	amte	st73 354e	ads	cn-	Normal Details	View	Operations Terminal Machine Management ~
	amte	st73 -ac1c	ecs	cn-	Normal Details	View	Operations Terminal Machine Management ×

3. (Optional)Select a project or enter the cluster name or machine name to search for the corresponding machine.

ltem	Description						
Hostname	Click the hostname to go to the Machine Details page.						
Status	The current status of the machine. Use T 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. Metric Metri Metric Metri Metri Metric Metric Metric Metri Metric Metric Metri						
Actions	 metrics. Click Operations to go to the Machine Details page. Click Terminal to log on to the machine and perform operations. You can select multiple machines and then click Batch Terminal in the upper-right corner to log on to multiple machines at a time. Click Machine Management to perform an out-of-band restart operation on the machine. 						

5.6.2.8. Monitoring center

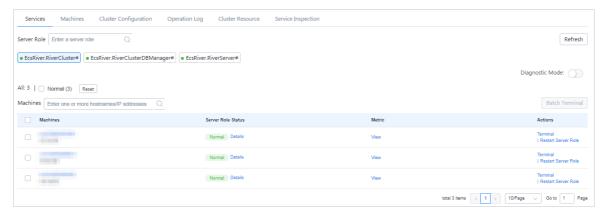
You can view the alert status, alert rules, and alert history in the monitoring center.

5.6.2.8.1. View the monitoring instance status

You can view the status of a monitoring instance after it is deployed.

Procedure

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



5.6.2.8.2. 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, 20:10:00	Search	
Service	Cluster	Instance	Alert Status	Alert Level	Alert Name	Alert Time	Actions
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	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	P1	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	P1	memo_cluster_host	11/23/19, 13:04:00 Lasted for 18 Days 7 Hours 6 Minutes 35 Seco nds	Show

- 4. (Optional)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.

ltem	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.

5.6.2.8.3. 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 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	•	Cluster A	All	•	Enter an alert name.	Search		
Service	Cluster	Aler	rt Name	Alert Conditions		Periods	Alert Contact	Status
yundun-semawaf		sema	nawaf_check_disk	\$Use>90		60	4	Running
yundun-semawaf		sema	nawaf_check_disk	\$Use>90		60	*	Running
yundun-semawaf		app_	_vip_port_check_serverrole	\$state!=0;\$state!=0		60	4	Running
yundun-semawaf		alert	t_ping_yundun-soc	<pre>\$rta_avg>500 \$loss_</pre>	max>80;\$rta_avg>400 \$loss_max>60	60		Running
yundun-consoleservice		chec	ck_auditLog_openapi	\$totalcount>9		300		Running
yundun-consoleservice		chec	ck_sas_openapi	Stotalcount>9		300		Running
yundun-consoleservice		chec	ck_aegis_openapi	Stotalcount>9		300	*	Running
yundun-consoleservice		chec	ck_secureservice_openapi	Stotalcount>9		300		Running
yundun-consoleservice		cons	soleservice_check_disk	long(\$size)>2097152	0	60		Running
yundun-consoleservice		chec	ck_aegis_openapi	Stotalcount>9		300	4	Running

- 4. (Optional)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.

ltem	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.

5.6.2.8.4. 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.

Service	All		•	Cluster All		-						
Hour	12 Hours	1 Day	1 Week	1 Month	3 Months	Custom	12/10/19, 20:16:00 ~ 1	12/11/19, 20:16:00				Sea
Service		Cluster		Alert Instanc	e		Status	Alert Level	Alert Name	Alert Time	Alert Contact	Actions
irds-conse	ole			service=drds-c	onsole,serverrol	e=drds-cons	sol ORestored	Restored	tianji_drds_prectrl_check_url	12/10/19, 20:38:13	4	Show
csTdc			288c	cluster	288c,serve	rrole=EcsTo	lc OAlerting	P4	ecs_server_compute-cpu_usa ge	12/10/19, 20:39:49	*	Show
csTdc			288c	cluster=	288c,serve	rrole=EcsTo	dc ORestored	Restored	ecs_server_compute-cpu_usa ge	12/10/19, 20:41:49	4	Show
so-syster	mMgr			service=aso-s	vsternMgr, server	role=aso-sy	ste OAlerting	P1	tianji_aso_auth_check_url	12/10/19, 21:48:28	-	Show
cs-houyi		ECS-HOUY	RE	cluster=		28a2,serv	err OAlerting	P4	ecs-houyi_ecs_regionmaster- unknow_error	12/10/19, 21:57:39	忠	Show
ecs-houyi		ECS-HOUY	RE	cluster=		-28a2,serv	err ORestored	Restored	ecs-houyi_ecs_regionmaster- unknow error	12/10/19, 22:08:39	4	Show

- 4. (Optional)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.

ltem	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 P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members that are notified when an alert is triggered.
Actions	Click Show to show the data before and after the alert time.

5.6.2.9. 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.

ks / My Tasks / Task Details							
Summary							Refresh
Task Status: Successful		Submission Time:	Dec 11, 2019, 20:25:3	2	Submitter: aliy	vuntest	
Duration: 12 minutes		Task Description:	RemoveMachine: ['iZh	n5i05w9770q3zm	qiltxdZ', 'iZh5i066934zp0of5l54mzZ	Z', 'iZh5i05w9770q3zmqiltxcZ', 'i	Zh5i05w9770q3z
Server Role All	~						
Included Clusters							
Clusters Q	Region T	Status		Progress		Start Time	Actions
-ac17	cn	Successful		🕢 Build -	— 🕢 Change	Dec 11, 2019, 20:25:32	View Version Differences Operation Log
						total 1 items < 1 > 1	Go to 1 P
E Change Details Clust	ers ac17	Ser	rvice Upgrade (8)	Machine Change	e (4)		
Server Role Q		Services T	Upgrade Type	Status T	Progress		Actions
rds-hbase.DbHBase# 🔗		rds-hbase	Configuration Change	Successful	🕑 Download — 🕑 Upgrade		Details
rds-hbase.Dblnit# @			Configuration				
		rds-hbase	Change	Successful	🥑 Download — 🕑 Upgrade		Details

5.6.2.10. Reports

5.6.2.10.1. 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.

- 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 27 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, 13:14:48	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. 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.

5.6.2.10.2. Add a report to favorites

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

- 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. (Optional)Enter the name of the report that you are about to add to favorites in the search box.
- 4. At the right of the report, click Add to Favorites in the Actions column.
- 5. In the displayed Add to Favorites dialog box, enter tags for the report.
- 6. Click Add to Favorites.

5.6.2.11. Tools

5.6.2.11.1. 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.
- 3. Select the operation scene according to actual situations.

Operation scene	Description	Action
Scene 1: NC Scale-out (with existing machines)	Scales out an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 2: Host Scale-out (with existing machines)	Scales out the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled out in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 3: NC Scale-in	Scales in an SRG of the worker type.	Select a target cluster and a target SRG. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.
Scene 4: Host Scale-in	Scales in the DockerHost#Buffer of an Apsara Infrastructure Management Framework cluster.	Select a target cluster. Select the machines to be scaled in in the left-side section and then click Select> to add them to the right-side section. Click Submit and then click Confirm in the displayed dialog box.

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.	

5.6.2.11.2. IDC shutdown

If you are about to maintain the data center or shut down all of the machines in the data center, you must shut down the data center.

Prerequisites

• Warning 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 data center with one click.

5.6.2.12. Appendix

5.6.2.12.1. Project component info report

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

ltem	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.

ltem	Description
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

5.6.2.12.2. 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

ltem	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.
Docker IP	The Docker IP address.

5.6.2.12.3. 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.

ltem	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.

ltem	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.

ltem	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

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.

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

Monitoring Status

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

ltem	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.

5.6.2.12.4. 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.

Operations and Maintenance Guide • Operations tools

ltem	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.

ltem	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.

ltem	Description
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.

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

5.6.2.12.5. 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 basic information of pending approval machines.

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

Item	Description
Actions	The approval button.

Machine Serverrole

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

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

Machine Component

Displays the hard disk information of pending approval machines.

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

5.6.2.12.6. Registration vars of services

This report displays values of all service registration variables.

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

5.6.2.12.7. 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.

5.6.2.12.8. 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.

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

5.6.2.12.9. 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

ltem	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

ltem	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

ltem	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.

ltem	Description
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.

5.6.2.12.10. 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.

ltem	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.

ltem	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.

ltem	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.

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

Service Inspector Information

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

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

5.6.2.12.11. 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.

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

5.6.2.12.12. 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.

Operations and Maintenance Guide · Operations tools

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.

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

5.6.2.12.13. Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

ltem	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

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

ltem	Description
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

5.6.2.12.14. 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.

5.6.2.12.15. 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.

ltem	Description
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.

ltem	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.

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

6.Products

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

6.1. Product list

On the **Product List** page, you can 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 the ASO console.
- 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.

6.2. ISV access configurations

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

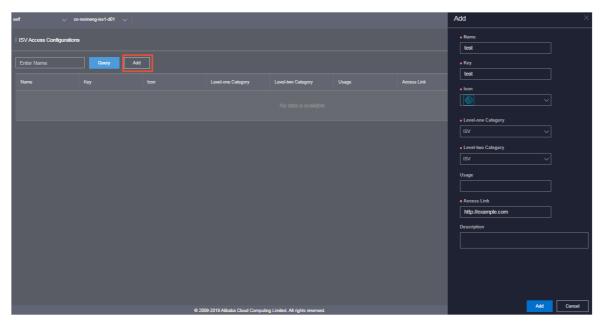
6.2.1. Configure the ISV access information

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

- 1. Log on to the ASO console.
- 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.
lcon	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.

6.2.2. Modify the ISV access information

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

- 1. Log on to the ASO console.
- 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 Query. Fuzzy search is supported.
- 4. Find the ISV whose access information is to be modified. Click Modify in the Actions column.



- 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.

6.2.3. Delete the ISV access information

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

Procedure

- 1. Log on to the ASO console.
- 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 Query. 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**.

7.Network operations

7.1. Apsara Network Intelligence

7.1.1. 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.

7.1.2. Log on to the Apsara Network Intelligence console

This topic describes how to log on to the Apsara Network Intelligence console.

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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

Log on to ASO

Log On								
<u>8</u>	Enter a user name							
F	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 your username and password.

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

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

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

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- It must be 10 to 20 characters in length.
- 4. Click Log On to go to the ASO console.
- 5. In the left-side navigation pane, click **Products**. On the right side of the page, click **Apsara Network Intelligence**.

7.1.3. Query information

You can enter an instance ID to query details of the instance.

- 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

ion					
Cluster	EIP Type		Gateway Type	SLB Mode	status
cn-qingdao-env8d-d01	intranet		classic	fnat	active
Proxies	Created At		Modified At	After WAF/Anti-DDoS Protection	Actions
		Ne	o data		
			Black Hole Threshold		
			None		
Status	Tunnel ID		Service Unit Name	Primary IDC/LVS Name	Secondary IDC/LVS Name
			** ·		
VIP OUT bit/s	VIP IN bit/s	VIP QPS	VIP CPS	Specifications	Instance Type
	co-opgiss-andid-d51 Proxie	en-engela-endel-601 internet Presies Created AL	an-spagae-anded-s01 intranet Reside Cented A Statu Taxand ID	or-opgiso-ended:01 intravet classic Passia Cented Al Modified Al Intravel Foodratic Read Mole Therebald Status TasselD Sanisa Unit Name Foodratic Foodratic Foodratic	$\begin{tabular}{ c c c c } \hline \begin{tabular}{ c c c c c } \hline \begin{tabular}{ c c c c c c c } \hline \begin{tabular}{ c c c c c c c } \hline \begin{tabular}{ c c c c c c c } \hline \begin{tabular}{ c c c c c c c c c c c c c c c c c c c$

Listener information

Click **Show** in the **Back-end Server/Health Check** column to view details on backend servers.

C Resures / Ski Jestera Datais											
Instance Information Listener Information											
Enter 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	
Ib-q8ckibpdtql	tcp	80	No	No	None	None	 running 	Show	2019-05-16 03:14:45	2019-05-16 03:14:56	
Ib-q8ckibpdtql	tcp	22	No	No	None	None	 running 	Show	2019-05-16 03:14:36	2019-05-16 03:14:56	
4										÷	

7.1.4. Manage cloud service instances

You can create a cloud service in a region or query the instance information of a region.

- 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.

7.1.5. Tunnel VIP

7.1.5.1. 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.

Create SLB Instance	Add Back-end Server to SLB Instance	Create Listener
	ter VPC	
	VPC ID : Enter a VPC ID to add a VM in ored Servers: 192.168.9.0.200	VPC (0) [Enter a VPC ID to add a VM in the VPC.

- 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.

7.1.5.2. 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.

Procedure

- 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.

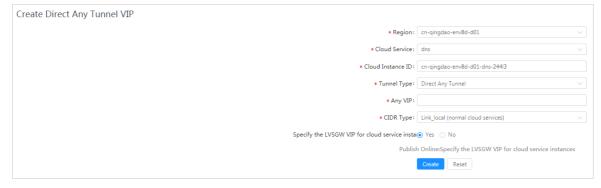
Tunnel VIP Management	Tunnel VIP Management											
• Report D :	1481		Doud Service.	gin .			Status Ranning					
										South		
Report	Dout Service	Coul Instance I D	S.J. Instance ID	LS 119	Same	Control As	Modified Ac	Hodfed by	Part C connec Sixity	Actions		
cr-singdes-envloi-d01	93	to singles and	10.1416/201 on all states and at	10.65	2	2218-06-03 10:17:38	2215-05-02 12:10:05	alyumat		Actors ~		

7.1.6. 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.

7.1.7. Leased line connection

7.1.7.1. 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:

- 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.

7.1.7.2. 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-qingo Search	dao-env8d-d01 v Reset	Access Point ID: ap-	devie en Ré Ré									
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 🔿 Busy 💿 Full 🔿 Unavailable	
* Access Point Location: AMTEST61	
* Access Point IDC Operator: Other	
Physical Region :	~
Modify Cancel	

7.1.7.3. 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.

? Note If Device ID is not set, the information about all devices in a region is queried.

5. Click Search.

Access Devices											
Rejon: Or-anglas-en-86-001 Declar B1 (259-10-4)											
Search	Reset										
Device ID	Region	Access Point ID	Device Status	Physical Location	Access Nethod	Device Name	Description	Created At	Modified At	Actions	
C94-VAH-VPC-C	ch-sinsdap en 64.401	ap-or-grigitas-modul-	mainthe	AMTERNI	stanToVstanRo	098-00-090-00-	C9A-VM-VPC-61-	2777-04-29 22:59:32	2019-04-29 22:50.32	Medity Shew Details	
		0.000			uting	1000					

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-G	
* 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 Add	10.48.	
Device Manufacturer:	Ruijie	
Device Model:	RG-S6220-	
Device Name:	CSW-VM-VPC	
Device Description:	CSW-VM-VPC-	
	Modify Cancel	

3. Click Modify.

7.1.7.4. 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.

Create Leased Line		×
Name:	The leased line name. It can be 2 to 128 characters in length and	d cannc
Description:	The leased line description. It can be 2 to 128 characters in leng	th and
* BID :	26842	
* UID :	EnterUID	
* Region:	cn-qingdao-env8d-d01	~
-	The region ID is used for managing access devices (which is cessarily the same as the attached region ID of the access of 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 VI orks 	PC netw
* Access Point ID :		
	Access Point ID	
Device Name:	Device names can be 2 to 256 characters in length and cannot s	start wit
Device Port:	CSW Port	
Bandwidth:	[2-10000]	Mbps
	The inbound interface bandwidth of the leased line. Unit: N alue range: [2-10000].	1bit/s. V
* Port Type:	Select	\sim
	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 redundant one and upload its ID. If you do so, Alibaba Cl cates a separate access device for higher availability. The leased line that you specify must exist and be in Allou onfirmed, or Enabled status. 	oud allo

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.

7.1.7.5. 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 on-premises 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.

Operations and Maintenance Guide • Network operations

Create VBR	×
* BID :	26842
* UID:	EnterUID
* Region:	cn-qingdao-env8d-d01 v
	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 ID Addr	Local Gateway IP Address
· Local Galeway IP Addie	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 addres
	s and peer IP address.
	 It is required when the interface status is not waiting. Only the VBR 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
	Create Cancel

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.

/BRs								Create VBR
* Region : cn	n-qingdao-env5b-d01 V	* BID:		* UID :	119			
VBR ID:								
S	earch Reset							
VBR ID	VLAN ID	VLAN Interface ID	Status	Routing Table ID	Local Gateway IP Address	Peer Gateway IP Address	Subne Actions	
vbr-	33	ri- f9rt	active	vtb- f9i	192.168.	192.168.	255.25	Actions 🔨
							Release	
-1/1							Modify Terminate	
							Show Details	

You can click **Release**, **Modify**, **Terminate**, or **Show Details** in the **Actions** column to manage a VBR.

7.1.7.6. Create router interfaces

After you create a VBR, you must create a pair of router interfaces to connect the VBR and VPC. The connection initiator must be 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 ◯ Double	
Name:	name of router interface	
Description:	description of router interface	
* Bid :	EnterBid	
* Uid:	EnterUid	
* Region:	cn-qingdao-env8d-d01	~
* Router Type:	● VRouter ○ VBR	
Zone:	SelectZone	
* Router ID :		
* Role:	InitiatingSide	
* Specifications:		~
Health Check Source IP:	EnterHealth Check Source IP	
Health Check Destination	EnterHealth Check Destination IP	
Skip Inventory Check:	🔿 Yes 💿 No	
	Next Cancel	

Operations and Maintenance Guide • Network operations

When the router interface state is Active, the interface is created.

Router Inte	rfaces									Create Router Interface
* Region : cr	n-qingdao-er	w5b-d01	~ * BID :	26		* UID :	11			
s	Search	Reset								
_										
Local Router ID		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-f5		VRouter	ri-f9n,	Active	None	Accepting Side	vbr-ft	VBR	Deactivate	Actions A
vrt-f9r 2		VRouter	ri-f9rij)	Inactive	None	Accepting Side	vbr-f9 g	VBR	Modify Attribute Modify Specificat	ion
-2/2									Show Details	′page ∨ Goto

7.1.7.7. 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:
 - i. Log on to the Apsara Network Intelligence console.

- ii. From the Products menu, choose Express Connect > Network Environment Management.
- iii. Choose Function Modules > Routing Tables.
- iv. Set search conditions such as Region, BID, UID, Router Type, Routing Table ID, and Router ID, and click Search to query routing tables.
- v. Click Add Route Entry in the Actions column corresponding to a routing table.
- vi. 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 a route destined for a destination VPC

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:	Ves 💿 No	
* Next Hop Type:	Instance	
	The next hot type. Valid values: Instance, Tunnel, HaVip, RouterInterfaceSet the value to RouterInterface for ECMP.	
* Next Hop ID:		
	The next hop interface ID for the route entry.	
	Create	

vii. 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.

7.1.8. 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.

Flow Management page

Flow Manager	nent							
Region:	Region: on-gingdao-env8d-d01		Row Name: common il and		UID:			
Step Status:	Success		Instance Id: m-q8mmin		Request Id:			
Time Range: Cus	tom V 2019-05-01 11:21:33 ~ 2019-0	6-03 11:21:33						Search
Enter filter conditions	i. Items per Page: 10 🗸							
	Service Type	Flow Name	Instanceld	Execution Status(Success	Ealed Demossion			Stopped) (Resend)
					0			0.11
	VPC	commonAsyncTask	ms		0	0	() handleResult	
	100	End Flo			C submit lask	C CONTRACTAGE	() handlenester	
			Res	end				

7.1.9. 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.

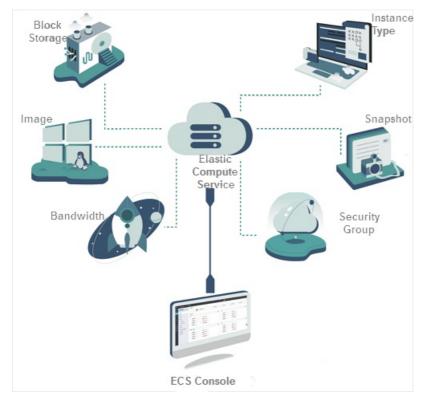
The value of Instance Port must be an integer value. You can specify multiple instance ports separated by commas (,). Example: 10,20,30. You can configure up to 10 instance ports.

8.Operations of basic cloud products 8.1. Elastic Compute Service (ECS)

8.1.1. 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 ECS instance.



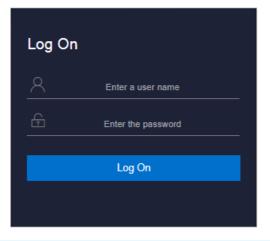
ECS instance

8.1.2. Log on to the Apsara Stack Operations console

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

- 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.

Log on to ASO



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

3. Enter the correct username and password.

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

8.1.3. ECS operations and maintenance

8.1.3.1. 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.

8.1.3.2. VM

8.1.3.2.1. 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.

8.1.3.2.2. 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 Apsara Stack Operations.
- 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 Mainte	nance Platform							
VMs Disks	Snapshots	Images	Security G	Groups				
Region Security Group	Status	ommas (,) to separate		Address mmas (,) to separate	AliUid			
Start Stop	Reboot Stop	and Migrate Mi	ore	Disks 🎝	Security Groups 🎝		Vie	W Clear
	Host	VM IP Address	Public IP Address	Region	CPU (C) Memory (M)	Disk Information	Internet Bandwidth	Status
	10.80%) #40.4118	-		2.000 arriv				Running

5. In the VM list, click a VM ID. You can view the VM information in the VM Details message that appears.

8.1.3.2.3. 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.

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

8.1.3.2.4. 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.

? Note 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 Apsara Stack Operations.
- 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 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**.

(?) Note 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.

8.1.3.2.5. 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 Apsara Stack Operations.
- 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.

8.1.3.2.6. 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 Apsara Stack Operations.
- 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

Operations and Maintenance Guide · Operations of basic cloud products

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.

8.1.3.2.7. 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 Apsara Stack Operations.
- 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.

8.1.3.3. Disks

8.1.3.3.1. Overview

In an ECS instance, cloud disks can be considered as physical disks. You can mount, detach, and create snapshots for disks.

8.1.3.3.2. 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 Apsara Stack Operations.
- 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. Specify the filter conditions and click View.Region is a required filter condition.

ECS O	ECS Operations and Maintenance Platform										
VMs Disks Snapshots Images Security Groups											
Region		(Disk ID				Ali	Jid			
) to separate multip	Enter onl	y one ID.					
Disk Type			Disk Status								
Select			Select							View	Clear
	Disk ID	AliUid	Disk ID	Release Auto Snapshot	Independent Disk	Disk Size	Disk Type	VM ID	Mount Point	Region	Disk Status
								1000	-	12^{mm}	In use

8.1.3.3.3. View snapshots

You can view the list of existing snapshots and their information in the Apsara Stack Operations console.

Procedure

- 1. Log on to Apsara Stack Operations.
- 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 whose snapshots you want to view, and choose **_____** > View Snapshot.

The information of all snapshots on the disk is displayed.

8.1.3.3.4. 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 Apsara Stack Operations.
- 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 to be mounted and choose **Mount**.
- 7. In the dialog box that appears, set VM ID and Operation Reason. Click OK.

8.1.3.3.5. 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 Apsara Stack Operations.
- 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 to be detached and choose **Detach**.
- 7. In the dialog box that appears, set Operation Reason. Click OK.

8.1.3.3.6. Create a snapshot

You can manually create a snapshot for a disk in the Apsara Stack Operations console.

Procedure

- 1. Log on to Apsara Stack Operations.
- 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

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	🗘 Detacl	h Take Sn	apshot							
_	d-bs005ma64vv otmptiff9	100000021	20211-19968	Yes	No	100G	SYSTEM	i-bs005ma64vv otmpwfe3p	/dev/xvda	cn-qingdao-env 12-d01	In use
	View Snapshot	A Take S	Snapshot								

7. In the dialog box that appears, set Snapshot Name, Snapshot Description, and Operation Reason. Click OK.

8.1.3.4. Snapshots

8.1.3.4.1. 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.

8.1.3.4.2. 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 Apsara Stack Operations.
- 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. Specify the filter conditions and click View.Region and AliUid are required filter conditions.

ECS Operations and Maint							
VMs Disks Region	Snapshots AliUid	Images	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

8.1.3.4.3. Delete a snapshot

You can delete snapshots that are no longer needed in the Apsara Stack Operations console.

Procedure

- 1. Log on to Apsara Stack Operations.
- 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 to be deleted and choose **Delete**.
- 7. In the dialog box that appears, set Operation Reason. Click OK.

8.1.3.4.4. 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 Apsara Stack Operations.
- 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**.

8.1.3.5. Images

8.1.3.5.1. 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.

8.1.3.5.2. 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 Apsara Stack Operations.
- 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 Images tab.
- 5. Specify the filter conditions and click View.Region is a required filter condition.

ECS Operations and Maintenance Platform					
VMs Disks	Snapshots I	nages Security Groups			
• Region	Image Type	AliUid	Image ID Use commas	s (,) to separate multiple	View Clear
Image ID	Image Name	Snapshot ID	Created At Im	наде Туре	Operating System Type
 3000.000000 			Sep 20, 2019, 18:26:22	-	Windows Server 2008

8.1.3.6. Security groups

8.1.3.6.1. 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.

8.1.3.6.2. 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.

? Note

- 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 Apsara Stack Operations.
- 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 Security Groups tab.
- 5. Specify the filter conditions and click View.Region is a required filter condition.

ECS Operations and Maintenance Platform				
VMs Disks Snapsł	nots Images Se			
	curity Group ID Use commas (.) to separate multiple	VM ID Enter only one ID.		View Clear
Security Group ID	Security Group Name	Created At	VPC ID	Region
+	West and some sectors			

8.1.3.6.3. Add security group rules

You can add rules to security groups based on your needs.

- 1. Log on to Apsara Stack Operations.
- 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 Security Groups tab.

- 5. On the tab that appears, set the filter conditions and click View.
- 6. Find the target security group and choose \rightarrow Add Rule.
- 7. In the dialog box that appears, configure the parameters.For more information about the parameter configurations, see Security group rule parameters.

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.
Operation Reason	Optional. Enter a reason for the operation.

8. Click OK.

8.1.4. VM hot migration

8.1.4.1. 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.

8.1.4.2. 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.

? Note VMs created in Apsara Stack versions earlier than V3.3 do not support hot migration. Hot migration becomes available after you restart the VMs.

8.1.4.3. 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_ch
eck_image] [no_check_load] [downtime]== Usage: == houyi_api.sh <function_name> [--help|-h] [name=
value]

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
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

Parameter description

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=valu
e] == Functions usage: == |- cancel_live_migrate_vm <region_id> <vm_name>

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

Parameter description

8.1.4.4. 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_cha nge_migration_status <vm_name> <region_id> <where>

Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
region_id	The ID of the region where the target VM is located.	N/A	N/A
where	The ID of the NC where the VM is located.	N/A	N/A

8.1.4.5. 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
 - \circ 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 can cel_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?

Hot migration 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_load=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_image=true to skip this inspection. Be aware of the risks involved.

8.1.5. Hot migration of disks

8.1.5.1. 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.

8.1.5.2. 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* .

8.1.5.3. 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(gm t_finish) < DATE_ADD(CURDATE(), INTERVAL -1 DAY)' command.</p>
- 3. Perform the following operations for each set of <task_id,src_pangu_path,dst_pangu_path>:
 - i. 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.
 - ii. Run the /apsara/deploy/bsutil delete-image --dir=\$src_pangu_path command on the host that runs the bstools role in the storage cluster.
 - iii. Run the /apsara/river/river_admin migrate recycle \$task_id command on the host that runs the river role in the storage cluster.

8.1.6. Upgrade solution

8.1.6.1. Overview

For both hot and cold migration of GPU and FPGA clusters, you must understand the limitations that apply to cluster upgrades.

8.1.6.2. 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 reformatted, resulting in data loss. These disks must be backed up before they can be migrated.

8.1.6.3. 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.

8.1.7. Disk maintenance of an instance

8.1.7.1. 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.

8.1.7.2. 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
- $\circ~$ Run the go2ncinfo 10.10.3.5 command to obtain the NC ID.
- NC ID: 21765-26

	:/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
\$ go2ncinfo 10.10.3.5	
nc_id:	21765-26
ip	10.10.3.5
hostname:	c43b07003.cloud.b07.amtest1221
biz_status:	free
priority:	8
health:	

- 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_device_id={mount_point}

{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.

Example:

go2hyapi query_vm_list format=json region_id=cn-neimeng-env10-d01 nc_id=21765-26 nc_storage_ device_id=/apsarapangu/disk1

[admin@ :/home/admin] [ECS-IO11-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
\$ go2hyapi query_vm_list format=json region_id=cn-neimeng-env10-d01 nc_id=21765-26 nc_storage_device
[ERROR] [2018-05-10 16:41:36] The function 'query_vm_list' doesn't exist!
= Usage: =
houyi_api.sh <function_name> [name=value]

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/tm p.log nc_id=21765-26 storage_device_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.
 - i. 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}
```

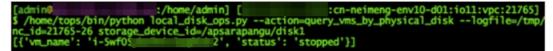
{vm_name} is the ID of the running VM obtained in the preceding step.

Example:

go2hyapi stop_vm vm_name=i-5wf05ykw7mic5aq65dv2



Wait until the VM status changes to Stopped.



- ii. If the VM is in the pending or stopped state, you do not need to shut it down.
- iii. 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_local_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_local_disks_by_physical_disk --logfile=/t mp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1

[admin@ :/home/admin] [:cn-neimeng-env10-d01:io11:vpc:21765]
\$ /home/tops/bin/python local_disk_ops.py --action=query_local_disks_by_physical_disk --logfile
[{'vm_name': 'i-5wf05y /2', 'disk_id': '1000-3388'}]

Only the local data disk with the ID 1000-3388 is associated.

- 5. Replace the damaged physical disk on the NC.
 - i. 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.

- ii. 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

iii. 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:

C4-3.NT12	B07	06	CVXKB7CD001	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.



- iv. Insert a new disk.
- v. 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/sdb: 60 Units = sectors of Sector size (logi I/O size (minimum Disk label type: Disk identifier:	of 1 * 512 cal/physic n/optimal): dos	= 512 bytes al): 512 byt 4096 bytes	es / 4096 by	tes	168 sectors
Device Boot /dev/sdb1 Partition 1 does	1 4	4294967295	2147483647+	ee	

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

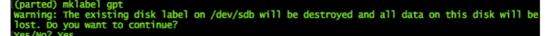
Run the parted command if the hard disk capacity is greater than 2 TB.

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.



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.

[root					:/root]
[root [#lsblk	grep	sdb			
sdb └─sdb1	8:16	0	5.5T	0 disk	
└─sdb1	8:17	0	5.5T	0 part	

c. Format the partition.

```
mkfs -t {filesystem_type} {device_name}
```

{filesystem_type} is the type of the file system to be formatted. {device_name} is the name of the partition to be formatted.

Example:

mkfs -t ext4 /dev/sdb1

[root@
#mkfs.ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem_label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
183144448 inodes, 1465130240 blocks
73256512 blocks (5.00%) reserved for the super user First data block=0
Maximum filesystem blocks=3613392896
44713 block groups
32768 blocks per group, 32768 fragments per group
4096 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
102400000, 214990848, 512000000, 550731776, 644972544
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

MOUNTPOINT
391b8 /boot
Ba6da /
5794d
1821e /apsarapangu
:0979 /apsara
/824d
db54e /apsarapangu/disk2
L4093 /apsarapangu/disk3
662cc /apsarapangu/disk4
20979 /absara 7824d db54e /apsarapang 14093 /apsarapang

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.

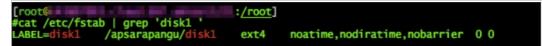
Device files in the */etc/fstab* directory are identified by their labels, so you must change the label of the new disk.

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.

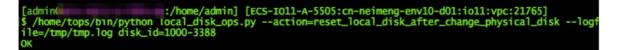
\$/home/tops/bin/python local_disk_ops.py --action=reset_local_disk_after_change_physical_disk

--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_local_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.

8.1.7.3. 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_physical_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.

8.1.8. Handle routine alarms

8.1.8.1. 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.

Description of metric types

Operations and Maintenance Guide · Operations of basic cloud products

Metric type	Function	Solution
BasicMonitors the basic performance of the host and the availability of the services on the host. This kind of availabilityavailabilitymetrics includes CPU, memory, and handle count.	the host and the availability of the	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.	
Connectivit y 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 application problem, forward the error information to the back-end R&D team for troubleshooting.

8.1.8.2. API proxy

This topic describes the metrics of API proxy.

Metric	Alert item	Description
check_apiproxy_dns	Database HA switchover occurs or not	Checks whether Server Controller database switchover occurs. If so, nginx will be reloaded automatically.
		Checks the connectivity to the Server Controller database.
check_apiproxy_con	check apiproxy conn ne	

n_new Metric	w Alert item	Description
		Checks the connectivity to the API Server:Checks whether the API Server is down.Checks the network connectivity.
check_apiproxy_pro c_new	check_apiproxy_proc_ne w	Checks the memory usage and CPU utilization for nginx and memcache processes.

8.1.8.3. API Server

The topic describes the metrics of the API Server.

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	
check_API Server_conn_new	Checks the connectivity between the API Server and Server Controller database.		
	Checks the connectivity between the API Server and TAIR.	Checks whether the corresponding component is down. If the corresponding component is down, fix the issue by taking necessary O&M measures. If the database is down, contact DBA to fix the issue.	
	Checks the connectivity between the API Server and RegionMaster.	Checks whether the VIP is connected to the corresponding component. If not, contact th network engineer to fix it.	
	Checks the connectivity between the API Server and the RMS.		
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. 	

8.1.8.4. RegionMaster

This topic describes the metrics of RegionMaster.

Metric description

Metric	Alert item	Description
check_regionmaster_proc	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage.
	rms_connectivity	Checks the connectivity to RMS.
shock regionmester work	regiondb_connectivity	Checks the connectivity to the houyiregiondb database.
check_regionmaster_work	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_errorlo	errorlog_for_db	Checks whether the SQL statements are properly
g	check_regionmaster_errorlog	executed.
check_workflow_master	Checks the operating state of the master in the workflow process.	-
check_workflow_worker	Checks the operating state of the worker in the workflow process.	-

8.1.8.5. RMS

This topic describes the metrics of RMS.

Metric	Alert item	Description
check_rms_proc	Checks the process status, CPU utilization, and memory usage of RMS.	-
check_rabbitmq_proc	Checks the process status, CPU utilization, and memory usage of the rabbitmq cluster.	-

Operations and Maintenance Guide · Operations of basic cloud products

Metric	Alert item	Description	
check_rabbitmq_status	Checks the number of queues, exchanges, and bindings in the rabbitmq cluster.	Follow the maintenance guide for the rabbitmq cluster.	
	Checks whether messages are accumulated.	If messages are accumulated, it will also check for the cause.	
check_rabbitmq_queues	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.	

8.1.8.6. PYNC

This topic describes the metrics that are monitored for PYNC.

Metric description

Metric	Alert item	Description
check_vm_start_f ailed	Checks the causes of a VM startup fault.	You do not need to handle it immediately. It is typically caused by custom images.
	Checks the CPU utilization and memory usage of PYNC.	-
	PYNC has too many open file handles.	-
	PYNC process count.	PYNC must have four processes.
check_pync	It has been long since pyncVmMonitor.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.

8.1.8.7. Zookeeper

This topic describes the metrics of Zookeeper.

Metric	Alert item	Description	
		The process does not exist.	
check_zookeeper_proc	proc	The memory usage or CPU utilization is too high.	

8.1.8.8. AG

This topic describes the metrics of AGs.

	Description	
apsara_90	/apsara disk usage.	
homeadmin_90	Usage of /home/admin.	
mem_85	Memory usage.	
cpu_98	CPU utilization.	
df_98	Disk usage of the root directory.	
check_ag_disk_usage	Disk usage.	
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. 	
	homeadmin_90 mem_85 cpu_98 df_98 check_ag_disk_usage check_recover_failed check_repeat_recovered check_continuous_nc_down	

Metric	Alert item	Description
check_ag_fhtd_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.

8.1.8.9. Server groups

This topic describes the metrics that are monitored for server groups.

Metric description

Metric	Alert item	Description
	pync_mem	Monitors the memory usage of PYNC.
	pync_cpu	Monitors the CPU utilization of PYNC.
check_pync	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.

8.1.9. Inspection

8.1.9.1. Overview

ECS inspection includes cluster basic health inspection and cluster resources inspection.

8.1.9.2. Cluster basic health inspection

8.1.9.2.1. Overview

Cluster basic health inspection includes monitoring inspection, inspection of basic software package versions, and basic public resources inspection.

8.1.9.2.2. Monitoring inspection

This topic describes basic monitoring inspections and connectivity monitoring inspections.

8.1.9.2.3. Inspection of basic software package versions

This topic describes the version inspections of Server Controller components, Apsara system, virtualization packages, and basic service packages.

8.1.9.2.4. 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 information:

\$ houyiregiondb

```
mysql>select name,os_type,version,path,oss_info from iso_resource where os_type! ="\G
```

Parameters in the command are as follows:

- *name*: 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.

8.1.9.3. Cluster resource inspection

8.1.9.3.1. Overview

Cluster resource inspection includes cluster inventory inspection and VM inspection.

8.1.9.3.2. 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_resourc e,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.

Note 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 Mappings of host states between the ECS database and Apsara Infrastructure Management Framework.

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 proactively locked.

Host states in ECS database	Cluster	Host state	Scenario
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.

8.1.9.3.3. 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.

8.2. Container Service

8.2.1. Components and features

8.2.1.1. Console

The Container Service console provides a user interface that serves as an entry for all operations on Container Service. It adopts the deployment mode that applies to standard Java applications in Alibaba Cloud. Each console instance contains a Tengine server and a Jetty container.

Command entry

- Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page that appears, locate the AcsControlCluster-A-201812 cluster, and click Cluster Configuration in the Actions column corresponding to the cluster. On the Cluster Configuration page that appears, locate Service Role, and find the CosConsoleAliyunCom service role and the corresponding host.
- In the middle of the left-side navigation pane, enter the hostname in the Server search box. Hover over the vertical dots next to the hostname and choose Terminal from the shortcut menu to log on to the host through a terminal session. Run the docker ps command to obtain the ID of the cos-console-aliyun-com container.
- Run the sudo docker exec -it container_id bin/bash command to access the container.
- Go to the specified directory to find Tengine and Jetty.

O&M commands

- Restart Tengine: /etc/rc.d/init.d/tengine restart
- Restart Jetty: /etc/init.d/jetty restart

Directory structure

- Root directory of Web applications: /alidata/www/
- WAR directory of applications: /alidata/www/wwwroot/cos-console-aliyun-com

Application log files

- The root directory that stores log files: /alidata/www/logs
- The path to Jetty: /alidata/www/logs/jetty
- The path to application log files: /alidata/www/logs/java/cos-console-aliyun-com/applog

8.2.1.2. Troopers

Troopers is used to create clusters and hosts and to manage their information in Container Service.

Troopers uses the Go language to compile code. A container only runs the Troopers process and does not use any daemons.

Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page that appears, locate the AcsControlCluster-A-201812 cluster, and click **Cluster Configuration** in the Actions column corresponding to the cluster. On the Cluster Configuration page that appears, locate Service Role, and find the host that corresponds to the Troopers service role.

In the middle of the left-side navigation pane, enter the hostname in the Server search box. Hover over the vertical dots next to the hostname and choose Terminal from the shortcut menu to log on to the host through a terminal session. Run the docker ps command to obtain the ID of the Troopers container.

Run the sudo docker exec -it container_id bin/bash command to access the container.

The directory structure is as follows:

- */usr/aliyun/acs/troopers*: the root directory of the application.
 - troopers: the main program of Troopers.
 - troopers.json: the configuration file of Troopers.
 - troopers.ym: the certificate encryption configuration information.
 - start.sh: the script used to start Troopers. If the Troopers process already exists, do not run the *start.sh* script.
- */opt/aliyun/install/check_health.sh*: the script for health checks.
- /usr/aliyun/acs/certs/control: the directory that stores the certificate Troopers uses to access the Region Controller (RC). You can use OpenSSL to verify the certificate.

Troopers log files are exported to stdout directly. No log files are stored in the container. To view log records, run the docker logs command outside the container.

8.2.1.3. Mirana

Command entry

Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. On the Cluster Operations page that appears, locate the AcsControlCluster-A-201812 cluster, and click **Cluster Configuration** in the Actions column corresponding to the cluster. On the Cluster Configuration page that appears, locate Service Role, and find the host that corresponds to the Mirana service role.

Log query

Log on to the Apsara Infrastructure Management Framework console. In the middle of the leftside navigation pane, enter a specified hostname in the Server search box. Hover over the vertical dots next to the hostname and choose Terminal from the shortcut menu to log on to the host through a terminal session. Run the docker ps command to obtain the ID of the Mirana container. Run the docker logs container_id command to view the log information.

The Mirana container is stateless. You can try to restart the container if the service is unavailable. Run the docker restart container_id command to restart a container.

Deployment mode

- A Mirana container is deployed in each cluster. The deploy mode of the Mirana container is similar to that of the Commander container.
- Mirana containers are deployed on control hosts and use HTTPS to provide external services.

The Kubernetes API certificate must be provided when a Troopers container is created.

Features

- Uses the Helm client to manage orchestration templates.
- Supports the blue-green deployment of APIs.

8.2.2. System restart

8.2.2.1. Restart a control node

A container control node is a Docker container where a service, such as CosConsoleAliyunCom, Troopers, or Etcd, is deployed. To restart a control node, perform the following operations:

Procedure

- Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page that appears, locate the AcsControlCluster-A-201812 cluster, and click Cluster Configuration in the Actions column corresponding to the cluster. On the Cluster Configuration page that appears, locate Service Role, and find the host where the control node is deployed. In the middle of the leftside navigation pane, enter the hostname in the Server search box. Hover over the vertical dots next to the hostname and choose Terminal from the shortcut menu to access the host through a terminal session.
- 2. Run the docker ps|grep [app] command to obtain the container ID.

[app] indicates the name of the application deployed in the container. You can obtain the container ID based on the application name.

3. Run the docker restart container_id command to restart the container.

8.3. Auto Scaling (ESS)

8.3.1. Log on to the Apsara Stack Operations

console

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

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.

Log on to ASO

Log On	
<u>8</u>	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.

8.3.2. Product resources and services

8.3.2.1. Application deployment

All the applications in the ESS Business Foundation System are stateless. You must restart the applications by running the docker restart command.

• 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.

8.3.2.2. Troubleshooting

This topic describes how to troubleshoot issues related to product resources and services.

Prerequisites

When issues related to Business Foundation System occur, you can submit tickets on the AliCloud Business Center Platform and check the status of related services 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. For more information, see Unavailable services and their impacts.

Unavailable services and their impacts

Service	Key impact
middleWare.dubbo	Deployment is affected, and the service is unavailable.
middleWare.tair	Deployment is affected, and the service is unavailable.
middleWare.metaq (message midddleware)	Deployment is affected.
middleWare.zookeeper	Deployment is affected, and the service is unavailable.
middleWare.jmenvDiamondVips	Deployment is affected, and the Diamond configuration item cannot be obtained.
ram.ramService (RAM)	The RAM user is unavailable.
webapp.pop (API Gateway)	The OpenAPI service is unavailable.

Service	Key impact
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 ApsaraDB for RDS association requests become invalid.
tianjimon (Monitoring System)	Some services are unavailable.

8.3.3. Inspection

8.3.3.1. Overview

ESS inspection monitors the basic health conditions of clusters.

The inspected basic health conditions include the following aspects:

- Monitoring inspection
- Basic software package version inspection

8.3.3.2. Monitoring inspection

The monitoring inspection includes the basic monitoring and connectivity monitoring inspection.

8.3.3.3. Basic software package version inspection

The basic software package version inspection includes the version inspection for trigger, coordinator, worker, and base services.

8.4. Resource Orchestration Service (ROS)

8.4.1. ROS component O&M

8.4.1.1. API Server

The API Server is used to receive ROS requests, send requests to RabbitMQ clusters, and send the responses returned by the Engine Server to callers. The API Server is used to connect the frontend and backend business.

• Components

The Engine Server and API Server share three servers, all of which are attached to a special SLB instance.

- O&M methods
 - The storage path of the API Server information is /home/admin /ros-server/bin/.

- The basic operation of the API Server: #/usr/local/ros-python/bin/python /home/admin/ros-ser vice/bin/ros-api {stop|status|--daemon}
 - Stop : stops the API Server.
 - Status : queries the status of the API Server.
 - --daemon : starts the API Server in daemon mode.
- Health criteria
 - Intrinsic availability: The CPU usage and system memory are within the normal range. The API Server is running normally.
 - Associated component availability: ROS is available.

8.4.1.2. Engine Server

The Engine Server is used to process stack requests. It shares the three servers with the API Server.

- O&M methods
 - The storage path of the API Server information is /home/admin /ros-server/bin/.
 - You can run the following command to perform operations on the Engine Server: /usr/local/ ros-python/bin/python /home/admin/ros-service/bin/ros-engine {stop|status|--daemon}
 - Stop : stops the Engine Server.
 - Status : queries the status of the Engine Server.
 - --daemon : starts the Engine Server in daemon mode.
- Health criteria
 - Intrinsic availability: The CPU usage and system memory are within the normal range. The API Server is running normally.
 - Associated component availability: ROS is available.

8.4.1.3. RabbitMQ clusters

RabbitMQ clusters are used to receive requests from the API Sever and responses from the Engine Server.

• Components

RabbitMQ clusters are composed of nodes.

RabbitMQ clusters are used for messaging. Nodes in the clusters use disks for non-persistent storage. Messages are written into the queues that correspond to the nodes. Nodes in a cluster can communicate with each other. Typically, to ensure data accuracy, the minimum number of working nodes is set to [Total number of nodes/2] rounded up. If data of nodes are inconsistent, the secondary nodes synchronize queue messages from the primary nodes.

• O&M methods

The storage path of the RabbitMQ information is /opt/rabbitmq-server/.

Common Rabbit MQ commands are as follows:

• You can run the following command to query the cluster status: sudo /usr/local/sbin/rabbitm g-server/sbin/rabbitmgctl cluster_status

[root@dataset_]
#/usr/local/sbin/rabbitmqctl cluster_status
Cluster status of node ros_rabbit@docker011165194088
[{nodes,[{disc,[ros_rabbit@docker011165194088]},
{ram,[ros_rabbit@docker011165194091]}]},
<pre>{running_nodes,[ros_rabbit@docker011165194091,ros_rabbit@docker011165194088]},</pre>
{cluster_name,<<"ros_rabbit@docker011165194088">>},
{partitions,[]}]

- Nodes : indicates the nodes in the cluster.
- Disc : indicates that the cluster uses disks for storage.
- Mem : indicates that the cluster uses memory for non-persistent storage.
- Running_nodes : indicates the information of the running nodes in the cluster.
- Partition: indicates the partitions of the cluster. If the value field is brackets [], the cluster has no partitions. If this parameter is not empty, the cluster nodes are divided into several partitions.
- You can run the following command to query the virtual hosts in a cluster: sudo /usr/local/s bin/rabbitmqctl list_vhosts



Typically, there are two virtual hosts. One is displayed as a forward slash (/), and the other is named based on the region where it resides.

- Health criteria
 - Intrinsic availability: The CPU usage and system memory are within the normal range. RabbitMQ is running normally, which indicates that clusters have no partitions, queues are properly processed, and messages are properly consumed.
 - Associated component availability: ROS is available.

8.4.1.4. Notify Server

The Notify Server is the proxy server for ECS instances that reside inside a VPC. It sends the execution status and information of operations on ECS instances to ROS.

Components

The Notify Server consists of three servers, all of which are attached to a Server Load Balancer (SLB) instance.

• O&M methods

For example, the virtual IP address of the SLB instance is 10.152.XX.XX, you can run curl http://10.152.XX.XX:80/health-check to check whether the Notify Server is running.

- Health criteria
 - Intrinsic availability: The CPU usage and system memory are within the normal range.
 - Associated component availability: ROS is available.

8.5. Object Storage Service (OSS)

8.5.1. Log on to the Apsara Stack Operations console

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

Prerequisites

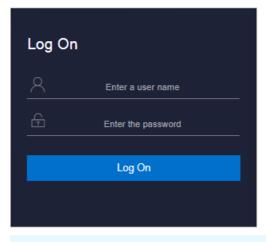
• The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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



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

3. Enter your username and password.

(?) Note Obtain the username and password for logging on to the ASO console from the deployment personnel or an administrator.

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

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

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- It must be 10 to 20 characters in length.
- 4. Click Log On to go to the ASO console.

8.5.2. OSS operations and maintenance

8.5.2.1. User data

8.5.2.1.1. 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 V alijuntest View
Product List	Bucket Basic Information User Data Overview	Data Monitoring
✓ ECS	View Procket	
ECS Operations and <	ECS Operations and Maintenance Platform	
Image Upload	Bucket Name:	al6a-005
RDS	DUCKH NAMH.	atis# 455
▼ OSS	User Account:	allyuntes(199999999)
User Data	Enterprise/Individual Name:	
Cluster Data	Application:	fie
✓ MPS User Configurations	BID:	26842
Batch Retranscoding	Current Capacity:	Standard: 1994/20173468, IA: 06, AR: 06
Apsara Distributed Fil	Storage Type:	standard
ISV Access Configur	Objects:	Standard: 2504947, U-C 0, AP: 0
	Log Service:	inactivated
	Location:	oss-cn-qingdao env40-401 a/Ossih)bridCluster A-20191028-eac5)
	Permissions:	private
	Anti-Leech Settings:	ConfiguredRutes [Empty Refer: Allowed] [Refer List: 0 Entries]
	CORS Settings:	0 Ruler: Configured
		A Philip Desidence #

8.5.2.1.2. 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

> Document Version:20200918

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.

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 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.

Products											
Products		Alibaba Cloud	d Account 🛛 🗸	< aligunia	HE1					View	
Product List	Bucket Basic Information				ata Monito	rina					
- ECS	Salact Data										
ECS Operations and	ECS Operations and Maintenance Platform	m	View								
Image Upload											
BMS	Total Storage Capacity				d Traffic		Total Outbou	ind Traffic		Total Charged Requ	lests
RDS	0B				0B			0B		0	K l
OSS	STD IA			Public	Internal	CDN	Public		CDN	QPS 1 Requests	QPS 2 Requests
User Data	08 08			Network	Network 0B	0B	Network	Network 0B		0Ten thousand	OTen thousand
Cluster Data				UD	08		08	00			
- MPS	• Show Storage Data Show Traffic Da	ta									
User Configurations	Bucket Name	Regi	ion			Total Storage Capacity	Standard	IA (occupied size) IA (charged)	Archive (occupied size)	Archive (charged)
Batch Retranscoding											
Monitor Video											

4. Set Date. Click OK. Click View, as shown in the following figure.

8.5.2.1.3. 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.

? Note By default, this metric is not displayed. You can select this metric from the Monitoring Items drop-down list.

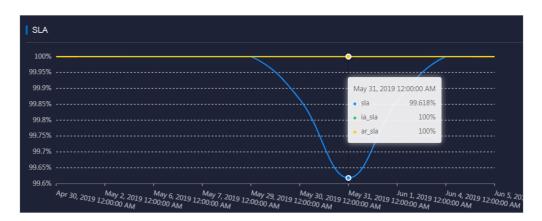
- Traffic: collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
- QPS: collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, 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.





• Move the pointer over the trend chart to display data at a specific point in time.

Data monitoring 2



8.5.2.2. Cluster data

8.5.2.2.1. Inventory monitoring

Metrics of inventory monitoring include the total capacity, available capacity, used capacity, backup ratio, and inventory usage.

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 Inventory Monitoring tab, you can view statistics by Apsara Distributed File System, metric data, or KV data usage.

Products	Cluster Data									
Product List		Bucket Statistics	Object Statistics	Data Monitoring	Resource Usage Ran	king				
✓ ECS ECS Operations and	Report Type: Storage Inventory Data Dimension: Apsara Distribute	d File System Data 🗸 Statistica	I Time: Jan 22, 2020 🛗	View				\$	Sampling Time: Jai	n 22, 2020, 14:24:11 C* Refresh
Image Upload							Data Increment(TB)			
RDS	Region T	Cluster ↓]*	Total Capacity(TB).↓h	Used Capacity(TB) ↓	Unused Capacity(TB) 1	Utilization 11				Actions
▼ OSS	cn-qingdao-env4b-d01	osshybridcluster-a-20191028-e ac5	505.39	40.25	465.14	7.96%		0.42	8.06	Show Details
User Data Cluster Data	cn-qingdao-env4b-d01	osshybridduster-a-20191028-e b52	519.83	26.84	492.99	5.16%	-0.02			Show Details
▼ MPS										<pre> Prev 1 Next > </pre>
User Configurations Batch Retranscoding Agears Distributed F.E HSV Access Configur	The marks: 1 Prevaining days of 2 provide the space of	peak incoment is calculated forced in when the Ayourt Datituded File Syste	n 190% of the ductor shorage: en utilization it 70%–15%, yellow w	hen the utilization is over 85%, and i	ed when Accara Debribuled File Sys	tem sopires in 30 days of the physica	il space of Apsara b	Distributed File Sys	tem is two linnes har	per han the OSS logical

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.
- $\circ~$ KV Data: includes the logic KV data, KV data in the recycle bin, and data increment (by

day, week, or month).

8.5.2.2.2. Bucket statistics

This topic describes how to collect statistics for the number of buckets 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 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	Resource Usage Ranking	
✓ ECS	Display Method: Report Statistics V Specify Time Range: 01/15	/2020 - 01/22/2020 📷 View		
ECS Operations and	Display meniod. Report Statistics V Speciny nine Range. Viria			
Image Upload	Region	Cluster	Active Users	Active Buckets
RDS	cn-qingdao-env4b-d01	o:=8-eb52		29
✓ OSS	cn-qingdao-env4b-d01	o:		1961
User Data	Total Number v	vilhout Duplicates		1988
▼ MPS				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 mo nths*, *six months*, or *one year*.
- 4. Click View.

8.5.2.2.3. Object statistics

This topic describes how to view the statistics for the number and trend of objects 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 **Object Statistics** tab, select **Current Overall Statistics** or **Growth Trend** to view object statistics.

Operations and Maintenance Guide · Operations of basic cloud products

Products	Cluster Data			
Product List	Inventory Monitoring Bucket Statistics 0	bject Statistics Data Monitoring Resource Usage Ran	king	
ECS Operations and	Display Method: Current Overall Statistics A			
RDS	Growth Trend Region	Cluster	Objects	
• OSS	cn-qingdao-env3b-d02	o; -285b	113991990null	
User Data	cn-gingdao-env3b-d02	os 867	11804840null	
Cluster Data	cn-gingdao-env4b-d01	oc eac5	40107104null	
• MPS	cn-qingdao-env4b-d01	oceb52	12554185null	
Batch Retranscoding		Total	178458119null	
Apsara Distributed Fil				✓ Prev 1 Next >
ISV Access Configur	ff Active User Comparison			
	Ten thousand			
	Ten thousand		·····	
	Ten thousand			
	Ten thousand			
	Ten thousand	/Ten thousand		
	Ten thousand			
	Ten thousandOsshybridcluster-a-20191030	osshybridcluster-a-2019107a	cn-qingdao-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 mo nths*, *six months*, or *one year*.
- 4. Click View.

8.5.2.2.4. Data monitoring

This topic describes how to collect statistics for each metric by cluster.

Context

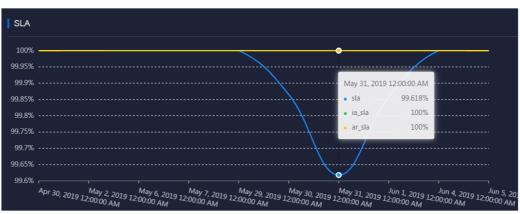
Cluster data metrics are similar to user data metrics except that the object of cluster data metrics is the data collected by cluster.

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.

- ? 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%.
 - 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.



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.

⑦ Note 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.



Data monitoring 2

• Move the pointer over the trend chart to display data at a specific point in time.



Data monitoring 2

8.5.2.2.5. 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

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 Resource Usage Ranking tab, select Report or Trend from the Display Mode dropdown list. Select a number from the Top drop-down list. Set Specify Time Range and Monitoring Items to view resource usage.

Products	Cluster Data						
Product List	Inventory Monitoring Bucket St	atistics Object Statistics	Data Monitoring	Resourc	e Usage Ranking		
← ECS	Display Method Top Specify Time	e Range					
ECS Operations and	Report V 10 V 01/22/2020	12:16:06 - 01/22/2020 14:31:06					
Image Upload	Monitoring Items (11/11) Total Requests × Request Errors × Public Inbour	of Traffic X. Public Outbound Traffic X. Inter	nal inhound Traffic X		nk Traffic X CDN Downlink Traffic X Storage	Canarity × Storage Increment × Storage F	Increment X
RDS	Total Holdeston Hindester Elling and Hispan					outering and an an and an and an	
▼ OSS	Total Requests-TOP10				Request Errors-TOP10		
User Data	Bucket	UD	Total Requests		Bucket	UD	Request Errors
Cluster Data							Request Errors
▼ MPS			3.46Ten thousand				6289
User Configurations			3.36Ten thousand				0
Batch Retranscoding			1.57Ten thousand				o
ISV Access Configur			1.12Ten thousand				o
			1.02Ten thousand				0
			4456				0
			4132				0
			2340				o
			2064				0
			© 2009-2019 Alibaba Cloud Computing Limited. All ri	ights reser	rved.		

- In report mode, you can view the top 10, 30, or 50 buckets by resource usage.
- $\circ~$ In trend mode, you can view the top 10 buckets by resource usage.
- 4. Click View.

8.5.3. Tools and commands

8.5.3.1. 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
```

8.5.3.2. 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 mod_nginx on mod_nginx on 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, mod_tcpx,mod_apache,mod_pcsw,mod_io,mod_percpu,mod_nginx,mod_tcprt

8.6. Table Store

8.6.1. Table Store Operations and Maintenance

System

8.6.1.1. 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.

8.6.1.2. User data

8.6.1.2.1. 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.

Operations and Maintenance Guide · Operations of basic cloud products

ilte	r: Instance	ID 🔻		Searc	:h					
egi	on: cn-qin;	pdao-envilid-d01	(APSARA_STACK) •	Cluster: ots-ssd-a-2018	1031-25ce •	Search				
в	atch Update								Add Instance	Refre
						Click to v	iew the	instance infor	mation pag	e.
	Region	Cluster	User ID	Instance ID	Instance Name	Туре	At	Description	Opera	
	-									

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.

ilte	: Instance ID	•	_qFtsUsYyICdR7stgs2	A. Sea	irch					
Regi	on:		٣	Cluster:		Ŧ	Search			
В	atch Update								Add Instance	Refres
	Region	Cluster	User ID	Instance ID	Instance Name	Туре	Modified At	Description	Opera	tion
	tri-gingdat- env68-d01	tionji-a- 25ee	1365544539690241	_qRirUnY/CdKTiqL2A	asootsins	INTERNAL	2018-12- 11		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.

🛱 Instar	nce asootsi	ns
Details	Tables	
Monitorin asootsins M		Click to view the instance monitoring page.
Endpoint		
Public Netw	ork:	
Private Netv	work	and the couple with the subconverting on
Table Stat	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.

Stance odps								
Details Tables								
Table Name	Max Version	TTL(s)		Write CU 👻	Partitions	Data Size ▼	Pangu Data Size ▲	Timestamp
ODPS_META_X_META_HISTORY	view the table	-1	0 0	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 Click to view the table monitoring page.	
Monitoring ODPS_META_X_META_HISTORY Monitoring	
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.

_								
De	etails Partitions							
are	ch: Worker 🔻			Search				
		Start			Pangu File Size	Data Size	Youchao	
ID	Partition ID	Key	End Key	Worker	^	Ŧ	Files 🔻	Timestamp
4	1891d981-771c-45af-b239- 84312b750ba9		\xfd\xfd\xfd\xf.	a36f01001.cloud.f01.amtest10	59.4MB	0B	9	2019-02-0

The available filtering conditions include:

- Worker (For more information, see the value in the Worker column.)
- Partition ID

8.6.1.3. Cluster management

8.6.1.3.1. Cluster information

You can obtain cluster information through cluster searches, cluster usage, and top requests.

Function description

• Clusters

Cluste	r Information			
Region: A	11	v	OCM Cluster Synchronization	Refresh
Status	Cluster	Region	Storage Type	Operation
	Click	to view the cluster infor	rmation page.	
using	ots-hy-a-20181217-2e46	cn-qingdao-env8	HYBRID	Delete
		an ainadan anu0	CCD	Delete
using	ots-ssd-a-20181031-25ce	cn-qingdao-env8	SSD	Delete
		an aire dan anu0		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.

Cluste	er Information			
Region:	All	•	OCM Cluster Synchronization	Refresh
Status	Cluster	Region	Storage Type	Operation
using	ots-hy-a-20181217-2e46	Click to view the cluster i cn-qingdao-env8.		Delete
using	ots-ssd-a-20181031-25ce	cn-qingdao-env8.	SSD	Delete
using	tianji-a-25ee	cn-qingdao-env8.	HYBRID	Delete

• Cluster details

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-20	181217-2e46
Overview Top R	lesource Usage
*Region: cn-qingdao-anvõ	did01(AFSANA_STACK) Cluster: db-hy-a-20181217-2x46 Switch Cluster
Region Description	APSARA_STACK
Region	cn-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 L	Jsage	Click-to	view the information p	hade of top request	c	
		ew the table infor	mation page.	Circle to				
Top Partit	ion <mark>s</mark> by F	Pangu File Size			Top Partitions by	Youchao Files		
	/		Click to view the p	oartition info	rmation page.			Мо
Table N	ame	Partition ID	Pangu File Siz	ze 🔻	• Table Name	Partition ID	Youchao Files 🔻	r
Top Table	s by Pan	gu File Size			Top Tables by You	uchao Files		
	📕 Click t	o view the instan	ce information p	age More				Mo
	Name	Table Name	Panou File S		Instance Name	Table Name	Youchao Files	

• 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	hy-a-	20181217-2	3e46						
Overviev	v T	op	Resource Us	age	Cli	ck to ma	nually collect	resourc	e usage in	formation
Collect	ed At:	~								Collect Data
Check	Result	:								
Storage	e Reso	urce l	Jsage							
	Total Disk Size Total File Size		ize	Recycle Bin Size		Table Size F		pace I	Disk Usage Ratio (%)	
									•	%
Gap Siz	ze	Hosts	Total/Master/O	TSServer/S	qlWorker	Ну	vbrid Deployment	Clus	ster Type	Scale-out Requirement
		111								
OTSSer	ver Re	sourc	e Usage							
Hosts	Failed Hosts		rg/Max CPU sage (%)	Increase CPU Cor			ncreased Hosts Due Excessive NetIn		Avg/Max NetOut (MB/s)	Increased Hosts Due to Excessive NetOut
		1			1			1		

8.6.1.4. Inspection center

8.6.1.4.1. 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

anji-a- Sao	2019- 02-02	64.46TB	6.21TB	3.25TB	1.64TB	1.32TB	48.80TB	24.31%	15日本語、Reach Safe Level in -1Days, Growth Rate:-35.27GB/Days 3日本語語 Reach Safe Level in -1Days, Growth Rate:-35.28GB/Days
	Date	Total Disk Size	Total File Size	Gap	Recycle Bin Size	Table Size	Free Space	Disk Usage Ratio (%)	Scale-out Requirement
Cluster Name						Abnorm	al Resource	e Usage	
									Collect

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.

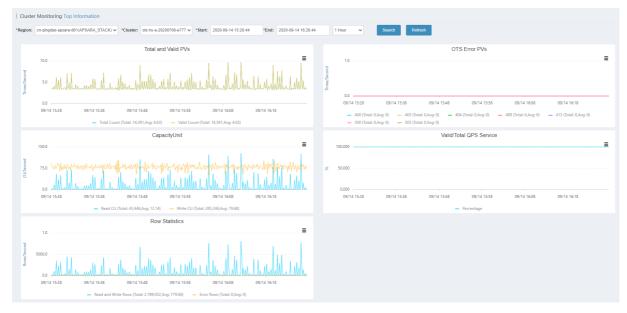
8.6.1.5. Monitoring center

8.6.1.5.1. Cluster monitoring

You can determine the service status of a cluster based on a series of metrics such as clusterlevel 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.



8.6.1.5.2. 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.

? Note The Instance field is required. Table a	and Operation fields are optional.
*Region: On-singdae-onalid-d01(APSARA_STACK) *Cluster: tanj-a-25es ODPS_META_X_CHANGE_LOGS * Operation: *Start: 2019-02-02 12:37:09 *End: 2019-02-02 13:37:09 1 Hout	▼ *Instance: odps Table: ▼ ur ▼ Search Refresh
QPS - ODPS_META_X_CHANGE_LOGS	Error QPS - ODPS_META_X_CHANGE_LOGS 1.0 0.0 02/02 12:37 02/02 12:57 02/02 13:07 02/02 13
CapacityUnit - ODPS_META_X_CHANGE_LOGS 15.0 7.5 0.0 0.2/02 12:37 Read CU (Totat: 1,017,Avg: 0.29)	Net Flow - ODPS_META_X_CHANGE_LOGS 11.7KB 2019-02-02 12:37:50 08 Internal Network Input (Total: 1,873,195,Avg: 541.39): 08 Internal Network Output (Total: 101,811,Avg: 29.43): 08 Internal Network Output (Total: 101,811,Avg: 29.43): 08 Internal Network Output (Total: 0,Avg: 0): 08 Internal Output User Data (Total: 0,Avg: 0): 08 Public Internal Output User Data (Total: 1,588,655,Avg: 459.15): 08 Internal Output User Data (Total: 0,Avg: 0): 08

Operations and Maintenance Guide · Operations of basic cloud products

RowsCount - ODPS_META_X_CHANGE_LO 24.0 2019-02-02 14:16:30 0.0	2/02 14:17	1B Satos OB	TableStorage - ODPS_META_X_CHANGE_LOGS Table 3213:28 02/02 13:38 02/02 13:48 02/02 13:58 02/02 14:08 02/02 14:18 Table 5ize (Avg:0)
Return Code	Times		Ratio
200	1,680		100.00%
Total	1,680		100.00%
SQLWorker Error Distribution (Rows) - ODPS_META_X_CHANGE_LOGS			
Return Code	Times		Ratio
0	1,680		100.00%
Total	1,680		100.00%

8.6.1.5.3. 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.

Region: m-qingda	e-enind-d01jA	SARA STAC	*Cluster	: Sanji-a-25e	• •	*Time: 20	19-02-02 12:40:15	9	2019-02-0	2 13:40:19	1 Hour	Y	
Monitoring Level:	Instance	•	*SortBy: Tot	al Requests	۲	"TopN: 100			Search				
op 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=m etric	1,643,542	73,033,4 06	0	0B	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	186,686	185,768	0	08	08	45.4MB	100.7MB	180,059	11,366	203,426 US	203,268 us	{"200":186686}	{"0":186686}

8.6.1.5.4. 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.

Operations and Maintenance Guide · Operations of basic cloud products

Region: m-engdad	-eni8d-d01(APSARA_STACK) *	*Cluster: tanji-a-25ec	* *Request ID:	Search
og Search Result				
og Search Result				

8.6.1.6. System management

8.6.1.6.1. 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 back-end tasks:

- View task details such as the specific parameters and running time of each task.
- Enable or disable a task.

Onte 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.

S 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/rur	n
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	

8.6.1.6.2. 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.

Operations and Maintenance Guide · Operations of basic cloud products

All Tasks	Host VIP/Net	Application	Resource Usage	Remote HTTP		
ime Range:	2019-02-02	To 2019)-02-02	Check		
Status	Name	Туре	Starte	ed At	Ended At	Operation
Abnormal	collect_water_level	Remot	e HTTP 2019-	02-02 06:00:00	2019-02-02 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remot	e HTTP 2019-	02-01 06:00:00	2019-02-01 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remot	e HTTP 2019-	01-31 06:00:00	2019-01-31 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remot	e HTTP 2019-	01-30 06:00:00	2019-01-30 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remot	e HTTP 2019-	01-29 06:00:00	2019-01-29 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remot	e HTTP 2019-	01-28 06:00:00	2019-01-28 06:00:10	View All View Exceptions
Abnormal	collect water level	Remot	e 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.

collect_water	_level task res	ult			
total 1 count, 0 e	execute success, /	1 execute fail, 1 e	execute warnin	ng	
Executelp	StartTime	EndTime	TaskResult "env: APSARA_STACK, inner task collect water	Warning env: APSARA_STACK, inner task collect water	IsSuccess
HTTP	Feb 2, 2019 2:00:00 AM	Feb 2, 2019 2:00:10 AM	[ots-hy-a-20181217-2e46,	level fail: Trigger collect water level fail, cluster list: [ots-hy-a-20181217-2e46, " ots-ssd-a-20181031-25ce]	fail

8.6.1.7. Platform audit

8.6.1.7.1. 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.

Operations and Maintenance Guide · Operations of basic cloud products

ime Range: 2018-12-31 00:00:00 To 2019-0	2-02 01:05:00 Add Condition Ope	rator 🔻		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	2019-01-18 13:42:54
/ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:42:53
ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:42:53
ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:39:38
ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:39:37
/ots/apsarastack/v1/user/instance_list.json?prev	get_user_instance_list	10.148.64.187	alivuntest	2019-01-18 13:36:08

8.6.2. 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

8.6.3. System roles

- TableStoreOCM
 - OCMInit: the OCM initialization tool used to create tables and bind POP APIs
 - $\circ~$ 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
 - $\circ~$ 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

8.6.4. Pre-partition a table

8.6.4.1. 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

Once You can also use the preceding methods to partition a table that already has data.

8.6.4.2. 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.

8.7. ApsaraDB for RDS

8.7.1. Architecture

8.7.1.1. System architecture

8.7.1.1.1. 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.

8.7.1.1.2. 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.

8.7.1.1.3. 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.

8.7.1.1.4. 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.

8.7.1.1.5. 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.

8.7.2. Log on to the Apsara Stack Operations

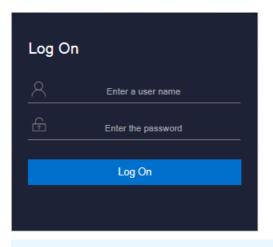
console

Prerequisites

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

Procedure

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



? 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.

? Note Obtain the username and password used to log on to ASO from the deployment personnel or the administrator.

- 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.

8.7.3. 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 **Instances**.

Instances



• View instance details

Click the ID of an instance to view details, as shown in Instance details. You can switch your service between primary and secondary instances and query history operations on this page.

? Note 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 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 AE -
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 💭 Note:	Saturday 🔽 Sunday
Create Single Database Backup	

Instance details

• View user information

On the Instance Management tab, click **User Information** in the **Actions** column corresponding to an instance, as shown in **User information**.

IISer	into	rmation
Obci		mation

User Information 5								
						User Info	mation:	
		Database Typ e	Instance Usa ge Type					
	CREATING	Redis	-	- %		%		- %
	CREATING	Redis	-	- *		- %		- *
	CREATING	Redis	-	- *	- 5	- %		- *
	CREATING	Redis	1000	- *		- %		
	CREATING	Redis	1000	- *	- 5	- %		- *
	CREATING	Redis	1000	- *		- %		- s
Contract Description	CREATING	Redis	10000	 - s	 - x	 - x		- s

• Create backups

For ApsaraDB RDS for MySQL instances, click **Create Backup** in the **Actions** column to view the backup information, as shown in **Backup information**. You can also click **Create Single Database Backup** on the Backup Information page to back up a single database.

	Backup 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:	day 🔽 Sunday
Create Single Database Backup	

Backup information

8.7.4. Manage hosts

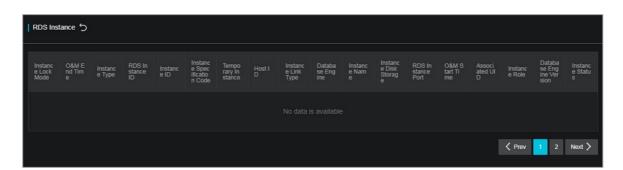
You can view and manage hosts.

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							
Instance Manag	ement Host N	lanagement					
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
	Normal offline	cn-qingdao-env8d-d0 1				5.6	MySQL
	Normal offline	cn-qingdao-env8d-d0 1				5.6	MySQL
	Normal offline	cn-qingdao-env8d-dD 1				5.6	MySQL

4. Click a hostname to go to the **RDS Instance** page. You can view all instances on this host.



8.7.5. Security maintenance

8.7.5.1. 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.

8.7.5.2. 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.

8.8. AnalyticDB for PostgreSQL

8.8.1. 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

> Document Version:20200918

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.

8.8.2. Architecture

Physical cluster architecture

The following figure shows the physical cluster architecture of AnalyticDB for PostgreSQL.

Physical cluster architecture

Instar	nce 1	Insta	nce 2	Instar	nce 3			
Coordinator Node (Primary)	Coordina (Seconda	tor Node Iry)	Compute (Primary)	Node 1	Compute (Seconda			
Compute Node 2 (Primary)	Compute (Seconda		Coordina (Primary)			Coordinator Node (Secondary)		
Compute Node 1 (Primary)	Compute (Seconda		Compute (Primary)	Node 2		Compute Node 2 (Secondary)		
Compute Node 2 (Primary)		Compute Node 3 (Secondary)		Node 4		Compute Node 4 (Secondary)		
Coordinator Node (Primary)		Coordinator Node (Secondary)		Node 1		Compute Node 1 (Secondary)		
Compute Node 2 (Primary)	Compute (Seconda							
Database Server 1		tabase rver 2		itabase erver 3		atabase erver 4		

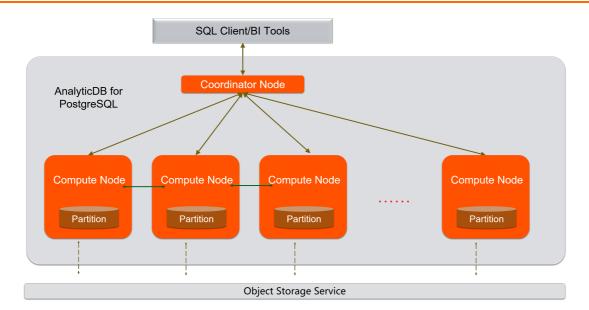
You can create multiple instances within a physical cluster of AnalyticDB for PostgreSQL. Each cluster includes two components: the coordinator node and the compute node.

- 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.

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.

8.8.3. Routine maintenance

8.8.3.1. 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.
 - i. 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. ii. 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_random('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.
 - i. 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.

ii. 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.

8.8.3.2. 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.

8.8.4. Security maintenance

8.8.4.1. 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.

8.8.4.2. Account password maintenance

Account passwords include the superuser password of AnalyticDB for PostgreSQL and the password of the host operating system.

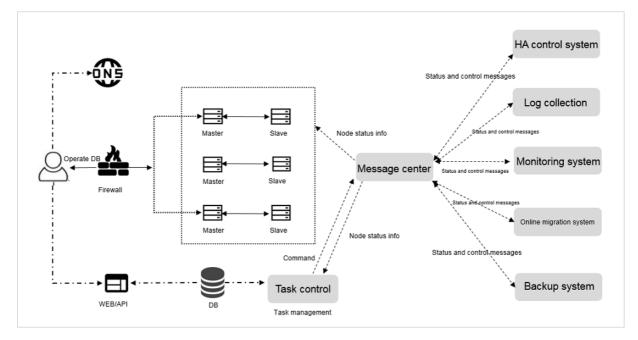
To ensure account security, use complex passwords and periodically change the passwords of systems and devices.

8.9. KVStore for Redis

8.9.1. O&M tools

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.



8.9.2. Architecture diagram

8.9.3. Architecture

8.9.3.1. Architecture

8.9.3.1.1. 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.

8.9.3.1.2. 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.

8.9.3.1.3. 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.

8.9.3.1.4. 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.

8.9.3.1.5. 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.

8.9.4. Log on to the Apsara Stack Operations console

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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

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

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

3. Enter your username and password.

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

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

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

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- It must be 10 to 20 characters in length.
- 4. Click Log On to go to the ASO console.

8.9.5. Instance management

You can view instance details, logs, and user information.

Procedure

- 1. Log on to Apsara Stack Operations.
- 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:
 - $\circ~$ 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.

8.9.6. Host management

Host management allows you to view and manage hosts.

Procedure

- 1. Log on to Apsara Stack Operations.
- 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							
Instance Managen	nent Host M	lanagement					
Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine
different sources and come			$\frac{1}{2} = 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1$				
	mand problem						
							-
CONTRACTOR OF T							
		-			-		-
					-		
estiv			wi				
		© 2009-2018 Alibaba (loud 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	ince 🖒															
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	O&M Start Time	Instance Role	Datab Engine Version	Instance Status
				-	-				-					-		-
		<u> 1997</u>		-	<u>.</u>				-			-				
				-	***				10.00							1946
			-	-					-					-		1995
									-					 2		ine
				62	009-2018 Alba	aba Cloud Con	nputina Limite	d. All rights res	erved.							

8.9.7. Security maintenance

8.9.7.1. Network security maintenance

Network security maintenance involves device security and network security.

Device security

Check network devices, and enable security management protocols and configurations 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.

8.9.7.2. 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.

8.10. ApsaraDB for MongoDB

8.10.1. Service architecture

8.10.1.1. System architecture

8.10.1.1.1. Backup system

Automatic backup

ApsaraDB for MongoDB supports both physical backup and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can initiate a temporary backup as required. The backup files are retained for seven days.

Log management

ApsaraDB for MongoDB generates operation logs and allows you to download them. You can use the operation logs for local incremental backup.

Data backtracking

ApsaraDB for MongoDB can use backup files and logs to generate a temporary instance for any time point within the past seven days. After verifying that the data in the temporary instance is correct, you can use the temporary instance to restore data to the specified time point.

Creating a temporary instance does not affect the running of the current instance.

Only one temporary instance can be created for each ApsaraDB for MongoDB instance at a time. A temporary instance is valid for 48 hours. You can create a maximum of 10 temporary instances for an ApsaraDB for MongoDB instance each day.

8.10.1.1.2. Data migration system

Database replication between instances

ApsaraDB for MongoDB allows you to easily migrate databases from one instance to another.

Data migration to or from ApsaraDB for MongoDB

ApsaraDB for MongoDB provides a professional tool and a migration wizard to help you migrate data to or from ApsaraDB for MongoDB.

Backup file download

ApsaraDB for MongoDB retains backup files for seven days. During this period, you can log on to the ApsaraDB for MongoDB console to download the backup files.

8.10.1.1.3. Monitoring system

Performance monitoring

ApsaraDB for MongoDB provides nearly 20 metrics for monitoring system performance, such as the disk capacity, IOPS, number of connections, CPU utilization, network traffic, transactions per second (TPS), queries per second (QPS), and cache hit rate. You can obtain such status information for an ApsaraDB for MongoDB instance within the past one year.

SQL auditing

> Document Version:20200918

The system records SQL statements and additional information sent to ApsaraDB for MongoDB instances, such as the IP addresses of connections, database names, access accounts, execution time, and number of records returned. You can use SQL auditing to locate problems and check instance security.

Threshold alerting

ApsaraDB for MongoDB provides short message service (SMS) notifications to indicate status or performance exceptions that occur in ApsaraDB for MongoDB instances.

These exceptions include instance locking, disk capacity, IOPS, connection quantity, and CPU exceptions. You can configure alert thresholds and up to 50 alert recipients (of which five are effective at a time). If a metric of an ApsaraDB for MongoDB instance exceeds a specific threshold, an SMS notification is sent to alert the recipients.

Web operation logging

The system logs all modification operations in the ApsaraDB for MongoDB console for administrators to check. These logs are retained for a maximum of 30 days.

8.10.1.1.4. Control system

If a host or an instance crashes, the ApsaraDB for MongoDB high-availability (HA) component fails services over within 30 seconds after the exception is detected. This guarantees that applications run properly and ApsaraDB for MongoDB is highly available.

8.10.1.1.5. Task scheduling system

You can use the ApsaraDB for MongoDB console or APIs to create or delete instances or switch instances between the intranet and Internet. All instance operations are scheduled, traced, and displayed as tasks.

8.10.2. ApsaraDB for MongoDB O&M overview

Apsara Stack Operations Console provides the following O&M features for ApsaraDB for MongoDB:

- Instance management: allows you to view instance details, instance logs, and user information.
- Host management: allows you to view and manage hosts.

8.10.3. Log on to the Apsara Stack Operations

console

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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

Enter a user name
Enter the password
Log On

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

3. Enter your username and password.

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

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

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

- It must contain uppercase and lowercase letters.
- $\circ~$ It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- It must be 10 to 20 characters in length.
- 4. Click Log On to go to the ASO console.

8.10.4. Manage ApsaraDB for MongoDB instances

This topic describes how to manage ApsaraDB for MongoDB instances. You can view instance details, logs, and user information.

Procedure

- 1. Log on to the Apsara Stack Operations console.
- 2. In the left-side navigation pane, choose Products > RDS to go to the RDS page. On the Instance Management tab that appears, you can perform these operations:
 - View the list of instances.

View the instances that belong to the current account on the **Instance Management** tab, as shown in **Instance list**.

Instance list

RDS										
Instance Management		Host Mana	agement							
Instance Name V Please		٥								
Instance Name	Availa	CPU Perfor	QPS Perfor	IOPS Perfor	Conne	Disk Usage	Instance Status	Database Type	Actions	
	Yes	1.6 %				0.5 %	Using	gpdb		
	Yes	1.5 %			4.5		Using	ppassql		
-	Yes	1.4 %				0.1 %	Using	ppassql		
	Yes	1.4 %				0.3 %	Using	gpdb		
	Yes	1.2 %				0.2 %	Using	gpdb		ation
-	Yes	1.11 %	5.62 %		0	1.3 %	Using	mysql	User Inform	ation

• View the details of an instance.

Click the ID of an instance to view its details, as shown in Instance details.

Instance details

Instance Details 5		
Instance Information		
	Instance Name:	CPU Performance: 0.4 %
	Active-Standby Delay: 0	QPS Performance: 0 %
	Connections: 0	IOPS Performance: 0 %
	Traffic: 0	Active Threads: 0
	Client Instance Level: P4	Instance Status: Using
	Database Version: 3.0	Link Type: Ivs
	Cluster:	Created At: 04/30/2019, 14:29:35
Network Details of Instance Host		
	Host IP Addresses:	Proxies:
	VIP IB_D List of SLB:	ECS-typed Dedicated Host of Client Instance: No
Network Details of Instance-Attache	d Host	
	Host IP Addresses: (MASTER) (SLAVE)	Proxies:
	VIP IB_ID List of SLB:	ECS-typed Dedicated Host of Client Instance: No

? Note On the Instance Details page, you can also perform active/standby switchovers and query historical operations.

• View user information.

Click User Information in the Actions column, as shown in User information.

User information

User Information 5								
								Information: 0_
Instance Name	Availability	CPU Performance	QPS Performance	IOPS Performance	Connections	Disk Usage	Instance Status	Database Type
	Yes					0.6 %	Using	mongodb
	Yes	78.8 %		0.4 %	0.5	51.1 %	Using	pgsql
	Yes	5.4 %			0.2	0.1 %	Using	gpdb

8.10.5. Host management

Host management allows you to view and manage hosts.

Procedure

- 1. Log on to the Apsara Stack Operations console.
- 2. On the Host Management tab of the RDS page, view information about all hosts.

RDS							
Instance Manage	Instance Management Host Management						
Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine
Photo and the second se	normal operation	(marine a	la mar			5.6	MySQL
100	normal operation	Lawrence .	la rear			5.6	MySQL
100 million and the	normal operation	pressore	la mare			5.6	MySQL
(22)	normal operation	(second s	la mare			5.6	MySQL
Part and a second	normal operation	[second	haranaa			3.0	MongoDB

3. Click a host name to go to the **RDS Instance** page. On this page, you can view all instances on this host.

RDS Instance	5								
Instance Lock Mode	O&M End Time	Instance Type	RDS Instance ID	Instance ID	Instance Specification Code	Temporary Instance	Host ID	Instance Link Type	Database Engine
	06:00	Primary Instance	417		dds.mongo.mid	No		lvs	MongoDB
	06:00	Primary Instance	420		dds.mongo.mid	No		lvs	MongoDB
	06:00	Primary Instance	1546		dds.mongo.mid	No		lvs	MongoDB

8.10.6. Security maintenance

8.10.6.1. Network security maintenance

Network security maintenance is aimed at ensuring device security and network security.

Device security

Check network devices, and enable security management protocols and configurations of devices.

Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner.

For more information about the security maintenance method, see the product document of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check Internet and intranet traffic and defend the network against abnormal behaviors and attacks.

8.10.6.2. Account password maintenance

Account passwords include the ApsaraDB for MongoDB system and device passwords.

To ensure account security, change the system and device passwords periodically, and use passwords that meet the complexity requirements.

8.11. Log Service

8.11.1. O&M methods

This topic describes two O&M methods of Log Service.

Log Service is deployed, operated, and maintained by using the Apsara Infrastructure Management Framework console. Log Service supports the following two O&M methods.

- Terminal: In the Apsara Infrastructure Management Framework console, you can use the Terminal to log on to the machine where Log service resides and view logs.
- Portal: The Portal provides a user interface for operating Log Service. The Portal complies with the standard Java applications of Alibaba Cloud.

Terminal

- Log on to the Apsara Stack Operations console. For more information, see Apsara Stack Main tenance Guide. You can navigate in the guide as follows: Operations of basic platforms > Apsara Stack Operations (ASO) > Log on to Apsara Stack Operations.
- 2. In the left-side navigation pane, click **Products**. The product list page is displayed.
- 3. Click**Apsara Infrastructure Management Framework** to go to the Apsara Infrastructure Management Framework console.
- 4. Choose Operations > Service Operations.
- 5. On the page that appears, find sls-backend-server in the Services column, and then click Operations in the Actions column.
- 6. On the Clusters tab, find the target cluster in the Clusters column, and then click Operations

in the Actions column.

7. On the Services tab, select the target server role, for example, sls-backendserver.ServiceTest#, and then click Terminal in the Actions column.

Services Machines Cluster Configuration O	peration Log Cluster Resource Service Inspection			
Server Role Enter a server role Q				Refresh
• sls-backend-server.DBPostHandler#	erver.DBPreHandler#	sls-backend-server.FuxiServiceSlsEtlFramework#	sls-backend-server.FuxiServiceSlsLoghubMaster4	*
• sls-backend-server.FuxiServiceSlsMeteringService#	sls-backend-server.FuxiServiceSlsPrestoWorker#	iServiceSIsQueryMaster# 📔 💿 sls-backend-server.FuxiSer	viceSlsQuotaServer# sls-backend-server.Fu	xiServiceSIsReplayWorker#
• sls-backer 1er.FuxiServiceSlsShennongWorker#	\bullet sls-backend-server.FuxiServiceSlsToolServiceWorker#) (\bullet sls-backend-server.FuxiServiceSlsToolServiceWorker#)	ver.InitSIsCluster#] • sIs-backend-server.Nginx#] • sI	s-backend-server.PackageManager#	ckend-server.RedisServer#
• sls-backend-server.ServiceTest# • sls-backend-serve	r.SlsConsole# sls-backend-server.SlsFastcgi# sls-backend-server.Sls	ImportOdpsScheduler# sls-backend-server.SlsLogtai	il# esls-backend-server.SlsScmc# esls-ba	ckend-server.SIsScmg#
sis-backend-server.SisTools# sis-backend-server.SisTools# All: 1 Normal (1) Reset	weby sh-backend-server.SkWebTools# sh-backend-server.ToolSe	rvice#		Diagnostic Mode:
Machines Enter one or more hostnames/IP addresses	Q			Batch Terminal
Machines	Server Role Status	Metric		Actions
0 vm/	Normal Details	View		Terminal 2 Restart Server Role

8. Log on to the machine through the Terminal and go to the corresponding file directory to view logs.

Portal

You can collect logs of your machine and send the logs to the Portal. Then, you can query, retrieve, and analyze these logs on the Portal.

- 1. Log on to the Apsara Infrastructure Management Framework console to obtain the endpoint of the Portal.
 - i. Log on to the Apsara Infrastructure Management Framework console. For more information, see Terminal.
 - ii. Choose Reports > All Reports. The All Reports page is displayed.
 - iii. In the report list, click Registration Vars of Services.
 - iv. In the Service column, click the 📃 icon, enter sls in the search box, and then click Apply

Filter.

Registration Vars of Services	a			≜ © 2
Service	Sen Sen	vice Registration	Cluster	Update Time
acs-acs_control	Contains •	acs.intra.env17.shuguang.com","cert.contr	And and a second second second	04/22/20, 11:42:54
acsTest		formance.acs-test.intra.env17.shuguang.co	and a finite of strength when	04/22/20, 13:31:11
ads-service	3 Apply Filter	env17.shuguang.com","ads_ag_ip":"10.17		04/22/20, 12:01:47
adsprecheck-service	10.000	CONTRACTOR OF TAXABLE		04/22/20, 14:19:27

v. Right-click the service in the Service Registration column of sls-backend-server, and then click Show More.

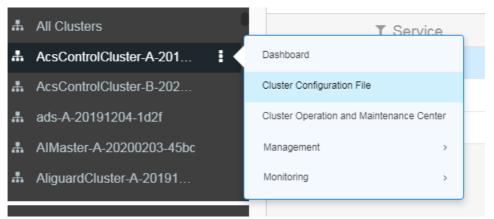
Registration Vars of Services			4 C Z
T Service	Service Registration	Cluster	Update Time
sls-backend-server	("cluster.name":"PublicBasicCluster-A-202003 Show More	P 00330-3588	04/22/20, 11:08:09
sis-common	{"cluster.name":"tianji-A-35fe","sls_admin_ak": Copy ,"s	11	04/22/20, 11:07:59

vi. On the Details page, find the endpoint of the Portal.

Formatted Value Original Value *Cluster.name": "PublicBasi *Sls_admin_skt': "20 *Sls_admin_skt': "20 *Sls_admin_sithout_owner_skt': "SAL *Sls_admin_without_owner_skt': "SaL *Sls_admin_stocked	tails	
<pre>"sls_admin_ak": "z8 "sls_admin_ak": "z8 "sls_admin_without_owner_ak": "SAI "sls_admin_without_owner_ak": "SAI "sls_admin_without_owner_sk": "slarid", "sls_cluster_name": "PublicBasicClu "sls_cluster_name": "PublicBasicClu "sls_configgerver.endpoint": "logtal.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_consle.endpoint": "sls.console.en-qingdao-env17-d01.inter.env17.shuguang.com", "sls_consle_vip": "1</pre>	Formatted Value 💿 Original Value	
<pre>"sls_admin_ak": "28 "sls_admin_ak": "28 "sls_admin_without owner_ak": "SAL "sls_admin_without owner_ak": "SAL "sls_duster_vip": "sls_conside.asiactick "sls_cluster_vip": "li"", "sls_configerver.endpoint": "logtal.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_conside_vip": "li"", "sls_conside_vip": "li", "sls_conside_vip": "li", "sls_conside_vip": "li", "sls_pon_topint": "data.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pon_endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pon_endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pon_endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pon_endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scat.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scat.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scat.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scat.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scat.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scated"</pre>		
<pre>"sls_admin_sk": "Oi "sls_admin_without_owner_ak": "SAL "sls_admin_without_owner_ak": "SAL "sls_admin_without_owner_sk": "saL "sls_cluster_wip": "PublicBasicClu "sls_cluster_wip": "lu ", " "sls_consigerver.endpoint": "logtail.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_console_vip": "lu ", " "sls_console_vip": "lu ", " "sls_console_vip": "lu ", " "sls_console_vip": "lu ", " "sls_console_vip": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.vec.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scm.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scm.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portel.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",</pre>	"cluster.name": "PublicBasi	
<pre>"sls_admin_without_owner_sk": "SAL "sls_admin_without_owner_sk": "SAL "sls_denin_without_owner_sk": "sls</pre>	"sls_admin_ak": "z8	
<pre>"sls_gadain_without owner_sk": "slsrid", "sls_cluster.name": "Public@asicClu "sls_cluster.name": "Public@asicClu "sls_configserver.endpoint": "logtail.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_console_vip": "l=, "sls_console_vip": "l=, "sls_console_vip": "l=, "sls_console_vip": "l=, "sls_pon_topint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pon_topint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pon_topint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pon_topint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pon_topint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pon_topint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scm.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_cm_scm.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",</pre>	"sls admin sk": "Oi	
<pre>"sls_cluster.name": "PublicBasicClu "sls_cluster_vip": "10"" "sls_contigserver.endpoint": "logtail.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_console_vip": "10" ", "sls_console_vip": "10" ", "sls_console_vip": "10" ", "sls_opo.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.vpc.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portel.endpoint": "scre.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scre.endpoint": "scre.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scre.endpoint": "scre.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scre.endpoint": "scre.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portel.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",</pre>	"sls_admin_without_owner_ak": "SAL	
<pre>"sls_cluster_vip": "1 "," "sls_configserver.endpoint": "logtail.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_console.endpoint": "sls.console.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_dtata.endpoint": "diata.endpindao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.vpc.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.tendpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.tendpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.tendpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.tendpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scated"</pre>	"sls_admin_without_owner_sk": "sle	
<pre>"sls_configserver.endpoint": "logtail.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_console.endpoint": "sls.console.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_data.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.tal.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_scm.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_portal.endpoint": "sls.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scm.endpoint": "sls.cn-qingdao-env17-d01</pre>	"sls_cluster.name": "PublicBasicClu	
<pre>"sls_console.endpoint": "sls.console.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_console_vip: "l= ", "sls_console_vip: "l= ", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop_vpc.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop_vpc.endpoint": "srs-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_ports.endpoint": "srs-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_ports.endpoint": "srs-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_wn_located"</pre>	"sls_cluster_vip": "1(",	
"sl= console_vip": "10 ", "sl=got_endpoint": "data.enqingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sl=pop.endpoint": "sls-cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "portal.en-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_wn_located"	"sls_configserver.endpoint": "logtail.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",	
<pre>"sls_data.endpoint": "data.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop_vpc.endpoint": "sls-vpc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "sorc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",</pre>	"sls_console.endpoint": "sls.console.cn-qingdao-env17-d01.inter.env17.shuguang.com",	
"sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_pop_vpc.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal endpoint": "portal cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_wn_located"	"sls_console_vip": "1 •,	
<pre>"sls_pop_vpc.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com", "sls_portal.endpoint": "portal.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc_cn-ninedao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_vm_located"</pre>	"sls_data.endpoint": "data.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",	
<pre>'sls_portal.endpoint": "portal.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmg.endpoint": "scmc.cn-ningdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_wn_located"</pre>	"sls_pop.endpoint": "sls.cn-qingdao-env17-d01.inter.env17.shuguang.com",	
"sls_scmc.endpoint": "scmc.en-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_scmg.endpoint": "scme ra.ningdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_wn_located"	"sls_pop_vpc.endpoint": "sls-vpc.cn-qingdao-env17-d01.inter.env17.shuguang.com",	
"sls_scmg.endpoint"' "scmg.cn-ningdao-env17-d01.sls-pub.inter.env17.shuguang.com", "sls_vm_located"	'sls_portal.endpoint": "portal.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",	
"sls_vm_located"	"sls_scmc.endpoint": "scmc.cn-qingdao-env17-d01.sls-pub.inter.env17.shuguang.com",	
	"sls_scmg.endpoint". "scmg_cn-gingdao-env17-d01.sls-pub.inter.env17.shuguang.com",	
"ele web vin": '	"sls_vm_located"	
SIS_RCD_VIP :	"sls_web_vip": '	

- 2. Log on to the Apsara Infrastructure Management Framework console to obtain the AccessKey pair of the Portal.
 - i. In the Apsara Infrastructure Management Framework console, choose 📩 > Cluster

Configuration File.



ii. Click kv.conf to obtain the AccessKey pair of the Portal.

lote] One page displays full information. Latest cluster configura	ation files int	tegrate the original service configuration and cluster configuration pages, and supports quick acces	is to files by file type. Do Not Show Aga
et's get started with more operations.			
File List @ 📄 Cluster 📕 Template	🗋 kv.	conf	
Create File	8	"HSM_CLIENI_KEY": "Empty", "HSM_SERVER_CERT": "Empty",	
Luster.conf	9 10	"HSM_SERVER_LIST": "Empty", "HSM_TLS":	
kv.conf	11 12	"REGION": " ROOT CERT": "Empty",	
	13	"ROOT_CERT_KEY_SLOT": "Empty",	
machine_group.conf	14	"account.acs.accesskey-id": "p	
norolling_config	15 16	"account.acs.accesskev-secret" "account.acs.id": "10	
	10	"account.acs.id : 10 "account.adminportal.accesskey-id": "7Wckp	
🕀 🗀 acs	18	"account.adminportal.accesskey-secret": "6	
plan.conf	19	"account.adminportal.id": "1	
	20	"account.ads.accesskey-id": "0fspJl	
🗅 🗀 services	21	"account.ads.access" "70",	
🕂 🗀 acs-acs control	22 23	"account.ads.id": ' "account.ads.passwc ,	
	25	"account.ads.user": "test10000	
🕀 🗀 hids-client	24	"account.all.accesskey-id": "	
	26	"account.all.accesskey-secret'	
🕂 🗀 os	27		and the second
🕀 🗅 tianji	28	"account.all.id": " ,	
	29	"account.all.user":	
🕀 🗀 tianji-dockerdaemon	30 31	"account.asm.id": "i	

- 3. Log on to the Apsara Infrastructure Management Framework console. Log on to the Portal by using the endpoint obtained in Step 1 and AccessKey pair obtained in Step 2.
- 4. Find the corresponding project and Logstore, and then query and analyze logs.

8.11.2. O&M

8.11.2.1. View logs on machines

InitSlsCluster#

- Startup log: /cloud/app/sls-backendserver/InitSlsCluster#/init_sls_cluster/current/log/start.log
- Service log: none

Nginx#

- Startup log: /cloud/app/sls-backend-server/Nginx#/nginx/current/log/start.log
- Service logs:
 - /apsara/nginx/logs/access.log
 - /apsara/nginx/logs/error.log
 - o /apsara/nginx/logs/fastcgi_agent_access.log
 - /apsara/nginx/logs/offline_access.log
 - o /apsara/nginx/logs/scmc_access.log
 - o /apsara/nginx/logs/scmc_err_log
 - /apsara/nginx/logs/scmc_op_log
 - o /apsara/nginx/logs/scmg_access.log
 - o /apsara/nginx/logs/scmg_err_log
 - o /apsara/nginx/logs/scmg_op_log
 - /apsara/nginx/logs/sls_console.log
 - /apsara/nginx/logs/web_access.log

PackageManager#

- Startup log: /cloud/app/sls-backendserver/PackageManager#/package_manager/current/log/start.log
- Service log: none

RedisServer#

- Startup log: /cloud/app/sls-backend-server/RedisServer#/sls_redis/current/log/start.log
- Service log: /var/log/redis/redis.log

SlsConsole#

- Startup log: /cloud/app/sls-backend-server/SlsConsole#/sls_console/current/log/start.log
- Service logs: /alidata/www/logs/

- o /alidata/www/logs/java/sls/
 - /alidata/www/logs/java/sls/dashboard.log
 - /alidata/www/logs/java/sls/debug.log
 - /alidata/www/logs/java/sls/error.log
 - /alidata/www/logs/java/sls/info.log
 - /alidata/www/logs/java/sls/reasons.log
 - /alidata/www/logs/java/sls/tairSave.log
- /alidata/www/logs/java/sls-service/applog
 - /alidata/www/logs/java/sls-service/applog/error.log
 - /alidata/www/logs/java/sls-service/applog/info.log
 - /alidata/www/logs/java/sls-service/applog/warn.log
- o /usr/share/jetty/logs/
 - /usr/share/jetty/logs/request.log
 - /usr/share/jetty/logs/stderrout.log

SlsFastcgi#

- Startup log: /cloud/app/sls-backend-server/SlsFastcgi#/sls_fastcgi/current/log/start.log
- Service logs:
 - o /apsara/fcgi_agent/FastcgiAgent.LOG
 - o /apsara/fcgi_agent/metering.LOG
 - /apsara/fcgi_agent/monitor.LOG
 - /apsara/fcgi_agent/ols_operation.LOG

SlsLogtail#

- Startup log: /cloud/app/sls-backend-server/SlsLogtail#/sls_ilogtail/current/log/start.log
- Service logs
 - Service log on Apsara Stack: /usr/local/ilogtail_private/ilogtail.LOG
 - Service log on on-premises machines: /usr/local/ilogtail/ilogtail.LOG

SlsScmc#

- Startup log: /cloud/app/sls-backend-server/SlsScmc#/sls_scmc/current/log/start.log
- Service logs:
 - o /var/www/html/SCMC/logs/scm_op_log
 - /var/www/html/SCMC/logs/scm_err_log

SlsScmg#

- Startup log: /cloud/app/sls-backend-server/SlsScmg#/sls_scmg/current/log/start.log
- Service logs:
 - /var/www/html/SCMG/logs/scm_err_log
 - /var/www/html/SCMG/logs/scm_op_log

SlsTools#

- Startup log: /cloud/app/sls-backend-server/SlsTools#/aliyun_log_cli/current/log/start.log
- Service log: none

SlsWeb#

- Startup log: /cloud/app/sls-backend-server/SlsWeb#/sls_web/current/log/start.log
- Service logs:
 - /apsara/sls/web/logs/access.log
 - /apsara/sls/web/logs/apidetail.log
 - /apsara/sls/web/logs/httpclient.log
 - /apsara/sls/web/logs/normal.log
 - /apsara/sls/web/logs/sysinfo.log
 - /apsara/sls/web/logs/worker.log

SlsWebTools#

- Startup log: /cloud/app/sls-backendserver/SlsWebTools#/sls_web_tools/current/log/start.log
- Service log: none

ToolService#

- Startup logs:
 - o /cloud/app/sls-backend-server/ToolService#/init_db/current/log/start.log
 - o /cloud/app/sls-backend-server/ToolService#/init_diamond/current/log/start.log
 - o /cloud/app/sls-backend-server/ToolService#/init_odps/current/log/start.log
 - o /cloud/app/sls-backend-server/ToolService#/init_pop/current/log/start.log
 - o /cloud/app/sls-backend-server/ToolService#/jdk_uploader/current/log/start.log
- Service log: none

SlsImportOdpsScheduler#

- Startup log: /cloud/app/sls-backendserver/SlsImportOdpsScheduler#/sls_import_odps_scheduler/current/log/start.log
- Service Logs: Job Scheduler service

FuxiServiceSlsConfigService#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsConfigService#/sls_config_service/current/log/start.log
- Service log: none

FuxiServiceSlsEtlFramework#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsEtlFramework#/sls_etl_framework/current/log/start.log
- Service log: none

FuxiServiceSlsLoghubMaster#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsLoghubMaster#/sls_loghub_master/current/log/start.log
- Service log: none

FuxiServiceSlsMeteringService#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsMeteringService#/sls_metering_service/current/log/start.log
- Service log: none

FuxiServiceSlsPrestoWorker#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsPrestoWorker#/sls_presto_worker/current/log/start.log
- Service log: none

FuxiServiceSlsQueryMaster#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsQueryMaster#/sls_query_master/current/log/start.log
- Service log: none

FuxiServiceSlsQuotaServer#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsQuotaServer#/sls_quota_server/current/log/start.log
- Service log: none

FuxiServiceSlsReplayWorker#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsReplayWorker#/sls_replay_worker/current/log/start.log
- Service log: none

FuxiServiceSlsShennongWorker#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsShennongWorker#/sls_shennong_worker/current/log/start.log
- Service log: none

FuxiServiceSlsToolServiceWorker#

- Startup log: /cloud/app/sls-backendserver/FuxiServiceSlsToolServiceWorker#/sls_tool_service_worker/current/log/start.log
- Service log: none

NGINX

Error log: /apsara/nginx/log/error.log

Operations and Maintenance Guide · Operations of basic cloud products

Error	Action
Bind Address Failed	Check the port listening information in <i>/etc/init.d/nginx.conf</i> .
open() failed	Check whether the item that you want to open exists in the static resource file.

Console

Error log: /alidata/www/logs/java/sls/error.log

Error	Action
SLS SDK Exception	No action is required.
Create Bean Failed	Check the dubbo settings in the console configurations of SlsConsole.

Service

Error log: /alidata/www/logs/java/sls-service/applog/error.log

Error	Action
Create Bean Failed	Check the dubbo settings in the service configurations of SlsConsole.
Invoke failed	Check the scmg settings in the service configurations of SlsConsole.

Query Job Scheduler service logs

1. In the startup log, find the rpc sql command.

For example, if the command is /apsara/deploy/pc_wrapper/rpc.sh spl EtlFramework, EtlFramework is the name of the Job Scheduler service.

am
re
he
cut
ер
002
nf
he
cut
he
cut
ер
002
nf

2. Find the Job Scheduler machine.

/apsara/deploy/rpc_wrapper/rpc.sh spl EtlFramework				
Partition WorkerName		LastUpdateTime	status	
66	EtlFrameworkPartitionRole@a34h11080	.cloud.h11.amtest87 Su	ın Jan 5 16:03:01 2020 load	
ed				
62	EtlFrameworkPartitionRole@a34h11080	.cloud.h11.amtest87 Su	ın Jan 5 16:03:01 2020 load	
ed				
111	EtlFrameworkPartitionRole@a34h1108	0.cloud.h11.amtest87 S	un Jan 5 16:03:01 2020 loa	
ded				
113	EtlFrameworkPartitionRole@a34h1108	0.cloud.h11.amtest87 S	un Jan 5 16:03:01 2020 loa	
ded				

3. Log on to the Job Scheduler machine without using a password.

ssh a34h11080.cloud.h11.amtest87

4. View the logs.

[root@a34h11078.cloud.h11.amtest87 /root] #ls /apsara/tubo/TempRoot/sys/EtlFramework/EtlFrameworkPartitionRole@a34h11078.cloud.h11. amtest87/etl_worker.LOG /apsara/tubo/TempRoot/sys/EtlFramework/EtlFrameworkPartitionRole@a34h11078.cloud.h11.amt est87/etl_worker.LOG

- /apsara/tubo/TempRoot/sys/: fixed directory
- EtlFramework: the service name obtained in Step 1.
- EtlFrameworkPartitionRole@a34h11078.cloud.h11.amtest87: the Job Scheduler machine name obtained in Step 2.
- etl_worker.LOG: the log name.

8.11.2.2. Use Log Service Portal to view logs

Project admin

Logstore	Log directory
metering	/tmp/metering_*.LOG metering.log
sls_service_error_log	/alidata/www/logs/java/sls- service/applog/error.log
sls_service_info_log	/alidata/www/logs/java/sls- service/applog/info.log
sls_console_error_log	/alidata/www/logs/java/slserror.log
sls_console_info_log	/alidata/www/logs/java/slsinfo.log
scmc_access_log	/apsara/nginx/logsscmc_access.log
scmc_err_log	/apsara/nginx/logs/scmc_err_log
scmc_op_log	/apsara/nginx/logs/scmc_op_log
sls_operation_agg_log	/apsara/fcgi_agent/metering_*.LOG
sls_operation_log	/apsara/fcgi_agent/ols_operation*.LOG
offline_scheduler_log	/apsara/sls/import_odps/scheduler/*.[Ll][Oo] [Gg]
sls_fastcgi_log	/apsara/fcgi_agent/FastcgiAgent*.LOG

Logstore	Log directory
trace_log	/apsara/shennong_agent/tracer/index_worker_t race.LOG
dispatch_worker_log	/apsara/tubo/TempRoot/sys/DispatchWorker/[[user@ip]]/log_dispatch_worker.LOG
etl_framework_log	/apsara/tubo/TempRoot/sys/EtlFramework/[[us er@ip]]/etl_worker.LOG
etl_golang_worker_log	/apsara/tubo/TempRoot/sys/EtlFramework/[[us er@ip]]/etl_golang_worker.LOG
fc_trigger_log	/apsara/tubo/TempRoot/sys/FcTriggerWorker/[[user@ip]]/fc_trigger.log
query_master_log	/apsara/tubo/TempRoot/sys/QueryMaster/[[use r@ip]]/query_master.LOG
sls_configservice_log	/apsara/tubo/TempRoot/sys/ConfigService/[[us er@ip]]/sls_config_service.LOG
sls_configservice_query_log	/apsara/tubo/TempRoot/sys/ConfigService/[[us er@ip]]/config_service_query.LOG
sls_consumergroup_log	/apsara/tubo/TempRoot/sys/QuotaServer/[[use r@ip]]/monitor.LOG
sls_index_status_log	/apsara/tubo/TempRoot/sys/ShennongWorker/[[user@ip]]/project_index_size.LOG
sls_indexworker_log	/apsara/tubo/TempRoot/sys/OlsIndexWorker/[[user@ip]]/ols_index_worker.LOG
sls_loghub_shard_status_log	/apsara/tubo/TempRoot/sys/LoghubMaster/[[us er@ip]]/loghub_master_meta.LOG
sls_loghubmaster_log	/apsara/tubo/TempRoot/sys/LoghubMaster/[[us er@ip]]/sls_loghub_master.LOG
sls_quotaserver_log	/apsara/tubo/TempRoot/sys/QuotaServer/[[use r@ip]]/quota_server.LOG
sls_quotausage_log	/apsara/tubo/TempRoot/sys/QuotaServer/[[use r@ip]]/charge.LOG
sls_replayworker_log	/apsara/tubo/TempRoot/sys/ShennongReplayW orker/[[user@ip]]/shennong_replay_worker.LOG
sls_shennongworker_log	/apsara/tubo/TempRoot/sys/ShennongWorker/[[user@ip]]/shennong_worker.LOG

Operations and Maintenance Guide · Operations of basic cloud products

Logstore	Log directory
worker_input_log	/apsara/tubo/TempRoot/sys/ShennongWorker/[[user@ip]]/shennong_worker_input.LOG

Project scmg

Logstore	Log directory
scmg_access_log	/apsara/nginx/logs/scmg_access.log
nginx_error_log	/apsara/nginx/logs/error.log
scmg_err_log	/apsara/nginx/logs/scmg_err_log
scmg_op_log	/apsara/nginx/logs/scmg_op_log
sls_portal_access_log	/apsara/sls/web/logsaccess.log
sls_portal_http_req	/apsara/sls/web/logshttpclient.log
sls_portal_sys_info	/apsara/sls/web/logssysinfo.log
sls_portal_normal	/apsara/sls/web/logsnormal.log
sls_portal_api_audit	/apsara/sls/web/logsapidetail.log

8.12. Apsara Stack Security

8.12.1. 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.

- 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.

8.12.2. Routine operations and maintenance of Server Guard

8.12.2.1. Check the service status

8.12.2.1.1. Check the client status

Check the following status information about the Server Guard client to verify that the client is running properly:

Client logs

Client logs are stored in the data directory under the directory of the Server Guard process file, for example, */usr/local/aegis/aegis_client/aegis_xx_xx/data*.

Client logs are saved by day, for example, data.1 to data.7

Client's online status

Run the following command to check the client's online status:

```
ps -aux | grep AliYunDun
```

Network connectivity

Run the following command to check whether the client has set up a TCP connection with the server:

netstat -tunpe |grep AliYunDun

Client UUID

Open the client log file data.x and check the character string following Currentuid Ret . This character string is the UUID of the current client.

Client processes

The Server Guard client has three resident processes: AliYunDun, AliYunDunUpdate, and AliHids.

When the client runs properly, all of the three processes run normally.

? Note On a Windows OS client, the AliYunDun and AliYunDunUpdate processes exist in the form of services. The service names are Server Guard Detect Service and Server Guard Update Service, respectively.

8.12.2.1.2. Check the status of Aegiserver

Context

To check the running status of Aegiserver, follow the following steps:

- 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:f9d292f54c58646b672a8533a0d78fba534d2 6d376a194034e8840c70d9aa0b3 "/bin/bash /startApp." 2 hours ago Up 2 hours 80/tcp, 7001/tcp, 80 05/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.con fig.file=/home/admin/aegiserverlite/.default/conf/logging.properties -Djava.util.logging.manager =org.apache.juli.ClassLoaderLogManager -server -Xms2g -Xmx2g -XX:PermSize=96m -XX:MaxPerm Size=384m -Xmn1g -XX:+UseConcMarkSweepGC -XX:+UseCMSCompactAtFullCollection -XX:CMSMa xAbortablePrecleanTime=5000 -XX:+CMSClassUnloadingEnabled -XX:+UseCMSInitiatingOccupancy Only -XX:CMSInitiatingOccupancyFraction=80 -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDum pPath=/home/admin/logs/java.hprof -verbose:gc -Xloggc:/home/admin/logs/gc.log -XX:+PrintGC Details -XX:+PrintGCDateStamps -Djava.awt.headless=true -Dsun.net.client.defaultConnectTimeo ut=10000 -Dsun.net.client.defaultReadTimeout=30000 -XX:+DisableExplicitGC -Dfile.encoding=UTF-8 -Ddruid.filters=mergeStat -Ddruid.useGloalDataSourceStat=true -Dproject.name=aegiserverlite -Dcatalina.vendor=alibaba -Djava.security.egd=file:/dev/./urandom -Dlog4j.defaultInitOverride=tr ue -Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_EQUALS_IN_VALUE=true -Dorg.apache.tomc at.util.http.ServerCookie.ALLOW_HTTP_SEPARATORS_IN_V0=true -Djava.endorsed.dirs=/opt/taoba o/tomcat/endorsed -classpath /opt/taobao/tomcat/bin/bootstrap.jar:/opt/taobao/tomcat/bin/to mcat-juli.jar -Dcatalina.logs=/home/admin/aegiserverlite/.default/logs -Dcatalina.base=/home/a dmin/aegiserverlite/.default -Dcatalina.home=/opt/taobao/tomcat -Djava.io.tmpdir=/home/admin /aegiserverlite/.default/temp org.apache.catalina.startup.Bootstrap -Djboss.server.home.dir=/ho me/admin/aegiserverlite/.default -Djboss.server.home.url=file:/home/admin/aegiserverlite/.defa ult start

5. Run the following command to perform the health check:

curl 127.0.0.1:7001/checkpreload.htm

If the response is "success", the service is normal.

6. View related logs.

- **Protocol logs:** View logs about upstream and downstream protocol messages between the server and client in */home/admin/aegiserver/logs/AEGIS_MESSAGE.log*.
- **Operation logs:** View abnormal stack information during operation in */home/admin/aegis erver/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*.

8.12.2.1.3. 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.

8.12.2.1.4. 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.

8.12.2.2. 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 Dete ct 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

8.12.3. Routine operations and maintenance of

Network Traffic Monitoring System

8.12.3.1. Check the service status

8.12.3.1.1. Basic inspection

The basic inspection feature of Network Traffic Monitoring System allows you to check whether the service has reached the final status.

- 1. Log on to the Apsara Infrastructure Management Framework console.
- 2. Choose **Operations > Project Operations**. On the page that appears, enter *yundun-advance* in the search bar. Click **Details** in the Actions column corresponding to the yundun-advance cluster.
- 3. On the page that appears, select **BeaverCluster**.
- 4. Check whether yundun-beaver-advance has reached the final status in Service Instances List.

8.12.3.1.2. Advanced inspection

The advanced inspection feature of Network Traffic Monitoring System allows you to check service status and features.

Procedure

Follow these steps to check the service status:

- 1. Log on to the Apsara Infrastructure Management Framework console.
- 2. Log on to two physical machines of Network Traffic Monitoring System.
 - i. Choose Operations > Project Operations.
 - ii. Enter *yundun-advance*, and click **Details** to go to the cluster O&M page.
 - iii. Select the BeaverCluster cluster.
 - iv. Select yundun-beaver-advance from Service Instances List, and click Details to go to the Service Instance Dashboard page.
 - v. Select BeaverAdvance# from Service Role List, and click Details to go to the Service Role Dashboard page.
 - vi. View Server Information, and use TerminalService to log on to two physical machines of Network Traffic Monitoring System.
- 3. Check the log status of Network Traffic Monitoring System. Run udo cat /var/log/messages . If any record is returned, the log function is normal.
- 4. Check the status of the mirrored traffic. Run sudo cat /proc/ixgbe_debug_info . If speed in the second-to-last row of the output is not 0, the traffic mirroring process is normal.
- 5. Check the protected 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 CIDR block in the output is supposed to be an SLB or EIP CIDR block in a classic network. If the CIDR block is connected to Network Traffic Monitoring System through CSWs, a CIDR block in a VPC is required.
- 6. Check the network connectivity between Network Traffic Monitoring System and the VM. Run ping *VMIP* to check the network connectivity. In the command, replace *VMIP* with an IP address in the CIDR block specified in 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-sched ule/conf/dj.conf
 Check whether the ip parameter in the aliguard_smart section in the output is the same as the ip parameter corresponding to the aliguard. \${global:internet-domain} domain.
- 9. Check the following logs:
 - DDoS alert logs

Run grep -A 10 -B 10 LIDS /var/log/messages to view the DDoS alert logs.

• TCP intercept command logs

Run grep add_to_blacklist.htm /var/log/messages to view the TCP intercept command logs.

• Outbound attack logs

Run grep zombie_new /var/log/messages to view the outbound attack logs.

8.12.3.2. Common operations and maintenance

8.12.3.2.1. 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 -r f /dev/shm/drv_setup_path

8.12.3.2.2. 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

8.12.3.2.3. Disable TCP blocking

Context

To disable TCP blocking for Network Traffic Monitoring System, follow the following steps:

- 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

8.12.3.2.4. 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

8.12.4. Routine operations and maintenance of

Anti-DDoS Service

8.12.4.1. Check the service status

8.12.4.1.1. 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.

8.12.4.1.2. 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.
 - i. Log on to the Apsara Infrastructure Management Framework console, and choose Operations > Project Operations.
 - ii. Enter yundun-advance, and click **Details** to go to the Cluster Operations page.
 - iii. Select AliguardCluster.
 - iv. Select yundun-aliguard from Service Instances List, and click Details to go to the Service Instance Dashboard page.
 - v. Select AliguardConsole# from Service Role List, and click Details to go to the Service Role Dashboard page.
 - vi. 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/Aliguar dDefender/bin/aliguard_defender_check , and check the output result.

? Note If a server of Anti-DDoS Service has just restarted, wait for three to five minutes before running the script to check the deployment status.

• If the message aliguard status check OK! appears, Anti-DDoS Service has been correctly deployed and the service status is normal, as shown in Check the status of Anti-DDoS Service.

Check the status of Anti-DDoS Service

1	<pre>[root@1]].cloud/home/admin]</pre>				
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 Reinjection route error message appears, the reinjection route is faulty.

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 BGP routing error message appears, the BGP protocol (for traffic routing) is faulty.

BGP routing error message

1 Error: bgp status error!

Troubleshooting: If BGP routing is faulty, troubleshoot the problem as follows:

- a. Use the ISW to check whether the BGP neighbor is in the normal status.
- b. Check whether the BGP route of the ISW contains a 32-bit attacked IP address of which the route is redirected to Anti-DDoS Service in the next hop.
- c. Check whether the route policy in the BGP configuration of the ISW is correctly configured.
- If the problem is caused by none of the above reasons, the core process is faulty. Contact Alibaba Cloud technical support.
- 3. Check the status of the NICs or optical modules of Anti-DDoS Service.

? Note Anti-DDoS Service has special requirements on optical modules. Only optical modules equipped with Intel X520 or Intel 82599 NICs can be used.

Run Ispci | grep Eth . If the command output contains four Intel 82599 NICs, the NICs are standard.

[roota cloud.am54 /root]
#lspci -v grep Eth
02:00.0 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
02:00.1 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
04:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
04:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
81:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
81:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2

8.12.4.2. Common operations and maintenance

8.12.4.2.1. 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/Aliguard Defender/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/Aligu ardDefender/bin/aliguard start
- 4. Run the service status check command five minutes after Anti-DDoS Service is restarted.

8.12.4.2.2. 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 <u>aliguard local</u> 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:

Anti-DDoS Service configuration example

Operations and Maintenance Guide · Operations of basic cloud products

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

i. 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.

- ii. Run the /home/admin/aliguard/target/AliguardDefender/bin/aliguard stop command to stop Anti-DDoS Service.
- iii. 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 Unbind NICs.

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

iv. In the /sys/bus/pci/drivers/ixgbe directory, bind the four NICs to the ixgbe driver for Linux, as shown in Bind NICs.

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

v. 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 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

vi. 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

vii. 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/targ et/AliguardDefender/bin/aliguard_rule -v 0.0.0.0 -d drop_icmp command to disable the drop_icmp policy.

- viii. 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.

8.12.5. Routine operations and maintenance of

Threat Detection Service

8.12.5.1. Check the service status

8.12.5.1.1. 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.

8.12.5.1.2. 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.
 - i. Choose Operations > Project Operations.
 - ii. Enter yundun-advance, and click Details to go to the Cluster Operations page.
 - iii. Select BasicCluster.
 - iv. Select yundun-sas from Service Instances List, and click Details to go to the Service Instance Dashboard page.

- v. Select SasApp# from Service Role List, and click Details to go to the Service Role Dashboard page.
- vi. 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-def ault.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/adm in/sas/logs/notify.log*.

8.12.5.2. 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

8.12.6. Routine operations and maintenance of WAF

8.12.6.1. Check the service status

8.12.6.1.1. 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 yundunsemawaf service instance is in final status.

? Note If the Final Status column for an instance is True, the instance has reached final status.

8.12.6.1.2. Advanced inspection

The advanced inspection feature of Web Application Firewall (WAF) focuses on the system status and service status.

- 1. Log on to the Apsara Infrastructure Management Framework console.
- 2. Log on to two WAF physical machines respectively.
 - i. In Apsara Infrastructure Management Framework, choose Operations > Project Operations.
 - ii. In the Fuzzy Search search box, enter *yundun-semawaf*. Click Details in the Actions column, and the Cluster Operations page is displayed.
 - iii. Click the SemaWafCluster cluster.
 - iv. In Service Instances, select yundun-semawaf, and click Details. The Service Instance Information Dashboard page is displayed.

- v. In Server Role List, select YundunSemawafApp#, and click Details. The Server Role Dashboard page is displayed.
- vi. In Machine Information, click Terminal to log on to two WAF physical machines respectively.
- 3. Check the system status.
 - i. Check the system logs. Run the dmesg -T |tail -30 command to check for exception logs.
 - ii. 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.

? Note The following check is based on the WAF installation directory, which is */home /safeline* by default.

- i. Run the cd /home/safeline command to open the installation directory.
- ii. 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.
- iii. 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.
- iv. Check the Redis service. Run the docker logs --tail 50 mgt-redis command to check for exception logs.

- v. 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_total parameter, which indicates the number of requests that have been processed within the last five seconds.
- vi. 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.
- vii. 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_pending 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.

8.12.7. Routine operations and maintenance of Sensitive Data Discovery and Protection

8.12.7.1. Check the service status

8.12.7.1.1. Basic inspection

During the basic inspection of Sensitive Data Discovery and Protection (SDDP), check whether the service has reached the final status.

- 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 yundunsddp service instance is in the final status.

8.12.7.1.2. 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.
 - i. Choose Operations > Project Operations.
 - ii. 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.
 - iii. In the cluster list, click the cluster name that starts with SddpCluster.
 - iv. 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							
Service Instance	Final Status	Expected Server Roles	Server Roles In Final	Server Roles Going O	Actions		
hids-client	True	1	1	0	Actions - Details		
0S	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		

v. 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	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

vi. In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.

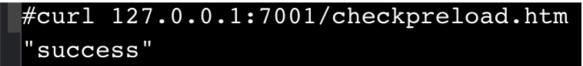
Achine Information									
Machi	IP	Mac	Mac	Serv	Serv	Curr	Targ	Error	Actions
a56g101	10	good		good P		2fb869ef	2fb869ef		Terminal Restart Details Machine System View Machine Operation
a56h1116	10	good		good P		2fb869ef	2fb869ef		Terminal Restart Details Machine System View Machine Operation

3. Log on to two Docker containers of the SddpService service, respectively. Run the sudo dock er 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 yundunsddp-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.acti
ve=cloud -server -Xms4g -Xmx4g -Xmn2g -XX:MetaspaceSize=256m -XX:MaxMetaspaceSize=512m -XX:MaxDirectMemorySize=1g
-XX:SurvivorRatio=10 -XX:+UseConcMarkSweepGC -XX:CMSMaxAbortablePrecleanTime=5000 -XX:+CMSClassUnloadingEnabled -X
X:CMSInitiatingOccupancyFraction=80 -XX:+UseCMSInitiatingOccupancyOnly -XX:+ExplicitGCInvokesConcurrent -Dsun.rmi.
dgc.server.gcInterval=2592000000 -Dsun.rmi.dgc.client.gcInterval=2592000000 -XX:ParallelGCThreads=4 -Xloggc:/root/
logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/root/logs
/java.hprof -Djava.awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun.net.client.defaultReadTime
out=30000 -DJM.LOG.PATH=/root/logs -DJM.SNAPSHOT.PATH=/root/snapshots -Dfile.encoding=UTF-8 -Dhsf.publish.delayed=
true -Dproject.name=yundun-sddp-service -Dpandora.boot.wait=true -Dlog4j.defaultInitOverride=true -Dserver.port=70
01 -Dmanagement.port=7002 -Dmanagement.server.port=7002 -Dpandora.location=/home/admin/yundun-sddp-service/target/
taobao-hsf.sar -classpath /home/admin/yundun-sddp-service/target/yundun-sddp-service -Dapp.location=/home/admin/yu
ndun-sddp-service/target/yundun-sddp-service -Djava.endorsed.dirs= -Djava.io.tmpdir=/home/admin/yundun-sddp-servic
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-req uest.log* file.
 - View system logs in the */home/admin/yundun-sddp-service/logs/service-stdout.log* file.

8.12.7.1.3. Advanced inspection: Check the status of the

SddpData service

This topic describes how to check the running status of the SddpData service.

- 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.
 - i. Choose Operations > Project Operations.
 - ii. 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.
 - iii. In the cluster list, click the cluster name that starts with SddpCluster.
 - iv. In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.
 - v. In the Server Role List section, find SddpData# and click Details in the Actions column to go to the Server Role Dashboard page.
 - vi. 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.

8.12.7.1.4. Advanced inspection: Check the status of the

SddpPrivilege service

This topic describes how to check the running status of the SddpPrivilege service.

- 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.
 - i. Choose Operations > Project Operations.
 - ii. 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.
 - iii. In the cluster list, click the cluster name that starts with SddpCluster.
 - iv. In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.
 - v. In the Server Role List section, find SddpPrivilege# and click Details in the Actions column to go to the Server Role Dashboard page.
 - vi. 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 doc ker 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 yundu n-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/service-stdout.log* file.

8.12.7.1.5. 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.
 - i. Choose Operations > Project Operations.
 - ii. 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.
 - iii. In the cluster list, click the cluster name that starts with SddpCluster.
 - iv. In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.
 - v. In the Server Role List section, find SddpLog# and click Details in the Actions column to go to the Server Role Dashboard page.
 - vi. 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 e xec -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-sdd p-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/application.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.

8.12.7.2. 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.

8.12.8. Routine operations and maintenance of Apsara Stack Security Center

8.12.8.1. Check service status

8.12.8.1.1. 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-secure console has reached the final status in Service Instances List.

8.12.8.1.2. 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.
 - i. Choose Operations > Project Operations.
 - ii. Enter yundun-advance, and click Details to go to the Cluster Operations page.
 - iii. Select BasicCluster.
 - iv. Select yundun-secureconsole from Service Instances List, and click Details to go to the Service Instance Dashboard page.
 - v. Select SecureConsoleApp# from Service Role List, and click Details to go to the Service Role Dashboard page.
 - vi. 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 \$(sud o 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*.

8.12.8.2. 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 secureconsole 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

8.12.9. Routine operations and maintenance of

secure-service

8.12.9.1. Check the service status

8.12.9.1.1. 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 *yundun-advance*, and click Details to go to the Cluster Operations page.
- 3. Select BasicCluster.
- 4. Check whether yundun-secureservice has reached the final status in Service Instances List.

8.12.9.1.2. 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.
 - i. Choose Operations > Project Operations.
 - ii. Enter yundun-advance, and click Details to go to the Cluster Operations page.
 - iii. Select BasicCluster.
 - iv. Select yundun-secureservice from Service Instances List, and click Details to go to the Service Instance Dashboard page.
 - v. Select SecureServiceApp# from Service Role List, and click Details to go to the Service Role Dashboard page.
 - vi. View Server Information, and click Terminal to log on to two physical machines, respectively.
- 3. Log on to two secure-service Docker containers, respectively. Run sudo docker exec -it \$(sudo docker ps | grep secureservice | awk '{print \$1}') bash .
- 4. Check the secure-service process status. Run ps aux |grep secure-service . If any record is returned, the secure-service process is normal.
- 5. Check the health status. Run curl 127.0.0.1:3010 . If OK is returned, the service is normal.
- 6. Run the following command to go to the Docker container:

sudo docker exec -it [imageId] /bin/bash

- 7. View related logs.
 - View the Server Guard logs in /home/admin/secure-service/logs/aegis-info.log.
 - View the error logs in */home/admin/secure-service/logs/Error*.
 - View the vulnerability analysis and scanning logs in */home/admin/secure-service/logs/le akage-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*.

8.12.9.1.3. 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/common-error.log*.
 - View the metaq messages received by the Dolphin service in */home/admin/dolphin/logs/ dolphin-message-consumer.log*.

Note Currently, only Threat Detection Service (TDS) sends messages to the Dolphin service.

• View the metaq messages sent by the Dolphin service in */home/admin/dolphin/logs/dolp hin-message-producer.log*.

Onte Currently, the Dolphin service sends messages only to TDS.

8.12.9.1.4. 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*.

8.12.9.2. 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

8.13. Apsara Stack DNS

8.13.1. 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 configuration 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.

8.13.2. Maintenance

8.13.2.1. 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.

8.13.2.2. 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
ospfd stopserviceospfd stopcommand to disable the OSPF service before you run the
service named stopservice named stopcommand 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.

8.13.2.3. 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.

8.13.3. DNS API

8.13.3.1. 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 > Machine Operations** in the Apsara Infrastructure Management Framework console.

Context

> Document Version:20200918

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.
 - i. Log on to the Apsara Infrastructure Management Framework console.
 - ii. In the top navigation bar, choose Tasks > Deployment Summary to open the Deployment Summary page.
 - iii. Click Deployment Details.
 - iv. On the **Deployment Details** page, find the dnsProduct project.
 - v. 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.

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 \bigcirc < dnsService \bigcirc
dts	Final 3 Days 23 Hours	A standardCluster < d < hids-client < + bindServerRole#
ecs	Final 1 Hour 24 Minutes	C «tianji (c) A dnsServiceDblnit# (c)
edas	Final 4 Days 21 Hours	o≤ tianji-dockerdae ⊘
elasticsearch	Final 11 Hours 57 Minutes	c
emr	Final 4 Days 21 Hours	c
ess	Final 3 Days 22 Hours	c

View API status

- 2. Obtain the IP addresses of servers where the API services are deployed.
 - i. Log on to the Apsara Infrastructure Management Framework console.
 - ii. In the top navigation bar, choose **Operations > Cluster Operations**.
 - iii. Click a cluster URL to open the Cluster Dashboard page.

iv. On the Cluster Dashboard page, choose Operations Menu > 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		

Cluster Operation and Maintenance Center

v. On the **Cluster Operation and Maintenance Center** page, view and obtain the IP addresses of servers that are deployed with the API service.

View the IP addresses of servers

Cluster Operation and Maintenance Center (Cluster: dnsc/Luster:A-20190827-4eb3) SR not in Final Status: N/A Running Tasks: No rolling is available. Head Version Submitted AI: 09/26/19, 10.29:12 Head Version Analysis0: done Service dnsService Service Role Scale-ont Scale-ont
Service Server Role drsServerRole# Total Machines 2 Expected Mac: 2 Scale-out Scale-out Machine: In: 0 out 0 Abnormal Machines 0 Ping Failed: 0 No Heartbeat: 0 Abnormal Services 0 TJ-Client: 0 Status Error: 0
Total Machines 2 Expedied Mac2 Scale out Scale out Scale out Machines in 0 out 0 out 0 out 0 Abnormal Machines Ping Failed 0 No Heartbeat 0 Status Error 0 Abnormal Scale out Abnormal Other Status
Orall Machines 2 Expected Mac2 Scale-out Scale-out medunes it of sr. in: 0 Court of out: 0 Anormal Machines 0 No Hearbeat 0 Status Error: 0 Anormal Services 0 Other SR: 1
Machine Search Supports multiple machine search. Q Reset
Machine Action Status Action Action Status Monitoring Statistics Actions
vm010012012075 Normal GOOD N/A N/A Error: 0 Warning: 0 Terminal Approval A 18 12 / 37 Tb 000D N/A N/A N/A Good: 7 Restart Server Role
wm010012016048 Normal GOOD N/A N/A Error: 0 Warning: 0 Terminal / Approval A 13:11:16:48 0cod: 7 Restart Server Role Restart Server Role Restart Server Role
Batch Terminal Items per Page [10 •] 🤘 🤘 🤘

- 3. Log on to the DNS API server. Run the curl http://localhost/checkpreload.htm command, and check whether the command output is "success".
 - i. Log on to the Apsara Infrastructure Management Framework console.
 - ii. In the top navigation bar, choose **Operations > Machine Operations**.
 - iii. Click Terminal in the Actions column of a server to log on to the server.

iv. 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".



8.13.3.2. 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.

8.13.4. DNS system

8.13.4.1. 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.
 - i. Log on to the Apsara Infrastructure Management Framework console.
 - ii. In the top navigation bar, choose Tasks > Deployment Summary.
 - iii. On the Deployment Summary page, click Deployment Details.
 - iv. On the **Deployment Details** page, find dnsProduct.
 - v. Click **Details** in the **Deployment Progress** column to check whether the bindServerRole# role is in its final state.

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 E	<u>ietails</u>
drds	Final 4 Days 7 Hours	C 🚠 dnsCluster-A-20 ⊘	of dnsService	ServiceTest#
dts	Final 3 Days 23 Hours	at standardCluster⊘ C		 bindServerRole# dnsServerRole#
ecs	Final 1 Hour 24 Minutes	c		 Instantion (and a second secon
edas	Final 4 Days 21 Hours	c	∞° tianji-dockerdae… (⊘ -the monitorSrDemo# ⊘
elasticsearch	Final 11 Hours 57 Minutes	с		
emr	Final 4 Days 21 Hours	с		
ess	Final 3 Days 22 Hours	c		

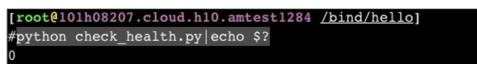
- 2. Obtain the IP addresses of the servers where DNS services are deployed.
 - i. Log on to the Apsara Infrastructure Management Framework console.
 - ii. In the top navigation bar, choose **Operations > Cluster Operations**.
 - iii. Click a cluster URL to go to the Cluster Dashboard page.
 - iv. On the Cluster Dashboard page, choose **Operations Menu > Cluster Operation and** Maintenance Center.

Cluster Operation and Maintenance Center

Cluster Dashboard	Operations Menu 👻	
	Change Machine	
Basic Cluster Information	Deploy Service	7 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	
Machines Not In Final Status	Service Final Status Query	
Machines Not in Filldi Status	Cluster Configuration	
Real/Pseudo Clone	Operation Logs	
Expected Machines		_

- v. 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.
 - i. Log on to the Apsara Infrastructure Management Framework console.
 - ii. Choose Operations > Machine Operations.
 - iii. Select a server and click Terminal to log on to the server.
 - iv. 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.

Verifying the server



8.13.4.2. 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.

8.13.4.3. 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.

8.13.5. 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.

8.13.6. 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.

9.Operations of middleware products 9.1. Enterprise Distributed Application Service (EDAS)

9.1.1. O&M overview

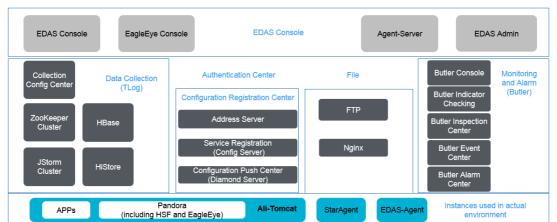
This topic describes the system architecture, component architecture, and O&M architecture of EDAS.

9.1.1.1. Architecture

This topic describes the system architecture and component architecture of EDAS. Familiarize yourself with this knowledge when performing O&M for EDAS.

System architecture

The system architecture of EDAS consists of the console, data collection center, configuration registry, authentication center, and file system. EDAS architecture shows the overall architecture of EDAS.



EDAS architecture

• EDAS console

The EDAS console is the only EDAS system component that you can use directly. You can implement resource management, application lifecycle management, maintenance control and service governance, three-dimensional monitoring, and digital operation in the console.

• Data collection system

It allows you to collect, compute, and store the runtime status, trace logs, and other information about clusters and instances where applications are deployed in EDAS in real time.

• Configuration registry

This is a central server that is used to publish and subscribe to HSF services (RPC framework) and to push distributed configurations.

• Authentication center

This system component controls permissions for user data to ensure data security.

• O&M system

EDAS uses Butler as its O&M system. Butler monitors the data of EDAS and triggers alarms when specified criteria are met. Butler provides routine inspection and alarm functions for all EDAS components.

• File system

This system component stores WAR packages and required components, such as JDK and Ali-Tomcat, uploaded by users.

Component architecture

Each system component of EDAS consists of one or more components. The following figure shows the component architecture of EDAS.

Component	Node type	Node quantity	Description
EDAS console	Control node	2	The console of EDAS. It provides the core functions of the PaaS platform, including resource management, application lifecycle management, service governance, and auto scaling.
EDAS admin	Control node	2	A background task service. It provides the instance synchronization and application health check functions.
EDAS server	Control node	2	The EDAS server synchronizes status information with EDAS Agent.
Cai-fs	Control node	2	A file server. It stores the EDAS Agent installation package and EDAS application packages.
EagleEye console	Control node	2	You can query and view service traces in the EagleEye console.

Operations and Maintenance Guide · Operations of middleware products

Component	Node type	Node quantity	Description
Cai-address	Control node	3	An address discovery service. It provides the address lists for DiamondServer and ConfigServer.
DiamondServer	Control node	3	A configuration management service. It provides configuration storage, query, and notification functions, and mainly stores database metadata and EDAS function switch configurations in EDAS.
ConfigServer	Control node	3	An RPC service registry. It is used to query and store the publishing and subscription data of services.

For information about other external components, such as Butler, DAuth, and TLog, see the corresponding O&M documents.

9.1.1.2. O&M architecture

This document describes the O&M architecture of EDAS. Before using this document, familiarize yourself with the system architecture of EDAS.

EDAS O&M is performed mainly through Butler and the CLI.		
O&M category	Description	O&M tool
Routine maintenance	It includes inspection and monitoring.	 You can use Butler to automatically inspect and monitor the containers and components of EDAS. You can use the CLI to manually inspect the containers and components of EDAS.
Power-off maintenance	 Check and determine the statuses of containers and components. Stop and start containers and components. 	CLI

EDAS O&M is performed mainly through Butler and the CLI.

O&M category	Description	O&M tool
Troubleshooting	This topic mainly describes how to handle the component availability and service continuity faults of EDAS.	 Butler and the CLI. Butler inspects and monitors components and reports errors. Determine the component status through Butler and retrieve related logs through Butler or the CLI. Handle faults in the CLI.

9.1.2. Overview of critical operations

Routine O&M for EDAS must be performed in strict accordance with the O&M guide. Failure to follow the O&M guide may cause risks to components and services.

O&M operations are classified into three levels: G1, G2, and G3. Operations vary by level. See the following table.

Level	Description
G1	L1 L2: Operations can be performed safely based on documented instructions, without having to apply for changes. Such operations will not affect the service.
G2	L1 L2: The onsite personnel must obtain confirmation from the product personnel before performing operations, which require applying for changes and following the documented instructions. Such operations will not affect the service.
G3	L1 L2: The onsite personnel must obtain confirmation from the product personnel and the customer before performing operations, which require applying for changes and following the documented instructions. Such operations may affect the service.

Definitions of operation levels

G3 is the highest level, which involves critical operations. See the following table.

A list of critical operations

Operation	Operation or Command
Check the AccessKeyId and AccessKeySecret	cat /home/admin/.spas_key/default

Operation	Operation or Command
Clear logs	 For EagleEye: find /home/admin/eagleeye/logs/ -name "*log.*" - exec rm {} For EDAS: find /home/admin/edas/logs/ - name "*log.*" -exec rm {}
Restart containers	docker start {containerld}

9.1.3. Maintenance preparation

This topic describes the logon portal, account, permissions, and tools required for maintenance.

Maintenance preparation

Name	Function	Description
Butler console	Configure and view the inspection, monitoring, and alarm settings of EDAS.	A domain name, such as butler.console.example.com. example indicates a root domain, which must be determined based on the actual environment.
Remote SSH logon tool (such as <i>MobaXterm</i> or <i>PuTTY)</i>	Log on to the instances where components are located.	The logon account must be assigned the corresponding permissions. We do not recommend that you use the root or admin account for logon.
Account	Log on to the console or an instance.	 Obtain the account and password for console logon from EDAS Customer Services. The account used to log on to the instances where EDAS components are located must be assigned the corresponding permissions. We do not recommend that you use the root or admin account for logon.

9.1.4. Routine maintenance

EDAS routine maintenance includes inspection and monitoring.

• Inspection is the process where a periodic dialing test is performed on URLs or ports to determine whether EDAS services are normal. Currently, inspections in HTTP, TCP, ping, and

JDBC modes are supported.

• Monitoring is the process where logs are collected from clients through TLog to summarize key metrics for measuring the system runtime status. Monitoring includes infrastructure monitoring and JVM monitoring.

9.1.4.1. Inspection

You can configure the inspection rules of the HTTP, TCP, ping, and JDBC types in Butler to inspect the components and services of EDAS. An inspection rule provides a response code that is used to check the configured alarm rule. The alarm content is configured in Alarm Description. You can also log on to the instances where components and services are deployed and run commands for inspection.

EDAS inspection itemslists the default inspection items of EDAS. You can create an inspection configuration as needed.

Inspected object	Description	Inspection method
ConfigServer	Request /configserver/serverlist to check whether ConfigServer is normal.	Check /configserver/serverlist.
DiamondServer	Check whether the API operation for querying the DiamondServer status is normal.	Check /diamond- server/diamond.
	Check whether the DiamondServer database is connected.	Check DB connection.
	Check whether the TLog service is normal.	Check /api/StageHealthCheck.
TLog	Check whether the TLog listening port is normal.	Check port 8080.
	Check whether the TLog database is connected.	Check DB connection.
	Check whether the edas- console service is normal.	Check /checkpreload.htm.
edas-console	Check whether the edas- console port is normal.	Check port 8080.
	Check whether the EDAS database is connected.	Check DB connection.
	Check whether the HiStore listening port is normal.	Check port 5029.

EDAS inspection items

Hispected object	Description	Inspection method
	Check whether HiStore is connected.	Check DB connection.
Redis	Check whether the Redis listening port is normal.	Check port 6379.
	Check whether the edas-admin service is running properly.	Check /index.
edas-admin	Check whether the edas-admin listening port is normal.	Check port 8080.

9.1.4.1.1. Component inspection

You can inspect EDAS components by using the inspection management function of Butler or the CLI.

9.1.4.1.1.1. Manual inspection

Butler cannot fully support inspection with complex logic. To address this issue, EDAS provides CLI-based inspection.

Manual inspection uses commands typical of Linux operating systems. Set specific parameters based on the actual environment.

9.1.4.2. Monitoring

You can monitor the containers, system, and services of EDAS by using the product metrics and the metric rules created in Butler.

Container monitoring

By default, the container status of each EDAS component is checked based on the monitoring script that is configured in an environment variable.

System monitoring

System monitoring includes infrastructure monitoring and JVM monitoring.

System metrics

Monitoring type	Metric	Threshold
Infrastructure monitoring	CPU usage	70%
	Memory usage	70%
	Disk usage	90%
	Young GC times	60

Monitoring type	Metric	Threshold
JVM monitoring	Full GC	5
	Old generation usage	90%
	Permanent generation usage	90%

Service monitoring

Service monitoring is configured by EDAS and uploaded to Butler through APIs for monitoring and alarm purposes. Service metrics are as follows:

Monitoring target	Туре	Monitoring description
HTTP service	QPS	Monitors the QPS of the HTTP service.
	Response Time	Monitors the I/O time consumed by the HTTP service.
HSF service	Service Provision QPS	Monitors the QPS of HSF service provisioning.
	Service Provision RT	Monitors the time spent on HSF service provisioning.
	Service Consumption QPS	Monitors the QPS of HSF service consumption.
	Service Consumption RT	Monitors the time spent on HSF service consumption.
Container	Heap Memory Usage	Measures how much heap memory is used by services.
	Off-Heap Memory Usage	Measures how much non-heap memory is used by services.

9.1.4.2.1. Monitoring logs

Logs are critical for EDAS 0&M. You can monitor logs to promptly locate runtime faults.

You can use Butler to monitor	r EDAS logs, including:	

Component	Log	Path
EDAS console	console.log	/home/admin/edas/logs
EDAS admin	admin.log	/home/admin/edas/logs
EDAS server	agent-server.log	/home/admin/edas/logs

EDAS also provides other component logs, including the infrastructure monitoring log, service monitoring log, container monitoring log, and JVM monitoring log. For more information, see the "Log reference" topic.

9.1.5. Troubleshooting

Faults may occur during EDAS usage. This topic describes the typical faults that may occur during O&M as well as their handling methods.

Fault classification

Currently, EDAS-related faults are classified into two categories:

- Component unavailability
- Service discontinuity

Fault locating

You can locate faults through inspection, monitoring, logs, and alarms.

9.1.5.1. Alarm handling

Butler inspects and monitors the status and metrics of each EDAS component during EDAS O&M. Alarms are triggered when inspection and monitoring are abnormal. This topic describes how to handle inspection and monitoring alarms.

9.1.5.1.1. CPU utilization alerts

The CPU utilization threshold is 70%. The CPU utilization is abnormal and an alert is triggered if it exceeds the threshold.

Possible causes

- High access concurrency
- Insufficient application instances

Impact on the system

Service performance is compromised.

Procedure

- 1. Open the Secure Sockets Layer (SSL) tool (MobaXterm Personal Edition).
- 2. Run ssh <Username>@<IP address of your Ark client> and enter your *password* to log on to the client.
- 3. Go to the target container and run top to check the CPU utilization of components.
 - If the CPU utilization is normal, no further action is required.
 - If the CPU utilization is abnormal, check the number of calls.
- 4. Run netstat -tnlp | grep -E "80|8080" | wc -l to check the call status of components.
 - If the number of calls is relatively large and meets the service status, scale out instances.
 - If the number of calls is not large, identify the cause of high CPU utilization by completing

the following steps.

? Note You can use the open-source script to locate and print the processes with high CPU utilization.

- a. Run top -Hp <Component process ID> to locate the processes with high CPU utilization or memory usage and convert them into the hexadecimal format.
- b. Go to the *jstack process id* path, open **ps.txt**, and identify the specific process class based on the hexadecimal thread ID.
- c. If Full GC occurs, check gc.log or run jstat -gcutil [pid] to check the corresponding GC log. Record *top/jstack file/full GC* and send the record to EDAS Customer Services.

9.1.5.1.2. Memory usage alarms

The memory usage threshold is 90%. Memory usage is abnormal if it exceeds the threshold.

Possible causes

- High access concurrency
- Insufficient application instances

Impact on the system

Service performance is compromised.

Procedure

- 1. Open the SSL tool (MobaXterm Personal Edition).
- 2. Run the command ssh <Username>@<IP address of the client> and enter your *password* to log on to the client.
- 3. Go to the target container and run top to check the memory usage of components.
 - If memory usage is within the normal range, check disk usage and JVM metrics.
 - If memory usage is abnormal, check whether this is caused by JVM.
- 4. Run jmap to check memory usage.
- 5. Run jstat -gcutil [pid] to check memory usage.
- 6. Run vmstat to analyze and collect statistics on virtual memory.

7.

Result

- Log on to the Butler console to check whether the related alarms are cleared.
- Log on to the target container by using the SSH tool and run top to check whether memory usage is normal (less than 90%).

9.1.5.1.3. Disk usage alarms

The disk usage threshold is 90%. Disk usage is abnormal if it exceeds the threshold.

Possible causes

- High access concurrency
- Insufficient application instances
- Insufficient disk space or no periodic disk cleanup

Impact on the system

Service performance is compromised.

Procedure

- 1. Open the SSL tool (MobaXterm Personal Edition).
- 2. Run the command ssh <Username>@<IP address of the client> and enter your *password* to log on to the client.
- 3. Go to the target container and run df -lh to check the disk usage of each directory and identify the directories with excessive and fast disk usage.
- 4. Run **iostat** to check the data write status. Perform disk cleanup if the *logs* directory occupies excessive disk space.
- 5. Run netstat to check the number of calls, view logs for call errors, and take relevant measures such as scale-out.

Result

- Log on to the Butler console to check whether the related alarms are cleared.
- Go to the target container by using the SSH tool and run df -lh to check whether disk usage is normal (less than 90%).

9.1.5.2. Service continuity exceptions

9.1.5.2.1. EDAS monitoring exceptions

This topic describes how to troubleshoot EDAS monitoring exceptions.

Symptoms

- No application data can be monitored in the EDAS console.
- Data monitoring in the EDAS console has a significant lag.
- The monitoring and alarm functions are ineffective.
- Traces cannot be queried.

Possible cause

The EDAS components and dependent components are abnormal.

Impact on the system

EDAS cannot monitor applications or services, or monitoring is inefficient.

Procedure

1. Log on to the Butler console. In the left-side navigation pane, click **Container Monitoring**. On the **Container List** page, check whether the related components (such as TLog, JStorm, and

HBase) are normal.

- If TLog is abnormal, log on to the instance where TLog is located, go to */home/admin/logs* /, and check the **tlogconsole.log** file. Then, troubleshoot the problem and restart TLog.
- If JStorm is abnormal, log on to the instance where JStorm is located and check the log for the data collection task, such as */home/admin/logs/tlog_eagleeye-worker-6801.log*. Troubleshoot the problem and restart JStorm.

? Note You need to restart each JStorm process. Otherwise, data cannot be written to HBase due to a connection error.

• If HBase is abnormal, log on to the instance where HBase is located and check the related error log. Troubleshoot the problem and restart HBase.

Result

Check whether EDAS monitoring returns to normal.

9.1.5.2.2. Excessive node logs

This topic describes how to troubleshoot the problem of excessive node logs for EDAS.

Symptoms

- Traces and system responses slow down, and Butler issues a disk usage alarm.
- The instance generates excessive log files.

Possible cause

A large amount of log files are not cleared from the disk in a timely manner, which affects system performance.

Impact on the system

Service nodes become less responsive.

Procedure

- 1. Log on to the EDAS component node to check logs.
 - Path to EagleEye logs: /home/admin/logs/eagleeye
 - Path to EDAS logs: /home/admin/edas/logs
- 2. Ensure that service logs are not printed on the preceding paths or that service logs have been backed up.
- 3. Clear backup logs by running find /home/admin/logs/ -name "*log.*" -exec rm {}; or find /home/admin/edas/logs/ -name "*log.*" -exec rm {};.

Result

Check whether the service nodes return to normal.

9.1.5.2.3. Console access failure

This topic describes how to troubleshoot EDAS console access failures.

Symptoms

The EDAS console cannot be accessed.

Possible causes

- The edas-console node is abnormal.
- An error occurs during DNS resolution.

Impact on the system

The EDAS console is unavailable.

Procedure

- 1. Troubleshoot the edas-console node errors.
 - If the EDAS console becomes accessible again, no further action is required.
 - If the EDAS console remains inaccessible, proceed with the next step.
- 2. Log on to the ECS instance where the edas-console node is located, go to */home/admin/ed as/logs*, and check console.log for the problem.

Result

The EDAS console becomes accessible again.

9.1.5.2.4. ECS instance import failure

This topic describes how to troubleshoot failures when importing ECS instances.

Symptoms

ECS instances cannot be imported.

Possible causes

- An Alibaba Cloud API call fails.
- Images cannot be converted.
- ECS instance registration fails.

Impact on the system

Service availability and reliability degrade.

Procedure

- 1. Log on to the EDAS console and import ECS instances manually.
- 2. If ECS instance registration fails, register the ECS instance by running edas init.
- 3. If image registration fails, log on to the ECS console and view the status.

9.1.5.2.5. TLog data collection errors

This topic describes how to fix TLog data collection errors.

Symptoms

- The application monitoring dashboard is inaccessible.
- Application and service monitoring is inaccessible, and alarms cannot be triggered.
- Infrastructure monitoring is inaccessible, auto scaling is ineffective, and alarms cannot be triggered.
- Traces cannot be queried.

Possible causes

Collection point: Each collection job of TLog is called a collection point.

Collection points are the basic units for task processing by TLog. The monitoring function of EDAS is provided by one or more collection points in TLog. When working properly, collection points are in the activated or running state.

If the collection point for a product encounters an error, the corresponding EDAS monitoring data or page shows an exception.

- The basic data of the monitoring dashboard corresponds to the collection point service group TLog and the collection point infrastructure.
- The service data of the monitoring dashboard corresponds to the collection point service group EagleEye and the collection point stats_logger_agg.
- The zoom-in (more than 30 minutes) function in infrastructure monitoring corresponds to the collection point service group TLog and the collection point infrastructure.
- Service monitoring corresponds to the collection point service group EagleEye and the collection point stats_logger_agg.
- Trace analysis and query corresponds to the collection point service group TLog and the collection point EagleEye.

Impact on the system

An application change fails.

Procedure

- 1. Identify the corresponding TLog collection point based on the abnormal function.
- 2. Check whether the collection point has been started properly.

- i. On the Collection Points page, locate the row that contains the collection point, and click **More** > **Manually Assign Task** in the Actions column. View the dialog box that appears.
 - If the number in the dialog box is greater than 0, the collection point has been started properly. You can go to the next step.
 - If the collection point is not started properly, return to the Collection Points page and click Edit/Deployment Process. On the page shown in the following figure, click Start and wait until "Operation successful" appears in the result. If "Operation successful" does not appear in the result, contact EDAS Customer Services and give feedback like "The xxx collection point does not start properly."
 - If "Operation successful" appears in the result, return to the Collection Points page, wait for three to five minutes, and click Manually Assign Task again. If the number in the dialog box that appears is greater than 0, the collection point has been started properly. Go to the next step.
- 3. If the collection point has been started, check whether the collection rules are correctly distributed (operation risk level: G1).
 - On the Collection Points page, check whether the distribution status is Active. If it is Inactive, click Activate and OK in sequence.
 - Click Collection Point Details to go to the Collector Status tab. If the status list is not empty, collection rules are distributed properly. Go to the next step.
 - If the status list on the Collector Status tab is empty, return to the Collection Points page to manually distribute collection rules as follows: Click Create Task by Rule under Collection Rule. If the number of created tasks in the dialog box shown in the following figure is greater than 0, click OK. Return to the Collection Points page and click Manually Assign Task and OK in sequence. If manual distribution is successful, a dialog box appears.
 - If the number of created tasks in Create Task by Rule is 0, contact EDAS Customer Services for troubleshooting and give feedback like "The xxx collection point has 0 created tasks in Create Task by Rule."
- 4. If the collection point has been started and collection rules are distributed properly but no data exists, perform troubleshooting as follows (operation risk level: G1):
 - Click Collection Point Details to go to the Collector Status tab. Check the data in the Last Collection Attempt column. Normally, the time in this column is less than 1 minute. If the time in this column is generally greater than 1 minute, contact EDAS Customer Services and give feedback like "On the Collector Status tab for the xxx collection point, the time in the Last Collection Attempt column is generally greater than 1 minute. The collection point must be scaled up."
 - On the Collector Status tab, check the data in the Status column. Normally, the states in this column are Normal or File Not Modified. If states such as File Not Found, No Permission, Connection Timeout, and SProxy Not Found appear in this column, contact EDAS Customer Services and give feedback like "On the Collector Status tab for the xxx collection point, the yyy state appears in the Status Column."
 - On the Collection Points page, click More > Perform Health Check. Then, contact EDAS Customer Services and provide the JSON content on the health check page to help engineers quickly locate the problem.

9.1.6. Log reference

You can check logs to view the status of each EDAS component or locate faults during O&M.

EDAS provides logs for the following components:

- EDAS console
- EDAS admin
- EDAS server
- Cai-fs
- DiamondServer
- ConfigServer
- Cai-address
- EagleEye console

EDAS archives and clears the logs for these components based on predefined policies.

9.1.6.1. EDAS console logs

The EDAS console is the console component of EDAS. It provides the core functions of the PaaS platform, including resource management, application lifecycle management, service governance, and auto scaling.

Log files

EDAS console logs

File	Description
console.log	The EDAS console log.
changeorder.log	The change order log.
openapi.log	The API log.
tengine.log	The TEngine log.
debug.log	The log that records the internal API calls of the EDAS console.

Path

- Logging path: *\${user.home}/edas/logs*
- Archive path: *\${user.home}/edas/logs/bak*

Format

openapi.log: %msg%n (print log information directly) others: %d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %-5level %logger{50}:%line - %msg%n (date and time, thread name, log level, class name: number of lines - specific log information)

Archiving policies

EDAS console log archiving policies

Operations and Maintenance Guide · Operations of middleware products

Log	Archiving policy
cons ole .log	 Maximum size: 100 MB The name of the new file takes the format console.{d}.log. Seven logs are retained.
changeorder.log	 A file is created every day. The name of the new file takes the format changeorder.{yyyy-MM-dd}.log. Logs from the last seven days are retained.
openapi.log	 A file is created every day. The name of the new file takes the format opanapi.{yyyy-MM- dd}.log. Logs from the last seven days are retained.
tengine.log	 A file is created every day. The name of the new file takes the format tengine.{yyyy-MM- dd}.log. Logs from the last seven days are retained.
debug.log	 Maximum size: 100 MB The name of the new file takes the format debug.{d}.log. Three logs are retained.

9.1.6.2. EDAS admin logs

The EDAS admin is a background task service that provides the instance synchronization and application health check functions.

File

EDAS admin logs

File	Description
admin.log	The scheduling task log.
tengine.log	The Tengine log.

Path

- Logging path: *\${user.home}/edas/logs*
- Archive path: *\${user.home}/edas/logs/bak*

Format

%d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %-5level %logger{50}:%line - %msg%n (date and time, thread name, log level, category name: number of lines - specific log information)

Archiving policy

> Document Version:20200918

EDAS admin log archiving policies

Log	Archiving policy
admin.log	 Maximum size: 100 MB The name of the new file takes the format admin.{d}.log. Three logs are retained.
tengine.log	 A file is created every day. The name of the new file takes the format tengine.{yyyy-MM- dd}.log. Logs from the last seven days are retained.

9.1.6.3. EDAS server logs

The EDAS server synchronizes status information with EDAS Agent.

File

EDAS server logs

File	Description
agent-server.log	The log for the instance where EDAS Agent is installed.
changeorder.log	The change order log.
tengine.log	The Tengine log.

Path

- Logging path: *\${user.home}/edas/logs*
- Archive path: *\${user.home}/edas/logs/bak*

Format

agent-server: %d{HH:mm:ss.SSS} [%thread] %-5level %logger{36}:%line - %msg%n(time, thread name, log level, category name: number of lines - specific log information) others: %d{yyyy-MMdd HH:mm:ss.SSS} [%thread] %-5level %logger{50}:%line - %msg%n (date and time, thread name, log level, category name: number of lines - specific log information)

Archiving policy

Archiving policy for EDAS server logs

Log

Archiving policy

Log	Archiving policy		
agent-server.log	 A file is created every day. The name of the new file takes the format agent-server{yyyy-MM-dd}.log. Logs from the last seven days are retained. 		
tengine.log	 A file is created every day. The name of the new file takes the format tengine.{yyyy-MM- dd}.log. Logs from the last seven days are retained. 		
changeorder.log	 A file is created every day. The name of the new file takes the format changeorder.{yyyy-MM- dd}.log. Logs from the last seven days are retained. 		

9.1.6.4. DiamondServer logs

A configuration management service. It provides configuration storage, query, and notification functions, and primarily stores database metadata and EDAS function switch configurations in EDAS.

File

DiamondServer logs

File	Description
diamondServer.log	The DiamondServer log.
fata.log	The most important system log, which records database service errors, "master db not found" messages, and other information.
dump.log	The log that records the dumping of configurations to the local device.

Path

- Logging path: *\${user.home}/admin/diamond/logs*
- Archive path: *\${user.home}/admin/diamond/logs*

Format

[%p] [%t] %d{MM-dd HH:mm:ss,SSS} [%c{1}] - %m%n (log information priority, log event thread name, logging time, log information category, and specific log information)

Archiving policy

The archiving policies vary depending on the version. For example, in the latest version 3.8.8, a log is 15 MB in size and 10 logs are retained.

9.1.6.5. Cai-fs logs

A file server. It stores the EDAS Agent installation package and EDAS application packages.

Log files

EDAS console logs

File	Description
efs-server.log	Cai-fs logs

Path

- Logging path: *\${user.home}/efs/logs*
- Archive path: *\${user.home}/efs/logs*

Format

%d{HH:mm:ss} [%thread] %-5level %logger{36} - %msg%n (time, thread name, log level, class name: number of lines - specific log information)

Archiving policies

EDAS admin log archiving policies

Log	Archiving policy
efs-server.log	 Maximum size: 100 MB The name of the new file takes the format efs-server.log.{d}. Three logs are retained.

9.1.6.6. ConfigServer logs

An RPC service registry. It is used to query and store the publishing and subscription data of services.

Log files

EDAS console logs

File	Description
cluster.log	The log that records cluster operations, such as merging tasks and connecting to or disconnecting from other instances in the cluster.
memory.log	The log that records memory statuses, including the total number of subscriptions and the amount of persistent data of an instance.
persistent.log	The data persistence log.

File	Description		
push.log	The data push log.		
http.log	The log that records the instance commands called over HTTP.		
monitor.log	The warning code log.		

Path

- Logging path: *\${user.home}/admin/configserver/log*
- Archive path: *\${user.home}/admin/configserver/log*

Format

%date %level %msg%n%n (logging time, log level, and specific log information)

Archiving policies

- A file is created every day.
- The name of the new file takes the format {module}.log.%d{yyyy-MM-dd}.log. Logs from the last 15 days are retained.

9.1.6.7. Cai-address logs

An address discovery service. It provides the address lists for DiamondServer and ConfigServer.

Log files

EDAS console logs

File	Description
access.log	All access logs
error.log	Error logs

Path

- Logging path: *\${user.home}/admin/cai/logs*
- Archive path: *\${user.home}/admin/cai/logs*

Format

"\$remote_addr \$request_time_usec \$http_x_readtime [\$time_local] \"\$request_method http://\$host\$request_uri\" \$status \$body_bytes_sent \"\$http_referer\" \"\$http_user_agent\" \"\$md5_encode_cookie_unb\" \"\$md5_encode_\$cookie_cookie2\" \"\$eagleeye_traceid\""; records the client IP address - request elapsed time - request header - request status - the number of bytes sent to the client - records the link from which the access request is received records information about the web browser of the client - performs MD5 on cookies to obtain fixed-length cookies - eagleeye trace id

Archiving policies

Log splitting is not performed.

9.1.6.8. EagleEye console logs

You can query and view service traces.

Log files

EDAS console logs

File	Description	
eagleeye-console.log	All console access logs	
eagleeye-sql.log	Trace query logs	

Path

- Logging path: *\${user.home}/admin/logs*
- Archive path: \${user.home}/admin/logs

Format

%d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %msg%n (time, thread name, and specific log information)

Archiving policies

Archiving policies of EagleEye console access logs

Log	Archiving policy
eagleeye-console.log	 Maximum size: 500 MB Retention period: 30 days
eagleeye-sql.log	 Maximum size: 200 MB Retention period: 15 days

9.1.7. Configuration reference

You need to complete basic configuration and optimization configuration during the EDAS O&M process.

The configuration during EDAS O&M is divided into component configuration and JVM configuration.

9.1.7.1. Component configuration

You can configure the basic settings of components by using configuration files.

Parameters

		1			
Component	Configuration file	Path	Configuration item	Description	Value
EDAS console	AS console config.proper ties	/home/admin /edas/conf/	dataSource.c onfig.URL	Database connection string	Standard database connection string, such as jdbc:mysql:// edastest.mys ql.rds.aliyunc s.com/edas? rewriteBatch edStatement s=true
			dataSource.c onfig.user	Username	Standard database username
			dataSource.c onfig.passwo rd	Password	Standard database password
EDAS admin	DAS admin	/home/admin /edas/conf/	dataSource.c onfig.URL	Database connection string	Standard database connection string, such as jdbc:mysql:// edastest.mys ql.rds.aliyunc s.com/edas? rewriteBatch edStatement s=true
		dataSource.c onfig.user	Username	Standard database username	
			dataSource.c onfig.passwo rd	Password	Standard database password
			dataSource.c onfig.URL	Database connection string	Standard database connection string
Redis console	redis.conf	/home/admin /redis- 2.8.17/src/	dataSource.c onfig.user	Username	Standard database username

			dataSource.c onfig.passwo rd	Password	Standard database password
			config.tlog.zk .servers	ZooKeeper connection string	Standard ZooKeeper address information, such as 192.168.1.2:21 81, 192.168.1.3:21 81, and 192.168.1.4:21 81
TLog console	tlog- cloud.propert ies	/home/admin /taobao- tomcat- production- 7.0.59.3/lib/	config.tlog.hb ase.zkServer s	ZooKeeper used by HBase	Standard ZooKeeper connection string, which is shared by default and is consistent with the preceding ZooKeeper
			config.tlog.hb ase.zkRootNo de	HBase root node	Default value: /hbase
			config.nimbus .host	JStorm nimbus node	IP address of the primary node
			config.edas.c onsole.url	EDAS admin address	EDAS admin domain name
HBase	hbase- site.xml	/home/admin /hbase{- Version}/conf /	hbase.rootdir	Host name of the primary instance	Note that the host name cannot be an IP address and must be bound in /etc/hosts if it cannot be resolved. All HBase-based applications must be bound to the host names of all HBase instances.

Operations and Maintenance Guide • Operations of middleware products

			hbase.zooke eper.quorum	ZooKeeper connection string	ZooKeeper connection string
			hbase.zooke eper.property .clientPort	ZooKeeper port	ZooKeeper port
			storm.zookee per.servers	IP addresses of all storm nodes	IP addresses of all storm nodes
Jstorm	storm.yaml	/home/admin /jstorm/conf/	nimbus.host	Primary node of JStorm nimbus	Primary node of JStorm nimbus
			supervisor.sl ots.port.cpu. weight	CPU weight	Number of CPUs occupied by each task
ConfigServer	/ ConfigServer confsrv.conf /		serverlist	Server list	A list of IP addresses separated with commas (,)
		/conf/	unitserverlist	Modular server list	The content is the same as that of serverlist.
DiamondServ er	config.proper ties	/home/admin /diamond/tar get/diamond. war/WEB- INF/classes/	openInnerInt erfaceFilter	Indicates whether to enable internal interface access verification.	The default value is false. Enter false to avoid failed verification because Address- Server is typically configured with a virtual IP address rather than a real one.
			OPEN_SPAS	Indicates whether to enable authenticatio n.	The value is true, which indicates that authenticatio n is enabled.

9.1.7.2. JVM configuration

You can optimize system performance through a JVM configuration.

The parameters vary slightly depending on the JDK versions.

Parameters

Name	Description	Applicable JDK version	Reference value
-Xms	Specifies the initial heap memory size for the JVM.	All JDK versions	4 GB
-Xmx	Specifies the maximum heap memory size for the JVM.	All JDK versions	4 GB
-Xmn	Specifies the size of the young generation.	All JDK versions	2 GB
-Xss	Specifies the stack size of each thread.	All JDK versions	2 MB
- XX:+UseCompressedO ops	Compresses common object pointers.	All JDK versions	-
-XX:SurvivorRatio	Specifies the ratio of Survivor to Eden.	All JDK versions	10
- XX:+UseConcMarkSwe epGC	collector for memory		-
- XX:+UseCMSCompactA tFullCollection	Instructs the CMS collector to compress the old generation upon full GC.	All JDK versions	-
- XX:CMSMaxAbortable PrecleanTime		All JDK versions	5000
- XX:+CMSClassUnloadi ngEnabled	Specifies that CMS GC is triggered after class unloading.	All JDK versions	-
- XX:CMSInitiatingOccup ancyFraction	Sets the threshold size of the old generation that triggers CMS GC.	All JDK versions	80

Operations and Maintenance Guide • Operations of middleware products

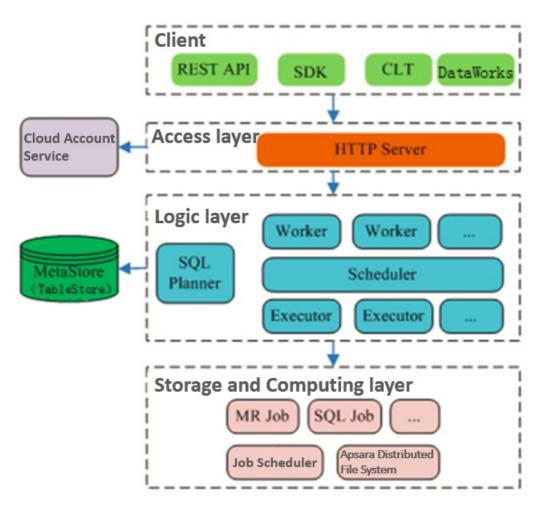
Name	Description	Applicable JDK version	Reference value
-XX:PermSize	Specifies the initial value of the permanent generation.	1.7 and earlier versions	196 MB
-XX:MaxPermSize	Specifies the maximum value of the permanent generation.	1.7 and earlier versions	256 MB
MetaspaceSize	Sets the threshold size of the allocated metadata space that triggers full GC.	1.8 and later versions	196 MB
MaxMetaspaceSize	Sets the maximum size of the allocated metadata space that triggers full GC.	1.8 and later versions	256 MB
-XX:+DisableExplicitGC	Disables System.gc().	All JDK versions	-
- XX:+HeapDumpOnOut OfMemoryError		All JDK versions	-
-XX:HeapDumpPath	Specifies the heap dump path.	All JDK versions	/home/admin/logs/oo mDump.log

10.Operations of big data products 10.1. MaxCompute

10.1.1. Concepts and architecture

MaxCompute architecture shows the MaxCompute architecture.

MaxCompute architecture



The MaxCompute service is divided into four parts: client, access layer, logic layer, and storage and computing layer. Each layer can be horizontally scaled.

The following methods can be used to implement the functions of a MaxCompute client:

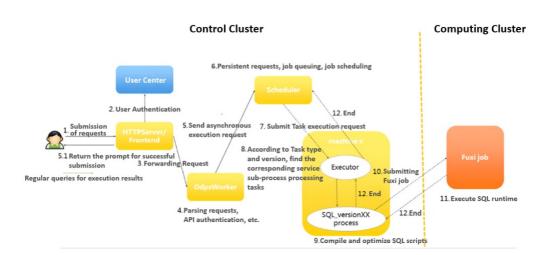
- API: RESTful APIs are used to provide offline data processing services.
- SDK: RESTful APIs are encapsulated within SDKs. SDKs are currently available in programming languages such as Java.
- **Command line tool (CLT):** This client-side tool runs on Windows and Linux. CLT allows you to submit commands to manage projects and use DDL and DML.
- DataWorks: DataWorks provides upper-layer visual ETL and BI tools that allow you to synchronize data, schedule tasks, and create reports.

The access layer of MaxCompute supports HTTP, HTTPS, load balancing, user authentication, and service-level access control.

The logic layer is at the core of MaxCompute and supports project and object management, command parsing and execution logic, and data object access control and authorization. The logic layer contains two clusters: control and compute clusters. The control cluster is designed to manage projects and objects, parse and start queries and commands, and control and authorize access to data objects. The compute cluster executes tasks. Both control and compute clusters can be horizontally scaled as needed. The control cluster has three roles: Worker, Scheduler, and Executor. These roles are described as follows:

- The Worker role processes all RESTful requests and manages projects, resources, and jobs. Workers forward jobs that need to launch Fuxi tasks (such as SQL, MapReduce, and Graph jobs) to the Scheduler for further processing.
- The Scheduler role schedules instances, splits instances into multiple tasks, sorts tasks that are pending for submission, and queries resource usage from FuxiMaster in the compute cluster for throttling. If there are no idle slots in Job Scheduler, the Scheduler stops processing task requests from Executors.
- The Executor role is responsible for launching SQL and MapReduce tasks. Executors submit Fuxi tasks to FuxiMaster in the compute cluster and monitor the operating status of these tasks.

When you submit a job request, the web server at the access layer queries the IP addresses of registered Workers and sends API requests to randomly selected Workers. The Workers then send these requests to the Scheduler for scheduling and throttling. Executors actively poll the Scheduler queue. If the necessary resources are available, the Executors start executing tasks and return the task execution status to the Scheduler. The following figure shows the MaxCompute job execution process.



MaxCompute job execution process

The following concepts are involved in the MaxCompute job execution process:

1. MaxCompute instance: the instance of a MaxCompute job. A job is anonymous if it is not defined. A MaxCompute job can contain multiple MaxCompute tasks. In a MaxCompute instance, you can submit multiple SQL or MapReduce tasks, and specify whether to run the tasks in parallel or serial mode. This scenario is rarely seen because MaxCompute jobs are

not commonly used. In most cases, an instance contains only one task.

- 2. MaxCompute task: a specific task in MaxCompute. Currently, there are almost 20 task types, such as SQL, MapReduce, Admin, Lot, and Xlib. The execution logic varies greatly depending on the task type. Different tasks in an instance are differentiated by their task name. MaxCompute tasks can run in the control cluster. Simple tasks such as metadata modification can run in the control cluster for their entire lifecycles. To run computing tasks, submit Fuxi jobs to the compute cluster.
- 3. Fuxi job: a computing model provided by the Job Scheduler module. A Fuxi job corresponds to a Fuxi service. A Fuxi job represents a task that can be completed, while a Fuxi service represents a resident process.
 - The DAG scheduling approach can be used to schedule Fuxi jobs. Each job has a job master to schedule its job resources.
 - For SQL, Fuxi jobs are divided into offline and online jobs. Online jobs evolve from the service mode jobs. An online job is also called a quasi-real-time task. An online job is a resident process that can be executed whenever there are tasks, reducing the time required to start and stop a job.
 - You can submit a MaxCompute task to multiple compute clusters. The primary key name of a Fuxi job is the cluster name followed by the job name.
 - The JSON plan for Job Scheduler to submit a job and the status of a finished job are stored in Apsara Distributed File System.
- 4. Fuxi task: a sub-concept of Fuxi job. Similar to MaxCompute tasks, different Fuxi tasks represent different execution logics. Fuxi tasks can be linked together as pipes to implement complex logic.
- 5. Fuxi instance: the instance of a Fuxi task. A Fuxi instance is the smallest unit that can be scheduled by Job Scheduler. During the actual execution process, a task is divided into many logical units to improve the processing speed. Different instances will run on the same execution logic but work with different input and output data.
- 6. Fuxi worker: an underlying concept of Job Scheduler. A worker represents an operating system process. A worker can be reused by multiple Fuxi instances, but a worker can only handle one instance at a time.
- ? Note
 - InstanceID: the unique identifier of a MaxCompute job. It is commonly used for troubleshooting. You can construct the LogView of the current instance based on the project name and instance ID.
 - Service master or job master: a primary node of the service or job type. The primary node is responsible for requesting and scheduling resources, creating work plans for workers, and monitoring workers across their entire lifecycles.

The storage and computing layer of MaxCompute is a core component of the proprietary cloud computing platform of Alibaba Cloud. As the kernel of the Apsara system, this component runs in the compute cluster independent of the control cluster. The architecture diagram illustrates only the major modules.

10.1.2. O&M commands and tools

10.1.2.1. Before you start

Before using MaxCompute O&M commands and tools, you must be aware of the following information:

During the MaxCompute O&M process, the default account is admin. You must run all commands as an admin user. You must use your admin account and sudo to run commands that require sudo privileges.

10.1.2.2. odpscmd commands

You can use the command line to perform operations and maintenance. You must log on to the command line tool before you can run commands. The specific procedure is as follows:

- 1. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose **Operations > Cluster Operations**. In the **Cluster** search box, enter odps to search for the expected cluster.
- 2. Click the cluster in the search result. On the Cluster Details page, click the Services tab. In the Services search box, search for odps-service-computer. Click odps-service-computer in the search result.
- 3. After you access the odps-service-computer service, select ComputerInit# on the Service Details page. In the Actions column corresponding to the machine, click Terminal. In the TerminalService window that appears, you can perform subsequent command line operations.

Console command directories and configurations

The MaxCompute client is located in the clt folder under the */apsara/odps_tools* directory of odpsag. The client configuration file is located in the conf directory under the clt folder. The access_id, access_key, end_point, log_view, and tunnel_point parameters are configured by default. You can use the ./clt/bin/odpscmd command to view information such as the version number in interactive mode. For example, run the HTTP GET /projects/admin_task_project/system; command to check the version information of MaxCompute.

Description of client command options

The following figure shows the client command options.

\$/apsara/odps tools/clt/bin/odpscmd -h	
Usage: odpscmd [OPTION]	
where options include:	
help config= <config_file> project=<prj_name> endpoint=<http: host:port=""> -u <user_name> -p <password> instance-priority=<priority> -M -k <n></n></priority></password></user_name></http:></prj_name></config_file>	(-h)for help specify another config file use project set endpoint user name and password priority scope(0-9) read machine readable data will skip begining queries and start from specified posi
ion	with puth polynomia duories and pours from shoetifed boot
<pre>-r <n> -f <"file_path;"> -e <"command;[command;]"> -C -y</n></pre>	set retry times execute command in file execute command, include sql command will display job counters will not submit jobs to fuxi master

Client command options

- -e: The MaxCompute client does not execute SQL statements in interactive mode.
- --project, -u, and -p: The client directly uses the specified values for the project, user, and

pass parameters. If you do not specify a parameter, the client uses the corresponding value configured in the conf file.

- -k and -f: The client directly executes local SQL files.
- --instance-priority: This option is used to assign a priority to the current task. Valid values: 0 to 9. A lower value indicates a higher priority.
- -r: This option indicates the number of times a failed command will be retried. It is commonly used in scripting jobs.

Commonly used SQL commands for O&M

The following table lists the commonly used commands.

Commonly used commands

Command	Description			
whoami;	Allows you to view your Apsara Stack tenant account and endpoint information.			
show p;	Allows you to view information about all instances that have been run.			
wait <instanceid>;</instanceid>	Allows you to re-generate the LogView and Fuxi job information of a task. To run this command, you must have owner permissions, and the LogView and Fuxi job information must be stored in the same project.			
kill <instanceid>;</instanceid>	Allows you to terminate specified instances.			
tunnel upload/download;	Allows you to test whether Tunnel is functioning.			
desc project <projectname> -extended;</projectname>	 Allows you to view the project usage. desc extended table: allows you to view table information. desc table_name partition(pt_spec): allows you to view partition information. desc resource \$resource_name: allows you to view project resource information. desc project \$project_name -extended: allows you to view cluster information. 			
export <project name=""> local_file_path;</project>	Allows you to export DDL statements of all tables in a project.			
create table tablename () ;	Allows you to create a table.			
<pre>select count(*) from tablename;</pre>	Allows you to search for a table.			
Explain	Allows you to create plans without submitting Fuxi jobs to view resources required for tasks.			
list	Allows you to list tables, resources, and roles.			

Command	Description
show	Allows you to view table and partition information.
	Allows you to remove all data from the MaxCompute recycle bin directly to the Apsara Distributed File System recycle bin.
purge	 purge table <tablename>: allows you to purge a single table.</tablename>
	 purge all: allows you to purge all tables from the current project.

10.1.2.3. Tunnel commands

The client provides Tunnel commands that implement the original functions of the Dship tool. Tunnel commands are mainly used to upload or download data.

Tunnel commands

Command	Description
tunnel upload	Allows you to upload data to MaxCompute tables. You can upload files or level-1 directories. Data can only be uploaded to a single table or table partition each time. The destination partition must be specified for partitioned tables.
tunnel download	Allows you to download data from MaxCompute tables. You can only download data to a single file. Only data in one table or partition can be downloaded to one file each time. For partitioned tables, the source partition must be specified.
tunnel resume	If an error occurs because of network or Tunnel service faults, you can resume file or directory transmission after interruption. This command only allows you to resume the previous data upload. Every data upload or download operation is called a session. Run the resume command and specify the ID of the session to be resumed.
tunnel show	Allows you to view historical task information.
tunnel purge	Purges the session directory. Sessions from the last three days are purged by default.

Tunnel commands allow you to view help information by using the Help sub-command on the client. The sub-commands of each Tunnel command are described as follows:

Upload

Imports data of a local file into a MaxCompute table. The following example shows how to use the sub-commands:

odps@ project_name>tunnel help upload;
usage: tunnel upload [options] <path> <[project.]table[/partition]></path>
upload data from local file
-acp,-auto-create-partition <arg> auto create target partition if not</arg>
exists, default false
-bs,-block-size <arg> block size in MiB, default 100</arg>
-c,-charset <arg> specify file charset, default ignore.</arg>
set ignore to download raw data
-cp,-compress <arg> compress, default true</arg>
-dbr,-discard-bad-records <arg> specify discard bad records</arg>
action(true false), default false
-dfp,-date-format-pattern <arg> specify date format pattern, default</arg>
yyyy-MM-dd HH:mm:ss
-fd,-field-delimiter <arg> specify field delimiter, support</arg>
unicode, eg \u0001. default ","
-h,-header <arg> if local file should have table</arg>
header, default false
-mbr,-max-bad-records <arg> max bad records, default 1000</arg>
-ni,-null-indicator <arg> specify null indicator string,</arg>
default ""(empty string)
-rd,-record-delimiter <arg> specify record delimiter, support</arg>
unicode, eg \u0001. default "\r\n"
-s,-scan <arg> specify scan file</arg>
action(true false only), default true
-sd,-session-dir <arg> set session dir, default</arg>
D:\software\odpscmd_public\plugins\ds
hip
-ss,-strict-schema <arg> specify strict schema mode. If false,</arg>
extra data will be abandoned and
insufficient field will be filled
with null. Default true
-te,-tunnel_endpoint <arg> tunnel endpoint</arg>
-threads <arg> number of threads, default 1</arg>
-tz,-time-zone <arg> time zone, default local timezone:</arg>
Asia/Shanghai
Example:
tunnel upload log.txt test_project.test_table/p1="b1",p2="b2"

Parameters:

• -acp: indicates whether to automatically create the destination partition if it does not exist.

No destination partition is created by default.

- -bs: specifies the size of each data block uploaded with Tunnel. Default value: 100 MiB (MiB = 1024 * 1024B).
- -c: specifies the local data file encoding format. Default value: UTF-8. If this parameter is not set, the encoding format of the downloaded source data is used by default.
- -cp: indicates whether to compress the local data file before it is uploaded to reduce network traffic. By default, the local data file is compressed before it is uploaded.
- -dbr: indicates whether to ignore dirty data (such as additional columns, missing columns, and columns with mismatched data types).
 - If this parameter is set to true, all data that does not comply with table definitions is ignored.
 - If this parameter is set to false, an error is returned when dirty data is found, so that raw data in the destination table is not contaminated.
- -dfp: specifies the DateTime format. Default value: yyyy-MM-dd HH:mm:ss.
- -fd: specifies the column delimiter used in the local data file. Default value: comma (,).
- -h: indicates whether the data file contains the header. If this parameter is set to true, Dship skips the header row and starts uploading data from the second row.
- -mbr: terminates any attempts to upload more than 1,000 rows of dirty data. This parameter allows you to adjust the maximum allowable volume of dirty data.
- -ni: specifies the NULL data identifier. Default value: an empty string ("").
- -rd: specifies the row delimiter used in the local data file. Default value: \r\n.
- -s: indicates whether to scan the local data file. Default value: false.
 - If this parameter is set to true, the system scans the source data first, and then imports the data if the format is correct.
 - If this parameter is set to false, the system imports data directly without scanning.
 - If this parameter is set to only, the system only scans the source data, and does not import the data after scanning.
- -sd: sets the session directory.
- -te: specifies the Tunnel endpoint.
- -threads: specifies the number of threads. Default value: 1.
- -tz: specifies the time zone. Default value: Asia/Shanghai.

Show

Displays historical records. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help show;
usage: tunnel show history [options]
show session information
-n,-number <ARG> lines
Example:
tunnel show history -n 5
tunnel show log
```

Parameters:

-n: specifies the number of rows to be displayed.

Resume

Resumes the execution of historical operations (only applicable to data upload). The following example shows how to use the sub-commands:

odps@ project_name>tunnel help resume; usage: tunnel resume [session_id] [-force] resume an upload session -f,-force force resume Example:

tunnel resume

Download

The following example shows how to use the sub-commands:

odps@ project_name>tunnel help download;						
usage: tunnel dowr	usage: tunnel download [options] <[project.]table[/partition]> <path></path>					
download d	ata to local file					
-c,-charset <arg></arg>	specify file charset, default ignore.					
	set ignore to download raw data					
-ci,-columns-index	<arg> specify the columns index(starts from</arg>					
	0) to download, use comma to split each					
	index					
-cn,-columns-name	<arg> specify the columns name to download,</arg>					
	use comma to split each name					
-cp,-compress <ar(< td=""><td>G> compress, default true</td></ar(<>	G> compress, default true					
-dfp,-date-format-	pattern <arg> specify date format pattern, default</arg>					
	yyyy-MM-dd HH:mm:ss					
-e,-exponential <ai< td=""><td>RG> When download double values, use</td></ai<>	RG> When download double values, use					
	exponential express if necessary.					
	Otherwise at most 20 digits will be					
	reserved. Default false					
-fd,-field-delimiter	<arg> specify field delimiter, support</arg>					
	unicode, eg \u0001. default ","					
-h,-header <arg></arg>	if local file should have table header,					
	default false					
-limit <arg></arg>	specify the number of records to					
	download					
-ni,-null-indicator <	ARG> specify null indicator string, default					
	""(emnty string)					

(cinpty string)
-rd,-record-delimiter <arg> specify record delimiter, support</arg>
unicode, eg \u0001. default "\r\n"
-sd,-session-dir <arg> set session dir, default</arg>
D:\software\odpscmd_public\plugins\dshi
p
-te,-tunnel_endpoint <arg> tunnel endpoint</arg>
-threads <arg> number of threads, default 1</arg>
-tz,-time-zone <arg> time zone, default local timezone:</arg>
Asia/Shanghai
usage: tunnel download [options] instance://<[project/]instance_id> <path></path>
download instance result to local file
-c,-charset <arg> specify file charset, default ignore.</arg>
set ignore to download raw data
-ci,-columns-index <arg> specify the columns index(starts from</arg>
0) to download, use comma to split each
index
-cn,-columns-name <arg> specify the columns name to download,</arg>
use comma to split each name
-cp,-compress <arg> compress, default true</arg>
-dfp,-date-format-pattern <arg> specify date format pattern, default</arg>
yyyy-MM-dd HH:mm:ss
-e,-exponential <arg> When download double values, use</arg>
exponential express if necessary.
Otherwise at most 20 digits will be
reserved. Default false
-fd,-field-delimiter <arg> specify field delimiter, support</arg>
unicode, eg \u0001. default ","
-h,-header <arg> if local file should have table header,</arg>
default false
-limit <arg> specify the number of records to</arg>
download
-ni,-null-indicator <arg> specify null indicator string, default</arg>
""(empty string)
-rd,-record-delimiter <arg> specify record delimiter, support</arg>
unicode, eg \u0001. default "\r\n"
-sd,-session-dir <arg> set session dir, default</arg>
D:\software\odpscmd_public\plugins\dshi
p
-te,-tunnel_endpoint <arg> tunnel endpoint</arg>
-threads <arg> number of threads, default 1</arg>

-tz,-time-zone <ARG> time zone, default local timezone:

Asia/Shanghai

Example:

tunnel download test_project.test_table/p1="b1",p2="b2" log.txt

tunnel download instance://test_project/test_instance log.txt

Parameters:

- -c: specifies the local data file encoding format. Default value: UTF-8.
- -ci: specifies the column index (starting from 0) for downloading. Separate multiple entries with commas (,).
- -cn: specifies the names of columns to be downloaded. Separate multiple entries with commas (,).
- -cp, -compress: indicates whether to compress the data file before it is uploaded to reduce network traffic. By default, a data file is compressed by it is uploaded.
- -dfp: specifies the DateTime format. Default value: yyyy-MM-dd HH:mm:ss.
- -e: allows you to express the values as exponential functions when you download Double type data. If this parameter is not set, a maximum of 20 digits can be retained.
- -fd: specifies the column delimiter used in the local data file. Default value: comma (,).
- -h: indicates whether the data file contains a header. If this parameter is set to true, Dship skips the header row and starts downloading data from the second row.

⑦ Note -h=true and threads>1 cannot be used together.

- -limit: specifies the number of files to be downloaded.
- -ni: specifies the NULL data identifier. Default value: an empty string ("").
- -rd: specifies the row delimiter used in the local data file. Default value: \r\n.
- -sd: sets the session directory.
- -te: specifies the Tunnel endpoint.
- -threads: specifies the number of threads. Default value: 1.
- -tz: specifies the time zone. Default value: Asia/Shanghai.

Purge

Purges the session directory. Sessions from the last three days are purged by default. The following example shows how to use the sub-commands:

```
odps@ project_name>tunnel help purge;
usage: tunnel purge [n]
force session history to be purged.([n] days before, default
3 days)
Example:
tunnel purge 5
```

10.1.2.4. LogView tool

10.1.2.4.1. Before you start

You must confirm the LogView process status before using LogView. If the process status is off, you must start the LogView process.

The procedure for querying the process status and starting the process is as follows:

- Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose Operations > Cluster Operations. In the Cluster search box, enter odps to search for the expected cluster.
- 2. Click the cluster in the search result. On the Cluster Details page, click the Services tab. In the Service search box, search for odps-service-console. Click odps-service-console in the search result.
- 3. After you access the odps-service-console service, select LogView# on the Service Details page. In the Actions column corresponding to the machine, click Terminal to open the TerminalService window.
- 4. Run the following command to find the Docker container where LogView resides:

docker ps|grep logview

5. Run the following commands to view the LogView process status:

ps -aux|grep logview

netstat -ntulp|grep 9000

6. If the process status is off, run the following command to start the process:

/opt/aliyun/app/logview/bin/control start

The following sections describe what is LogView and how to use LogView to perform basic operations.

10.1.2.4.2. LogView introduction

LogView is a tool for checking and debugging a job submitted to MaxCompute. LogView allows you to check the running details of a job.

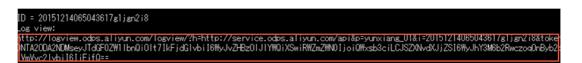
LogView functions

LogView allows you to check the running status, details, and results of a job, and the progress of each phase.

LogView endpoint

Take the odpscmd client as an example. After you submit an SQL task on the client, a long string starting with logview is returned.

A long string starting with logview



Enter the string with all carriage return and line feed characters removed in the address bar of the browser.

Composition of a LogView string

A LogView string consists of five parts, as shown in the following figure.

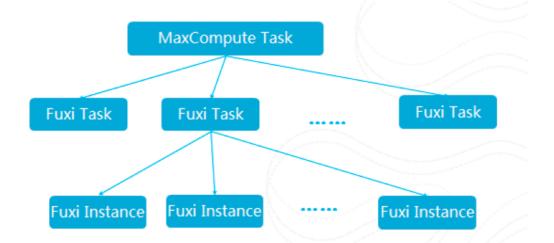
Composition of a LogView string



10.1.2.4.3. Preliminary knowledge of LogView

For complex SQL queries, you must have an in-depth knowledge of the relationships between MaxCompute tasks and Fuxi instances before you can understand LogView.

In short, a MaxCompute task consists of one or more Fuxi jobs. Each Fuxi job consists of one or more Fuxi tasks. Each Fuxi task consists of one or more Fuxi instances.



Relationships between MaxCompute tasks and Fuxi instances

The following figures show the relevant information in LogView.

MaxCompute Instance

MaxCompute Instance

Operations and Maintenance Guide • Operations of big data products

ODPS Instance							
URL	Project	InstanceID	Owner	StartTime	EndTime	Status	Sourcester
http://service.odps.aliyun.co	yunxlang_01	20151214065043617g	ALIYUN\$traini	2015-12-14 14:5	2015-12-14 14:5	Terminated	
						50	
		-				sole_select_qu	Source for: 20151214065043617c11m218
Node XML: [console	_select_qu	ery_task_145007	5843613]		×		xml version="1.0" encoding="UTF-8"?
7(10.10.52.38/ali-87 <property> <name>guid<!--1<br--><value>69f5682 </value></name></property> <property> <name>uuid<!--1</td--><td></td></name> idata.userage (315n)","odp 21-a782-45b Vame></property>		- ent":"CLT(0.17.3 : 9	a2149c); Wi format":"Hun 5d6	nanReadable"}			<pre>chois void in the set of the</pre>

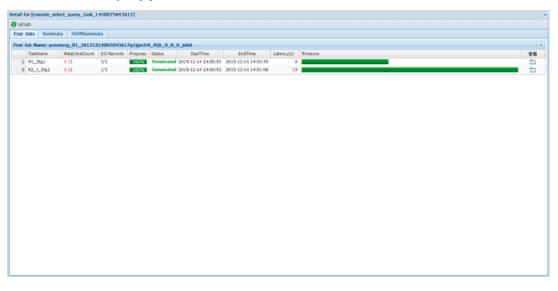
MaxCompute Task

MaxCompute Task

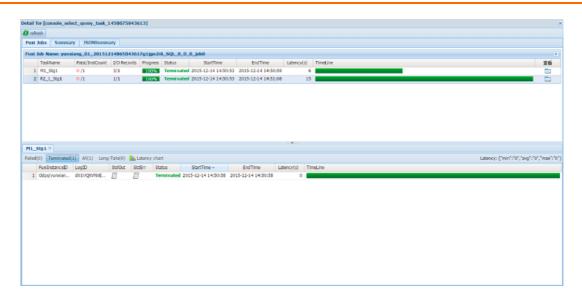
ODP5 Tasks								
Name	Type	Status	Result	Detail	StartTime	EndTime		TimeLine
console_select_query	SQL	Success			2015-12-14 14:50:43	2015-12-14 14:51:14	3	
Result for [console	_select	_query_ta	sk_1450	07584	3613]		×	
++								
_c0 ++								Defaul for Journele_selent_poory_teal_1430073413013]
3								Ø shuh Fruit John Semmany Holedownany
++								Paul Mit Rame ywnoang 61, 20131214943001817g1gechill, 102, 6, 9, 340 Tealfame fast heferau (2014eosh) Pageo Salar SarTine beffine (2014eos) Tealtre
								1 H0_806 0/1 3/1 MNNN Terrenaled 2029-12-04 14:30:30 000-12-04 04:30:30 0
								2 42_1_901 0/1 1/1 MANN 1055104195910 2015034195910 15

Task Detail - Fuxi Job

Task Detail - Fuxi Job(1)



Task Detail - Fuxi Job(2)



1000 CV/*

Task Detail - Summary

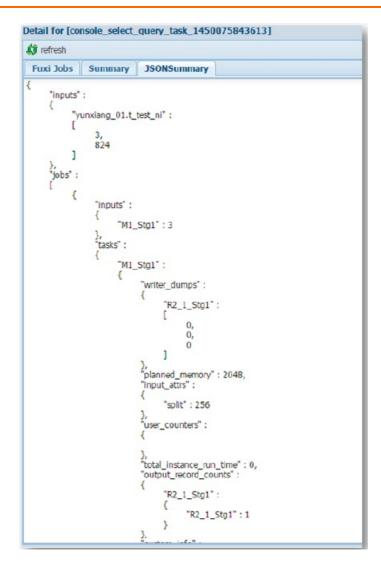
Task Detail - Summary

Fuxi Job	s Summary JSONSummary
	ost: cpu 0.00 Core * Min, memory 0.00 GB * Min
nputs:	inc. 01 h hat als 2 (024 h has)
yunx	iang_01.t_test_ni: 3 (824 bytes)
	ne: 15.000
	de: fuxi job
41 Sto1:	
	nce count: 1
run t	ime: 6.000
insta	nce time:
	min: 0.000, max: 0.000, avg: 0.000
inpu	records:
outo	input: 3 (min: 3, max: 3, avg: 3) ut records:
outp	R2_1_Stg1: 1 (min: 1, max: 1, avg: 1)
write	r dumps:
	R2 1 Stg1: (min: 0, max: 0, avg: 0)
2_1_Stg1	
	nce count: 1
	ime: 15.000
insta	nce time:
1 march	min: 0.000, max: 0.000, avg: 0.000
inpu	records: input: 1 (min: 1, max: 1, avo: 1)
outo	ut records:
outp	R2 1 Stg1FS 940124: 1 (min: 1, max: 1, avg: 1)
read	r dumps:
	input: (min: 0, max: 0, avg: 0)

odilitat y
resource cost: cpu 0.00 Core * Min, memory 0.00 GB * Min
inputs:
yunxiang_01.t_test_ni: 3 (824 bytes)
outputs:
Job run time: 15.000
Job run mode: fuxi job
M1_Stg1:
instance count: 1
run time: 6.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:
input: 3 (min: 3, max: 3, avg: 3)
output records:
R2_1_Stg1: 1 (min: 1, max: 1, avg: 1)
writer dumps:
R2_1_Stg1: (min: 0, max: 0, avg: 0)
R2_1_Stg1:
instance count: 1
run time: 15.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:
input: 1 (min: 1, max: 1, avg: 1)
output records:
R2_1_Stg1FS_940124: 1 (min: 1, max: 1, avg: 1)
reader dumps:
input: (min: 0, max: 0, avg: 0)

Task Detail - JSONSummary

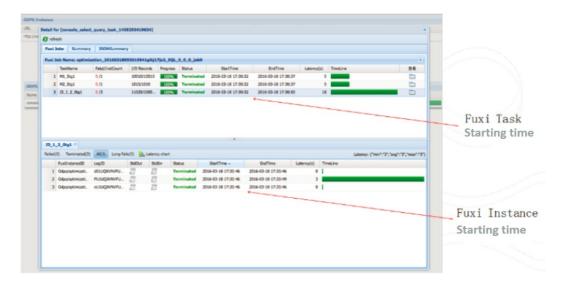
Task Detail - JSONSummary



10.1.2.4.4. Basic operations and examples

View each point in time in the life cycle of a job.

View each point in time in the life cycle of a job



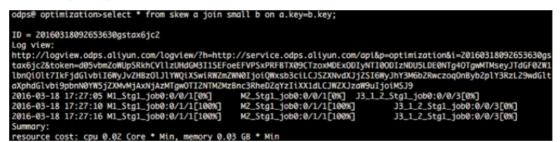
View the time it takes for Job Scheduler to schedule an instance.

	Netail for [console_sele	d_query_task_145	8293419634]							×	
	A) refresh										
	Fuxi Jobs Summar	y 350NSummar	γ								
	Fuxi Job Name: optim	ization_201603180	193019841gikj	17je2_9QL	0_0_job0					(A)	
	TaskName	Fatal/InstCount	1/O Records	Progress	Satus	StartTime	EndTime	Latency(s)	TimeLine	24	
	1 M1_Stg1	n/ 0	10010/10010	100%	Terminated	2016-03-18 17:30:32	2016-08-18 17:30:37	5			
s	2 M2_Stg1	0/1	1010/1010	100%	Terminated	2016-03-18 17:30:32	2016-08-18 17:30:37	5			
	3 33_1_2_9kg1	0,/3	11020/1000.	100%	Terminated	2016-03-18 17:30:32	2016-08-18 17:30:50	18			Fuxi Instanc
15-	33_1_2_Stg1 *	D AICD Long-Te	le(0)	wy chart					Latence: Chai	10137/10001037/100001037	Starting time - Fuxi Task Starting tim
P	FueiInstanceID	LegID	201. A. C.		tus	StartTime .	EndTine Lu	stency(s) Tir			July Starting time
п	1 Odps/optimizati.	d01UQKVNVFU	3 5	Te	rminated 20	16-03-18 17:30:46 2	016-03-18 17:30:46	0			=
	2 Odps/optimizati.	PUSUQININVEU	5 8	j Te	rminated 20	16-03-18 17:30:46 2	016-03-1817:30:49				Everi a da a da di
	3 Odps/optimizati.	eUSUQRVIN/FU.	3 8	j Te	rminated 20	116-03-18 17:30:46 2	046-03 1817:30.46	0			Fux i scheduli
I											takes time

View the time it takes for Job Scheduler to schedule an instance

View the polling interval.

View the polling interval



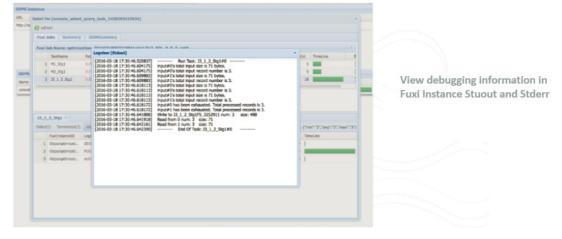
After a MaxCompute instance is submitted, odpscmd polls the execution status of the job at a specified interval of approximately 5s.

Check for data skews

Check for data skews

13 -	for [console_select	(very_task_145	8295419634]							×	
-	Jobs Summary	350NSummar	γ.								
Pasi	Job Name: optimiz	rien_201603180	193019841glkj1	7je2_9QL	0_0_0_0400					1	
	TaskNerre	Fatal/InstCount	1/0 Records	Progress	Satus	StartTime	EndTime	Latency(s)	Timbine	24	2
	M1_\$1p1	0/1	10110/10010	100%	Terminated	2016-03-18 17:30:3	32 2016-03-18 17:3	0.37	5		Different instances in the same I
	2 M2 Sto1	0/1	1000/1010	100%	Terminated	2016-03-18 17:30:3	12 2016-03-18 17:3	0:37	5		A. I. I. I.I. A. I. I.I. I.
	33_1_2_91g1	0/3	11020/1000	100%	Terminated	2016-03-18 17:30:3	2016-03-18 17:3	0:50 1	8		task should run for similar times
											💻 🛛 In this example, data skew occu
Palled	2_Stg1 = (0) Terminated(3) FuelInstanceID Odgs/optimizati Odgs/optimizati	AIC3) Long-Tai logED s01UQX/WVFU NJ1UQX/WVFU	laten SkOut Ski	Err Sa Ter	minated 20	StatTime = 06-03-18 17:30:46 16-03-18 17:30:46	EndTime 2014-03-18 17:30 46 2014-03-18 17:30 49	0	Ineline	latency: (hein195,16g1/95,16an195)	
	Odps/optimizati		ля	Ter	minated 20	18-02-18-12-30-46	2014-03-18 17:30	0			Click on stdout to see the amoun
											data processed, which can accur determine the data skew, that is amount of data processed betwe different instances varies greatly

View the UDF and MR debugging information



View the UDF and MR debugging information

View the task status - Terminated

Steve the task status - Terminated

 Steve
 The status - Terminated

 Steve
 The steve

 Steve
 Steve</

10.1.2.4.5. Best practices

Locate LogView based on the instance ID

After you submit a job, you can press Ctrl+C to return to odpscmd and perform other operations. You can run the wait <instanceid>; command to locate LogView and obtain the job status.

Locate LogView based on the instance ID

odpsē optimization-select * from skew a join skew2 b on a.key=b.key; ID = 20160318095028941gopbx61c2
Log view: http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=cptimization&i=20160318095028941gopbx6jc2&token=U02EU1RNbGhmRE5 jbHNIN2gwY016MjfobjhRPSxPRFBTK09CTzoxMDExODJyNTI00D1zNDU5LDE0NTg4OTK0MjkseyJTdGF0ZW1lbnQi01t71KfjdGlvbi16WyJvZHBz01JlYWQiXSwiRWZmZWN0IjoiQWxsb 3ciLCJSXWxMjJ32SMU9JhY3M65D2RvczoqOnByb2plY3RzL29wdGltoXphdGlvbi9pbnN0YW5jZXWxMjAxNj4zMTgwOTUwMjg5NDFh3BieDZqYZIiXX1dLCJWZXJzoW9uJjoiMSJ9 2016-03-18 17:50:40 NL_stg1_job0:0/0/1[0%] M2_stg1_job0:0/0/1[0%] J3_L2_Stg1_job0:0/0/3[0%]
2016-03-18 17:50:45 ML_5tgl_job0:0/1/1[100%] M2_Stgl_job0:0/1/1[100%] J3_1_2_Stgl_job0:0/0/3[0%] Instance running background. Use 'kill 20160318095028941gopbx6jc2' to stop this instance. Use 'wait 20160318095028941gopbx6jc2' to get details of this instance. odps@ optimization+wait 20160318095028941gopbx6jc2;
<pre>ID = 20160318095028941gopbx6jc2 Log view: http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=cptimization&i=20160318095028941gopbx6jc2&token=NVFFc1g2VIFSNmx 2TINGew91L2QvWU&z0UhFP5xPRFB1X09CTzoxM0ExDDJyNT100D1zNDUSLDE0NTg4OTK0NTcseyJTdGF0ZW1LDnqiOlt7IkFjdGlvb16NyJvZHBz0LJ1YNQiXSntRWZmZNN0Ijo1QWxsb 3c1LCJSXXNvdJjZS10MyJhT3M6b2Ruczoq0nBgb2p173RzL29mdGltaDphdGlvb19ppnN0YN5JZXNvMjAxNjazNTgAOTUMNjg5NDFnb3BieDZqYZ1iXX1dLCJWZXJzaW9uIjo1MSJ9 2016-09-18 17:50:8 ML_5tg1_job0:0/1/1[100%] M2_5tg1_job0:0/1/1[100%] J3_1_2_5tg1_job0:0/0/3[0%] Instance running background. Use 'will 20160318095028941gopbx6jc2' to stop this instance. Use 'will 20160318095028941gopbx6jc2' to et details of this instance.</pre>

Locate running tasks

After you exit the control window, you can run the show p; command to locate currently running tasks and historical tasks.

Locate running tasks

StartTime	RunTime	Status	InstanceID	Owner		Query	
2016-09-18 16:27:04	7s	Success	20160918082704275guto17jc2	ALIYUNŞ	liyun.com	select	from dual;

10.1.2.5. Apsara Bigdata Manager

Apsara Bigdata Manager (ABM) supports O&M on big data products from the perspective of business, services, clusters, and hosts. You can also upgrade big data products, customize alert configurations, and view the O&M history in ABM.

On-site Apsara Stack engineers can use ABM to easily manage big data products through actions such as viewing resource usage, checking and handling alerts, and modifying configurations.

For more information about how to log on to Apsara Bigdata Manager, see related documentation.

10.1.3. Routine O&M

10.1.3.1. Configurations

MaxCompute configurations are stored in the */apsara/odps_service/deploy/env.cfg* directory in odpsag. The configuration file contains the following content:

odps_worker_num=3 executor_worker_num=3 hiveserver_worker_num=3 replication_server_num=3 messager_partition_num=3

You can modify these parameter values based on your requirements and start the corresponding MaxCompute services based on the configured values. For more information, see *Restart a MaxCompute service*.

If you add xstream_max_worker_num=3 at the end of the configuration file, XStream will be started with three running workers.

10.1.3.2. Routine inspections

- 1. On the Cluster Operations page in Apsara Infrastructure Management Framework, check whether all machines have reached the desired state.
 - i. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose **Operations > Cluster Operations**. In the **Cluster** search box, enter odps to search for the expected cluster.
 - ii. Based on the information in the Status, Machine Status, and Server Role Status columns, check whether all machines have reached the desired state. The following figure shows that some machines have not reached the desired state.
 - iii. Click the exceptions in the Machine Status and Server Role status columns to view the exception details.
- 2. Go to the */home/admin/odps/odps_tools/clt/bin/odpscmd -e* directory and run the following command:

select count(*) from datahub_smoke_test;

odps@ odps_smoke_test>select count(*) from dual;
odpse odps_smoke_test/select count(*) from dual;
ID = 20180420061754827g78x7i
Log view:
http://logview.cn-hangzhou-env6-d01.odps.aliyun-inc.com:9000/logview/?h=http://s
180420061754827g78x7istoken=aEVmNTF1dm5GMnF0V1BSWjViZE0r0WRERnZFPSxPRFBTX09CTzox
SwiRWZmZWN0IjoiQWxsb3ciLCJS2XNvdXJjZSI6WyJhY3M6b2RwczoqOnByb2p1Y3RzL29kcHNfc21va
39
Job Queueing.
Summary:
resource cost: cpu 0.00 Core * Min, memory 0.00 GB * Min
inputs:
odps_smoke_test.dual: 1 (1408 bytes)
outputs:
Job run time: 0.000
Job run mode: service job
Job run engine: execution engine
M1:
instance count: 1 run time: 0.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:
TableScan REL5136522: 1 (min: 1, max: 1, avg: 1)
output records:
StreamLineWrite REL5136523: 1 (min: 1, max: 1, avg: 1)
R2 1:
instance count: 1
run time: 0.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:
StreamLineRead_REL5136524: 1 (min: 1, max: 1, avg: 1)
output records:
ADHOC_SINK_5136527: 1 (min: 1, max: 1, avg: 1)
_c0 !

As shown in the following figure, fuxi job is running. The command output indicates that the cluster functions properly.

odps@ odps_smoke_test*select count(*) from datahub_smoke_test
>;
ID = 20180420065305115gv5pf9d
Log view:
<pre>http://logview.cn-beijing-bgm-d01.odps.bgm.com:9000/logview/?h=http://servio 80420065305115gv5pf9d&token=VS9hRzc4RjAzeXJ2bmRF0UtyYnNWSXFkNW0wPSxPRFBTX090</pre>
iI6WyJvZHBzOlJlYWQiXSwiRWZmZWN0IjoiQWxsb3ciLCJSZXNvdXJjZSI6WyJhY3M6b2Rwczoq
UzMDUxMTVndjVwZjlkIl19XSwiVmVyc2lvbiI6IjEifQ==
2018-04-20 14:53:10 M1 Stg1_job0:0/0/1[0%] R2_1_Stg1_job0:0/0/1[0%]
2018-04-20 14:53:15 M1_Stg1_job0:0/1/1[100%] R2_1_Stg1_job0:0/0/1[0%]
2018-04-20 14:53:20 M1_Stg1_job0:0/1/1[100%] R2_1_Stg1_job0:0/1/1[100%]
2018-04-20 14:53:25 M1_Stg1_job0:0/1/1[100%] R2_1_Stg1_job0:0/1/1[100%]
Summary:
resource cost: cpu 0.00 Core * Min, memory 0.00 GB * Min
inputs:
odps_smoke_test.datahub_smoke_test: 10 (745 bytes)
outputs: Job run time: 10.000
Job run mode: fuxi job
Ml_Stg1:
instance count: 1
run time: 5.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:
input: 10 (min: 10, max: 10, avg: 10)
output records:
R2_1_Stgl: 1 (min: 1, max: 1, avg: 1) writer dumps:
R2_1_Stg1: (min: 0, max: 0, avg: 0)
R2_1_Stg1:
instance count: 1
run time: 10.000
instance time:
min: 0.000, max: 0.000, avg: 0.000
input records:

- 3. Run the following commands to check whether the following workers exist and whether they have been restarted recently:
 - i. r swl Odps/MessagerServicex

¢r swl Odps/MessagerServicex WorkerName	LastUpda	te	Time		pid	planned	loaded	unloaded
MessageServerRole@101h05215.cloud.h07.amtest1284	Mon Apr	9	16:49:03	2018	24697	1	1	
MessageServerRole@101h11210.cloud.h13.amtest1284	Mon Apr	9	16:48:37	2018	15149	1	1	
MessageServerRole@101h08109.cloud.h09.amtest1284	Mon Apr	9	16:49:03	2018	23586	1	1	0

ii. r swl Odps/OdpsServicex

WorkerName	LastUpda	te	Time		pid	planned	loaded	unloaded
RecycleWorker@101h08114.cloud.h09.amtest1284	Mon Apr	9	17:05:42	2018	52905			
DdpsWorker@101h08114.cloud.h09.amtest1284	Mon Apr	9	17:05:42	2018	52904			
DdpsWorker@101h11010.cloud.h11.amtest1284	Mon Apr	9	17:04:06	2018	4454			
ExecutorWorker@101h08114.cloud.h09.amtest1284	Mon Apr	9	17:05:42	2018	52903			
ExecutorWorker@101h11010.cloud.h11.amtest1284	Mon Apr	9	17:04:22	2018	6524			
SchedulerWorker@101h08114.cloud.h09.amtest1284	Mon Apr	9	17:05:47	2018	53609	0	0	
WorkflowWorker@101h08114.cloud.h09.amtest1284	Mon Apr	9	17:05:48	2018	53610	0	0	0

iii. r swl Odps/HiveServerx

WorkerName	Las	tUpd	ate	Time			pid	planned	loaded	unloaded
AuthServer@101h08114.cloud.h09.amtest1284	Tue	Apr	10	18:05:54	2018		23585			
HiveServer@101h11010.cloud.h11.amtest1284	Mon	Apr	9	17:03:07	2018		1696	1	1	
HiveServer@101h08114.cloud.h09.amtest1284	Tue	Apr	10	18:06:02	2018		23587	2	2	
CatalogServer@101h08114.cloud.h09.amtest1284	Tue	Apr	10	18:05:55	2018	1	23586	1	1	0

iv. r swl Odps/QuotaServicex

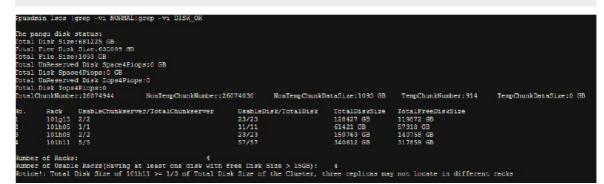
<pre>\$r swl Odps/QuotaServicex</pre>													
WorkerName		LastUpda	te:	Time			pid		planned		loaded		unloaded
QuotaWorkerRole@101h08114.cloud.h09.amtest1284	I.	Mon Apr	9	16:55:32	2018	I	32814	I	0	1	0	T	0

v. r swl Odps/ReplicationServicex

WorkerName	LastUpd	ate	Time		pid	planned	loaded	unloaded
ReplicationServer@101h05215.cloud.h07.amtest1284	Mon Apr	9	16:49:12	2018	26594			
ReplicationServer@101h11210.cloud.h13.amtest1284	Mon Apr	9	16:48:51	2018	26859			
ReplicationServer@101h11215.cloud.h13.amtest1284	Mon Apr	9	16:49:18	2018	3453			
ReplicationMaster@101h11010.cloud.h11.amtest1284	Mon Apr	9	16:50:21	2018	34315	0	0	0

4. Run the following command to check for errors:

puadmin lscs |grep -vi NORMAL|grep -vi DISK_OK



- 5. Run the following commands to check the data integrity:
 - i. puadmin fs -abnchunk -t none

```
$puadmin fs -abnchunk -t none
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type FoundTime
```

ii. puadmin fs -abnchunk -t onecopy

```
$puadmin fs -abnchunk -t onecopy
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type FoundTime
```

iii. puadmin fs -abnchunk -t lessmin

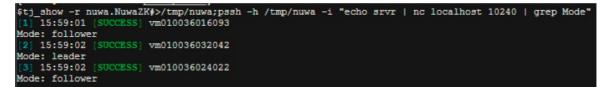
```
$puadmin fs -abnchunk -t lessmin
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type FoundTime
```

6. Log on to the machine where Apsara Name Service and Distributed Lock Synchronization System resides.

echo srvr | nc localhost 10240 | grep Mode

Example:

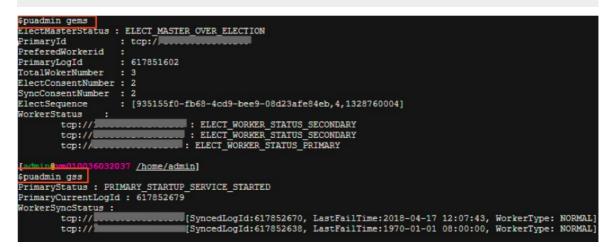
tj_show -r nuwa.NuwaZK#>/tmp/nuwa;pssh -h /tmp/nuwa -i "echo srvr|nc localhost 10240|grep Mode"



7. Run the following commands to check whether Apsara Distributed File System functions properly:

puadmin gems

puadmin gss



8. Perform daily inspections in Apsara BigData Manager to check disk usage.

10.1.3.3. Shut down a chunkserver, perform maintenance,

and then clone the chunkserver

Prerequisites

- A customer has asked to fix a faulty instance of odps_cs and clone a new one.
- You must inform the customer that this operation will temporarily render a chunkserver in the cluster unavailable, but will not affect the overall operation of the service.
- All MaxCompute services have reached the desired state and are functioning properly.
- All services on the OPS1 server have reached the desired state and are functioning properly.
- You must ensure that the disk space available is sufficient for data migration triggered when a node goes offline.
- If the primary node exists on the machine to be brought offline, you must ensure that services are switched from the primary node to the secondary node.

Procedure

1. In Apsara Infrastructure Management Framework, find **ComputerInit#** in the odps-servicecomputer service of the odps cluster, and open the corresponding TerminalService window. Run the following commands to check the data integrity of Apsara Distributed File System:

puadmin abnchunk fs -t none

-- Check for any missing files. If no output is displayed, no files are missing.

puadmin abnchunk fs -t onecopy

-- Check whether each file has only one copy. If no output is displayed, each file has only one copy

puadmin abnchunk fs -t lessmin

-- Check whether the number of files is smaller than the minimum number of backups. If no output

is displayed, the number of files is smaller than the minimum number of backups.

- 2. Add the machine to be shut down to a Job Scheduler blacklist.
 - i. Run the following command to enable the blacklisting function of Job Scheduler (ignore this step if the function has been enabled):

/apsara/deploy/rpc_caller --Server=nuwa://localcluster/sys/fuxi/master/ForClient --Method =/fuxi/SetGlobalFlag --Parameter={\"fuxi_Enable_BadNodeManager\":false}

ii. Run the following command to check the hostnames in the existing blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

iii. Run the following command to add the machine to be shut down to the blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster add \$hostname

iv. Run the following command to check whether the machine to be shut down is already included in the blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

3. Shut down the machine, perform maintenance, and then restart the machine.

⑦ Note Do not compromise the system during maintenance.

4. Run the following commands to remove the Job Scheduler blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster remove \$hostname /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

5. Set the status of rma to pending for the faulty machine.

i. Log on to the OPS1 server. Set the status of the rma action to pending for the faulty machine. The hostname of the faulty machine is m1.

Run the following command:

curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1" -d '{"action_name":"rma", "action_status":"pending"}'

The command output is as follows:

```
{
    "err_code": 0,
    "err_msg": "",
    "data": [
        {
            "hostname": "m1"
        }
    ]
}
```

ii. Run the following command to configure the audit log:

```
curl "http://127.0.0.1:7070/api/v5/AddAuditLog?object=/m/m1&category=action" -d
'{"category":"action", "from":"tianji.HealingService#", "object":"/m/m1", "content": "{\n
\"action\": \"/action/rma\",\n \"description\": \"/monitor/rma=error, mtime:
1513488046851649\",\n \"status\": \"pending\"\n}\n"}'
```

The mtime parameter, which represents action_description@mtime, is set to 1513488046851649 in the example. Set the parameter to the current system time when you configure the audit log. Run the following command to query the mtime value:

```
curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?
hostname=m1&attr=action_name,action_status,action_description@mtime"
```

The command output is as follows:

```
{
    "err_code": 0,
    "err_msg": "",
    "data": {
        "action_description": "",
        "action_description@mtime": 1516168642565661,
        "action_name": "rma",
        "action_name@mtime": 1516777552688111,
        "action_status": "pending",
        "action_status@mtime": 1516777552688111,
        "hostname": "m1",
        "hostname@mtime": 1516120875605211
    }
}
```

- 6. Wait for approval.
 - i. Wait until the status of the rma action becomes approved or doing on the machine. Check the action status.

Run the following command to obtain the machine information:

curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1"

Command output:

A large amount of information is returned. You can locate the following keyword: "action_status": "pending".

ii. Check the SR approval status on the machine. pending indicates that the SR is being approved. approved, doing, or done indicates that the SR has been approved. If no action was taken, the SR was not approved.

Run the following query command:

curl http://127.0.0.1:7070/api/v5/GetMachineInfoPackage? hostname=m1&attr=sr.id,sr.action_name,sr.action_status

Command output: A large amount of information is returned. You can also view items in the doing state on the webpage.

7. Shut down the machine when the status of rma becomes approved or doing. After the maintenance is completed, start the machine.

? Note If you need to clone the machine after the maintenance is completed, proceed with the next step. Otherwise, skip the next step.

- 8. Clone the machine.
 - i. After the maintenance is completed, run the following command to clone the machine on the OPS1 server:

curl "http://127.0.0.1:7070/api/v5/SetMachineAction? hostname=m1&action_name=rma&action_status=doing" -d '{"action_name":"clone", "action_status":"approved", "action_description":"", "force":true}'

The command output is as follows:

```
{
"err_code": 0,
"err_msg": "",
"data": [
{
"hostname": "m1"
}
]
}
```

- ii. Access the clone container. Run the following commands to check the clone status and confirm whether the clone operation takes effect.
 - a. Run the following command to query the clone container:

docker ps|grep clone

The command output is as follows:

18c1339340ab reg.docker.god7.cn/tianji/ops_service:1f147fec4883e082646715cb79c3710f7 b2ae9c6e6851fa9a9452b92b4b3366a ops.OpsClone__.clone.1514969139

b. Run the following command to log on to the container:

docker ps|grep clone

c. Run the following command to query the clone task:

/home/tops/bin/python /root/opsbuild/bin/opsbuild.py acli list --status=ALL -n 10000 | vim -

9. Run the following command to restore the machine status:

curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1&action_name=rma" -d '{"action_name":"rma","action_status":"done", "force":true}'

10. Check the machine status through the command or Apsara Infrastructure Management Framework. If the status is GOOD, the machine is normal.

Run the following command to check the machine status:

curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1&attr=state,hostname"



- 11. Check whether the cluster has reached the desired state. Ensure that all services on the machine being brought online have reached the desired state.
- 12. Run the following commands to remove the Job Scheduler blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster remove \$hostname

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

10.1.3.4. Shut down a chunkserver for maintenance

without compromising the system

Prerequisites

Check that all MaxCompute services have reached the final status and are functioning properly.

Procedure

1. In Apsara Infrastructure Management Framework, locate **ComputerInit#** in the odps-servicecomputer service of the odps cluster, and open the corresponding TerminalService window. Run the following commands to check the data integrity of Apsara Distributed File System:

puadmin abnchunk fs -t none

-- Check for any missing files. If no output is displayed, no files are missing.

puadmin abnchunk fs -t onecopy

-- Check whether each file has only one copy. If no output is displayed, each file has only one copy

puadmin abnchunk fs -t lessmin

-- Check whether the number of files is smaller than the minimum number of backups. If no output

is displayed, the number of files is smaller than the minimum number of backups.

- 2. Add the machine to be shut down to a Job Scheduler blacklist.
 - i. Run the following command to enable the blacklisting function of Job Scheduler (ignore this step if the function has been enabled):

/apsara/deploy/rpc_caller --Server=nuwa://localcluster/sys/fuxi/master/ForClient --Method =/fuxi/SetGlobalFlag --Parameter={\"fuxi_Enable_BadNodeManager\":false}

ii. Run the following command to check the hostnames in the existing blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

iii. Run the following command to add the machine to be shut down to the blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster add \$hostname

iv. Run the following command to check whether the machine to be shut down is already included in the blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

3. Shut down the machine for maintenance and then restart the machine.

? Note Do not compromise the system during maintenance.

4. Run the following commands to remove the Job Scheduler blacklist:

/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster remove \$hostname /apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get

Expected results

During the shutdown of Pangu_chunkserver, Apsara Distributed File System will keep trying to read data, and SQL tasks will remain in the running state. The tasks are completed after seven to eight minutes, or after the machine resumes operation.

10.1.3.5. Adjust the virtual resources of the Apsara system

in MaxCompute

Prerequisites

All MaxCompute services have reached the desired state and are functioning properly.

Procedure

- 1. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose **Operations > Cluster Operations**. In the **Cluster** search box, enter odps to search for the expected cluster.
- 2. Click the cluster in the search result. On the Cluster Details page, click the **Cluster Configuration** tab. In the left-side file list, find the role.conf file in the fuxi directory.

ile List 😧	E Cluster	Template	B	role.conf
Create File			1	MachineGroups:
T Create File			2	BigGraphInstance:
Alustar sanf			3	- a56d03127.cloud.d04.amtest73
Lange Cluster.conf			4	- a56d05125.cloud.d06.amtest73
kv.conf			5	- a56d07026.cloud.d07.amtest73
KV.COIII			6	GraphInstance:
machine group.c	onf		7	- a56d03007.cloud.d03.amtest73
			8	- a56d03008.cloud.d03.amtest73
blan.conf			9	- a56d05021.cloud.d05.amtest73
			10	- a56d05121.cloud.d06.amtest73
🗄 🗀 services			11	- a56d07022.cloud.d07.amtest73
			12	- a56d07107.cloud.d08.amtest73
🕀 🗀 alicpp			13 14	 a56d07108.cloud.d08.amtest73 a56d07122.cloud.d08.amtest73
🕀 🗀 apsaralib			14	- a56d0/122.cloud.d08.amtest/3 OdpsCommonInstance:
C C apourano			16	- a56d05007.cloud.d05.amtest73
🗄 🗀 apsarasecurity	V		17	- a56d05007.cloud.d05.amtest73
			18	OdpsSpecialInstance:
🕀 🗀 bigdata-sre			19	- a56d05007.cloud.d05.amtest73
			20	- a56d05008.cloud.d05.amtest73
🕀 🗀 disk-driver			21	RTInstance:
🖃 🗀 fuxi			22	- a56d03007.cloud.d03.amtest73
			23	- a56d03008.cloud.d03.amtest73
dependence	y.conf		24	- a56d05021.cloud.d05.amtest73
			25	- a56d05121.cloud.d06.amtest73
role.conf			26	- a56d07022.cloud.d07.amtest73
ton on t			27	- a56d07107.cloud.d08.amtest73
tag.conf			28	- a56d07108.cloud.d08.amtest73
template.co	onf		29	- a56d07122.cloud.d08.amtest73
emplate.co	2111		30	SInstance:
🕀 🗀 tianji			31	- a56d03007.cloud.d03.amtest73
			32	- a56d03008.cloud.d03.amtest73

3. Adjust the machine tags on the right and click Preview and Submit.

Adjust machine tags

Create File	1 MachineGroups: 2 BigGraphInstance:
	3 - a56d03127.cloud.d04.amtest73
E cluster.conf	4 - a56d05125.cloud.d06.amtest73
kv.conf	5 - a56d07026.cloud.d07.amtest73
RV.COM	6 GraphInstance:
hachine group.conf	7 - a56d03007.cloud.d03.amtest73
	8 - a56d03008.cloud.d03.amtest73
blan.conf	9 - a56d05021.cloud.d05.amtest73
	10 - a56d05121.cloud.d06.amtest73 11 - a56d07022.cloud.d07.amtest73
E services	12 - a56d07022.cloud.d07.amtest73
🕀 🗀 alicpp	13 - a56d07108.cloud.d08.amtest73
C C anopp	14 - a56d07122.cloud.d08.amtest73
🕀 🗀 apsaralib	15 OdpsCommonInstance:
	16 - a56d05007.cloud.d05.amtest73
🕀 🗀 apsarasecurity	17 - a56d05008.cloud.d05.amtest73
	18 OdpsSpecialInstance:
	19 - a56d05007.cloud.d05.amtest73
🕀 🗀 disk-driver	20 - a56d05008.cloud.d05.amtest73
	21 RTInstance:
😑 🗀 fuxi	22 - a56d03007.cloud.d03.amtest73
	23 - a56d03008.cloud.d03.amtest73
dependency.conf	24 - a56d05021.cloud.d05.amtest73 25 - a56d05121.cloud.d06.amtest73
Fole.conf	26 - a56d07022.cloud.d07.amtest73
TOIC.COM	27 - a56d07107.cloud.d08.amtest73
tag.conf	28 - a56d07108.cloud.d08.amtest73
	29 - a56d07122.cloud.d08.amtest73
template.conf	30 SInstance:
🕀 🗅 tianji	31 - a56d03007.cloud.d03.amtest73
ta 🖂 uanji	32 - a56d03008.cloud.d03.amtest73
⊞ ⊡ user	33 - a56d05021.cloud.d05.amtest73
	34 - a56d05121.cloud.d06.amtest73
version.conf	35 - a56d07022.cloud.d07.amtest73

4. In the **Confirm and Submit** dialog box that appears, enter the change description and click **Submit**.

nange Desc	ription:		
Difference	File 0: services/fuxi/role.conf		▼ Previous File Next F
🗈 servi	ces/fuxi/role.conf		
-	00 -1,74 +1,75 00		
1	MachineGroups:	1	MachineGroups:
2	BigGraphInstance:	2	BigGraphInstance:
3	- a56d03127.cloud.d04.amtest73	3	- a56d03127.cloud.d04.amtest73
4	- a56d05125.cloud.d06.amtest73	4	- a56d05125.cloud.d06.amtest73
5	- a56d07026.cloud.d07.amtest73	5	- a56d07026.cloud.d07.amtest73
6	GraphInstance:	6	GraphInstance:
7	- a56d03007.cloud.d03.amtest73	7	- a56d03007.cloud.d03.amtest73
8	- a56d03008.cloud.d03.amtest73	8	- a56d03008.cloud.d03.amtest73
9	- a56d05021.cloud.d05.amtest73	9	- a56d05021.cloud.d05.amtest73
10	- a56d05121.cloud.d06.amtest73	10	- a56d05121.cloud.d06.amtest73
11	- a56d07022.cloud.d07.amtest73	11	- a56d07022.cloud.d07.amtest73
12	- a56d07107.cloud.d08.amtest73	12	- a56d07107.cloud.d08.amtest73
13	- a56d07108.cloud.d08.amtest73	13	- a56d07108.cloud.d08.amtest73
14	- a56d07122.cloud.d08.amtest73	14	- a56d07122.cloud.d08.amtest73
15	OdpsCommonInstance:	15	OdpsCommonInstance:
16	- a56d05007.cloud.d05.amtest73	16	- a56d05007.cloud.d05.amtest73
17	- a56d05008.cloud.d05.amtest73	17	- a56d05008.cloud.d05.amtest73
		18	
18	OdpsSpecialInstance:	19	OdpsSpecialInstance:
19	- a56d05007.cloud.d05.amtest73	20	- a56d05007.cloud.d05.amtest73
20	- a56d05008.cloud.d05.amtest73	21	- a56d05008.cloud.d05.amtest73
21	RTInstance:	22	RTInstance:
22	- a56d03007.cloud.d03.amtest73	23	- a56d03007.cloud.d03.amtest73
23	- a56d03008.cloud.d03.amtest73	24	
24	- a56d05021.cloud.d05.amtest73	25	- a56d05021.cloud.d05.amtest73
25	- a56d05121.cloud.d06.amtest73	26	- a56d05121.cloud.d06.amtest73

5. The cluster starts rolling and the changes start to take effect.

⑦ Note You can check the task status in the operation log. If the changes take effect, the status becomes Successful.

6. After the changes are made, run the r ttrl command in the TerminalService window to

Submit

confirm the changes.

10.1.3.6. Restart MaxCompute services

Procedure

- 1. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, choose **Operations > Cluster Operations**. In the **Cluster** search box, enter odps to search for the expected cluster.
- 2. Click the cluster in the search result. On the Cluster Details page, click the Services tab. In the Service search box, search for odps-service-computer. Click odps-service-computer in the search result.
- 3. After you access the odps-service-computer service, select ComputerInit# on the Service Details page. In the Actions column corresponding to the machine, click Terminal. In the TerminalService window that appears, you can perform subsequent command line operations.
- 4. Run the following command to obtain the number of machines:

tj_show -r fuxi.Tubo#

5. Divide the number of machines by 3 to obtain the workernum value.

ONOTE The workernum value ranges from 1 to 3.

6. Modify workernum in *vim /apsara/odps_service/deploy/env.cfg*.

odps_worker_num = 2

executor_worker_num = 2

hiveserver_worker_num = 2

replication_server_num = 2

messager_partition_num = 2

-- The values here are used as an example. Set these values as needed.

7. Restart Hive and MaxCompute.

/apsara/odps_service/deploy/install_odps.sh restart_hiveservice

-- Restart Hive.

/apsara/odps_service/deploy/install_odps.sh restart_odpsservice

-- Restart MaxCompute.

r swl Odps/OdpsServicex

r swl Odps/HiveServerx

-- Check the service update status and time after restart.

8. Restart the messager service.

- cd /apsara/odps_service/deploy/; sh install_odps.sh pedeploymessagerservice
- -- Restart the messager service.
- r swl Odps/MessagerServicex
- -- Check the service update status and time after restart.
- 9. Restart the quota service.
 - cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployquotaservice
 - -- Restart the quota service.
 - r swl Odps/QuotaServicex
 - -- Check the service update status and time after restart.
- 10. Restart the replication service.
 - cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployreplicationservice
 - -- Restart the replication service.
 - r swl Odps/ReplicationServicex
 - -- Check the service update status and time after restart.
- 11. Restart the service mode.
 - r plan Odps/CGServiceControllerx >/home/admin/servicemode.json
 - r sstop Odps/CGServiceControllerx
 - r start /home/admin/servicemode.json
 - -- Restart the service mode.
 - r swl Odps/CGServiceControllerx
 - -- Check the CGServiceControllerx service update status and time after restart.

10.1.4. Common issues and solutions

10.1.4.1. View and allocate MaxCompute cluster resources

This topic describes how to view the storage and computing resources in a MaxCompute cluster. This topic also describes the quota group-related concepts, relationships between a quota group and a MaxCompute project, and quota group division policies.

Resources that can be allocated to projects in a MaxCompute cluster

• Storage resources: The total sum of storage resources available in a MaxCompute cluster is limited and can be calculated based on the number of compute nodes in the entire cluster. The storage capacity in a MaxCompute cluster is managed through Apsara Distributed File System. You can run Apsara Distributed File System commands to view the total storage capacity, such as the current storage usage statistics. The following metrics are available for

measuring storage resources:

Storage capacity metric: indicates the total size of files that can be stored in a cluster. You can calculate the total file size in a cluster based on the following formula: Total file size in a cluster = Number of machines * (Size of a single disk * (Number of disks on a single machine - 1)) * System security level * System compression ratio/Number of distributed replicas.

? Note

- Based on the standard TPC-H test data set, the ratio of the original data size to the compressed data size is 3:1. The ratio varies depending on the characteristics of business data.
- Typically, three replicas are stored in a distributed manner.
- Security level: The default value is 0.85 in the MaxCompute system. You can set a custom security level as required. For example, when the business data increases rapidly and reaches 85% of the total storage quota, the security level is low. You must scale out the system as required or delete unnecessary data.

How to view the storage capacity of a MaxCompute cluster

Run the puadmin lscs command on the cluster AG. The total disk size, total free disk size, and total file size are displayed at the end of the command output.

Capacity information

The pa	angu disk status:
Total	Disk Size:681225 GB
Total	Free Disk Size:635921 GB
Total	File Size:997 GB
Total	UnReserved Disk Space4Piops:0 GB
Total	Disk Space4Piops:0 GB
Total	UnReserved Disk Iops4Piops:0
Total	Disk Iops4Piops:0

⑦ Note Parameters:

- Total Disk Size: the total amount of physical space. Each file is stored in three copies. The logical space is one third the size of the physical space.
- Total Free Disk Size: the total size of available disks, excluding recycle bins on chunkservers.
- Total File Size: the total amount of physical space used by Apsara Distributed File System files, including the /deleted/ directory.

Run the following command on the cluster AG to view the storage capacity used by all projects:

```
pu ls -l pangu://localcluster/product/aliyun/odps/
```

Example:

pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -A 4

-- View the capacity used by a single project, such as adsmr.

Project capacity information

\$pu ls -l pa	ngu://localcluster/product/aliyun/odps/ grep_adsmr_A_4
pangu://loca	lcluster/product/aliyun/odps/adsmr/
Length	: 551267930
FileNumber	: 570
DirNumber	: 143
Pinned	: 0

⑦ Note Parameters:

- Length: the logical length used by a project. The physical length required is three times the logical length.
- FileNumber: the number of files used.
- DirNumber: the number of directories used.
- File size metric: The total size of files that can be stored in a cluster is limited based on the memory capacity of PanguMaster. The existence of a large number of small files or an improper number of files in a cluster can also affect the stability of the cluster and its services.

The Apsara Distributed File System index files, including the information of Apsara Distributed File System files and directories, are stored in the PanguMaster memory. Each file in PanguMaster corresponds to a file node. Each file node uses XXX bytes of memory, each level of directory uses XXX bytes of memory, and each chunk uses XXX bytes of memory. A large file is split into multiple chunks in Apsara Distributed File System. Therefore, the factors that affect PanguMaster memory usage include the number of files, directory hierarchy, and number of chunks.

If the size of the original files in Apsara Distributed File System is large, the memory usage of PanguMaster is relatively low. When a large number of small files exist, the memory usage of PanguMaster is relatively high.

We recommend that you perform the following operations to reduce the memory usage of PanguMaster:

- Reduce or even delete empty directories which occupy memory, and reduce the number of directory levels.
- Do not create directories. A directory is created automatically when you create a file.
- Store multiple files in a directory. However, a maximum of 100,000 files can be stored.
- Decrease the length of file names and directory names to reduce the memory usage and network traffic in PanguMaster.

 Reduce the number of small tables and files. We recommend that you use Tunnel to upload and commit MaxCompute tables only when the table data size reaches 64 MB.

The following figure shows the numbers of files that can be stored in Apsara Distributed File System for different PanguMaster memory capacities.

Numbers of files that can be stored for different PanguMaster memory capacities

48G memory	Upper limit of total number of files : 87.5 million
96G memory	Upper limit of total number of files:175 million
128G memory	Upper limit of total number of files : 233 million

How to view the number of files stored in a MaxCompute cluster

Run the pu quota command on the cluster AG to view the total number of files stored in a MaxCompute cluster.

Total number of files

\$pu quota
quota under pangu://localcluster/
EntryNumber Limit:unlimited
Used:16632877
Used(excluding hardlink):16632712
FileNumber Limit:unlimited
Used:8594596
Used(excluding hardlink):8594431
FilePhysicalLength Limit:unlimited
Used:1415115960895
Used(excluding hardlink):1414395196936
FileLogicalLength Limit:unlimited
Used:467814050981
Used(excluding hardlink):467573796328

This example uses the adsmr project to demonstrate how to view the number of files. Run the following command on the cluster AG to view the number of files for a single project in a MaxCompute cluster:

Number o	of files for a single project
· ·	angu://localcluster/product/aliyun/odps/ grep adsmr -A 4 alcluster/product/aliyun/odps/adsmr/ : 551267930
FileNumber DirNumber Pinned	

- FileNumber: the number of files used.
- DirNumber: the number of directories used.
- FileNumber + DirNumber = Number of files for the current project.
- Computing resources: CPU and memory are typically referred to as computing resources in a MaxCompute cluster. The total amount of computing resources is calculated based on the following formula: Total amount of computing resources = (Number of CPU cores + Memory size of each machine) * Number of machines. For example, each machine has 56 CPU cores. One core on each machine is used by the system. The remaining 55 cores are managed by the distributed scheduling system and are scheduled for use by the MaxCompute service. The memory (aside from the chunk of memory for system overhead) is allocated by Job Scheduler. Typically, 4 GB of memory is allocated per CPU core in each MaxCompute task. The ratio varies depending on MaxCompute tasks.

How to view computing resources

• Run the r ttrl command on the cluster AG to view all computing resources.

All computing resources

\$r ttrl total tubo in cluster=13						
detail table for every machir	ne:					
Machine Name		CPU		Memory		Other
.cloudamtest128	34 1	6,300		170,453		GraphInstance:8 RTInstance:4 SInstance:99
.cloudamtest128	34	6,300		234,014		BigGraphInstance:99
.cloudamtest128	34	6,300		170,453		GraphInstance:8 RTInstance:4 SInstance:99
.cloudamtest128	34	6,300		170,453		ElasticSearchInstance:5
.cloudamtest128	34	6,300		234,014		BigGraphInstance:99
.cloudamtest128	34	6,300		170,453		
.cloudamtest128	34	6,300		170,453		GraphInstance:8 RTInstance:4 SInstance:99
.cloudamtest128	34	6,300		170,453		OdpsSpecialInstance:20 OdpsCommonInstance:20
.cloudamtest128	34	6,300		170,453		ElasticSearchInstance:5
.cloudamtest128	34	6,300		170,453		ElasticSearchInstance:5
.cloudamtest128	34	6,300		234,014		BigGraphInstance:99
.cloudamtest128	34	6,300		170,453		OdpsSpecialInstance:20 OdpsCommonInstance:20
.cloudamtest128	34	6,300		170,453		GraphInstance:8 RTInstance:4 SInstance:99
Total	1	81,900	I	2,406,572	1	NA

(?) Note In the command output, the domain name, total CPU capacity (Unit: U. 100 U = 1 core), and total memory (Unit: MB) of each Tubo machine, as well as the role of each Tubo machine in Job Scheduling System are listed in four columns.

• Run the r tfrl command on the cluster AG to view the remaining computing resources.

Remaining computing resources

Machine Name	for ev	very machine:	-	Manager	Other
			CPU	Memory	Other
	oud.	.amtest1284	5,025	150,990	GraphInstance:8 RTInstance:4 SInstance:81
	oud.	.amtest1284	6,090	226,874	BigGraphInstance:98
.cle	oud.	.amtest1284	5,285	153,634	GraphInstance:8 RTInstance:4 SInstance:83
.cl	oud.	.amtest1284	6,100	68,521	ElasticSearchInstance:3
.cle	oud.	.amtest1284	6,190	227,850	BigGraphInstance:98
.cl	ud.	.amtest1284	6,200	169,453	
.clo	oud.	.amtest1284	5,035	150,450	GraphInstance:8 RTInstance:4 SInstance:83
.clo	oud.	.amtest1284	4,600	131,565	OdpsSpecialInstance:15 OdpsCommonInstance:12
.cl	oud.	.amtest1284	6,200	104,921	ElasticSearchInstance:4
.clo	oud.	.amtest1284	6,000	67,521	ElasticSearchInstance:3
.cl	oud.	.amtest1284	5,790	218,634	BigGraphInstance:97
.cl	oud.	.amtest1284	5,400	133,089	OdpsSpecialInstance:20 OdpsCommonInstance:13
.cle	ud.	.amtest1284	5,485	157,634	GraphInstance:8 RTInstance:4 SInstance:87
Total			73,400	1,961,136	NA

Note In the command output, the domain name, total CPU capacity (Unit: U. 100 U = 1 core), and total memory (Unit: MB) of each Tubo machine, as well as the role of each Tubo machine in Job Scheduling System are listed in four columns.

• Run the r cru command on the cluster AG to view the resources used by all running jobs in MaxCompute.

pr cru WorkItemName		CPU		Memory		VirturlResource
Ddps/DiskDriverService		280		13,600		0
Ddps/odps_elasticsearch_elasticsearch_mdu_es_demo_20170509064623398g2q8q9d		200		1,024		0
Ddps/CGServiceControllerx		1,980		66,660		{'SInstance': 60}
Ddps/ReplicationServicex		200		2,000		{'Odps5pecialInstance': 1}
Ddps/OdpsServicex		1,400		45,128		{'OdpsSpecialInstance': 4, 'OdpsCommonIns
ance': 7} Jdps/HiveServerx		850		37,864		{'OdpsCommonInstance': 4}
Ddps/XStreamServicex		14,070		146,370		0
Ddps/QuotaServicex		100		1,024		{'OdpsSpecialInstance': 1}
Ddps/MessagerServicex		300		3,092		0
m/sm used resource		1,000		11,192		0
otal Planned Resource 5, 'OdpsCommonInstance': 11}	I	20,380	I	327,954	I	{'SInstance': 60, 'OdpsSpecialInstance':

Resources used by all running jobs

(?) Note The name, total CPU capacity, total memory of each job, as well as the number of Fuxi instances started in the role of each job in Job Scheduling System are listed in four columns.

How to allocate project resources in a MaxCompute cluster

• Storage resource allocation: Based on the characteristics of a project, the space size and file size limit are configured when you create the project.

If the following error messages are displayed, the file size limit of the project has been exceeded. In this case, you must organize the data in the project by deleting unnecessary table data or increasing the storage resource quota.

Error messages

018-03-16 18:24:46 1:0:383:log.txt 3% 15 bytes 0 bytes/s
ava.util.concurrent.ExecutionException: java.io.IOException: RequestId=2018031618244658a751640003a1fa, ErrorCode=InternalServerError, ErrorMes
quota not enough.
at java.util.concurrent.FutureTask\$Sync.innerGet(FutureTask.java:222)
at java.util.concurrent.FutureTask.get(FutureTask.java:83)
at com.aliyun.odps.ship.upload.DshipUpload.uploadBlock(DshipUpload.java:152)
at com.aliyun.odps.ship.upload.DshipUpload.upload(DshipUpload.java:101)
at com.aliyun.odps.ship.DShip.runSubCommand(DShip.java:73)
at com.aliyun.odps.ship.DShipCommand.run(DShipCommand.java:99)
at com.aliyun.openservices.odps.console.commands.InteractiveCommand.run(InteractiveCommand.java:225)
at com.aliyun.openservices.odps.console.commands.CompositeCommand.run(CompositeCommand.java:50)
at com.aliyun.openservices.odps.console.ODPSConsole.main(ODPSConsole.java:62)
aused by: java.io.IOException: RequestId=2018031618244658a751640003a1fa, ErrorCode=InternalServerError, ErrorMessage=Storage quota not enough.
at com.aliyun.odps.tunnel.io.TunnelRecordWriter.close(TunnelRecordWriter.java:72)
at com.aliyun.odps.ship.upload.BlockUploader.doUpload/BlockUploader.java:166)
at com.aliyun.odps.ship.upload.BlockUploader.upload(BlockUploader.java:95)
at com.aliyun.odps.ship.upload.DshipUpload\$1.call(DshipUpload.java:139)
at com.aliyun.odps.ship.upload.DshipUpload\$1.call(DshipUpload.java:136)
at java.util.concurrent.FutureTask\$Sync.innerRun(FutureTask.java:303)
at java.util.concurrent.FutureTask.run(FutureTask.java:138)
at java.util.concurrent.ThreadPoolExecutor\$Worker.runTask(ThreadPoolExecutor.java:886)
at java.util.concurrent.ThreadPoolExecutor\$Worker.run(ThreadPoolExecutor.java:908)
at java.lang.Thread.run(Thread.java:662)
aused by: RequestId=2018031618244658a751640003a1fa, ErrorCode=InternalServerError, ErrorMessage=Storage quota not enough.
at com.aliyun.odps.tunnel.io.TunnelRecordWriter.close(TunnelRecordWriter.java:70)
9 more RROR: TunnelException - ErrorCodewLocal Error, ErrorMessagewBlock ID:0 Failed.
NACK: TURNELEXCEPTION - ErrorCode=Local Error, ErrorNessage=Block 1010 Falled,

Notice The sum of the storage capacity of all projects cannot exceed the total allowable storage capacity of a service. Similarly, the total file size of all projects cannot exceed the total allowable file size. Therefore, you must properly allocate the storage space and file size limit by project and make timely adjustment based on your business requirements.

- Computing resource allocation: division of quota groups.
 - What is a quota group?

A MaxCompute cluster allows you to divide computing resources into different quota groups, and schedule them as required. A quota group represents a certain amount of CPU and memory resources. MinQuota and MaxQuota are used for CPU and memory configurations. MinQuota is the minimum quota allowed for the quota group, and MaxQuota is the maximum quota allowed for the quota group. For example, MinCPU=500 indicates that the quota group has been assigned at least 500/100=5 cores. MaxCPU=2000 indicates that the quota group has been assigned at least 2000/100=20 cores.

MaxCompute uses a FAIR scheduling policy and a first-in-first-out (FIFO) scheduling policy by default. The difference between the FAIR and FIFO scheduling polices lies in the keys by which tasks in waiting queues are sorted. If each schedule unit has its own priority, both FAIR and FIFO scheduling policies allocate high-priority schedule units first. If all schedule units share the same priority, the FIFO scheduling policy sorts the schedule units by the time when they are submitted. The earlier they are submitted, the higher priority they have. The FAIR scheduling policy sorts the scheduling units by the slotNum allocated to them. The smaller the slotNum is, the higher priority they have. For the FAIR policy group, this can basically ensure that the same amount of resources are assigned to schedule units with the same priority.

You can run the r quota command on the cluster AG to view quota group settings.

\$r quota									
Account Alias	SchedulerType	Strategy	Init0	uota	ScaledQuota	[ScaleRatio	Runtime	UsageInf	
				[CPU:31500					CPU:480
			Stati	c	CPU:31500	CPU:37800	CPU:1008	Used	
				Mem:852265					Mem: 9040
9242 odps_quota	Fair	NoPreenp	t						
				CPU:100					CPU:489
			[Min		Mem:852265	Men: 1022718	Mem:21408	Availab]	e
				Mem:1024					Mem:10280

View quota group settings

You can run the following command on the cluster AG to create and modify a quota as needed:

sh /apsara/deploy/rpc_wrapper/rpc.sh setquota -i \$QUOTAID -a \$QUOTANAME -t fair -s \$max_cp u_quota \$max_mem_quota -m \$min_cpu_quota \$min_mem_quota

Note The command with \$QUOTAID is used to modify a quota. The command without \$QUOTAID is used to create a quota.

Create a quota

<pre>0 /home/tops/b) quotatest connecting td connected Method=SetAcd Parameter=[{' }, "returnRe' : 852265},"d in", "quota": {"d ir", "quota": {"d ir", "quota": 18900 turnResource' roups": false pt", "accound </pre>	in/python se o nuwa://loc. countQuota "scaleRatio" sourceType": canFreemptOtt y": "NoFreem CFU": 100, " : {"CFU": 18 alse, "alias alse, "alias alse, "alias , "Memory": ", "schedule e, "canBeFree LId": 9249}, }, "quota":	t_quota alclust "Retur herGrou pot", "a Memory" 900, "M ": "es_ 702042} rType": emptedB {"alia	_gro er/s nRes ps": ccou : 10 emor quot , "m "Fa yoth s":	up.py 9251 ys/fuxi/ma 7800, "Mem ource", "s false, "c ntId": 924 24), "retu y": 511359 a", "strat inQuota": ir", "guot erGroups": "guotatest	<pre>a -i 9251 -a (quotatest 50) ster/ForClien(chedulerType" anBePreempted) 2}, {"scaleRat rnResourceType" }, "canPreemp egy": "NoPreet ("CPU": 100, ' a": {"CFU": 11 false, "alias" ", "scheduler" ": 50000}, "ad</pre>	20 50000 500 t : "Fair", "qu ByOtherGroups : "ReturnRe CotherGroups" "Memory": 102 3900, "Memory s": "biggrap s": "biggrap	5000 fair - : {"CPU": 10 ota": {"CPU" ": false, "a 18900, "Mem source", "sc : false, "ca tId": 9243], 4], "returnR ": 702042], _quota", "st , "minQuota"	1 -1 0, "Memor : 31500, lias": "o ory": 511 hedulerTyy mBePreemp {"scaleR esourceTy "canPreem rategy":	y": 1024 "Memory" dps_quot 359], "m pe": "Fa tedByOth atio": { pe": "Re tedByOtherG "NoPreem
Account Alies	SchedulerType	Strategy	(Init)	juota	ScaledQuota	(ScaleRatio	Runtime	(UsageIn	fo
				ICP0:5000					ICPU:0
1			IStati		1CP0:5000	ICPT: 5000	ICPU:0	IUaed	1
1				(Hen: 50000					(Mem: 0
19251 Iquotatest	IFeir	INoPreemp							
1				[CP0:500					ICPU:0
1			IMin		Mem: 50000	IMem: 50000	IMem: 0	(Availab)	le!
				[Mem: 5000					[Mem:0

Modify a quota

Fair, reempt e, sc guota, ceType falae, 1024), emptedB fraceId	"quota": {"Cf ; "accountId": chedulerType": "strategy": " "alrategy": " "seturnResour byOtherGroups": =0 xgLevel-ALL	<pre>io": {"CFU": 5000, U": 2000, "Memory" 9251}, {"scaleRat "Fair", "quota": { NoPreempt", "accou trae", "achedulerTy quota", "atrategy" ceType": "Returnke</pre>	: 20000), io": ("CPU "CPU": 315 stid": 924 pe": "Fair : "NoPreem source",	"canPre ": 3780 00, "Me 2), ("s ", "quo pt", "a schedul	emptOtherGroups": 0, "Memory": 1022 mory": 052265), " caleRatio": {"CPU ta": {"CPU": 1890 ccountId": 9243), erType": "Fair",	: false, "canBePreem 718), "mimQuota": ["canPreemptOtherGrou J": 18900, "Memory": 00, "Memory": 511359 0, ["acaleRatio": {"C	2000), "returnResou ptedByOtherGroups": "CUT: 100, "Memory" 511359), "minQuota" 511359), "minQuota" 900, "Memory": 70204 : 9243)]	false, "alias": "g ': 1024], "returnRe cemptedByOtherGrou ': {"CPU": 100, "Me Snoups": false, "ca ': 702042}, "minDuo	<pre>puotatest", "at: sourceType": "ps": false, "al mory": 1024}, nBePreemptedBy ta": {"CFO": 10</pre>	rategy": "Noi ReturnResourd Liss": "odps "returnResour OtherGroups": D0, "Memory":
(A charges and	tlAlias	SchedulerType	IStrategy	Inito	uota	IScaledQuota	IScaleRatic	IRuntime	(VsageIn:	ťo
l		in an								
1		1			[CP0:2000					ICPU: 0
1		1		 Stati	(CP0:2000	 CPU: 2000	I ICF0:5000	I ICPU:0	। ।एउक्रत	CPU:0
		 		 Stati		। । СРО: 2000 ।	 CP0:5000 	1 1090±0 1	l IVaed I	CPU:0
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	 quotatest	l l l lFair	 NoPreempt		0		 CP0:5000 			1
			 NoPreempt 		0					1
			 NoPreempt 		Cl					

• How to divide quota groups

To divide quota groups correctly, you must understand the relationship between a MaxCompute project and a quota group.

You can select the quota group to which a project belongs upon project creation or modify the quota group after project creation.

Resources in a quota group can be used by all running tasks of all projects in this quota group. Therefore, the project tasks in the same quota group may be affected during peak hours. That is, one or several large tasks may take up all resources in the quota group, while other computing tasks can only wait for resources.

For example, in the following two figures, the first figure shows that a lot of jobs are waiting for resources (in red box). However, a lot of cluster resources are left unused. You can check the quota usage. In the second figure, quota 9243 is only allocated with 5000U, all of which are in use. The CPU quota for 9243 is used up, but there are still pending tasks in 9243. In this case, even if there are unused cluster resources, the tasks under this quota cannot have resources allocated to them.

drin&dccker192168000187 [home_/admin] crujsdd y, yg sort yz yz pjodpszawi sort yz yz yz pjodpszawi sort yz yz yz yz sort yz sort yz yz yz sort sort sort yz yz yz sort sort sort sort yz yz sort sort sort sort yz sort sort sort yz sort <	103400 {'SInstance': 50}	Contraction of the second second
--	----------------------------	----------------------------------

Jobs waiting for resources

Quota used up

quota							Incolematio	laura fine .	JusageInfo	
		I schedu l er Type		matte	CPU:42000	cru:12370	CPU:42000	CPU:0		CPUID
			NoPreempt		Hem:1293336 CPU:100		Men:1293336	Hem:0	Available	CPU:0 Hem:0
			Ī	static	CPU: 5000	CPU:1561	CPU: 5000	сри:5000	used	CPU: 5000 Hem:103400
1243	kalfa	Fair	horreempt	#in	CPU:100 Hem:100	Hem:164506	wem:620886	Hen: 620686	Available	CPU:0 Mem: 517486
				static	CPU:42000 Mem:1293336	CPU:12370	CPU:42000	CPU:100	used	CPU:109 Mem:2068
244	phq	Fair	air NoPreempt	sto	CPU:100 Mem:100	Hem: 342565	Hem:1293336	Hem:2068	Available	CPU:0 Mem:0
				static	CPU:42000	CPU:12370	CPU:42000	CPU:0	used	CPU:0 Mem:0
9245	The	rair	Fair NoPrempt	min	CPU:100	Hem: 342565	Mem:1293336	Hem:0	Available	CPU:0 Hem:0

You must divide quota groups based on the following general principles:

- You must plan quota groups in a way that they do not mutually interfere with each other in a large resource pool, and avoid overly fine-grained division of resource groups. For example, some large tasks cannot be scheduled due to quota group limits, or occupy a quota group for an extended period of time, which affects other tasks in the group.
- You must consider the configured MinQuota and MaxQuota when dividing quota groups.
- You can oversell the resources in your cluster, that is, the sum of MaxQuotas of all quota groups can be greater than the total amount of cluster resources. However, the oversell ratio cannot be too high. If the oversell ratio is too high, a quota group with a running project may perpetually occupy a large amount of resources.
- When dividing quota groups, you must consider the priorities of tasks, task execution duration, amount of task data, and characteristics of computing types.
- Properly configure quota groups for peak hours. We recommend that you configure a separate quota group for tasks that are important and time-consuming.
- The division of quota groups and the selection and configuration of projects are conducted based on a resource pre-allocation policy, which needs to be adjusted in a timely manner, based on actual requirements.

10.1.4.2. Common issues and data skew troubleshooting

Scenario 1: how to determine whether a job has stopped running due to insufficient resources

Symptom: The job does not progress as expected.

```
Symptom
```

2016-01-29	13:52:09	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:14	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:19	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:24	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:29	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:34	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:39	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:44	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:49	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:54	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:52:59	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:04	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:09	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:15	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:20	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:25	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:30	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:35	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:40	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:45	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:50	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]
2016-01-29	13:53:55	M1_Stg1_job0:0/0/5[0%]	R2_1_Stg1_job0:0/0/1[0%]

Cause: The issue is typically caused by insufficient resources. You can use LogView to determine the status of job resources (task instance status).

- Ready: indicates that instances are waiting for Job Scheduler to allocate resources. Instances can resume operation after they obtain the necessary resources.
- Wait: indicates that instances are waiting for dependent tasks to complete.

The task instances in the Ready state shown in the following figure indicate that there are insufficient resources to run these tasks. After an instance obtains the necessary resources, its status changes to Running.

M1_	Stg1 🖲				
Faile	d(0) Ready(5)	All(5) Long-Tails	(0) 🚹 La	atency char	t
	FuxiInstanceID	IP & Path	StdOut	StdErr	Status
1	Odps/odps_s			<u> </u>	Ready
2	Odps/odps_s				Ready
3	Odps/odps_s				Ready
4	Odps/odps_s				Ready
5	Odps/odps_s				Ready

Solution:

- If there are insufficient resources during peak hours, you can reschedule the tasks to run during off-peak hours.
- If the computing quotas are insufficient, check whether the quota group of the project has sufficient computing resources.
- If computing resources in the cluster are occupied for long periods of time, you can develop a computing quota allocation policy to scale the quota as necessary.
- We recommend that you do not run abnormally large jobs to prevent the jobs from occupying resources for extended periods of time.
- You can enable SQL acceleration, so that you can run small jobs without requesting resources from Job Scheduler.
- You can use the First-In First-Out (FIFO) scheduling policy.

Scenario 2: how to find the root cause of a job that has been running for an extended period of time

Symptom: The MaxCompute job execution progress has remained at 99% for a long period of time.

Cause: The running time of some Fuxi instances in the MaxCompute job is significantly longer than that of other Fuxi instances.

Cause analysis

rafiash										
uni Jobs Summar	ry JSONSumma	iry .								
uxi Job Name: cdo_l	boss_201507051	90608445go	uty7sb1	SQL 0 0 0	Odoj					
Taskhlame	Fatal/InstCount	1/O Records	Progress	Status	StartTime	EndTime	Latenció	0 TimeLine		26
1 M8_5tp3	0 /1	1604/1604	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:06:5	8	12		C
2 M2_Stg1	0 /19	28565726	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:07:5	4	1.8		C 2
3 M1_5tg1	0 /2	607317/6	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:07:2	7	41		
4 M5_Stg2	0 /1	143659/1	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:06:5	3	7		E 13
5 33 1 2 Stg1	0 /21	29173043	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:08:4	0 1	54		
6 16_3_5_Stp2	0 /11	750976/6	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:09:1	5 2	29		0
7 39_6_8_Stp3	0 /11	608921/6	100%	Terminated	2015-07-06 03:06:4	46 2015-07-06 03:10:1	9 3	23		20
						1.00				
G_1_2_Stg1 * M										
iled(0) Terminated	(11) Al(11) La	ng-Tais(1)	Latency d	hart					Latency: {"min":"2', "avg	167,7780
FuxInstanceID	IP & Path	StoOut Sto	dEn Sta	Cus	StartTime -	EndTime	atency(s)	TimeLine		
1 Odps/cdo_bo	1. C.	AT AT	Te	minated 201	5-07-06 03:08:47	2015-07-06 03:09:15	28			

Further analysis: Analyze the job summary in LogView, and calculate the difference between the max and avg values of input and output records of a slow task. If the max and avg values differ by several orders of magnitude, it can be initially determined that the job data is skewed.

```
Further analysis
```

Solution: If there are slow Fuxi instances on a particular machine, check whether a hardware failure has occurred on the machine.

Scenario 3: How to improve the concurrency of MaxCompute jobs

Fault locating: The concurrency of Map tasks depends on the following factors:

• Split size and merge limit.

Map takes a series of data files as inputs. Larger files are split into partitions based on the odps.sql.mapper.split.size value, which is 256 MB by default. An instance is started for each partition. However, starting an instance requires resources and time. Small files can be merged into a single partition based on the odps.sql.mapper.merge.limit.size value and be processed by a single instance to improve instance utilization. The default value of odps.sql.mapper.merge.limit.size is 64 MB. The total size of small files merged cannot exceed this value.

• Instances cannot process data across multiple partitions.

A partition is mapped to a folder in Apsara Distributed File System. You must run at least one instance to process data in a partition. Instances cannot process data across multiple partitions. In a partition, you must run instances based on the preceding rule.

Typically, the number of instances for Reduce tasks is 1/4 of that for Map tasks. The number of instances for Join tasks is the same as that for Map tasks, but cannot exceed 1,111.

You can use the following methods to increase the number of concurrent instances for Reduce and Join tasks:

set odps.sql.reducer.instances = xxx

set odps.sql.joiner.instances = xxx

Scenarios that require higher concurrency:

• A single record only contains a small amount of data.

Because a single record contains a small amount of data, there are many records in a file of the same size. If you split data into 256 MB chunks, a single Map instance needs to process a large number of records, reducing concurrency.

• Dump operations occur in the Map, Reduce, and Join stages.

Based on the preceding job summary analysis, the displayed dump information indicates that the instance does not have sufficient memory to sort data in the Shuffle stage. Improving concurrency can reduce the amount of data processed by a single instance to the amount of data that can be handled by the memory, eliminate disk I/O time consumption, and improve the processing speed.

• Time-consuming UDFs are used.

The execution of UDFs is time-consuming. If you execute UDFs concurrently, you can reduce the UDF execution time of an instance.

Solution:

• You can decrease the following parameter values to improve the concurrency of Map tasks:

```
odps.sql.mapper.split.size = xxx
odps.sql.mapper.merge.limit.size = xxx
```

• You can increase the following parameter values to improve the concurrency of Reduce and Join tasks:

```
odps.sql.reducer.instances = xxx
odps.sql.joiner.instances = xxx
```

Note: Improving concurrency will result in a greater amount of resources being consumed. We recommend that you take cost into account when improving concurrency. An instance takes an average of 10 minutes to complete after optimization, improving overall resource utilization. We recommend that you optimize jobs in critical paths so that they consume less time.

Scenario 4: how to resolve data skew issues

Different types of data skew issues in SQL are resolved in different ways.

• GROUP BY data skew

The uneven distribution of GROUP BY keys results in data skew on reducers. You can set the anti-skew parameter before executing SQL tasks.

```
set odps.sql.groupby.skewindata=true
```

After this parameter is set to true, the system automatically adds a random number to each key when running the Shuffle hash algorithm and prevents data skew by introducing a new task.

DISTRIBUTE BY data skew

Using constants to execute the DISTRIBUTE BY clause for full sorting of the entire table will result in data skew on reducers. We recommend that you do not perform this operation.

• Data skew in the Join stage

Data is skewed in the Join stage when the Join keys are unevenly distributed. For example, a key exists in multiple joined tables, resulting in a Cartesian explosion of data in the Join instance. You can use one of the following solutions to resolve data skew in the Join stage:

- When a large table and a small table are joined, use MapJoin instead of Join to optimize query performance.
- Use a separate logic to handle a skewed key. For example, when a large number of null values exist in the key, you can filter out the null values or execute a CASE WHEN statement to replace them with random values before the Join operation.
- If you do not want to modify SQL statements, configure the following parameters to allow MaxCompute to perform automatic optimization:

```
set odps.sql.skewinfo=tab1:(col1,col2)[(v1,v2),(v3,v4),...]
set odps.sql.skewjoin=true;
```

• Data skew caused by multi-distinct

Multi-distinct syntax aggravates GROUP BY data skew. You can use the GROUP BY clause with the COUNT function instead of multi-distinct to alleviate the data skew issue.

• UDF OOM

Some jobs report an OOM error during runtime. The error message is as follows: FAILED: ODPS-0123144: Fuxi job failed - WorkerRestart errCode:9,errMsg:SigKill(OOM), usually caused by OOM(out of memory) . You can fix the error by configuring the UDF runtime parameters. Example:

odps.sql.mapper.memory=3072;

set odps.sql.udf.jvm.memory=2048;

set odps.sql.udf.python.memory=1536;

The related data skew settings are as follows:

set odps.sql.groupby.skewindata=true/false

Description: allows you to enable GROUP BY optimization.

set odps.sql.skewjoin=true/false

Description: allows you to enable Join optimization. It is effective only when odps.sql.skewinfo is set.

set odps.sql.skewinfo

Description: allows you to set detailed information for Join optimization. The command syntax is as follows:

set odps.sql.skewinfo=skewed_src:(skewed_key)[("skewed_value")]

src a join src_skewjoin1 b on a.key = b.key;

Example:

set odps.sql.skewinfo=src_skewjoin1:(key)[("0")]

-- The output result for a single skewed value of a single field is as follows: explain select a.key c1, a.

value c2, b.key c3, b.value c4 from src a join src_skewjoin1 b on a.key = b.key;

set odps.sql.skewinfo=src_skewjoin1:(key)[("0")("1")]

-- The output result for multiple skewed values of a single field is as follows: explain select a.key c1, a .value c2, b.key c3, b.value c4 from src a join src_skewjoin1 b on a.key = b.key;

Scenario 5: how to configure common SQL parameters

Map settings

set odps.sql.mapper.cpu=100

Description: allows you to set the number of CPUs used by each instance in a Map task. Default value: 100. Valid values: 50 to 800.

set odps.sql.mapper.memory=1024

Description: allows you to set the memory size of each instance in a Map task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

set odps.sql.mapper.merge.limit.size=64

Description: allows you to set the maximum size of control files to be merged. Unit: MB. Default value: 64. You can set this variable to control the inputs of mappers. Valid values: 0 to Integer.MAX_VALUE.

set odps.sql.mapper.split.size=256

Description: allows you to set the maximum data input volume for a Map task. Unit: MB. Default value: 256. You can set this variable to control the inputs of mappers. Valid values: 1 to Integer.MAX_VALUE.

Join settings

set odps.sql.joiner.instances=-1

Description: allows you to set the number of instances in a Join task. Default value: -1. Valid values: 0 to 2000.

set odps.sql.joiner.cpu=100

Description: allows you to set the number of CPUs used by each instance in a Join task. Default

value: 100. Valid values: 50 to 800.

set odps.sql.joiner.memory=1024

Description: allows you to set the memory size of each instance in a Join task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

Reduce settings

set odps.sql.reducer.instances=-1

Description: allows you to set the number of instances in a Reduce task. Default value: -1. Valid values: 0 to 2000.

set odps.sql.reducer.cpu=100

Description: allows you to set the number of CPUs used by each instance in a Reduce task. Default value: 100. Valid values: 50 to 800.

set odps.sql.reducer.memory=1024

Description: allows you to set the memory size of each instance in a Reduce task. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

UDF settings

set odps.sql.udf.jvm.memory=1024

Description: allows you to set the maximum memory size used by the UDF JVM heap. Unit: MB. Default value: 1024. Valid values: 256 to 12288.

set odps.sql.udf.timeout=600

Description: allows you to set the timeout period of a UDF. Unit: seconds. Default value: 600. Valid values: 0 to 3600.

set odps.sql.udf.python.memory=256

Description: allows you to set the maximum memory size used by the UDF Python API. Unit: MB. Default value: 256. Valid values: 64 to 3072.

set odps.sql.udf.optimize.reuse=true/false

Description: When this parameter is set to true, each UDF function expression can only be calculated once, improving performance. Default value: true.

set odps.sql.udf.strict.mode=false/true

Description: allows you to control whether functions return NULL or an error if dirty data is found. If the parameter is set to true, an error is returned. Otherwise, NULL is returned.

MapJoin settings

set odps.sql.mapjoin.memory.max=512

Description: allows you to set the maximum memory size for a small table when running MapJoin. Unit: MB. Default value: 512. Valid values: 128 to 2048.

set odps.sql.reshuffle.dynamicpt=true/false

Description:

- Dynamic partitioning scenarios are time-consuming. Disabling dynamic partitioning can accelerate SQL.
- If there are few dynamic partitions, disabling dynamic partitioning can prevent data skew.

Scenario 6: how to check the storage usage of a single project

Launch the MaxCompute console as a project owner and run the desc project <project_name>extended; command to view the following information.

Storage information

odps@ odps_smoke_test>desc project od	
Name	odps_smoke_test
Description	
Owner	ALIYUN\$odpsadmin@aliyun.com
CreatedTime	Fri Dec 25 00:43:06 CST 2015
Properties:	
odps.table.lifecycle	optional
odps.function.strictmode	false
odps.table.drop.ignorenonexistent	false
odps.instance.priority.level	3
odps.task.sql.write.str2null	false
odps.instance.priority.autoadjust	false
odps.table.lifecycle.value	37231
odps.task.sql.outerjoin.ppd	false
odps.optimizer.mode	hbo
odps.instance.remain.days	30
READ_TABLE_MAX_ROW	10000
Extended Properties:	
tempDataLogicalSize	3642
tempDataPhysicalSize	10926
ableLogicalSize	20530
usedQuotaPhysicalSize	4162347
esourcePhysicalSize	4043403
cempResourcePhysicalSize	0
ableBackupPhysicalSize	38016
volumePhysicalSize	0
volumeLogicalSize	0
failoverPhysicalSize	8412
ableBackupLogicalSize	12672
ailoverLogicalSize	2804
cempResourceLogicalSize	0
ablePhysicalSize	61590
usedQuotaLogicalSize	1387449
resourceLogicalSize	1347801

The preceding figure shows the capacity-related storage information of the project. The relationship between the physical and logical values of the related metrics is: Physical value of a metric = Logical value of the metric * Number of replicas.

10.2. DataWorks

10.2.1. Basic concepts and structure

10.2.1.1. What is DataWorks (base)?

DataWorks, also known as base, is a visual workflow development platform that applies MaxCompute as its compute and storage engine. This platform is integrated with a hosted scheduling system, an administration system, and a synchronization system that can handle massive data. You can schedule your tasks by specifying a particular time and task relationships. You can also use the monitoring and management tools to ensure the punctual and accurate execution of millions of tasks. In addition, you are provided with a global overview of each workflow in the form of a directed acyclic graph (DAG).

10.2.1.2. Functions of base

Data collection

The data synchronization feature enables you to synchronize tables in a source database to a destination database using the data synchronization feature provided by base. Tables can be synchronized between heterogeneous data sources.

Data analysis

Write Shell, MapReduce (MR), or SQL code, and then submit the code to MaxCompute for computing.

Workflow

In base, you can combine task nodes of different types into a workflow. A workflow can contain data sync nodes, SQL nodes, Shell nodes, and MR nodes.

Task scheduling

You can run the tasks periodically with different cycles.

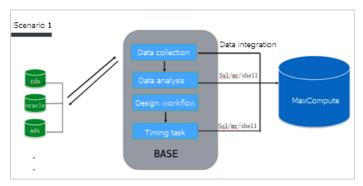
10.2.1.3. Introduction to data analytics

Scenario 1: data synchronization and analysis

Scenario 1 shows a typical scenario of data analytics.

- 1. You can collect data from various databases, and send the data to MaxCompute by using DataWorks.
- 2. You can log on to DataWorks, create SQL, MapReduce, and shell nodes, and commit the nodes to MaxCompute for data analysis.
- 3. You can use DataWorks to synchronize the analysis results from MaxCompute to the databases from which you collect data.

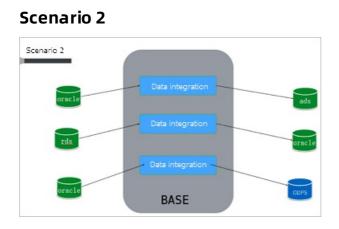
Scenario 1



Scenario 2: data synchronization

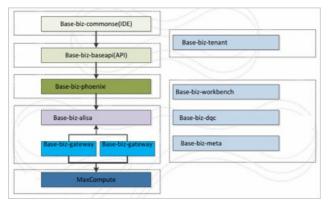
DataWorks supports data synchronization between various databases. You can synchronize data by using DataWorks.

> Document Version:20200918



10.2.1.4. DataWorks architecture in Apsara Stack V3

DataWorks architecture



Services shown in DataWorks architecture play an important role for node scheduling and running. You can perform all O&M operations for DataWorks of Apsara Stack V3 in Apsara Infrastructure Management Framework. The following figure shows the services in DataWorks.

DataWorks services

Operations and Maintenance Guide • Operations of big data products

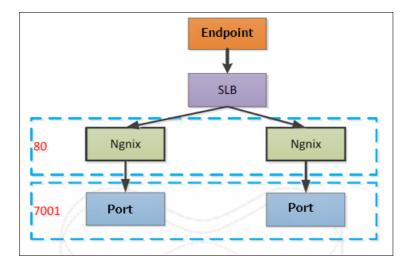
ALL Clusters	Service		
BasicCluster-A-20171113	Group		
BigdataCloudCluster-A-2	BaseBizAlisa#		
	BaseBizBaseapi#		
	BaseBizCdp#		
	BaseBizCdpGateway#		
	BaseBizCdpGatewayWithNc#		
	BaseBizCommonbase#		
Machine Service	BaseBizConsole#		
q	BaseBizDbinit#		
/ base-baseBizApp	BaseBizDfs#		
	BaseBizDmc#		
BaseBizAlisa#	BaseBizDqcexecutor#		
BaseBizBaseapi#	BaseBizDqcsupervisor#		
BaseBizCdp#	BaseBizGateway# BaseBizGatewayWithNc#		
BaseBizCdpGateway#			
BaseBizCdpGatewayWith	BaseBizMetaservice#		
BaseBizCommonbase#	BaseBizPhoenix#		
BaseBizConsole#	BaseBizSso#		
BaseBizDbinit#	BaseBizTenant#		
BaseBizD/s#	BaseBizWkbench#		
BaseBizDmc#	Redis1#		
BaseBizDqcexecutor#	Redis2#		
BaseBizDqcsupervisor#	ServiceTest#		
Base8izGateway#	pgsql-aurora#		
BaseBizGatewayWithNc#	pgsql-master#		
BaseBizMetaservice#	pgsgl-slave#		

All services in DataWorks are deployed in Docker containers. You can log on to a host and run the docker ps command to view the containers in which the services are deployed.

OUTADIER 10	3442	*/kis/bah -< /start.	COMMAND	OREATED	STATUS	PORTS
d6790e678d0			*/bin/bash -c /start.*	3 months ago	Up 3 months	22/trp, 0.0.0.0:8014->60/trp
fefdib9d989		-c /start.*/bin/bash -c /start.	"/bin/besh -c /start."	3 months ago	Lp 3 months	22/trp, 0.0.0.0:0013->00/trp
6875832f3d9	"/but/bash	-c /start. */bin/bash -c /start.	"/bin/besh -c /start."	3 months ago	Up 3 months	22/trp, 0.0.0.0:0012->00/trp
\$5e46a)175F	"/bin/bash	< /start. */bin/bash -c /start.	"/bin/besh -c /start."	3 months ago	Lp 3 months	22/trp, 0.0.0.0:0011->00/trp
SiceBide43e	*/bin/bash	< /start. */bin/bash < /start.	"/bin/bash -c /start."	3 months ago	Up 3 months	22/trp, 0.0.0.0:0010->00/trp
\$125458034		< /start. */hin/bash < /start.	"/bin/bash -c /start."	3 months ago	Up 3 months	22/trp, 0.0.0.0:8009->80/trp
bdb34825c3		< /start. "/sin/bash < /start.	"/bin/bash -c /start."	3 months ago	lip 3 months	22/trp, 0.0.0.1:8008->80/trp
cf1a4f41332		< /start. "/bin/bash < /start.	*/bin/bash -c /start.*	3 months ago	tip 3 months	22/trp, 0.0.0.0:8007->80/trp
97142750478		c /start. */bin/bash -c /start.	"/bin/bash -c /start."	l months ago	tp 3 months	22/tcp, 0.0.0.0:006->60/tcp
13f2b90dffa	*/tin/bash	-c /start. */bin/bash -c /start.	"/bin/bash -c /start."	3 months ago	tp 3 months	22/tcp, 0.0.0.1:005->00/tcp
67176000115	*/bin/bash	< /start. */bin/bash -c /start.	"/bin/bash -c /start."	3 months ago	tp 3 months	22/trp, 88/trp, 8.8.8.8.8004->7001/trp
fde287158fa	*/hin/hash ->	htert. "Ainhah - Atert.	"/bin/bash -c /start."	3 months ago	Up 3 months	22/trp, 0.0.0.0:8003->80/trp
117538:9254		c /start. */bin/bash -c /start.	"/bin/bash -c /start."	3 months ago	Up 3 months	22/trp, 0.0.0.8:8002->80/trp
ac48/152e5e5			"/bin/bash -c /start."	3 months ago	Up 3 months	22/trp, 0.0.0.8:8001->80/trp
485abd55bb1		: /start. */bin/bash -c /start.	"/bin/sh -c /start.sh"	3 months ago	Up 3 months	22/trp, 0.0.0.0:80-x80/trp
lae97cb4af49		c /start. */bin/bash -c /start.	"/bin/bash -c /start."	3 months ago	Up 3 months	22/trp, 6379/trp, 0.0.0.0:36379->36379
00488701935	-/11/5esh	< /start. */bis/bash < /start.	"/bin/bash -c /start."	3 months ago	Up 3 months	22/tcp, 0.0.0.0:16379->16379/tcp, 6379
51354658586	*/bin/bash -	c /start. */bin/bash -c /start.	"/docker-entrypoint.s"	3 months ago	Lp 3 months	0.0.0.0:5432->5432/tcp
\$33925c2157	*/bin/bash	-c /start. */bin/bash -c /start.	"/intrypaint.sh mjoqi"	3 months ago	Up 3 months	0.0.0.0:3306->3306/tcp
A4886:13354		< /start. */bis/bash -< /start.	"/bin/sh -c "/usr/sbi"	3 months ago	Up 3 months	80/trp, 0.0.0.0:1022-x22/trp
i6d78e4adLod	59811		"/swarm joinadvert"	4 months ago	lip 4 months	2375/tcp

Architecture shows the architecture of each service except gateway.

Architecture



10.2.1.5. Directory of each service

base-biz-gateway

This service receives tasks from the development platform and the scheduling system, and proceeds to run the tasks.

- logs: The directory stores operational logs of the gateway service.
- taskinfo: The directory stores the code and logs of tasks.
- target: The directory is the home directory of the gateway service, which includes service code, scripts for starting and stopping the service, and configuration files.

cdp

This service handles data synchronization tasks.

- logs: The directory stores operational logs of the cdp service.
- conf: The directory stores configuration files of the cdp service.
- bin: The directory stores the script for starting the service.

The directory structure of other services

The following example shows the directory structure of the alisa service.

- logs: The directory stores operational logs.
- conf: The directory stores configuration files.
- bin: The directory stores the script for starting the service.

10.2.2. Common administration tools and

commands

10.2.2.1. Find the container that runs the service

In Apsara Infrastructure Management Framework V3, select **base** from the **project** drop-down list, and then select BasicCluster.

Double-click **baseBizApp** in the lower part of the left-side navigation pane to view all services.

You can find the VM host that runs the service by double-clicking the service name. All services are deployed in containers. Therefore, you can run the docker exec -it [container ID] bash command to enter the container.

10.2.2.2. Cluster resource list

In Apsara Infrastructure Management Framework, select base from the project drop-down list. Select BasicCluster from the project list, move the pointer over the More icon next to BasicCluster, and select Dashboard from the menu to go to the Cluster Dashboard page.

On the Cluster Dashboard page, you can find the cluster resource list.

The Result column of the cluster resource list contains the details of each application. You can obtain the database logon information of a service from the Result column.

10.2.2.3. Commands to restart services

Enter the container that runs the service as an admin user, and then run the following commands to restart services.

Note Only admin users can run the following commands to restart the service.

- To restart the base-biz-cdp service, run the /home/admin/cdp_server/bin/appctl.sh restart command.
- To restart the base-biz-gateway service, run the /home/admin/alisatasknode/target/alisataskn ode/bin/serverctl restart command.
- To restart other services, run the /home/admin/base-biz-[application name]/bin/jbossctl restart command.

For example, to restart the base-biz-alisa service, run /home/admin/base-biz-alisa/bin/jbossctl r estart .

10.2.2.4. View logs of a failed node

Log on to the DataWorks console. Click the DataWorks icon in the upper-left corner and select **Operation Center** from the menu.

On the **Dashboard** page of **Operation Center**, you can view the statistics of the running status of nodes and node instances. Click **Failed** in the upper-left corner to view the list of nodes that failed to run.

In the failed node list that appears, find the target node and choose **More** > **View Runtime Log** in the Actions column to view the runtime log of the node.

10.2.2.5. Rerun a task

If you want to rerun a failed task, select the task in the Administration console and click Rerun.

10.2.2.6. Terminate a task

If you want to terminate a running task, select the task in Administration, and then click Terminate.

⑦ Note Only running tasks can be terminated.

10.2.2.7. Filter tasks in the administration center

You can choose Administration > Task List and filter the tasks to maintain.

10.2.2.8. Commonly used Linux commands

top: You can run this command to view the system load.

The load average section shows the average system load over the last 5, 10, and 15 minutes. The system is overloaded if any of the average load divided by the number of logical CPUs is greater than five.

du: You can run this command to view the file size.

Run the du -sh [file name] command to view the size of the file. Run the du -sh * command to list the sizes of all files in the current directory.

ps: You can run this command to view system processes.

Run the ps -ef command to view all processes that are running in the system.

grep: You can run this command to print lines which match a specified string.

Run the following command to print log file lines that match a specified string.

grep ["string"] [file_name]

Run the following command to print first few lines in a log file.

grep -C [NUM] ["string"] [file_name]

Onte C is uppercase, and NUM is the number of lines you want to print.

Run the following command to print last few log file lines that match a specified string.

grep -A [NUM] ["string"] [file_name]

kill: You can run this command to terminate a process.

Run kill -9 [process ID] to terminate the process.

Docker commands

docker ps -a : You can run this command to list all containers.

docker logs [container ID] : You can run this command to veiw the container logs.

docker exec -it [container ID] bash : You can run this command to enter the container.

10.2.2.9. View the slots usage of each resource group

Scenario: When a large amount of tasks are waiting for resources, you need to view the slots usage of each resource group.

Log on to the alisa database. In the Cluster Resource list, find and right-click the base-biz-alisa service that has a type of db, and then click Show More. The database logon address, username, and password are displayed. Connect to the database using MySQL statements.

Run the following command to view the top 10 longest running tasks.

select task_id,gateway,slot,create_time from alisa_task where status=2 order by create_time limit 10;

Run the following command to view the top 10 tasks that occupy most slots.

select task_id,gateway,slot,create_time from alisa_task where status=2 order by slot desc limit 10;

Run the following command to view the number of tasks that run in each slot. You can learn which tasks occupy a large number of slots.

select slot,count(*) from alisa_task where status=2 group by slot;

Run the following command to view the slots usage of each resource group.

select exec_target,sum(slot) from alisa_task where status=2 group by exec_target;

Run the following command to view the status of each gateway node. If either the live value or the active_type value of a node is 1, the server does not work properly.

select * from alisa_node;

10.2.3. Process daily administration operations

10.2.3.1. Daily check

10.2.3.1.1. Check the service status and the basic

information of the servers

Log on to the Apsara Infrastructure Management Framework console and select base from the project drop-down list. Hover over the vertical dots next to BasicCluster, and then click Dashboard. On the Dashboard page that appears, check whether servers are in GOOD status and services are in the desired status. If you find any issues, troubleshoot specific servers and services or contact an O&M engineer.

The blue column indicates the number of servers in GOOD status. If an orange column appears, errors occur on some servers.

10.2.3.1.2. Check the postgres database

- 1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster. On the Service tab in the lower part of the left-side navigation pane, find baseBizApp.
- 2. Double-click baseBizApp, and then double-click psql-master.
- 3. Open the terminal window of the VM server.
- 4. Run the docker ps|grep master command to view the container ID.
- 5. Run the docker exec -it [container ID] bash command to enter the container.
- 6. Run the psql -h127.0.0.1 -Uphoenix_prod -ddpphoenix -p3320 command, and enter the password pgsql to connect to the postgres database. Run the following statement in the database.

select to_char(to_timestamp(next_fire_time/1000), 'YYYY-MM-DD HH24:MI:SS') from qrtz_triggers;

View the result.

If the result contains 00:00:00 of the current day, the service is running properly. If not, ask for Alibaba Cloud technical support.

7. Run the following command in the database.

select pid ,(now() - xact_start) as time , state,query from pg_stat_activity where state ! = 'idle' ord er by time desc;

In the result, if the stat value is active, the service is running properly. If not, contact Alibaba Cloud Customer Support.

10.2.3.1.3. Check the status of each gateway server

- 1. In the Apsara Infrastructure Management Framework console, open the dashboard page of BasicCluster.
- 2. In the cluster resource list, find and right-click the **base-biz-alisa** service that has a type of db, and then click **Show More**. The database logon address, username, and password are displayed.
- 3. Connect to the database using MySQL statements and run the following statement.

Select * from alisa_node;

In the result that is returned, if either the active_type value or the live value is -1 or 0, the service does not run properly. In this case, contact Alibaba Cloud Customer Support.

10.2.3.1.4. Check the case test report

1. Log on to the Apsara Infrastructure Management Framework console, and enter base in the search box on the Service (S) tab.

- 2. In the search result, select base-baseBizApp to open the Dashboard page of the service instance.
- 3. In the Service Monitoring List, click Details.

If the Failed Cases tab contains any record, contact Alibaba Cloud Customer Support.

10.2.3.2. View logs of the services

Logs of the gateway service are stored in /home/admin/alisatasknode/logs/alisatasknode.log .

Logs of the cdp services are stored in /home/admin/cdp_server/logs/cdp_server.log .

Logs of other services are stored in /home/admin/base-biz-[service name]/base-biz-[service name].log .

For example, the logs of the base-biz-phoenix service are stored in /home/admin/base-biz-phoenix/base-biz-phoenix.log .

10.2.3.3. Scale out the node cluster that runs the base-biz-

gateway service

Prerequisites

Check whether the current environment meets the requirements for scale-out, such as disk space, file ownership and permissions, file execution path, software version, and any other necessary scale-out conditions.

- Before you scale out the BasicCluster cluster, make sure that it reaches the desired state and functions as expected.
- Save a screenshot of the key initial configurations for the cluster.
- Check for IP address conflicts. If you want to use a new buffer cluster for the scale-out, make sure that the IP addresses that Deployment Planner assigns to the servers in the cluster are not used in the current environment. This can avoid exceptions arising from IP address conflicts after the scale-out.
- Check the clone_mode parameter.

(?) Note Apsara Infrastructure Management Framework of V3.3 and later versions supports cloning protection. Before scaling out the cluster, you need to set the clone_mode parameter to normal. After the scale-out process is complete, you need to set this parameter to block.

Choose Apsara Infrastructure Management Framework > Operations > Cluster Operations > Global Clone Switch.

In the Global Clone Switch dialog box that appears, select normal, and then click OK.

Procedure

Add a buffer cluster

⑦ Note You can use idle servers in an existing buffer cluster for the scale-out operation, without adding a new buffer cluster. This method is applicable if the host, memory, CPU, and disk size of the idle servers match those of current servers that run the base-biz-gateway service. In this case, start from moving the idle servers to the default cluster.

In the scale-out procedure, use the actual parameter values and IP addresses instead of the specific parameter values in this guide.

? Note When you plan to scale out the cluster with Deployment Planner, make sure that the name of the new buffer cluster is different from that of any existing buffer cluster.

1. Copy and paste _tianji_imports to the /apsarapangu/disk3/u_disk/ directory of the ops1 server, and run the following command in the tianji_zhuque_sdk directory.

./tianji_zhuque_exchanger.py import --skip_packages -o \${desired state in the Apsara Infrastruct ure Management Framework}-c tianji_dest.conf

- 2. Log on to the Apsara Infrastructure Management Framework. In the left-side navigation pane, locate the buffer cluster in the cluster list. Then, move the pointer over the More icon next to the buffer cluster, and select Cluster Operations and Maintenance Center from the shortcut menu to view the status of servers in the buffer cluster.
- 3. Run the following commands on the ops1 server to check scale-out information by calling API operations.

cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current

ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.con

f clt2.conf

./tianji_clt machinestatus -c buffer --config clt2.conf

Scale in the buffer cluster

Note You can use the default cluster to scale out the cluster that runs base-biz-cdp and base-biz-gateway services.

1. Make sure that the value of the scalable tag value is true for the new buffer cluster.

2. Log on to the ops1 server, and then run the following commands to scale in the buffer cluster.

cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current

ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.con

f clt2.conf (You can ignore the message indicating that the symbolic link already exists.)

Scale-in command

./tianji_ops_tool.py contract_nc -c [buffer cluster name] -l [hostname of the server to be removed], [hostname of the server to be removed],.... --config clt2.conf -s [SRG name]

Parameters

- -c: the name of the buffer cluster that you scale in, which starts with buffer-cluster. This parameter is required.
- -l: a list of server hostnames that are included in the scale-in operation. Separate multiple hostnames with commas (,). This parameter is required.
- -s: the name of the SRG where the servers reside. You can find the SRG name in the machine_group.conf file of the buffer cluster. This parameter is required. If you want to remove the server, use this method to find the SRG name of the server.-config: the tianji_clt configuration file. This parameter is required.

? Note Chinese characters are not supported in the command line.

3. Check whether the operation takes effect in the Apsara Infrastructure Management Framework.

On the Cluster Operations page, make sure that the servers are removed.

4. Run the following commands to view the scaling information by calling API operations.

cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current

- ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.con
- f clt2.conf (You can ignore the message indicating that the symbolic link already exists.)

./tianji_clt machinestatus -c default --config clt2.conf

5. On the Cluster Configuration page of the buffer cluster, check whether the server is deleted from the machine_group.conf file. If the server still exists in the machine_group.conf file, delete the server, and then submit a rolling task.

Add servers to the BasicCluster cluster, and specify the SRG name where these servers reside.

- 1. Check whether the clone mode for the BasicCluster cluster is set to Real Clone.
- 2. Run the following commands to perform scaling. A rolling task is triggered after running the command.

cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current

ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.con
f clt2.conf (You can ignore the message indicating that the symbolic link already exists.)

To add servers to the cluster that runs the base-biz-gateway service, run the following command:

./tianji_ops_tool.py expand_nc -c [the name of a BasicCluster cluster] -s BaseGwGroup -l [machine 1,machine2] --config clt2.conf

To add servers to the cluster that runs the base-biz-cdpgateway service, run the following command:

./tianji_ops_tool.py expand_nc -c [the name of a BasicCluster cluster] -s BaseCdpGwGroup -l [mach ine1,machine2] --config clt2.conf

Parameters

- $\circ~$ -c: the name of a BasicCluster cluster. The name starts with BasicCluster.
- -l: a list of server hostnames that are included in the scale-out operation. Separate multiple hostnames with commas (,).

(?) Note Chinese characters are not supported in the command line.

3. You can run the following command to call an API operation to check the cluster to which the servers belong and the uplink information of the servers. This process may take a few minutes.

curl http://127.0.0.1:7070/api/v3/column/m. *?m.id=[machine hostname]

4. Log on to the OpsClone container, and run the following command to view the clone status:

/home/tops/bin/python /root/opsbuild/bin/opsbuild.py acli list --status=ALL -n 10000 | vim -

5. Check the rolling task status in the Apsara Infrastructure Management Framework.

Export the file that contains the information of desired state

After you complete the scale-out, export the file that contains the information of recent desired state to Deployment Planner. This ensures the success of subsequent scale-in and scale-out operations.

Verify the scale-out operation

1. View the heartbeat log.

Open the terminal of the added server, log on to the gateway container, and then run the t ail -f /home/admin/alisatasknode/logs/heartbeat.log command.

If the heartbeat log is updated every five seconds, the heartbeat function is running as expected.

2. Query the database.

In the Apsara Infrastructure Management Framework, open the dashboard page of the BasicCluster cluster. In the cluster resource list, find the base-biz-alisa service of the db type, right-click the result field, and then click Show More. You can find the database logon credentials. Connect to the database by using a MySQL command, and run the alisa_node; command. The information of all gateway servers is displayed.

Check the values of the live field and the active_type field for the added server. If both the two values are 1, the server is added.

3. Verify that the server reaches the desired state on the Cluster Operation and Maintenance Center page.

10.2.3.4. Scale in the base-biz-gateway cluster

Prerequisite

If a server in the base-biz-gateway cluster fails, you can repair and restart the server to redeploy the server.

If you want to remove a healthy server from the base-biz-gateway cluster, follow the instructions in this topic.

? Note Before removing a healthy server, perform an on-site check to guarantee that the following conditions are met:

- No business applications are running on the server.
- The hostname of the server is correct.

Procedure

Perform checks before the scale-in

1. Perform an on-site check.

Collect the detailed information of the server to be removed and the cluster that contains the server.

2. Make sure that the value of the scalable tag is true for the service resource group (SRG) of the server to be removed. If the value is false, change it to true and submit a rolling task.

Log on to Apsara Infrastructure Management Framework. In the left-side navigation pane, choose **BasicCluster > Cluster Configuration File > machine_group.conf.** In this file, verify that the value of the scalable tag is true for the SRG of the server to be removed.

Stop the base-biz-gateway service

- 1. Log on to the server to be removed and run the ps -ef|grep gateway command to obtain the container ID of the base-biz-gateway service.
- 2. Run the docker exec -it [container ID] bash command to enter the container.
- 3. Switch to the admin account and run the /home/admin/alisatasknode/target/alisatasknode/bin /servervtl stop command.
- 4. Run the ps -ef|grep java command to check whether any process is running on the server. If any process is running, run the kill -9 [process ID] command to terminate the process.
- 5. Delete the program directories from the server.

Clean up the disks of the server. Skip this step if you want to clone the server.

#rm -rf /home/admin/*

#rm -rf /opt/taobao/tbdpapp/

Move servers from the base-biz-gateway cluster to the default cluster in Apsara Infrastructure Management Framework

1. Log on to the ops1 server and run the following commands to remove a server from the base-biz-gateway cluster:

cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current

 $ln-s/cloud/data/bootstrap_controller/BootstrapController\#/bootstrap_controller/tianji_dest.con$

f clt2.conf (After you run this command, if a message appears indicating that a symbolic link alrea dy exists, proceed with the next command.)

./tianji_ops_tool.py contract_nc -c [clusterName] -l [machineList] --config tianji_clt.conf -s [SRGnam e]

The parameters are described as follows:

- -c: Required. Set this parameter to the name of the base cluster to be scaled in. To obtain the cluster name, choose Operations > Cluster Operations in the top navigation bar and select base from the Project drop-down list.
- -l: Required. Set this parameter to the hostname of the server to be removed. Separate multiple hostnames with commas (,).
- -s: Required. Set this parameter to the SRG name of the server to be removed. Find the machine_group.conf file among the configuration files of the base cluster. In this file, find the SRG of the server to be removed.
- -config: Required. Set this parameter to tianji_clt.conf.
- 2. After you run the preceding command, check whether the scale-in operation succeeds in Apsara Infrastructure Management Framework.

Go to the Cluster Operation and Maintenance Center of the base cluster.

- 3. On the **Cluster Operation and Maintenance Center** page, check the number of servers that are being removed.
- 4. Click the number next to Machine: in: to identify the status of the servers that are being removed.

If the scale-in operation succeeds, the number of servers that are being removed decreases to zero. Otherwise, check the server status on this page.

You can follow the preceding steps to scale in a node cluster by moving servers to the default cluster in Apsara Infrastructure Management Framework. The following section describes how to remove servers from Apsara Infrastructure Management Framework.

Remove servers from Apsara Infrastructure Management Framework

- 1. In the top navigation bar, choose **Operations > Machine Operations**.
- 2. On the Machine Operations page that appears, click Machine Online/Offline in the upperright corner.
- 3. In the Machine Online/Offline dialog box that appears, click Remove Machine.
- 4. On the Remove Machine tab, search for the server to be removed by hostname in the leftside Enter Machine List section. You can only remove servers in the default cluster.
- 5. Confirm the information of the server and click Clear Machines to remove it.

Verify the server removal result

1. Check whether the server is moved to the default cluster in Apsara Infrastructure Management Framework.

In the top navigation bar, choose **Operations > Machine Operations**. On the Machine Operations page that appears, search for the target server by hostname and check whether it is in the default cluster.

2. Check whether the server is removed from the default cluster.

In the top navigation bar, choose **Operations > Machine Operations**. On the Machine Operations page that appears, search for the server by hostname. If you cannot find the server in the search results, the server is removed.

3. To check whether the server is removed from the default cluster, run the following command on the ops1 server to call the GetMachineInfo operation:

curl http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=\$hostname

10.2.3.5. Restart the base-biz-alisa service

Procedure

- 1. Click C on the top of the left-side navigation pane in the Apsara Infrastructure Management Framework. Select base from the project drop-down list, and select BasicCluster from the cluster list.
- 2. Then, double-click base-baseBizApp to view the base-baseBizApp service list. Find and double-click BaseBizAlisa to view servers that run the base-biz-alisa service.
- Select one of the servers, and choose More > Terminal to open the Terminal Service page. In the upper part of the left-side navigation pane, click Add and then run the docker ps|grep ali sa command to obtain the container ID.
- 4. Run the docker exec -it [container ID] bash command to enter the container.
- 5. Switch to the admin account and run the /home/admin/base-biz-alisa/bin/jbossctl restart command to restart the service.

If you see NGINX start Done in the command output, the base-biz-alisa service is restarted.

10.2.3.6. Restart the base-biz-phoenix service

Procedure

- 1. Click C on the top of the left-side navigation pane in the Apsara Infrastructure Management Framework. Select base from the project drop-down list, and select BasicCluster from the cluster list. Then, double-click base-baseBizApp to view the base-baseBizApp service list. Find and double-click BaseBizPhoenix to view servers that run the base-biz-phoenix service.
- Select one of the servers, and choose More > Terminal to open the Terminal Service page. In the upper part of the left-side navigation pane, click Add and then run the oenix command to obtain the container ID.
- 3. Run the docker exec -it [container ID] bash command to log on to the container.
- 4. Switch to the admin account and run the /home/admin/base-biz-phoenix/bin/jbossctl restart command to restart the service.

If you see NGINX start Done in the command output, the phoenix service is restarted.

10.2.3.7. Restart base-biz-tenant

Procedure

- 1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster. In the lower part of the left-side navigation pane, double-click BaseBizTenant in the service list, and then the host that runs the service appears.
- 2. Open the terminal window of the vm host, and run docker ps|grep phoenix to find the container ID.
- 3. Run docker exec -it [container ID] bash to enter the container.
- 4. Switch to the admin account and run /home/admin/base-biz-tenant/bin/jbossctl restart to restart the service.

After you run the command, if the status is **OK** and the command output ends with **NGINX** start **Done**, the tenant service is restarted successfully.

10.2.3.8. Restart base-biz-gateway

Procedure

- 1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and then select BasicCluster from the search result.
- 2. On the Service tab in the lower part of the left-side navigation pane, double-click basebaseBizApp, double-click BaseBizCdpGateway, and then the host that runs the service appears.
- 3. Open the terminal window of the host, and run the docker ps|grep gateway command to find the container ID.
- 4. Run the docker exec -it [container ID] bash command to enter the container.
- 5. Switch to the admin account, and run the /home/admin/alisatasknode/target/alisatasknode/bi n/serverctl restart command to restart the service.
- 6. After the service is restarted, run the ps -ef|grep java command to check whether the process is started.

Note This method can only be used where the gateway service is deployed in a Docker container.

For the service deployed on a physical server

If the service is deployed on a physical server, use the following method to restart the service.

1. In the Apsara Infrastructure Management Framework console, open the Dashboard page of BasicCluster. In the cluster resource list, find and right-click the base-biz-alisa service that has a type of db, and then click Show More. The database logon address, username, and password are displayed.

- 2. Run the select * from alisa_node; command in the database to view the information of all gateway servers, and use the node IP address to find and maintain the gateway server.
- 3. In the terminal window of the server, switch to the admin account, and then run the /home/ admin/alisatasknode/target/alisatasknode/bin/serverctl restart command.

10.2.3.9. Restart the base-biz-api service

Procedure

- 1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster. On the Service tab in the lower part of the left-side navigation pane, double-click baseBizApp, double-click BaseBizCdpGateway, and then the host that runs the service appears.
- 2. Open the terminal window of the host, and run the docker ps|grep gateway command to find the container ID.
- 3. Run the docker exec -it [container ID] bash command to enter the container.
- 4. Switch to the admin account, and run the /home/admin/alisatasknode/target/alisatasknode/bi n/serverctl restart command to restart the service.
- 5. After the service is restarted, run the ps -ef|grep java command to check whether the process is started.

? Note The above method can only be used where the gateway service is deployed in a Docker container.

10.2.3.10. Restart the base-redis service

Procedure

- 1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list and select BasicCluster.
- 2. On the Service tab in the lower part of the left-side navigation pane, double-click basebaseBizApp, and you can find redis1 and redis2.
- 3. Open the terminal window of the VM host, and run the docker ps|grep redis command to find the container ID.
- 4. Run the docker exec -it [container ID] bash command to enter the redis container.
- 5. Run the following commands to restart the redis service.

/etc/init.d/ssh restart

/etc/init.d/redis-sentinel restart

10.2.3.11. Restart DataWorks Data Service

Procedure

- 1. In the Apsara Infrastructure Management Framework console, search for dataworksdataservice on the S tab.
- 2. Hover over the vertical dots next to BasicCluster, and then click Operations to open the Operations page to view the details of dataworks-dataservice.
- 3. Click the service instance name to open Service Instance Dashboard, and then find Service Role List.
- 4. If you want to restart the server, select BaseBizDataServiceServer#. If you want to restart the Web application, select BaseBizDataServiceWeb#. Hover over the vertical dots next to the service name, and then click **Details** to open the Service Role Dashboard page, and then find the virtual machine in the Server Information area.
- 5. Open the terminal window of the VM host, and run the docker ps|grep dataservice command to find the container ID.
- 6. Run the docker exec -it [container ID] bash command to enter the container.
- 7. Switch to the admin account, and run the /home/admin/data-service-web/bin/jbossctl restart command to restart the service.

If you are restarting the server, run the /home/admin/data-service-server/bin/jbossctl restart command.

8. After you run the command, if the status is **OK** and the command output displays [OK] --SUCCESS at the end, the dataservice service is restarted successfully.

10.2.3.12. Restart DataWorks Data Management

Procedure

- 1. In the Apsara Infrastructure Management Framework console, select base from the project drop-down list, and select BasicCluster.
- 2. In the lower part of the left-side navigation pane, double-click BaseBizAlisa in the service list, and then find BaseBizDmc.
- 3. Double-click BaseBizTenant, and then the host that runs the service appears.
- 4. Open the terminal window of the vm host, and run docker ps|grep dmc to find the container ID.
- 5. Run docker exec -it [container ID] bash to enter the container.
- 6. Switch to the admin account.
- 7. Run /home/admin/base-biz-dmc/bin/jbossctl restart to restart the service.
- 8. After you run the command, if the status is **OK**, and the command output ends with NGINX start Done, the Data Management (DMC) service is restarted successfully.

10.2.4. Common issues and solutions

10.2.4.1. Nodes remain in the Pending (Resources) state

Symptom

After you log on to the DataWorks console and click **Operation Center** in the upper-right corner of the console, the following issue occurs on the **Dashboard** page that appears: The instances of many recurring nodes remain in the Pending (Resources) state for a long period of time.

Causes

The issue may occur due to any one of the following four reasons:

- A gateway server is overloaded or offline and its status value is -1 in the database.
- The slots that handle concurrent jobs are fully occupied.
- The disk on a gateway server is full.
- The system time of servers in the base cluster is out of sync with the time of the Network Time Protocol (NTP) server.

Solutions

To resolve this issue, follow these steps:

- Check the status of a gateway server in the database.
 - i. Log on to the database that hosts the base-biz-alisa service. In Apsara Stack V3, you can find the database endpoint from the resource list of the base cluster in Apsara Infrastructure Management Framework.
 - ii. Run the select * from alisa_node; command to check the values of the active_type and live fields.

If the value of the live field is -1, the server is offline. If the value of the active_type field is -1, the server is overloaded.

? Note In either case, use SSH to connect to the gateway server and then check the server load and heartbeat.

Run the tail -f/home/admin/alisatasknode/logs/heartbeat.log command to check the heartbeat of the gateway server.

If the heartbeat log is updated every five seconds, the heartbeat is normal. Otherwise, check the configuration files for an error.

Run the top command to display the load of the gateway server.

The status of the server becomes -1 in the database as a result of the high load. In this case, check whether the CPU and memory are overloaded. You can find out the high-load processes in the output of the top command.

You can run the ps -ef[greppid command to view processes of the specified node and identify which process causes the high load. To terminate a process, run the kill -9 [process ID] command. After the load drops, check whether the status of the server resumes to 1.

• Check whether the slots that handle concurrent jobs are fully occupied.

Log on to the database that hosts the base-biz-alisa service and run the following statements:

select group_name,max_slot from alisa_group where group_name like '%default%'; select exec_target,sum(slot) from alisa_task where status=2 group by exec_target;

Compare the query results of the two statements.

- The first statement returns the maximum number of slots that can be assigned in each resource group.
- The second statement returns the number of slots that are occupied in each resource group.

If the query results of the two statements are the same or almost the same, all resource groups run out of slots. In this case, if a large number of nodes are running, the subsequent nodes do not run until the preceding nodes are completed.

Run the following statement to list the top 10 nodes that require the longest runtime:

select task_id,gateway,slot,create_time from alisa_task where status=2 and create_time>current_ti me order by create_time desc limit 10;

Log on to the gateway server and run the ps -ef|grep task_id command.

? Note Replace task_id in this command with one of the node IDs that are returned by the preceding SELECT statement. You can obtain the node name from the command output.

Then, you can troubleshoot the node. If required, run the kill -9 command to terminate the node and release resources immediately. Otherwise, new nodes can start only after the existing nodes are completed.

• Check whether the disk on a gateway server is full.

Log on to the gateway server and run the df -h command to check whether the disk attached to /home/admin is full. If the disk is full, run the du-sh command to identify the files in the /home/admin directory that consume a large amount of space. You can manually remove some large log files from the /home/admin/alisatasknode/taskinfo/ directory.

- Check the system time of servers in the base cluster against the time of the NTP server.
 - i. Log on to the database that hosts the base-biz-alisa service and run the select now(); command to view the current time of the database.
 - ii. Check the system time of servers in the base cluster against the time of the database.
 - iii. Run the date command on the servers to check whether the system time of each server is synchronized with the time of the database. If the time difference is greater than 30 seconds, the base-biz-alisa service may fail. In this case, synchronize the system time of servers in the base cluster with the time of the NTP server.

(?) Note In Apsara Stack V3, you can find the servers of the base cluster in the service list in the Apsara Infrastructure Management Framework console and follow the proceeding steps to resolve the issue.

• Rename the phoenix folder to change it to a .bak file and restart the base-biz-alisa service.

If the issue persists after you perform the preceding steps, run the following command on the gateway server:

cd /home/admin/alisatasknode/taskinfo/prevDay/phoenix/

? Note Replace prevDay in this command with the date of the previous day in the format YYYYMMDD, for example, 20180306.

In this directory, run the mkdir test command. If the error message "Cannot create directory too many links" appears, the issue occurs because the number of subdirectories in the directory has reached the maximum and you cannot create more subdirectories. To resolve this issue, follow these steps:

- i. Rename the /home/admin/alisatasknode/taskinfo/20180306/phoenix directory as /home/admin/alisatasknode/taskinfo/20180306/phoenix.bak.
- ii. Run the following command to restart the base-biz-alisa service:

sudo su admin -c "/home/admin/alisatasknode/target/alisatasknode/bin/serverctlrestart"

? Note This is a rare problem which tends to occur when a gateway server uses the third extended (ext3) file system.

10.2.4.2. An out-of-memory (OOM) error occurs when

synchronizing data from an Oracle database

Description

During the data synchronization from an Oracle database to MaxCompute or other platforms, an java.lang.OutOfMemoryError: Java heap space error is displayed in the task log.

Cause

This issue is often caused by a large volume of data in the data synchronization task, which causes a JVM OOM error.

Solution

Set a low fetchsize value.

Use MySQL statements to connect to the cdp database, and modify the template configuration of the Oracle reader plug-in by changing the fetchsize value from 1024 to 128. Run the following statement:

```
update t_plugin_template set template=replace(template,'1024','128') where name='oracle' and type='r eader';
```

Rerun the task after the fetchsize value is changed. To reset the fetchsize value, run the following statement:

update t_plugin_template set template=replace(template,'128','1024') where name='oracle' and type='r eader';

10.2.4.3. A task does not run at the specified time

Description

A periodic task does not run, and no data is displayed in the overview.

Solution

1. Check whether periodic scheduling is enabled in this workspace.

On the Workspace Configuration page in Workspace Management, ensure that the periodic scheduling is enabled.

2. If it is enabled, check whether the phoenix service runs properly.

Connect to the phoenix database and run the following statement.

select to_char(to_timestamp(next_fire_time/1000),'YYYY-MM-DDHH24:MI:SS') from qrtz_triggers;

If the output contains 00:00:00 of the next day, the service is running properly. If not, you need to check whether the time of the two base-biz-phoenix containers are different.

If the two containers have the same system time, you need to switch to the admin account and run the /home/admin/base-biz-phoenix/bin/jbossctl restart command to restart the phoenix service, and then check the time again.

3. After the time is corrected, you can run tasks that failed to run on the previous day.

Run the following command in either of the phoenix containers. Note that you can run this command only once.

curl -v -H "Accept:application/json"-H "Content-type: application/json"-X POST -d'{"opCode":11,"o pSEQ":12345,"opUser":"067605","name":"SYSTEM","bizdate":"2017-04-2300:00:00","gmtdate":"2017 -04-2400:00:00"}' http://localhost:7001/engine/2.0/flow/create_unified_daily

Note bizdate refers to the previous day, and gmtdate refers to the current day. Modify the command if needed before running it.

10.2.4.4. The test service of base is not in the desired status

- 1. On the S tab, select base-baseBizApp.
- 2. Select the cluster in the lower part of the left-side navigation pane, and then open the dashboard.
- 3. View the report of service monitoring.

Analyze the causes of the failed test based on the log.

10.2.4.5. The Data Management page does not display the number of tables and the usage of tables

Description

The Data Management page is blank.

Solution

- 1. Log on to the Apsara Infrastructure Management Framework console, select odps from the project drop-down list, and then open the HybridOdpsCluster dashboard page.
- 2. Find the accesskey type base_admin service in the Cluster Resource area.
- 3. Right-click the result field, and click Show More to view the username and the password.
- 4. Log on to DataWorks.

Note To log on to DataWorks, enter the domain name of base in the browser. By default, the domain name is ide.[your Apsara Stack second-level domain].

5. Select the base_meta workspace, and go to Administration.

Rerun all failed tasks, and then check whether the Data Management page is displayed properly. If the task fails again, contact Alibaba Cloud Customer Support.

10.2.4.6. Logs are not automatically cleaned up

Description

Logs are not cleaned up automatically because of an error.

Solution

Follow the following steps to clean up the logs manually.

- 1. Establish a terminal session to the VM.
- 2. Run the following command to clean up real-time analysis logs.

find /home/admin/cai/logs/cronolog/ -mtime +7 -type f -exec rm -rf {} \;
find /home/admin/dw-realtime-analysis/logs/ -mtime +7 -type f -exec rm -rf {} \;

3. Run the following command to clean up base-biz-diide application logs.

find /home/admin/cai/logs/cronolog/ -mtime +7 -type f -exec rm -rf {} \;
find /home/admin/base-biz-diide/logs/ -mtime +7 -type f -exec rm -rf {} \;

4. Run the following command to clean up base-biz-cdp application logs.

find /home/admin/cai/logs/cronolog/ -mtime +7 -type f -exec rm -rf {} \;
find /home/admin/base-biz-cdp/logs/ -mtime +7 -type f -exec rm -rf {} \;

10.2.4.7. The real-time analysis service is not in the desired

status

Description

The real-time analysis service is not in the desired status.

Solution

- 1. On the S tab, select dataworks-realtime.
- 2. Open the dashboard page of the cluster in the lower part of the left-side navigation pane.
- 3. View the report of service monitoring.

View the log to find out what caused the failed test.

10.3. Realtime Compute

10.3.1. Job status

10.3.1.1. Overview

StreamCompute allows you to view the real-time running information and instantaneous values of a job. You can also determine whether a job is running properly and whether the job performance meets expectations based on the job status.

10.3.1.2. Task status

A task can be in one of the following seven statuses: created, running, failed, completed, scheduling, canceling, and canceled. You can determine whether a job is running properly based on the task status.

10.3.1.3. Health score

To help you quickly locate job performance issues, Realtime Compute offers a health check feature.

If the health score of a job is lower than 60, lots of data has been piled up on the current task node and data processing performance needs to be optimized. To optimize the performance, you can enable automatic resource configuration or manually reconfigure the resources. You can optimize the performance based on your business requirements.

10.3.1.4. Job instantaneous values

Job parameters

Name	Description
Consumed compute time	Indicates the computing performance of a job.

Operations and Maintenance Guide \cdot Operations of big data products

Name	Description
Input TPS	Indicates the number of data blocks that are read from the source per second. For Log Service, multiple data records can be included in a log group and the log group functions as the basic unit of measurement for data. In this scenario, the number of blocks indicates the number of log groups that are read from the source per second.
Input RPS	Indicates the number of data records that are read from the source table per second.
Output RPS	Indicates the number of data records that are written into result tables per second.
Input BPS	Indicates the data transmission rate per second, which is measured in bytes per second.
CPU usage	Indicates the CPU usage of the job.
Start time	Indicates the start time of the job.
Running duration	Indicates the duration during which the job has been running.

10.3.1.5. Running topology

A running topology shows the execution of the underlying computational logic of Realtime Compute. Each component corresponds to a task. Each dataflow starts with one or more sources and ends in one or more result tables. The dataflows resemble arbitrary directed acyclic graphs (DAGs). For more efficient distributed execution, Realtime Compute chains operator subtasks together into tasks if possible. Each task is executed by one thread.

Chaining operators together into tasks provides the following benefits:

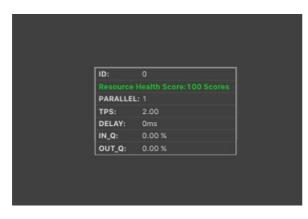
- Reduces the thread-to-thread handover.
- Reduces the message serialization and deserialization.
- Reduces the data handover in the buffer zone.
- Increases overall throughput while decreasing latency.

An operator indicates the computational logic, and a task is a collection of multiple operators.

View mode

The underlying computational logic is visualized in a view, as shown in View mode, to offer you a more intuitive display.

View mode



You can view the detailed information about a task by moving the pointer over the task. Parameter description describes the task parameters.

Parameter	Description
ID	The task ID in the running topology.
PARALLEL	The parallelism, which is the number of operator subtasks.
CPU	The CPU usage of a parallelism.
МЕМ	The memory usage of a parallelism.
TPS	The amount of data read from the inputs, which is measured in blocks per second.
LATENCY	The compute time consumed on the task node.
DELAY	The processing delay on the task node.
IN_Q	The percentage of input queues for the task node.
OUT_Q	The percentage of output queues for the task node.

Parameter description

You can also click a task node to access its details page. On this page, you can view its subtasks, as shown in Task details page.

Task details page



The Curve Charts tab provides curve charts to show the metrics of each task, as shown in Curve charts for task metrics.

Curve charts for task metrics

Operations and Maintenance Guide • Operations of big data products

Overview Cu	urve Charts	Failover	Checkpoints	s JobM	lanager	TaskExecutor	Data Lineag	e Prop	erties and Parameters			
Last 10 Minutes									Refresh Auto Refresh			
✓ OverView												
			Fa	ailover 💿							Delay 💿	
1.00												
0.80												
0.60												
0.40												
0.20									0 ms			
09:27:40 0												
			Input TPS	of Each So	urce 💿					Data Ou	utput of Each Sink 💿	
2.50 bps 2.00 bps 1.50 bps 1.00 bps 0.50 bps 0 bps 09:28:00												

List mode

In addition to the view mode, Realtime Compute also allows you to view each task in the list mode, as shown in List mode.

List mode



Parameter description describes the task parameters.

Parameter	Description
ID	The task ID in the running topology.
Name	The name of the task.
Status	The status of the task.
INQ max	The maximum percentage of input queues for the task node.
OUTQ max	The maximum percentage of output queues for the task node.
RecvCnt sum	The total amount of data that is received by the task node.
SendCnt sum	The total amount of data that is sent from the task node.
TPS sum	The total amount of data that is read from the inputs per second.
Delay max	The longest processing delay on the task node.
Task	The status of each parallelism on the task node.
StartTime	The start time of the task node.
Durations(s)	The running duration of the task node.

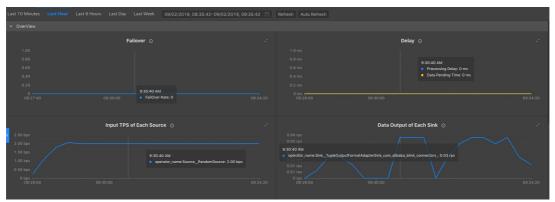
Parameter description

10.3.2. Curve charts

10.3.2.1. Overview

On the Curve Charts tab of the Realtime Compute development platform, you can view the key metrics of a job. This allows you to easily analyze the performance of a job. Currently, we are working on intelligent and automatic diagnosis by developing in-depth intelligent analysis algorithms based on the job running information.

Curve Charts ta	b
-----------------	---



? Note

- The metrics shown in this figure are displayed only when the job is in the running status.
- The metrics are asynchronously collected in the background, which results in delays. The metrics can be collected and displayed only after a job has been running for more than 1 minute.

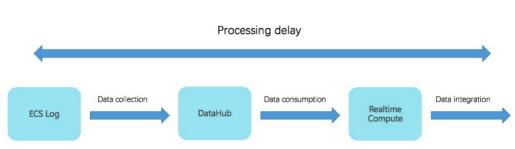
10.3.2.2. Overview

Failover rate

The failover rate indicates the percentage of the number of times that errors or exceptions occur on the current job. The failover rate curve allows you to easily analyze the issues of the current job.

Processing delay

The processing delay refers to the time interval between the current processing time and the time of reading data in the Realtime Compute service. If the time of reading data is not specified, the upstream DataHub or LogHub assigns the system timestamp to the data. The processing delay shows the timeliness of Realtime Compute end-to-end processing. For example, if the current processing time is 05:00 and the timestamp of the stored data is 01:00, the data to be processed was stored at 01:00, which is 4 hours earlier than the current processing time. In this scenario, the processing delay is 4 hours. The processing delay is used to monitor the data processing progress. If the source data fails to flow into DataHub because of certain faults, the processing delay increases accordingly. If the source data fails to enter DataHub because of certain faults, the processing delay.



Processing delay

The processing delay can be categorized into the following three types:

- Shortest delay: indicates the shortest processing delay of shards among data sources.
- Longest delay: indicates the longest processing delay of shards among data sources.
- Average delay: indicates the average processing delay of shards among data sources.

Input TPS of each source

Realtime Compute collects statistics about the streaming data inputs of each Realtime Compute job to help you easily view the input transactions per second (TPS). The input TPS describes the amount of data that is read from the source table, which is measured in blocks per second. Unlike the TPS, records per second (RPS) indicates the number of data records parsed based on the data blocks that are read from the source table.

For example, in Log Service, N log groups are read per second and M log records are parsed based on the N log groups. In this example, the input TPS is N, and the output RPS is M.

Data outputs of each sink

Realtime Compute collects statistics about data outputs of each Realtime Compute job to help you easily view the output RPS.

⑦ Note The outputs show all data outputs rather than streaming data outputs.

As an administrator, if you find that no data output is detected, you must check whether data inputs from the upstream exist. You also need to check whether data outputs in the downstream exist.

Input RPS of each source

> Document Version:20200918

Realtime Compute collects statistics about the streaming data inputs of each Realtime Compute job to help you easily view the input data records per second. As an administrator, if you find that no data output is detected, you must check whether data inputs from the source exist.

Input BPS of each source

Realtime Compute collects statistics about the streaming data inputs of each Realtime Compute job to help you easily view the input data bytes per second (BPS). The input BPS indicates the amount of data that is read from the source table per second.

CPU usage

The CPU usage describes the CPU resources consumed by a Realtime Compute job. Realtime Compute provides the following two metrics to reflect the CPU usage:

- The number of CPUs that you have applied for.
- The CPU usage of the current job at the specified time, which is shown in the curve chart.

Memory usage

The memory usage describes the memory resources consumed by a Realtime Compute job. Realtime Compute provides the following two metrics to reflect the memory usage:

- The size of memory space that you have applied for.
- The memory usage of the current job at the specified time, which is shown in the curve chart.

Dirty data from each source

Realtime Compute allows you to view the dirty data from each source through the corresponding curve chart.

10.3.2.3. Advanced view

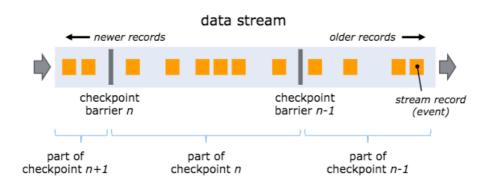
Realtime Compute offers a fault tolerance mechanism to consistently recover the state of data streaming applications. The central part of the fault tolerance mechanism is drawing consistent snapshots of the distributed data stream and the state. These snapshots act as consistent checkpoints to which the system can fall back when a failure occurs.

One of the core concepts of distributed snapshots is the barrier. Barriers are inserted into data streams and flow together with the data streams to the downstream. Barriers never overtake records, and the dataflow is strictly in line. A barrier separates the records in the data stream into two sets of records.

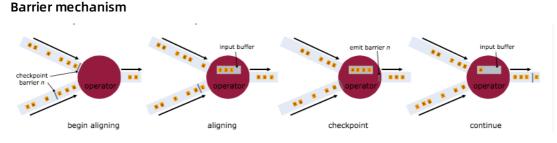
- One set of records is sorted into the current snapshot.
- The other set of records is sorted into the next snapshot.

Each barrier carries the ID of the snapshot that covers the records before the barrier. Barriers are a lightweight mechanism. They do not interrupt the flow of the stream. Multiple barriers from different snapshots can be in the stream at the same time. This means that multiple snapshots may be created concurrently.

Barriers



Stream barriers are injected into the dataflow at the stream sources. The point where the barrier for snapshot n is injected is the position in the source stream, up to which the snapshot covers the data. This point is indicated by Sn. The barriers then flow downstream. When an intermediate operator has received a barrier for snapshot n from all of its input streams, it emits a barrier for snapshot n into all of its outgoing streams. When a sink operator has received the barrier n from all of its input streams, it acknowledges that snapshot n to the checkpoint coordinator. A sink operator is the end of a streaming directed acyclic graph (DAG). After all sinks have acknowledged a snapshot, the snapshot is considered completed.



Checkpoint parameters

Checkpoint Duration

This parameter indicates the time spent on saving the state for each checkpoint. The duration is measured in milliseconds.

CheckpointSize

This parameter indicates the state size of a checkpoint, which is measured in MiB.

• checkpointAlignmentTime

This parameter indicates the time spent on receiving and acknowledging the barrier n from all incoming streams. When a sink operator (the end of a streaming DAG) has received the barrier n from all of its input streams, it acknowledges that snapshot n to the checkpoint coordinator. After all sinks have acknowledged a snapshot, the snapshot is considered completed. The time consumed by the acknowledgement is included in the checkpoint alignment time.

- CheckpointCount
- Get

This parameter indicates the longest time that a subtask spends on performing a GET operation on the RocksDB within a specified period.

• Put

This parameter indicates the longest time that a subtask spends on performing a PUT operation on the RocksDB within a specified period.

• Seek

This parameter indicates the longest time that a subtask spends on performing a SEEK operation on the RocksDB within a specified period.

• State Size

This parameter indicates the state size of a job. If the size increases excessively fast, you need to check and resolve potential issues.

• CMS GC Time

This parameter indicates the garbage collection (GC) time that is consumed by the underlying container that runs the job.

• CMS GC Rate

This parameter indicates how often the garbage collection is performed in the underlying container that runs the job.

10.3.2.4. Processing delay

Top 15 source subtasks with the longest processing delay

This metric describes the processing delays of each parallelism of a source.

10.3.2.5. Throughput

Task Input TPS

This indicates the data inputs of all tasks for the job.

Task Output TPS

This indicates the data outputs of all tasks for the job.

10.3.2.6. Queue

Input Queue Usage

This indicates the input data queues of all tasks for the job.

Output Queue Usage

This indicates the output data queues of all tasks for the job.

10.3.2.7. Tracing

The available parameters for advanced users are as follows:

• Time Used In Processing Per Second

This parameter indicates the time that a task spends on processing the data of each second.

• Time Used In Waiting Output Per Second

This parameter indicates the time that a task spends on waiting for outputs of each second.

• TaskLatency

This parameter indicates the computing delay of each task for a job. This delay is indicated by the interval between the time when data enters a task node and the time when data processing is completed on the task node. You can view the delay from the corresponding curve chart.

• WaitOutput

This parameter indicates the time that a task spends on waiting for outputs. You can view the waiting time from the corresponding curve chart.

• WaitInput

This parameter indicates the time that a task spends on waiting for inputs. You can view the waiting time from the corresponding curve chart.

• Source Latency

This parameter indicates the delay of each parallelism for a data source. You can view the delay from the corresponding curve chart.

10.3.2.8. Process

Process MEM Rss

You can view the memory usage of each process from the curve chart.

Memory NonHeap Used

You can view the non-heap memory usage of each process from the curve chart.

CPU Usage

You can view the CPU usage of each process from the curve chart.

10.3.2.9. JVM

Memory Heap Used

This indicates the Java Virtual Machine (JVM) heap memory usage of the job.

Memory NonHeap Used

This indicates the JVM non-heap memory usage of the job.

Threads Count

This indicates the number of threads for the job.

GC (CMS)

This indicates how often garbage collection (GC) is performed for the job.

10.3.3. FailOver

> Document Version:20200918

On the FailOver tab of the Realtime Compute development platform, you can check whether the job is running properly.

Latest FailOver

On the Latest FailOver tab, you can view the running errors of the job.

FailOver History

On the FailOver History tab, you can view the previous running errors of the job.

10.3.4. CheckPoints

Realtime Compute offers a fault tolerance mechanism to consistently recover the state of data streaming applications. The central part of the fault tolerance mechanism is drawing consistent snapshots of the distributed data stream and the state. These snapshots act as consistent checkpoints to which the system can fall back when a failure occurs.

Completed Checkpoints

On this tab, you can view the checkpoints that have been created. Parameter description describes the parameters for the created checkpoints.

Parameter description

Parameter	Description
ID	The ID of the checkpoint.
StartTime	The start time when the checkpoint is created.
Durations (ms)	The time that is spent on creating the checkpoint.

Task Latest Completed Checkpoint

On this tab, you can view the detailed information about the latest checkpoint. Parameter description describes the parameters for the latest checkpoint.

Parameter description

Parameter	Description
SubTask ID	The ID of the subtask.
State Size	The state size of the checkpoint.
Durations (ms)	The time that is spent on creating the checkpoint.

10.3.5. JobManager

After a Realtime Compute cluster is started, one JobManager and one or more TaskManagers are started. A client submits jobs to the JobManager, and the JobManager assigns the tasks of jobs to TaskManagers. During task execution, TaskManagers report the heartbeats and statistics to the JobManager. The TaskManagers exchange the data streams.

Similar to Storm Nimbus, a JobManager schedules jobs and functions as a coordinator to create checkpoints for tasks. A JobManager receives resources, such as jobs and JAR files, from a client. Then, the JobManager generates an optimized execution plan based on these resources and assigns tasks to TaskManagers.

10.3.6. TaskExecutor

After a Realtime Compute cluster is started, one JobManager and one or more TaskManagers are started. A client submits jobs to the JobManager, and the JobManager assigns the tasks of jobs to TaskManagers. During task execution, TaskManagers report the heartbeats and statistics to the JobManager. The TaskManagers exchange the data streams.

The number of slots is specified before a TaskManager is started. A TaskManager executes each task in each slot, and each task can be considered as a thread. A TaskManager receives tasks from the JobManager, and then establishes a Netty connection with its upstream to receive and process data.

TaskExecutor shows the detailed information about each TaskManager.

10.3.7. Data lineage

On the Data Lineage tab of the Realtime Compute development platform, you can view the dependencies of a job, including its relationship with its source table and result table. The topology on this tab allows you to easily and clearly view the complex dependencies of a job.

Data sampling

Realtime Compute provides the data sampling feature for source tables and result tables of jobs. The data to be sampled is the same as the data on the Development page. The data sampling feature allows you to check data at any time on the Administration page to facilitate fault locating. In the topology, click the button on the right side of the table name to enable the data sampling feature.

10.3.8. Properties and Parameters

The Properties and Parameters page provides detailed information about the current job, including the current running information and running history.

Job Code

On this tab page, you can preview the SQL job. You can also click **Edit Job** to go to the **Development** page.

Resource Configuration

On this tab page, you can view the resources that have been configured for the current job, including the CPU, memory, and parallelism.

Properties

> Document Version:20200918

On this tab page, you can view the basic running information of the current job. Job properties describes the basic job properties that are displayed on this tab page.

Job properties

No.	Field and Description
1	Job Name: indicates the name of the job.
2	Job ID: indicates the ID of the job.
3	Referenced Resources: indicates the resources that are referenced by the job.
4	Execution Engine: indicates the engine of the job.
5	Last Operated By: indicates the user who last operates the job.
6	Action: indicates the action that is last performed.
7	Created By: indicates the user who creates the job.
8	Created At: indicates the time when the job is created.
9	Last Modified By: indicates the user who last modifies the job.
10	Last Modified At: indicates the time when the job is last modified.

Running Parameters

On this tab page, you can view the underlying checkpoints, start time, and running parameters of the job.

History

On this tab page, you can view the detailed information about all versions of the job, including the start time, end time, and the user who operates the job.

Parameters

On this tab page, you can view additional job parameters, such as the separator used in the debugging file.

10.3.9. Improve performance by automatic

configuration

Background

To improve user experience, the Realtime Compute team offers the automatic configuration feature.

This feature optimizes the configuration of resources and parallelism for each operator of a job when the operators, data sources, and data sinks of Realtime Compute jobs are running properly. The automatic configuration feature also helps to globally improve job performance and handle issues, such as low throughput and data piling up on the upstream nodes.

This feature can optimize job performance in the following scenarios, but cannot address the performance bottlenecks of Realtime Compute jobs. To address the performance bottlenecks, contact the technical support team of Apsara Stack or your administrator.

- The performance of data sources or sinks needs to be improved.
 - Data sources. For example, the partitions of a DataHub source table are insufficient or the message queue (MQ) throughput is low. In this scenario, you need to increase the partitions of the data source table.
 - Data sinks. For example, an ApsaraDB for RDS deadlock occurs.
- The performance of user-defined extensions (UDXs) needs to be improved, such as userdefined functions (UDFs), user-defined aggregation functions (UDAFs), and user-defined table functions (UDTFs).

Improve the performance of a new job

- 1. After you write the SQL statements and the statements pass the syntax check, click **Publish**. The **Publish New Version** dialog box appears.
 - If you select Automatic CU Configuration, the automatic configuration algorithm determines the number of compute units (CUs) based on the system default configuration to optimize resource configuration. If automatic configuration is performed for the first time, the algorithm determines the number based on empirical values. We recommend that you perform automatic configuration after a job has been running for more than 10 minutes. In most cases, the resources are optimally allocated after you perform automatic configuration three to five times.
 - If you select Use Latest Manually Configured Resources, the latest saved resource configuration is used, no matter whether the resources are configured automatically or manually.

Note We recommend that you select Automatic CU Configuration. If you are performing automatic configuration for the first time, use the default number of CUs.

- 2. After you configure the resources for the job, click **Next** to check the data, and then click **Publish** to publish the job. Note that the default number of CUs is used.
- 3. Start the job.

The following section uses an example to describe how to improve job performance by using the automatic configuration feature. In this example, the default number of CUs for the job is 71. Note that the job must run for more than 10 minutes before automatic configuration is performed.

4. Improve the performance by using the automatic configuration feature.

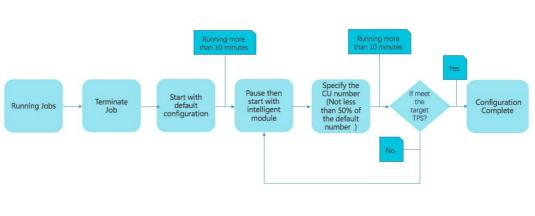
? Note Optimize the resource configuration. In this example, you can specify 40 CUs and select automatic configuration. You can increase or decrease the number of CUs based on the job running information. We recommend that you set the number of CUs to a value that is greater than or equal to 1 and 50% of the default number of CUs. For example, if the default number of CUs is 71, we recommend that you set the number of CUs to a value that is greater than or equal to 35.5 (71 CUs × 50% = 35.5 CUs). If the specified CUs cannot meet the throughput requirements of the job, you can increase the number of CUs. We recommend that you increase the number of CUs by more than 30% each time. For example, if 10 CUs were last specified, you can specify 13 CUs. If the result does not meet your needs, you can perform automatic configuration for several times and increase or decrease the number of CUs based on the job running information.

5. View the result of performance improvement.

Note If you are performing automatic configuration on a new job, do not select Use Latest Manually Configured Resources. Otherwise, an error message is displayed.

Improve the performance of an existing job

The following figure shows the procedure for improving the performance of an existing job.



Before performing automatic configuration on an existing job, check whether stateful operations are involved. This is because the saved state information of a job may be cleared

during the automatic configuration process. If a job is changed, for example, an SQL statement is modified or the Realtime Compute version is changed, the automatic configuration may fail. The reason is that these changes may lead to

changed, the automatic configuration may fail. The reason is that these changes may lead to topology changes, which further results in certain issues. These issues include: 1. Curve charts do not display the latest data. 2. The state cannot be used for fault tolerance. In this scenario, resource configuration cannot be optimized based on the job running history, and an error occurs while performing automatic configuration. To perform automatic configuration on a job that has been changed, perform steps 1 to 5 from the previous section on the changed job.

To perform automatic configuration on an existing job, perform steps 1 to 5 for a new job, and resume the job with the latest configuration.

Restrictions

Procedure

The result of automatic configuration may be compromised in the following scenarios:

- • The target job runs only for a short period. In this scenario, the useful information collected during data sampling is limited. This reduces the accuracy of the results calculated based on the automatic configuration algorithm. We recommend that you perform automatic configuration after the curves, such as Input RPS of Each Source, have been stable for 2 to 3 minutes.
 - The target job has encountered a failover. This reduces the accuracy of the results calculated based on the automatic configuration algorithm. We recommend that you check and handle failovers before performing automatic configuration.
 - Only a small amount of data is available for the target job. This reduces the accuracy of the results calculated based on the automatic configuration algorithm. We recommend that you trace more historical data.
 - The configuration obtained by using the automatic configuration feature is not always better than that from the last time. If the automatic configuration feature cannot meet your needs for improving the job performance, manually configure the resources.

Recommendations

- Before performing automatic configuration on a job, ensure that the job has been running stably and properly for more than 10 minutes. This helps to collect accurate job running information for the automatic configuration algorithm.
- You may need to perform automatic configuration for three to five times before the job performance is significantly improved.
- Before performing automatic configuration on a job, you can specify the start offset to read data from the past or even pile up large amounts of data for a job. This allows you to easily and quickly view performance improvement results.

Method for determining the effectiveness of automatic configuration

The automatic configuration feature for Realtime Compute is enabled based on a JSON configuration file. After performing automatic configuration, you can view the JSON configuration file to check whether this feature is running properly. You can view the JSON configuration file on either of the following tabs:

- Configuration Comparison tab under Properties on the Development page.
- **Resource Configuration** tab under **Properties and Parameters** on the **Administration** page.

The configurations in the JSON file are described as follows:

```
"autoConfig" : {
    "goal": { // The goal of automatic configuration.
    "maxResourceUnits": 10000.0, // The maximum number of CUs for a Blink job. The value cannot be
modified, and you can ignore this item when checking whether the feature is running properly.
    "targetResoureUnits": 20.0 // The number of CUs, which you have specified.
    },
    "result" : { // The results of automatic configuration. We recommend that you pay special attention
to this item.
    "scalingAction" : "ScaleToTargetResource", // The action of automatic configuration. *
    "allocatedResourceUnits" : 18.5, // The total resources.
    "allocatedCpuCores" : 18.5, // The total CPU cores.
    "allocatedMemoryInMB" : 40960 // The total memory size.
    "messages" : "xxxx" // We recommend that you pay special attention to the displayed messages.*
    }
}
```

- The InitialScale value of the scalingAction parameter indicates that automatic configuration is performed for the first time. The ScaleToTargetResource value of the scalingAction parameter indicates that automatic configuration is not performed for the first time.
- If no message is displayed, the automatic configuration feature is running properly. If certain messages are displayed, you need to analyze the messages and handle the issues. Messages are categorized into the following two types:
 - Warning: Messages of this type indicate that the feature is running properly, but you need to pay attention to potential issues, such as insufficient partitions of source tables.
 - Error or exception: Messages of this type indicate that the automatic configuration has failed. The following error message is usually displayed: Previous job statistics and configuration will be used. The automatic configuration for a job fails in either of the following two scenarios:
 - The job or Realtime Compute version has been modified. In this scenario, the previous running information cannot be used for automatic configuration.
 - The "xxxException" message is displayed. This message indicates that an error occurred while performing automatic configuration. You can analyze the error based on the job running information and logs. If the available information cannot help you to analyze the error, contact our technical support and development teams.

Error messages

IllegalStateException:

If the following error messages are displayed, the state cannot be used for fault tolerance. To resolve this issue, terminate the target job, clear its state, and then specify the start offset to re-read the data.

If you cannot migrate the target job to a backup node and you are concerned that online business may be interrupted, click **Properties** on the right side of the **Development** page, roll back the target job to the earlier version, and then specify the start offset to re-read the data during off-peak hours. java.lang.lllegalStateException: Could not initialize keyed state backend.

at org.apache.flink.streaming.api.operators.AbstractStreamOperator.initKeyedState(AbstractStreamOperator.java:687)

at org.apache.flink.streaming.api.operators.AbstractStreamOperator.initializeState(AbstractStream Operator.java:275)

at org.apache.flink.streaming.runtime.tasks.StreamTask.initializeOperators(StreamTask.java:870)

at org.apache.flink.streaming.runtime.tasks.StreamTask.initializeState(StreamTask.java:856)

at org.apache.flink.streaming.runtime.tasks.StreamTask.invoke(StreamTask.java:292)

at org.apache.flink.runtime.taskmanager.Task.run(Task.java:762)

at java.lang.Thread.run(Thread.java:834)

Caused by: org.apache.flink.api.common.typeutils.SerializationException: Cannot serialize/deserialize the object.

at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSecondaryState.deserializeStateE ntry(AbstractRocksDBRawSecondaryState.java:167)

at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRestoreOperation.restoreRawState Data(RocksDBIncrementalRestoreOperation.java:425)

at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRestoreOperation.restore(RocksDBI ncrementalRestoreOperation.java:119)

at com.alibaba.blink.contrib.streaming.state.RocksDBKeyedStateBackend.restore(RocksDBKeyedStateBackend.java:216)

at org.apache.flink.streaming.api.operators.AbstractStreamOperator.createKeyedStateBackend(Ab stractStreamOperator.java:986)

at org.apache.flink.streaming.api.operators.AbstractStreamOperator.initKeyedState(AbstractStreamOperator.java:675)

... 6 more

Caused by: java.io.EOFException

at java.io.DataInputStream.readUnsignedByte(DataInputStream.java:290)

at org.apache.flink.types.StringValue.readString(StringValue.java:770)

at org.apache.flink.api.common.typeutils.base.StringSerializer.deserialize(StringSerializer.java:69)

at org.apache.flink.api.common.typeutils.base.StringSerializer.deserialize(StringSerializer.java:28)

at org.apache.flink.api.java.typeutils.runtime.RowSerializer.deserialize(RowSerializer.java:169)

at org.apache.flink.api.java.typeutils.runtime.RowSerializer.deserialize(RowSerializer.java:38)

at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSecondaryState.deserializeStateE ntry(AbstractRocksDBRawSecondaryState.java:162)

... 11 more

10.3.10. Improve performance by manual configuration

10.3.10.1. Overview

You can manually configure resources to improve job performance using one of the following methods:

- Optimize resource configuration. You can modify the resources to improve the performance by reconfiguring parameters, such as parallelism, core, and heap_memory.
- Improve performance based on job parameter settings. You can specify the job parameters such as miniBatch to improve the performance.
- Improve upstream and downstream data storage based on parameter settings. You can specify related parameters to optimize the upstream and downstream storage for a job.

More details about these three methods are described in the following sections. After parameters are reconfigured to improve the performance of a job, the corresponding job must be re-published and started or resumed to apply the new configuration. The detailed process is provided in the following section.

10.3.10.2. Optimize resource configuration

Problem analysis

- 1. The percentage of input queues at task node 2 has reached 100%. Large amounts of data have piled up at task node 2, which results in the piling up of output queues at task node 1 in the upstream.
- 2. You can click task node 2 and find the subtask where the percentage of input queues has reached 100%. Then, click View TaskExecutor Logs to view the detailed information.
- 3. On the TaskExecutor page, you can view the CPU and memory usage. You can increase the number of CPU cores and expand the memory based on the current usage to handle the large amounts of data that have piled up.

Performance improvement

- 1. On the Development page of the StreamCompute development platform, click Properties.
- 2. Click Configure Resources to enter the page for editing resources.
- 3. Find the group (if any) or operator that corresponds to task node 2. You can modify the parameters of one operator or multiple operators in one group at a time.
 - Modify the parameters of multiple operators in a group.
 - Modify the parameters of an operator.
- 4. After modifying the parameters, click **Apply and Close the Page** in the upper-right corner of the page.

? Note

If the resources of a group have increased but the performance is not improved, you need to separately analyze each operator in the group and find the abnormal operators. Then, you can modify the resources for the abnormal operators for performance tuning. To separately analyze each operator in a group, click the target operator and change the value of its chainingStrategy parameter to HEAD. If the value is already set to HEAD, click the next operator and change the value of its chainingStrategy parameter are as follows:

- ALWAYS: indicates that operators are chained into a group.
- NEVER: indicates that operators are not chained.
- HEAD: indicates that operators are separated from a group.

Principles and recommendations

You can modify the following parameters:

- parallelism
 - Source

Set the parallelism parameter based on the number of source table partitions. For example, if the number of sources is 16, set the parallelism parameter to 16, 8, or 4. Note that the maximum value is 16.

• Operators

Set the parallelism parameter based on the estimated queries per second (QPS). For tasks with low QPS, set the parallelism parameter for the operators to the same value as that for the sources. For tasks with high QPS, set the parallelism parameter to a larger value, such as 64, 128, or 256.

• Sinks

Set the parallelism parameter for the sinks to a value that is two or three times the number of downstream sink partitions. However, if the specified parallelism limit is exceeded, a write timeout or failure occurs. For example, if the number of downstream sinks is 16, the maximum value of the parallelism parameter for sinks is 48.

• core

This parameter indicates the number of CPU cores. The default value is 0.1. Set this parameter based on CPU usage. We recommend that you set this parameter to a value whose reciprocal is an integer. The recommended value is 0.25.

• heap_memory

This parameter indicates the heap memory size, whose default value is 256 MB. The value is determined based on the actual memory usage. You can click GROUP on the resource editing page to modify the preceding parameters.

• For the task nodes that use the GROUP BY operator, you can configure the state_size parameter.

This parameter specifies the state size. The default value is 0. If the operator state is used, set the state_size parameter to 1. In this case, the corresponding job requests extra memory for this operator. The extra memory is used to store the state. If the state_size parameter is not set to 1, the corresponding job may be killed by YARN.

? Note

- The state_size parameter must be set to 1 for the following operators: GROUP BY, JOIN, OVER, and WINDOW.
- General users only need to focus on the core, parallelism, and heap_memory parameters.
- For each job, we recommend that you assign 4 GB memory for each core.

10.3.10.3. Improve performance based on job parameter

settings

The miniBatch parameter can be used to optimize only GROUP BY operators. During the streaming data processing of Flink SQL, the state is read each time a data record arrives for processing, which consumes large amounts of high I/O resources. After the miniBatch parameter is set, the state is read only once for data records with the same key, and the output contains only the latest data record. This reduces the frequency of reading state and minimizes the data output updates. The settings of the miniBatch parameter are described as follows:

1. The allowed delay for a job.

blink.miniBatch.allowLatencyMs=5000

2. The size of a batch.

blink.miniBatch.size=1000

10.3.10.4. Optimize upstream and downstream data

storage based on parameter settings

In Realtime Compute, each data record can trigger read and write operations on source and result tables. This brings considerable challenges for upstream and downstream data storage performance. To address these challenges, you can set batch size parameters to specify the number of data records that are read from a source table or written into a result table at a time. The following table describes the available batch size parameters.

Parameter description

Object	Parameter	Description	Value
DataHub source table	batchReadSize	The number of data records that are read at a time.	Optional. Default value: 10.

Operations and Maintenance Guide • Operations of big data products

Object	Parameter	Description	Value
DataHub result table	batchSize	The number of data records that are written at a time.	Optional. Default value: 300.
Log Service source table	batchGetSize	The number of log groups that are read at a time.	Optional. Default value: 10.
ApsaraDB for RDS result table	batchSize	The number of data records that are written at a time.	Optional. Default value: 50.

Note To complete batch data read and write settings, add the above parameters to the parameter list WITH in DDL statements for the corresponding data storage. For example, add batchReadSize='500' to the parameter list WITH in DDL statements for the DataHub source table.

10.3.10.5. Apply new configuration

After resources are reconfigured for a job, you must restart or resume the job to apply the new configuration. Perform the following operations:

- 1. Publish the job of the new version. In the Publish New Version dialog box, select Use Latest Configuration.
- 2. Suspend the job.
- 3. Resume the job. In the Resume Job dialog box, select **Resume with Latest Configuration**. Otherwise, the resource configuration cannot take effect.
- 4. After resuming the job, choose Administration > Overview > Vertex Topology to check whether the new configuration has taken effect.
- ? Note

We do not recommend that you terminate and restart a job to apply the new configuration. After a job is terminated, its status is cleared. In this case, the computing result may be inconsistent with the result that is obtained if you suspend and resume the job.

10.3.10.6. Concepts

• Global

isChainingEnabled: indicates whether the chaining is enabled. Use the default value (true).

- Nodes
 - id: specifies the unique ID of a node. The ID is automatically generated and does not need to be changed.
 - uid: specifies the UID of a node, which is used to calculate the operator ID. If this parameter is not specified, the value of id is used.

- pact: specifies the type of a node, such as the data source, operator, and data sink. Use the default value.
- name: specifies the name of a node, which can be customized.
- slotSharingGroup: Use the default value.
- chainingStrategy: specifies the chaining strategy. The options include HEAD, ALWAYS, and NEVER. Use the default value.
- parallelism: specifies the number of parallel subtasks. The default value is 1. You can increase the value based on the data volume.
- core: specifies the number of CPU cores. The default value is 0.1. The value is configured based on the CPU usage. We recommend that you set this parameter to a value whose reciprocal is an integer. The recommended value is 0.25.
- heap_memory: specifies the heap memory size. The default value is 256 MB. Set this parameter based on the memory usage.
- direct_memory: specifies the JVM non-heap memory size. We recommend that you use the default value (0).
- native_memory: specifies the JVM non-heap memory size for the Java Native Interface (JNI). The default value is 0. The recommended value is 10 MB.
- Chain

A Flink SQL task is a directed acyclic graph (DAG) that contains many nodes, which are also known as operators. Some upstream and downstream operators can be combined to form a chain when they are running. The CPU capacity of a chain is set to the maximum CPU capacity among operators in the chain. The memory size of a chain is set to the total memory size of operators in the chain. For example, after node 1 (256 MB, 0.2 cores), node 2 (128 MB, 0.5 cores), and node 3 (128 MB, 0.25 cores) are combined to form a chain, the CPU capacity of the chain is 0.5 cores and the memory is 512 MB. The prerequisite for chaining operators is that the operators to be chained must have the same parallelism settings. However, some operators cannot be chained, such as GROUP BY operators. We recommend that you chain operators to improve the efficiency of network transmission.

10.4. Quick BI

10.4.1. Introduction to O&M and tools

10.4.1.1. Introduction to operations and maintenance

Quick BI Operations and Maintenance (O&M) Guide provides step-by-step instructions to explain the O&M process for Quick BI. With the guide, You can perform daily operations, such as monitoring and maintaining Quick BI, and detecting, troubleshooting, and resolving issues. These operations can help ensure that Quick BI is available, stable, and secure.

You can use the Apsara Infrastructure Management Framework to troubleshoot the unavailability issues of Quick BI.

10.4.1.2. Troubleshoot Quick BI issues by using the Apsara

Infrastructure Management Framework

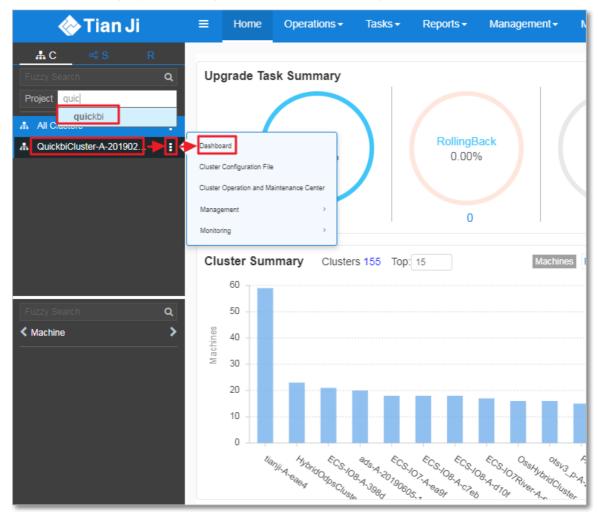
The Apsara Infrastructure Management Framework is a tool that allows you to perform O&M tasks on Quick BI. You can use the Apsara Infrastructure Management Framework to troubleshoot the service unavailability issue of Quick BI.

Prerequisites

Log on to the Apsara Infrastructure Management Framework.

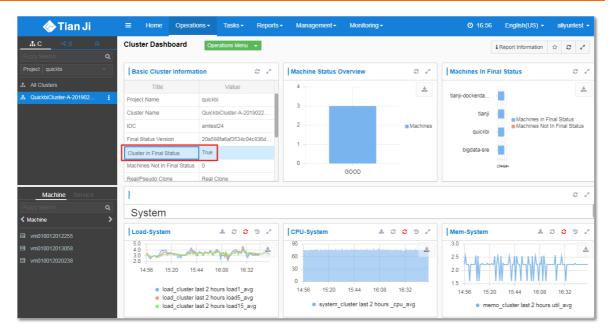
Procedure

1. Find the Quick BI project in the Apsara Infrastructure Management Framework.



2. On the Dashboard page, view the cluster status of Quick BI. Check whether the Quick BI cluster is at the desired state. If the cluster is at the desired state, the system works as expected. If the cluster is not at the desired state, go to the next step.

Operations and Maintenance Guide · Operations of big data products



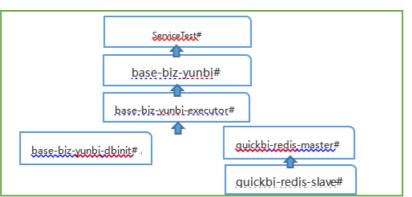
3. On the **Dashboard** page, find the **Service Instances** section, and view the service instance details of Quick BI.

🚸 Tian Ji	≡ Home Operat	ions	rts∓ Management∓ Mo	nitoring -	Ø 16:58 Engl	ish(US) - aliyuntest
<mark>≜C</mark> ≪S R	Cluster Dashboard	Operations Menu 👻			i Report	Information 🏠 😂 🖍
	tcp_cluster	last 2 hours retran_avg	1 ~ 100 / More			« < 1 / More >
All Clusters	Service Instances					Ø
QuickbiCluster-A-201902	Service Instance	Final Status	Expected Server Roles	Server Roles In Final	Server Roles Going O	Actions
	bigdata-sre	True	1	1	0	Actions - Details
	OS	True	-	-	-	Actions - Details
	quickbi	True	6	6	0	Actors - Details
	tianji	True	1	1	0	Actions - Details
	tianji-dockerdaemon	True	1	1	0	Actions - Details

🚸 Tian Ji	≡ Home Op	erations - Task	s - Reports - M	anagement - Monito	ring -	O 16:59	English(US) -	aliyuntest 👻
♣C ≪S R Fuzzy Search Q	Service Instance	Information Dash	operation	s Menu 👻		i	Report Information	\$ Ø 2
All Clusters								
A QuickbiCluster-A-201902								
	Server Role List				0 Z			
	Server Role	Current Status	Expected Machi	Machines In Fin	Machines Goin	Rolling Task Sta	Time Used	Actions
	base-biz-yunbi#	In Final Status	2	2	0	no rolling		Details
	base-biz-yunbi-dbinit#	In Final Status	1	1	0	no rolling		Details
	base-biz-yunbi-exec	In Final Status	2	2	0	no rolling		Details
	quickbi-redis-master#	In Final Status	1	1	0	no rolling		Details
Machine Service	quickbi-redis-slave#	In Final Status	2	2	0	no rolling		Details
	ServiceTest#	In Final Status	1	1	0	no rolling		Details
K Machine >								

4. If a service instance is not at the desired state, you need to follow these steps to troubleshoot the issue.

Dependencies exist between service roles. If an upstream service role has not reached the desired state, the downstream service role cannot reach the desired state. We recommend that you first troubleshoot the upstream service role. The following figure Relationship between Quick BI service roles shows the relationship between Quick BI service roles.



Relationship between Quick BI service roles

For example, if the base-biz-yunbi-executor# service role do not reach the desired state, the base-biz-yunbi# and ServiceTest# service roles cannot reach the desired state. You must first ensure that the base-biz-yunbi-executor# service role reaches the desired state. After the base-biz-yunbi-executor# service role reaches the desired state, the base-biz-yunbi# and ServiceTest# service roles will enter the desired state one by one excluding unexpected issues.

10.4.2. Routine maintenance

10.4.2.1. Introduction to Quick BI components

You can use container monitoring and periodical detection to check whether service roles related to Quick BI components are at the desired state. You can use these methods to manage and maintain Quick BI. This topic describes Quick BI operations and maintenance (O&M) components, related service roles, and the description about each component.

Component	Service role	Description	
Database initialization components	base-biz-yunbi-dbinit#	Allows you to initialize Quick BI metadata. The service role must be at the desired state before Quick BI can run as expected.	
Cache components	quickbi-redis-master#	Allows you to cache Quick BI data to improve query	
cache components	quickbi-redis-slave#	performance.	
Runtime components	base-biz-yunbi-executor#	Allows you to perform operations, such as retrieving table metadata and data from data sources.	
Web service components	base-biz-yunbi #	Provides Web services. The service role provides Web services that allow frontend clients to visit Quick BI Web pages.	

Quick BI O&M components, related service roles, and the description of each component

Component	Service role	Description
Automated testing components	ServiceTest#	Allows you to check the availability of Quick BI by running batch test cases.

? Note When you deploy or update Quick BI, the ServiceTest# service role is automatically started.

10.4.2.2. Database initialization components

This topic describes how to troubleshoot issues when you perform container monitoring on database initialization components.

In the Apsara Infrastructure Management Framework, you need to check whether the base-bizyunbi-dbinit# service role is at the desired state.

? Note The service role that is related to database initialization components must be at the desired state before Quick BI is running as expected. If the check result indicates that the service role is not at the desired state, we recommend that you contact Quick BI Technical Support.

10.4.2.3. Cache components

This topic describes how to detect and troubleshoot issues when you perform container monitoring on cache components.

Container monitoring

In the Apsara Infrastructure Management Framework, you need to check whether the **quickbi-redis-master#** and **quickbi-redis-slave#** service roles are at the desired state.

? Note You can also check the redis process. If the redis process exists, it means that the preceding service roles are at the desired state.

Quick BI is unavailable if the check result indicates that the linked service roles are not at the desired state. Cause: The redis process is interrupted or not started.

Solution: You need to restart the linked service roles. You need to restart the **quickbi-redis-master#** service role and then restart the **quickbi-redis-slave#** service role.

Periodical detection

You can check the service availability based on the exit status that is returned after you run the /checkRedis.sh script. Quick BI is available if the value of the exit status is 0. Otherwise, Quick BI is unavailable. You can use the preceding script to check whether the redis process exists. The redis process exists if the value of the returned exit status is 0. Otherwise, the redis process does not exist. The detection interval is one second.

10.4.2.4. Runtime components

This topic describes how to detect and troubleshoot issues when you perform container monitoring on runtime components.

Container monitoring

In the Apsara Infrastructure Management Framework, you need to check whether the **base-biz-yunbi-executor#** service role is at the desired state.

Quick BI is unavailable if the check result indicates that the linked service role is not at the desired state. Cause: The runtime component process is interrupted or not started.

Solution: You need to restart the base-biz-yunbi-executor# service role.

Periodical detection

You can visit http://container:7001/checkpreload.htm at regular intervals to call the HTTP service. Quick BI is available if a status code of 200 is returned. Otherwise, Quick BI is unavailable. The detection interval is one second.

? Note The container in the preceding HTTP link is a variable. You must replace the variable with an IP address that is used by the base-biz-yunbi# service role.

10.4.2.5. Web service components

This topic describes how to detect and troubleshoot issues when you perform container monitoring for Web service components.

Container monitoring

Check whether the base-biz-yunbi# service role is at the desired state.

Quick BI is unavailable if the check result indicates that the linked service role is not at the desired state. Cause:

- The Java process is interrupted or not started. Symptom: You cannot visit http://container:7001/checkpreload.htm.
- No HTTPS certificate is issued and port 443 is inaccessible. Symptom: You cannot visit https://container/checkpreload.htm.

? Note The container in the preceding link is a variable. You must replace the variable with an IP address that is used by the **base-biz-yunbi#** service role.

Solutions:

- If the Java process is interrupted or not started, you need to restart the base-biz-yunbi# service role.
- If no HTTPS certificate is issued, you need to restart the **base-biz-yunbi#** service after the HTTPS certificate is issued.

Periodical detection

You can visit https://container/checkpreload.htm at regular intervals to call an HTTPS service. Quick BI is available if a value of 200 is returned. Otherwise, Quick BI is unavailable. The detection interval is five minutes.

? Note The container in the preceding HTTPS link is a variable. You must replace the variable with an IP address that is requested by the base-biz-yunbi# service role.

10.4.2.6. Automated testing components

This topic describes how to detect and troubleshoot executor issues when you perform container monitoring on automated testing components.

Container monitoring

In the Apsara Infrastructure Management Framework console, check whether the ServiceTest# server role is at desired state.

If the server role is not at desired state, a service error occurs. Causes:

• The service is unavailable. Symptom: You cannot visit https://container/checkpreload.htm or log on to the Quick BI console.

Note "container" in the link is a variable. You must replace it with the IP address that is used by the base-biz-yunbi# server role.

• The service is available but an error is detected. Symptom: You can log on to the Quick BI console and search data. However, a logon error is reported in the Apsara Infrastructure Management Framework console. You can view the error message provided in the Description column.

Solutions:

- If the service is unavailable, check whether other server roles are at desired state. If they are not, handle the problem.
- If the service is available but an error is detected, contact Quick BI technical support and provide error information.

Periodical detection

You can execute test cases at regular intervals to check the availability of Quick BI. A service is available if the linked server role is at desired state. Otherwise, the service is unavailable. The detection interval is 30 minutes.

10.5. Graph Analytics

10.5.1. Operations and maintenance tools and

logon methods

10.5.1.1. Log on to Apsara Bigdata Manager

This topic describes how to log on to Apsara Bigdata Manager.

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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

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

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

3. Enter your username and password.

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

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

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

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- It must be 10 to 20 characters in length.
- 4. Click Log On to go to the ASO console.
- 5. In the left-side navigation pane, choose **Products** > **Apsara Bigdata Manager** to go to the homepage of Apsara Bigdata Manager.

Operations and Maintenance Guide \cdot Operations of big data products

	lanager ABM ፡		🗹 Monit	cor 웹 O&M 🕸 Management	
		C	Dashboard		
	∨ Jul 8, 2019, 15:22:31	~ Jul 8, 2019, 17:22:31 📋	×		
Overview					
Products 💠	∀ ок 🕏	중 Critical 🛊	∵ Warning 💠	♡ Exception ≑	
				Total Items: 9 < 1 >	50 / page $ arsigma$
> MaxCompute					
> DataWorks					
> StreamCompute					
> DataHub					

6. On the homepage of Apsara Bigdata Manager, click the icon in the upper-left corner and choose I+ to go to the operations page for Graph Analytics.

C 🕽 Apsara Bigdata	Manager I+ 🖀		器 O&M	ø Management	
	Services Clusters Host				
iplus v	bigdata-sre.Agent# Overview Server				
歳 bigdata-sre.Agent#	СРО /	DISK			
 Å iplus-iplus_biz.lplus Å iplus-iplus_biz.lplus Å iplus-iplus_biz.lplus Å iplus-iplus_biz.lplus Å iplus-iplus_biz.lplus 	40 30 20- 10- 10- 10- 10- 10- 10- 10- 10- 10- 1	10 8- 6- 4- 2- 0, 2019, 14:27:00 Jul 8,	2019, 15:03:00	Jul 8, 2019, 15:39:00	Jul 8, 2019, 16:15:0
🙏 tianji-dockerdaemo	LOAD	MEMORY			
, ≰ tianji.TianjiClient#	2 2 2 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4	39.1k 29.3k 19.5k 9.77k 18, 2019, 14:27:00 Jul	8, 2019, 15:04:00	Jul 8, 2019, 15:41:00	Jul 8, 2019, 16:1:

10.5.1.2. Log on to Apsara Infrastructure Management

Framework

This topic describes how to log on to Apsara Infrastructure Management Framework. Apsara Infrastructure Management Framework supports operations and maintenance (O&M) management for Graph Analytics.

Prerequisites

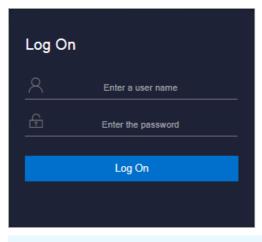
• The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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



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

3. Enter your username and password.

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

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

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

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

🔶 Tian Ji	Home Operations - Tasks - Reports - Ma	nagement - Monitoring -		O 10:33 English(US) -
C «S R ty Search Q act ALL ~	Announcement: The new version of Cluster Operations at	nd Maintenance Portal has been launched to suppo	rt operation and maintenance functions and task track	ing. Welcome to use it. Experience Immediately
(Clusters i	Rommg 100.00%	* Proceed 0.00%	,	Most-used Reports Registration Vars of Services PLst Thermonuter XOB Instance Meric Info Relationship of Service Dependency Machine Power On or Off Statuses of Clusters
Search Q section requires <u>double-</u> Double-click a cluster ine and a service chvely, and check ine loas. Choose	Cluster Summary Clusters 3 Top 3			Error Summary Rate of Abnormal Machines: 0.00% Machine 0.055 OS Errors: 0.00% HW Errors: Rate of Abnormal ServerRole Instances: 2.16% Attrix: 2.10%
g all machine tions from the <u>menu</u> , yet started, the data switching for associating reports	o tianji-paas-A-18	BigdataSaaS-A-20	modi-cluster-for	Abnormal: 16 Total: 741

10.5.1.3. Log on to the Graph Analytics container

You can log on to the Graph Analytics container through Apsara Infrastructure Management Framework to perform operations and maintenance.

Procedure

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. In the left-side **Project** drop-down list, enter or select *iplus* to display Graph Analytics clusters.
- 3. Select a Graph Analytics cluster. On the Services page, double click iplus-iplus_biz > IplusBizBackend#. Click the More icon next to vmxxxxxxxx and then select Terminal in the menu that appears. The TerminalService page appears.

Operations and maintenance are typically performed on the virtual machines of IplusBizBackend# and IplusBizBackendControl#. You can use the same method to open the virtual machine where the IplusBizBackendControl# service is deployed.

TerminalService page

TerminalService terminal service to reflect shell to web		Help Hello!
terminal service to reliect shell to web		neio:
 iplus 	"il vm) — — — ×	
d vmt1000000111100	⊕ [admin@vm /home/admin]	<u>ــــــــــــــــــــــــــــــــــــ</u>
	Ş	
	· · · ·	
Virtual AG		
Ad		-

The left-side navigation pane on the **TerminalService** page displays the virtual machine selected by you (vmxxxxxxxxx).

4. In the left-side navigation pane on the TerminalService page, click vmxxxxxxxxx, and

the command-line tool appears on the right-side of the page.

5. Run the docker ps|grep *iplus* command to query the docker ID in the Graph Analytics cluster.

0	uerv	the	docker ID	
~	~~.,			

[admin@vm	/home	/admin]		
\$docker ps grep ip	lus			
bc000	reg.doc	ker.	/ice_images/fr	ontend:aeed8a997e1
508810471f302b		"/bin/sh -c 'sudo sh "	2 days ago	Up 2 days
s_biz.IplusBizFron	tendip	lus-biz-frontend.153120	2126	
ad97d	reg.doc	ker.	/imore/imore_b	ackground_service:
c05393b231d593a656	f28ccd6	"sh /home/admin/ccbin"	3 days ago	Up 3 days
s_biz.IplusBizBack	endipl	us-biz-backend.15311539	75	

The query results of this sample display two docker IDs, which indicates that the **IplusBizBackend#** service is running on two containers.

6. Run the docker exec -ti dockerID bash command to log on to the docker container.

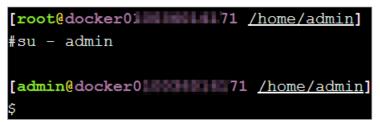
Enter the docker ID of the container you need to log on to in *dockerID*.

Log on to the docker container



7. The root account is used by default. You can use the **su - admin** command to switch to the admin account.

Switch to the admin account



10.5.2. Operations and maintenance

10.5.2.1. Operations and maintenance based on BigData

Manager

10.5.2.1.1. View and handle cluster alerts

The IT administrator must focus on Graph Analytics alerts and fix them in time, especially Warning alerts and Critical alerts.

Prerequisites

> Document Version:20200918

Your have obtained an Apsara BigData Manager account and the password with Graph Analytics O&M permissions.

Step one: View cluster alerts

- 1. Log on to Apsara BigData Manager.
- 2. Click the 🗃 icon in the upper-left corner. Click **Big Data Application** > I+.
- 3. On the page that appears, click **O&M > Clusters**.

C-) Apsara Bigdata	Manager I+ 🗃	眼 O&M 彩 Management 💮 🔤
	Services Clusters Hos	ts
Search by keyword. Q	iplus Overview Health Status	
▼ 昂 cn-	CPU 🧳	DISK 2
届 iplus	30 25 20 15 10 5 4 4 10 5 4 10 5 4 10 5 4 10 5 4 10 5 10 5	10 8- 6- 4- 2- 0 1, 2019, 14:34:00 Jul 8, 2019, 15:10:00 Jul 8, 2019, 15:46:00 Jul 8, 2019, 16:22:0
	LOAD 🧭	MEMORY 2
Recently Selected	2 1 1 1 1 2019, 1434.00 hul 8, 2019, 15.09:00 hul 8, 2019, 15:44:00 hul 8, 2019, 16:19:00	34.2k 29.3k 24.4k 19.5k 14.6k 9.77k 4.88k 0 1 k, 2019, 14.34.00 Jul 8, 2019, 15:11.00 Jul 8, 2019, 15:48.00 Jul 8, 2019, 16:21

4. On the **Clusters** page that appears, select a cluster in the left-side navigation pane, and then click **Health Status** on the right side. The **Health Status** tab page for the cluster appears.

		Services Clusters	Hosts		
Search by keyword. Q	Applications a statement	Overview Health Sta	itus		
▼ 器 cn-					
	Checker 🜲	♡ Source 矣	⊽ Critical 🚖	∵ Warning 🗢	ଟ Actions 🖕 େ ଟ
	+ bcc_check_ntp	tcheck			Details
	+ bcc_tsar_tcp_checker	tcheck			Details
	+ bcc_kernel_thread_count_checker	tcheck			Details
	+ bcc_network_tcp_connections_checker	tcheck			Details
	+ bcc_disk_usage_checker	tcheck			Details
	+ bcc_host_live_check	tcheck			Details
	+ bcc_process_thread_count_checker	tcheck			Details
	+ bcc_check_load_high	tcheck			Details
					< 1 >
Recently Selected					

The Health Status tab page displays all health check items of the current cluster. You need to focus on the check items with Critical and Warning alerts.

Step two: View hosts in the alert status and the alert causes

You can view the history of a check item and the check results.

1. On the Health Status tab page, click the Plus sign (+) in front of a check item that has alerts to view all hosts.

Checker					
Checker 🜲		∀ Source 🖨	중 Critical 🗢 🖓 W	Varning 🜲 🛛 🖓 Exception 🜲	♥Actions \$♥
- bcc_check_ntp		tcheck			
Host 🔺		∀ Status ≜	∵ 🛛 🛛 🕁 🖉 🖉 🖉 🖉	ত্ব Status Updated At ≜	ত Actions ≜ ত
a56		WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	
a56	2	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	
a56		WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	
a56		WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	
a56		WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	
a56		WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	
a56		WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	
a56		WARNING	Jul 8, 2019, 09:25:03	Jul 4, 2019, 18:55:07	
a56		WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:07	
a56		WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:06	
			Тс	otal Items: 32 < 1 2 3 4 >	10 / page \vee Goto

2. Click a host. In the dialog box that appears, click **Details**, and the cause of the alert appears on the right side.

a56History Status						
Status 🚖 🖓	7 Status Updated At 🜲	♀ Actions ↓	A	1562549106 sync=0 offset=0.001994		
WARNING	Jul 4, 2019, 18:55:10	Details				

Step three: View solutions and handle alerts

Apsara BigData Manager provides a solution for each alert to help you handle the alert quickly.

1. On the Health Status page, click Details of a check item that has alerts to view the corresponding solution.

Operations and Maintenance Guide · Operations of big data products

Name:	bcc_disk_usage_checker	Source:	tcheck
Alias:	Disk Usage Check	Application:	bcc
Type:	system	Scheduling:	Enable
Data Coll	ection: Enable		
	ection: Enable xecution Interval: 00/5***?		
	xecution Interval: 0 0/5 * * * ?		
Default E Description	xecution Interval: 0 0/5 * * * ?		red when the usage exceeds 80% and a critical alert is eted. Logrorate is not working. Fix:
Default E Description This check triggered	xecution Interval: 0 0/5 * * * ? on: ker checks the storage usage by using this	er operations. Old log data is not dele	-
Default E Descripti This check triggered 1. Lo <u>c</u>	xecution Interval: 0 0/5 * * * ? on: ker checks the storage usage by using this when the usage exceeds 90%. Reason: Use	er operations. Old log data is not dele executing this command: df -lh	eted. Logrorate is not working. Fix:

2. Handle the alerts based on the procedure described in Fix.

When you handle host alerts, you may need to log on to the host to perform related operations. For more information about how to log on to the host, see Step four: Log on to a host.

Step four: Log on to a host

1. On the Health Status tab page, click the Plus sign (+) of a check item.

Check	er								
	Checker 🜲	₽ S	Source 🜲	Critical 😄 🛛 🖓 V	Warning 🜲		Å	Actions 🔶	Å
-	bcc_check_ntp		tcheck						
	Host 🔺		Status 🔺	Last Reported At 🔺		Status Updated At 🔺		Actions 🔺	Å
	a56		WARNING	Jul 8, 2019, 09:25:07		Jul 4, 2019, 18:55:10			
			WARNING	Jul 8, 2019, 09:25:05		Jul 4, 2019, 18:55:09			
			WARNING	Jul 8, 2019, 09:20:07		Jul 4, 2019, 18:55:08			
			WARNING	Jul 8, 2019, 09:20:09		Jul 4, 2019, 18:55:08			
			WARNING	Jul 8, 2019, 09:20:33		Jul 4, 2019, 18:55:08			
			WARNING	Jul 8, 2019, 09:20:03		Jul 4, 2019, 18:55:07			
	a56	١	WARNING	Jul 8, 2019, 09:25:07		Jul 4, 2019, 18:55:07		Refresh	

2. Click the Logon icon of a host. The TerminalService page occurs.

TerminalService terminal service to reflect shell to web	Helo!
. 1 1 a56 💮 🕀	
	Welcome To Terminal service
	reiminal service
Virtual AG	

3. On the **TerminalService** page, select a host on the left side. You can log on to the host directly without entering the username and password.

TerminalService terminal service to reflect shell to web	
 Ignitigature i 20050-0.0 	al a56 ×
d a56 €	[admin@a56 /home/admin] \$]

10.5.2.1.2. View cluster performance metrics

IT administrators must regularly check and record server operation metrics of Graph Analytics for future troubleshooting. When a high resource consumption is detected, the IT administrator must take immediate measures to identify the cause and fix the issue.

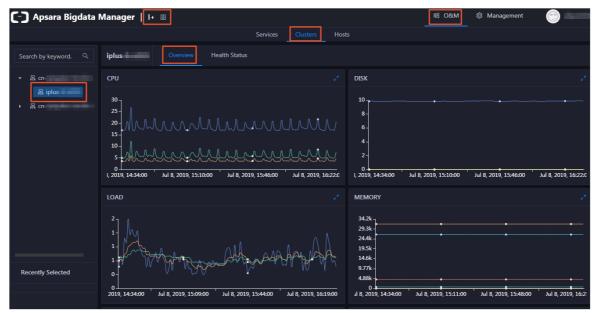
Prerequisites

Your have obtained an Apsara BigData Manager account and the password with Graph Analytics O&M permissions.

Procedure

1. Log on to Apsara BigData Manager.

- 2. Click the 🔳 icon in the upper-left corner. Click **Big Data Application** > I+.
- 3. On the page that appears, click **O&M > Clusters**.
- 4. On the **Clusters** tab page, select a cluster in the left-side navigation pane, and then click **Overview**. The **Overview** tab page of the cluster appears.



On the **Overview** page, you can view the usage trends of CPU, memory, disk, load, package, TCP, and disk root of the cluster.

10.5.2.1.3. View server operation metrics

IT administrators must regularly check and record server operation metrics of Graph Analytics for future troubleshooting. When a high resource consumption is detected, the IT administrator must take immediate measures to identify the cause and fix the issue.

Prerequisites

Your have obtained an Apsara BigData Manager account and the password with Graph Analytics O&M permissions.

Procedure

- 1. Log on to Apsara BigData Manager.
- 2. Click the 🗃 icon in the upper-left corner. Click **Big Data Application** > I+.
- 3. On the page that appears, click **O&M > Hosts**.
- 4. On the Hosts tab page, select a host in the left-side navigation pane, and then click **Overview**. The **Overview** tab page of the host appears.

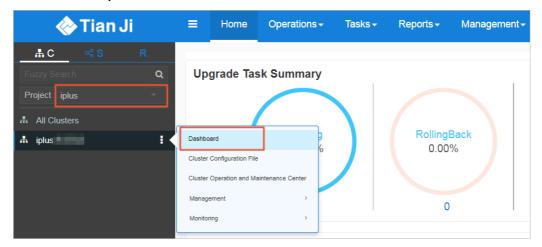
	Services Clusters	s Hosts
Search by keyword. Q	Vm Health Status	
Servers vm	Root Disk Usage	© 1-Minute Load © 1.9
vm (10	/tmp Usage 0 %	5-Minute Load 1 15-Minute Load 1
00 20 VII 00	(5) Total 22.4 %	
vm a < 4 / 4 >	System 5.2 % User 8.4 %	
	СРИ	Z DISK Z
Recently Selected		
vm (10.		6- 4-
vm 🖾 ao	5	0 • \bullet \bullet = \bullet \bullet = \bullet \bullet \bullet = \bullet \bullet = \bullet \bullet =\bullet \bullet =\bullet \bullet =\bullet \bullet \bullet =\bullet \bullet =

The **Overview** tab page displays the root disk usage, CPU usage, disk, memory, load, package, and TCP.

10.5.2.2. Operations and maintenance based on Apsara

Infrastructure Management Framework

- 1. Log on to Apsara Infrastructure Management Framework.
- 2. Enter iplus in the search box to search for the iplus cluster, as shown in Search for the iplus cluster.

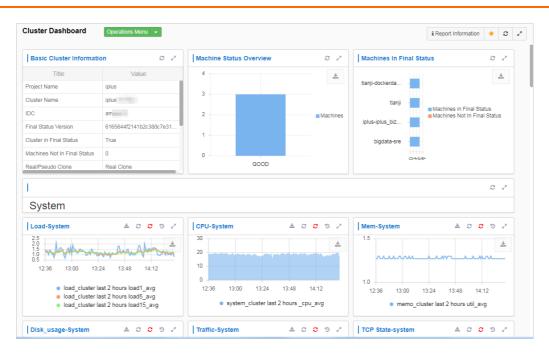


Search for the iplus cluster

3. Move the mouse pointer to the More icon next to the iplus cluster, and select **Dashboard** from the drop-down list. The Cluster Dashboard page appears, as shown in Cluster Dashboard page.

Cluster Dashboard page

Operations and Maintenance Guide • Operations of big data products



4. In the Service Instances list, click Details for iplus-iplus_biz. The Service Instance Information Dashboard page appears, as shown in Service instance list.

Service instance list

Service Instances						
Service Instance	Final Status	Expected Server Roles	Server Roles In Final S	Server Roles Going Off	Actions	
bigdata-sre	True	1	1	0	Actions - Details	
iplus-iplus_biz	True	5	5	0	Actions Details	
OS	True	-	-	-	Actions - Details	
tianji	True	1	1	0	Actions - Details	
tianji-dockerdaemon	True	1	1	0	Actions - Details	

You can restart any role in the server role list, as shown in Server role list. Typically, you only need to restart IplusBizBackendControl# and IplusBizBackend#.

✓ Notice

You must restart IplusBizBackendControl# and IplusBizBackend# in the following sequence:

- Restart IplusBizBackendControl# first, and then IplusBizBackend#.
- After you restart IplusBizBackendControl#, you must restart IplusBizBackend# within 10 minutes.

Other roles can be restarted in any order.

Server role list

Server Role List							0.
Server Role	Current Status	Expected Machi	Machines In Fin	Machines Going	Rolling Task Stat	Time Used	Actions
lplusBizBackend#	In Final Status	3	3	0	no rolling		Details
IplusBizBackendContr	In Final Status	2	2	0	no rolling		Details
lplusBizDbinit#	In Final Status	1	1	0	no rolling		Details
IplusBizFrontend#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

5. Select a role, and click Details. On the Server Role Dashboard page that appears, click **Restart** in the Actions column, as shown in **Restart server roles**.

Server Role Dashboard	Operations Menu 👻	
	Restart Server Role	
Server Role Summary	0 2	Machine Final Status Overview
Title	Value	
Project Name	iplus	
Cluster Name	iplus	
Service Instance	iplus-iplus_biz	
Server Role	lplusBizBackend#	
In Final Status	In Final Status	REACH_FINAL: 100.00%

10.5.2.3. Operations and maintenance based on the Graph

Analytics container

10.5.2.3.1. View instances

Restart server roles

By viewing and examining instances, you can know the running status of instances and fix the problematic instances, for example, perform a switchover or clear logs.

View Java running instances

Log on to the Graph Analytics container, and run the ps -ef|grep java|grep iplus command. If the progress shown in View Java running instances exists, Administration Console is in the normal status.

View Java running instances

Sps = e1[grep]ava[grep 1plus] admin 2578 1 0 hulps - 00:24:36 lava -server -xms1800m -xms1800m -xms600m -xss256k -xx:Permsize_512m -xx:MaxPermsize_512m -xx:Heaphump0routofMemorvError -xx:Heaphump2rbu/come/admin/l
ogs -XX:+UseConcMarkSweepGC -XX:+UseParNewGC -XX:CMSFullGCsBeforeCompaction=5 -XX:+UseCMSCompactAtFullCollection -XX:+CMSClassUnloadingEnabled -XX:+DisableExplicitGC -verbose:gc -XX:+PrintGCDetails -X
X:+PrintGCTimeStamps -Dfile.encoding=UTF-8 -jar /home/admin/iplus_pack/iplus-control.warspring.config.location=/home/admin/iplus_pack/config/application-control.yml
admin 27322 1 0 Julo5 ? 00:23:50 java -server -Xms5000m -Xmx5000m -Xmx1024m -Xss256k -XX:PermSize=512m -XX:HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/home/admin/
logs -XX:+UseConcMarkSweepGC -XX:+UseParNewGC -XX:+UseParNewGC -XX:+DisableExplicitGC -verbose:gc -XX:+PrintGCDateStamp
s -XX:+PrintGCDetails -XX:+PrintHeapAtGC -Xlogqc:/home/admin/logs/gc.log -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.authenticate=false -Dcom
.sun.management.jmxremote.port=9999 -Dfile.encoding=UTF-8 -jar /home/admin/iplus_pack/iplus.warspring.config.location=/home/admin/iplus_pack/config/application-service.yml

View node instances

Log on to the Graph Analytics application server, and run the ps -ef|grep node command. If the process shown in View node instances exists, the node service of Graph Analytics is normal.

View node instances

<pre>\$ps -ef grep node</pre>	
admin 7974 1 0 19:12 pts/0	00:00:00 node /home/admin/i3-admin/target/i3-admin/admin-patch.jsharmony 00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
admin 7991 7974 0 19:12 pts/0	00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.jsharmony undefined	
admin 7996 7974 0 19:12 pts/0	00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.jsharmony undefined	-
admin 7997 7974 0 19:12 pts/0	00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.jsharmony undefined	-
admin 8002 7974 0 19:12 pts/0	00:00:00 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-admin/target/i3-admin/lib/ai
o.jsharmony undefined	-
admin 14876 1 0 Aug16 ?	00:00:00 node /home/admin/i3-web/target/i3-web/dispatch.jsharmony
admin 14887 14876 0 Aug16 ?	00:02:20 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony	
admin 14892 14876 0 Aug16 ?	00:02:18 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony	
admin 14893 14876 0 Aug16 ?	00:02:19 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony	
admin 14898 14876 0 Aug16 ?	00:02:20 /usr/local/node-v4.4.2-linux-x64/bin/node /home/admin/i3-web/target/i3-web/index.js -
-harmony	

In the preceding information, i3-web indicates that Analytics Workbench is in a normal status, and i3-admin indicates that Administration Console is in a normal status. If Administration Console is not released, the i3-admin process may not exist.

10.5.2.3.2. Log files

Graph Analytics application log:

The log files of Graph Analytics are stored in the /home/admin/logs directory.

A 100-GB data disk is mounted to the /home/admin/logs directory. Log files will increase with the execution time, which requires automatic cleanup. Two cleanup policies are available:

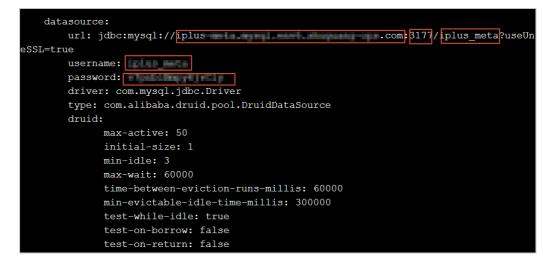
- Policy one: Time-based cleanup. The disk automatically deletes the log files that were created two weeks ago.
- Policy two: Cleanup based on the log size in the directory. If the log files occupy more than 80% of the total data disk space, the disk automatically deletes the earliest log files.

10.5.2.3.3. Database logs

Database logs record the execution information of i3-related programs, mainly the SQL statements. This information includes the execution time, whether the statements have been executed successfully, and whether an exception has occurred.

- 1. Log on to the Graph Analytics container.
- 2. Run the **cat** */home/admin/iplus_pack/config/application-service.yml* command to view the database information in *application-service.yml*.

View database information



- 3. Run the mysql -h\${db_host} -P\${db_port} -u\${db_user} -p\${db_password} -D\${db_name} command to log on to the database.
- 4. Query the latest SQL statement executed by Graph Analytics and the time track.

SELECT * from i3eye_time_trace WHERE main_time_trace_id in (SELECT max(main_time_trace_id) from i3eye_time_trace);

5. View the SQL statements executed within the last hour.

select * from i3eye_time_trace where name like 'com.alibaba.iplus.common.dal.manual%' and (gmt_ create < now() and gmt_create > date_sub(now(), interval 1 hour));

6. View the SQL statements that have errors within the last hour.

select * from i3eye_time_trace where complete = 0 and name like 'com.alibaba.iplus.common.dal.m
anual%' and (gmt_create < now() and gmt_create > date_sub(now(), interval 1 hour));

10.5.2.3.4. Stop the service

Use admin Log on to the Graph Analytics container, run the start script, and run the following ps commands to view processes:

- View Java process: ps -ef|grep java
- View node process: ps -ef|grep node

You can stop a service by killing the corresponding thread.

10.5.2.3.5. Restart the service

Use admin Log on to the Graph Analytics container and run the startup script:

- Directly start iplus, i3-web, and i3-admin: iplus-deploy.sh start
- Start iplus only: iplus-deploy.sh start_iplus
- Start i3web only: iplus-deploy.sh start_i3web

• Start i3admin only: iplus-deploy.sh start_i3admin

10.5.3. Security maintenance

10.5.3.1. Network security maintenance

Network security maintenance handles the device security and the network security.

Device security

- Check network devices, and enable security management protocols and configurations of the devices.
- Check for new versions of the network device software and update to a more secure version in a timely manner.
- For more information about the security maintenance methods, see the product documentation of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network situations to detect public and internal network traffic and protect the network against attacks and unusual activities.

10.5.3.2. Account password maintenance

Account passwords include the Graph Analytics system password and the device password.

To ensure account security, you must change the system and device passwords periodically, and use passwords with high complexity.

10.5.4. Troubleshooting

10.5.4.1. Fault response mechanism

The IT administrator must establish a fault emergency response mechanism, so that the service can be recovered quickly after a fault or security accident occurs.

10.5.4.2. Troubleshooting methods

After a system fault is detected during routine maintenance, the IT administrator can read the Operations and Maintenance part of this documentation for reference.

If the fault cannot be fixed, collect the fault information, including the system information and fault symptoms, contact Alibaba Cloud technical support engineers, and troubleshoot the fault under the guidance of the engineers.

After the fault is fixed, the IT administrator must analyze the causes, review the troubleshooting process, and make improvements.

10.5.4.3. Common failure troubleshooting

Insufficient disk space

Possible cause: The log size in the Graph Analytics system is too large.

Solution: Monitoring logs are usually stored in the /home/admin/logs directory. You can delete earlier logs to free up space.

Machine maintenance or downtime

Possible cause: The hardware is damaged or the warranty of the machine is expired.

Solution: Reinstall Graph Analytics.

Suspicious processes

Possible cause: If the process fails to start automatically or is terminated unexpectedly, view the logs in /home/admin/logs to identify the cause.

Solution: Restart Graph Analytics.

10.5.4.4. Hardware troubleshooting

Disk failure

Solution: Graph Analytics supports cluster deployment. Therefore, you can directly end all Graph Analytics threads, replace the hard drive, and then start the threads again.

Failures requiring server shutdown, including memory, MPU, CPU, and power supply failures

Solution:

Repairs involving server shutdown:

- If you can access the system, you can follow the service stop procedure to disable the Graph Analytics service on the server.
- If you cannot access the system, you must force the server to shut down.

10.6. Machine Learning Platform for AI

10.6.1. Query server and application information

10.6.1.1. Apsara Stack Machine Learning Platform for AI

10.6.1.1.1. Query server information

Machine Learning Platform for AI is deployed based on Apsara Infrastructure Management Framework. Its application information and database information can be found by accessing the corresponding Apsara Infrastructure Management Framework address. This topic describes how to query server information.

Procedure

1. Open Chrome and ensure that you can access internal services through the network.

2. Enter the username and password to log on to the homepage of Apsara Infrastructure Management Framework.

Notice To avoid logon failures, make sure that your network is connected and the hosts have been bound.

- 3. Click the C and search for pai. Hover over the dots next to PaiCluster-20170630-c34b, and choose Dashboard from the shortcut menu.
- 4. Query the server information for an application, such as the server where PaiDmscloud runs.
 - i. Find the service instance and click Details. The instance detail page appears.
 - ii. Find the role list and click **Details**. The role detail page appears.
 - iii. The IP address of the server is displayed in the server information list. You can click **Terminal** to manage the server on the terminal management page.

10.6.1.1.2. Log on to a server

Machine Learning Platform for AI is deployed based on Apsara Infrastructure Management Framework. Its application information and database information can be found by accessing the corresponding Apsara Infrastructure Management Framework address. This topic describes how to log on to a server.

Context

Each module is deployed on two servers with the same application package and configuration. You can log on to the back-end server through the server IP address and perform operations.

Procedure

- 1. Ensure that the network is connected and the IP address of the jump server has been obtained.
- 2. Log on to the jump server.
- 3. Switch to the root account.
- 4. All applications are deployed by using a Docker container. You can run the following command to view the current container:

sudo docker ps

5. Run the following command to go to the container:

sudo docker exec -ti container_id /bin/bash

The application log is stored in the */home/admin/logs/\${app}* path.

10.6.1.1.3. Query configurations

Prerequisites

Log on to the server of an application and go to the application container to view the configuration of the application.

Procedure

1. View the application configuration in the */home/admin/{app}/target/exploded/BOOTINF/cl asses/application.yml* file.

Note In the preceding file path, {app} indicates the component name, such as paidms.

2. View the application log in the */home/admin/pai-dms/* path.

The pai-dms.log, err_pai-dms.log, java.log, and access.log files store the application log, error log, framework log, and access log, respectively.

- 3. Log on to a database.
 - i. Query the database information of modules from the Dashboard cluster information of Apsara Infrastructure Management Framework. Find the corresponding result column and click More from the shortcut menu to obtain db_host, db_port, db_name, db_password, and db_user of the application.
 - ii. Run the following command to connect to the database through a MySQL client:

mysql -h\$db_host -P\$db_port -u\$db_user -p\$ db_password -D\$ db_name

10.6.1.1.4. Restart an application service

The application structures and directories of the PaiCap, PaiDmscloud, and PaiJcs modules are almost the same. You can restart an application service in either of the following ways:

• Log on to the container and run the following command to restart the service:

sudo -u admin /home/admin/pai-dms/bin/appclt.sh restart

• Run the following command on the server to restart the container:

sudo docker restart \$container_id

Run the following command to check whether the service is restarted:

curl localhost/status.taobao

10.6.1.2. Online model service

10.6.1.2.1. Query online model service information

Check the online model service status

Online model services are deployed in the Kubernetes cluster. Log on to the master node in the Kubernetes cluster and run the following command to query the service deployment status:

kubectl get pod -n eas-system

If no errors occur, all pods in the STATUS column display *Running*.

If not, run the following command to perform troubleshooting:

kubectl describe pod \${pod_name} -n eas-system

View the online model service configurations

- 1. Log on to the homepage of Apsara Infrastructure Management Framework.
- 2. Click the C tab and search for pai. Hover over the dots next to the PAI cluster, and choose Dashboard from the shortcut menu.
- 3. Search for the *eas-sentinel* role and log on to the VM from the terminal.
- 4. Run the docker ps |grep eas-sentinel command to view the ID of the container for the sentinel.
- 5. Run the docker logs \${sentinelcontainerid} command to view the output log, which contains the configuration information of the online model service.

10.6.1.2.2. Log on to the online model service container

Prerequisites

Ensure that the network is connected and the IP address of the jump server has been obtained.

Procedure

- 1. Log on to the jump server.
- 2. Switch to the root account.
- 3. All applications are deployed with a container. Run the following command to log on to the current pod: kubectl exec -ti \${pod_name} -n \${pod_namespace} bash

10.6.1.2.3. Restart a pod

Procedure

- 1. Log on to the master node in the Kubernetes cluster.
- 2. Run the kubectl get command to find the corresponding *pod name*.
- 3. Run the following command to restart the pod: kubectl delete \${pod_name}

10.6.1.3. GPU cluster and task information

10.6.1.3.1. Query GPU cluster information

Prerequisites

You must deploy the deep learning service before querying the GPU cluster information. Deep learning tasks are performed in the GPU cluster. You can log on to ApsaraAG of the GPU cluster to query the GPU cluster status.

Procedure

- 1. Log on to the homepage of Apsara Infrastructure Management Framework.
- 2. Click the C tab and search for *PAIGPU*. Move the pointer over the dots next to the deployed GPU cluster. Log on to the cluster O&M center.
- 3. Select *pai-deep_learning* from the Service drop-down list and *ApsaraAG#* from the Service Role drop-down list. Log on to the VM from the terminal.

4. Run the r ttrl command to view all GPU workers in the current GPU cluster.

If the Other column displays FUXI_GPU:200, the worker has two GPUs. If the column displays FUXI_GPU:800, the worker has eight GPUs.

10.6.1.3.2. Query GPU task information

Procedure

- 1. Perform steps 1 through 3 in Query GPU cluster information and log on to ApsaraAG of the GPU cluster.
- 2. Run the r al command to view the running tasks.
- 3. Run the r wwl WorkItemName
 command to view the status of a task and the allocated

 resources.
 WorkItemName
 : specifies the values in the first column displayed by the r al command.
- 4. Run the r cru command to view the resources allocated to the current cluster, including CPU, memory, and FUXI_GPU resources.
- 5. 🗘 Notice Use caution when performing this step.

Run the r jstop WorkItemName command to stop a Fuxi task. WorkItemName : specifies the values in the first column displayed by the r at command.

10.6.2. Maintenance and troubleshooting

10.6.2.1. Machine Learning Platform for AI maintenance

10.6.2.1.1. Run ServiceTest

After ServiceTest is run, the automated test case is executed.

- Log on to the homepage of Apsara Infrastructure Management Framework and choose Tasks
 > Deployment Summary from the top navigation bar. The Deployment Summary page
 appears.
- 2. On the **Deployment Summary** page, click **Deployment Details**. The Deployment Details page appears.
- 3. Move the pointer over the row in which the project name is PAI. Click **Details**, and click **ServiceTest#** to go to the server list page.
- 4. On the machine learning list page, click Terminal to access TerminalService.
- 5. Run the sudo docker ps -a command to find the ServiceTest instance of PAI, as shown in the following figure.

ServiceTest instance

pai	Final 21 Hours 19 Minutes	Cluster: 4 / 4 Service: 18	/ 18 Role: 23 / 23	Total: 21 Done:	21 0	0	*
officer .	Final 21 Hours 20 Minutes	C 🚠 AlgoMarketClust ⊘	👒 bigdata-sre 🔗	♣ PaiAlgoinit#	0	0	
rgántási	Final 21 Hours 20 Minutes	C AlinkCluster-A-2 ⊘ L EASCluster-A-20⊘	≪ os ⊘ ≪ pai-pai_service ⊘		0	0	
ram.	Final 21 Hours 20 Minutes	C 🚠 PaiCluster-A-20 🔗	📽 tianji 🔗	♣ PaiFront#	0	0	25
10	Final 1 Hour 7 Minutes	c	≪ tianji-dockerdae ⊘		0	0	26
101	Final 21 Hours 20 Minutes	c			0	0	24
10	Final 21 Hours 18 Minutes	c			0	0	*
els:	Final 11 Hours 48 Minutes	C			0	0	

6. Run the sudo docker restart e90f70353031 command to restart the ServiceTest service, as shown in the following figure.

 State
 PORTS
 NAMES
 COMMAND
 CREATED

 1997/0333311
 Lnc.com/ldst-pa1/pa1-web-test:tobl3083230eebc495/5186835eeta1
 bbc97
 "sh /usr/local/smokin"
 10 days ag

The test case is executed when the service_test service is restarted. After the execution, you can view the log information.

- 7. Run the sudo docker logs e90f70353031 --tail 1000 command to view the log. Only the last 1,000 rows are displayed.
- 8. After the test case is executed, the testing results for all algorithms are displayed, as shown in the following figure.



- PASS: The algorithm is running properly.
- SKIP or FAIL: The algorithm fails.

10.6.2.1.2. Common faults and solutions

10.6.2.2. Online model service maintenance (must be

activated separately)

Node maintenance

Online model service nodes are Kubernetes nodes. You can run the kubectl get node command to view all nodes in a cluster. A healthy node is in the Ready state. When a node is not in the Ready state, the one of the following errors may have occurred:

• Node failures

There are many reasons that may cause a node to fail. Typically, a node fails when the kernel crashes or the disk does not have sufficient space. If the node can be restarted properly, it rejoins the cluster after it is restarted. If the node cannot be restarted properly, contact the corresponding ECS support personnel.

• Docker daemon exceptions

A Docker daemon exception rarely occurs. Docker daemon exceptions are typically caused by storage issues. Run the systemctl restart docker command to restart the Docker daemon.

Online model service maintenance

- A service cannot be created or deleted.
 - If Error 500 is returned while an operation is called, the configurations of the eas-ui component are incorrect. Contact Apsara Stack delivery engineers.
 - If a creation or deletion operation is called but no response is returned in a timely manner, the jobworker of the service does not work properly. Check whether the KVStore for Redis service in the cluster is normal. If not, restart the pod for KVStore for Redis.
- The system fails to read the monitoring data.

Check whether the influxdb-0 pod under *eas-system* is created properly. If the pod is not in the running state, an influxdb out of memory error has occurred. You can expand the influxdb-0 memory.

Service maintenance

• Service creation failures.

The request is sent but the service creation result displays Failed. A model error has caused a crash. The system then fails to create the model. Check whether the model code contains any null pointers or has any other problems.

• The system fails to obtain the monitoring data.

Check whether the influxdb-0 of each service is normal. The service cannot be created because a persistent volume cannot be created. Check whether the Apsara Stack environment has sufficient disk space. If influxdb-0 runs properly but you cannot obtain the monitoring data, restart the influxdb-0 pod.

10.6.2.3. GPU cluster maintenance (deep learning must be

activated separately)

Node maintenance

A deep learning node is a server where a GPU cluster runs.

- 1. Perform steps 1 through 3 in Query GPU cluster information and log on to ApsaraAG of the GPU cluster.
- 2. Run the r ttrl command to view all nodes that support deep learning tasks.
- Node failures

There are many reasons that may cause a node to fail. Typically, a node fails when the kernel crashes or the disk does not have sufficient space. If the node can be restarted properly, it rejoins the cluster after it is restarted. If the node cannot be restarted properly, contact the corresponding service support team.

• Docker daemon exceptions

A Docker daemon exception rarely occurs. Docker daemon exceptions are typically caused by storage issues. Run the systemctl restart docker command to restart the Docker daemon.

Service maintenance

Failure to allocate resources to a task

Perform the following steps for troubleshooting:

- 1. Perform steps 1 through 3 in Query GPU cluster information and log on to ApsaraAG of the GPU cluster.
- 2. Run the r quota command to view the quota information of the GPU cluster.
- 3. Run the r cru command to view the resources allocated to each task in the current cluster.
- 4. Run the ral command to view all tasks submitted to the cluster.
- 5. Run the r wwl WorkItemName command to view the status of a specific task.
 - If only ChildMaster is displayed, no resources are allocated to the worker.
 - If worker name is displayed but no hostname is displayed, service resuming is pending or has failed. Log on to the server of the ChildMaster and locate the error. You can also contact the service support team.
- 6. Run the r ttrl command to check the value of FUXI_GPU in the Other column. If the value is 200, the worker has two GPUs. If the value is 800, the worker has eight GPUs.
- 7. Log on to a GPU worker in the worker list obtained in Step 3 over SSH. Run the nvidia-smi command to view the GPU status. If an exception occurs, contact the relevant service support personnel.

10.7. E-MapReduce (EMR)

10.7.1. Methods for logging on to O&M platforms

10.7.1.1. Log on to Apsara Infrastructure Management

Framework

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

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

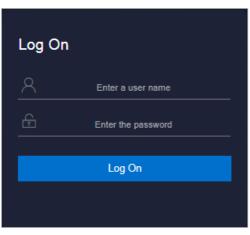
• A browser is available. We recommend that you use the Google Chrome browser.

Context

Apsara Infrastructure Management Framework is a cluster monitoring and management tool. It displays basic cluster information as well as machine and server monitoring information. It monitors system load, CPU, memory, disk, and transmission metrics and the status of service instances. This helps you promptly detect exceptions and take actions to troubleshoot the exceptions. After you log on to Apsara Infrastructure Management Framework, you can perform command-line operations on cluster machines.

Procedure

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



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

3. Enter your username and password.

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

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

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

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

🐟 Tian Ji 💦 🗧	Home Operations - Tasks -	Reports - Management - Monitoring -		Ø 10:33 English(US) -
I ALL -		uster Operations and Maintenance Portal has been launched to :	support operation and maintenance functions and task tra	
Musters I	Upgrade Task Summary	RelingBack 0.00% 0 0		Most-used Reports Registration Vars of Services IF Lat Thermometer XDB Instance Metric Info Relationship of Service Dependency Machine Power On or Off Statuses of Clusters
Scarch Q inc 5 sction requires <u>double-</u> Double-click a cluster te and a service	Cluster S Top 3		Machines Error Alerts OS Errors HW Errors	Error Summary Rate of Abnormal Machines: 0.00% Machine
tively, and check e logs. Choose all machine ons from the <u>menu</u> . et started. te data switching	0 tianji-paas-A-18.	BigdstaSaaS-A-20	mook-aluster-for	Alerts: 2.16% Abnormat: 16 Totat 741

10.7.2. Routine maintenance

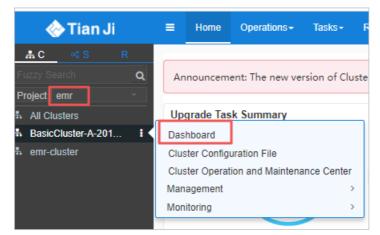
10.7.2.1. O&M on Apsara Infrastructure Management

Framework

This topic describes how to perform O&M on Apsara Infrastructure Management Framework.

Procedure

- 1. Log on to the Apsara Infrastructure Management Framework console.
- 2. In the left-side navigation pane, click the C tab and select emr from the Project drop-down list.



- 3. Move the pointer over the **i** icon next to an EMR cluster and select **Dashboard** to go to the Cluster Dashboard page.
- 4. In the Service Instances section, click Details in the Actions column that corresponds to emrservice to go to the Service Instance Information Dashboard page.

Service Instances					0 /
Service Instance	Final Status	Expected Server Roles	Server Roles In Final Status	Server Roles Going Offline	Actions
emr-service	True	8	8	0	Actions - Details
os	True		-		Actions - Details
tianji	True	1	1	0	Actions - Details
tianji-dockerdaemon	True	1	1	0	Actions - Details

5. In the Server Role List section, click **Details** in the Actions column that corresponds to the server role you want to restart to go to the Server Role Dashboard page.

Server Role List							e 2
Server Role	Current Status	Expected Machines	Machines In Final Status	Machines Going Offline	Rolling Task Status	Time Used	Actions
EcmAdmin#	In Final Status	2	2	0	no rolling		Details
EcmServer#	In Final Status	2	2	0	no rolling		Details
EmrConsoleAliyunCom#	In Final Status	2	2	0	no rolling		Details
EmrDbInit#	In Final Status	1	1	0	no rolling		Details
EmrMainversionInit#	In Final Status	1	1	0	no rolling		Details
EmrMiddlewareInit#	In Final Status	1	1	0	no rolling		Details
EmrPopserver#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

You can restart every role in this list.

6. Click the downward arrow next to Operation Menu and select **Restart Server Role**. In the message that appears, click OK.

Server Role Dashboard	Operat	ions Menu 🝷		
	Rest	art Server Role		
Server Role Summary				0 Z
Title			Value	
Project Name		emr		
Cluster Name		Basic		
Service Instance		emr-service		
Server Role		EcmAdmin#		
In Final Status		In Final Status		
Expected Machines		2		
Actual Machines		2		
Machines Not Good		0		
Machines with Role Status No	ot Good	0		

10.7.3. Troubleshooting

10.7.3.1. Troubleshooting methods

If you detect a system fault during routine maintenance, read the Routine Maintenance part of this documentation for reference.

If you fail to rectify the fault, collect related information such as system information and fault symptoms and contact Alibaba Cloud technical support for help.

After you rectify a fault, analyze its causes, review the troubleshooting process, and make improvements.

11.Apsara Asapi Management system 11.1. Apsara Asapi Management system overview

This topic describes the features and infrastructure of the Apsara Asapi Management system (Asapi).

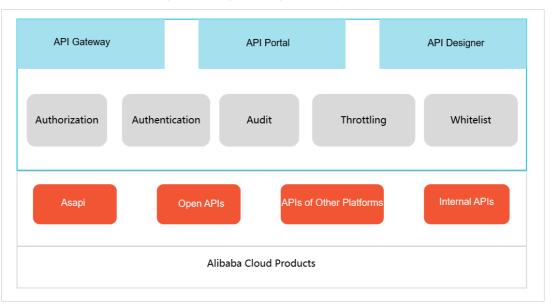
The Apsara Asapi Management system is a platform that manages operations and maintenance APIs and SDKs in the Apsara Stack environment in a unified manner. This system also manages API and SDK versions.

Most Apsara Stack products provide APIs for tenants. Only a few Apsara Stack products provide APIs for operations and maintenance. To address the business needs at the operations and maintenance side and meet custom development requirements of users such as developing their own operations and maintenance consoles or obtaining operations and maintenance data, Alibaba Cloud provides the Apsara Asapi Management system.

Features of the Apsara Asapi Management system are as follows:

- Provides APIs at the system level and typical APIs for resource usage, monitoring, and alerting.
- Manages APIs, including querying, editing, testing, and removing information about APIs.
- Provides an API designer to customize API flows based on existing APIs, which facilitates custom business.
- Manages versions and relationships between these versions. These versions include Apsara Stack versions, product versions, SDK versions, and API versions.
- Supports SDKs. The Apsara Asapi Management system provides SDKs for Java and Python to call operations and maintenance APIs.

Basic architecture of the Apsara Asapi Management system shows the basic architecture of the Apsara Asapi Management system.



Basic architecture of the Apsara Asapi Management system

The Apsara Asapi Management system contains the following components:

- api-server: contains operations and maintenance APIs and provides the APIs for SDKs so that SDKs can be used to call the APIs.
- API Portal: the operations and maintenance console to manage Asapis.
- api-node: the API designer.

11.2. Log on to the Apsara Asapi Management platform

You can log on to the Apsara Asapi Management platform through Apsara Infrastructure Management Framework.

Prerequisites

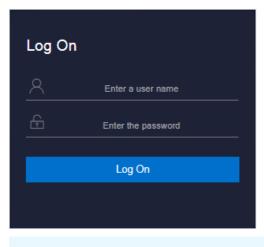
• The URL of the ASO console, and the username and password used for logging on to the console are obtained from the deployment personnel or an administrator.

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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



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

2. Enter your username and password.

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

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

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

- It must contain uppercase and lowercase letters.
- It must contain digits.
- It must contain special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- It must be 10 to 20 characters in length.
- 3. Click Log On to go to the ASO console.
- 4. In the left-side navigation pane of the Apsara Stack Operations console, click **Products**. On the Product List page that appears, click **Apsara Infrastructure Management Framework**.
- 5. In the left-side navigation pane of the Apsara Infrastructure Management Framework console, click **Reports**.
- 6. On the All Reports page that appears, search for Registration Vars of Services. Click Registration Vars of Services.

🐟 Tian Ji		Reports - Management - Mon
🛔 C 👒 S R		
Fuzzy Search Q	All Reports Favorites	
Project ALL	Fuzzy Search Registration Vars of Services	Q
All Clusters		
AcsControlCluster-A-202	Report 🗹	Group 🗹 Status
🚠 ads-A-20200416-661b		
AlgoMarketCluster-A-202	Registration Vars of Services	Tianji OPublished
AliguardCluster-A-20200		
🚓 amtest70		
A ansCluster-A-20200416-6		
👬 asaCluster-A-20200416-6 🚦		

7. On the Registration Vars of Services page that appears, click the 📃 icon next to Service. In

the dialog box that appears, search for asapi from the drop-down list.

- 8. Right-click the Service Registration column corresponding to asapi. Select Show More from the drop-down list.
- 9. In the **Details** dialog box that appears, the value of **asapi.gateway.endpoint** is the logon URL of the Apsara Asapi Management platform.

11.3. Manage APIs

The Apsara Asapi Management system provides APIs of various products in the Apsara Stack environment. You can upload, query, edit, test, and delete APIs. You can also use the API designer to customize APIs.

11.3.1. Register APIs

An API can be defined in an XML file. Each API corresponds to one XML file. You can upload an XML file to register an API in the Apsara Asapi Management system.

Context

The following fields must be defined in an XML file:

- API name
- namespace (or product name)
- API type
- Parameters

When you upload APIs, you can upload one or more XML files simultaneously. You must compress the XML files into a ZIP file before you upload these files.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.
- 3. In the upper-right corner, click **Upload API**. Select the XML or ZIP file to be uploaded. After you upload the XML or ZIP file, you can view the uploaded APIs in the API list on the **APIs** page.

11.3.2. Modify information about APIs

You can modify basic information and specific parameters of an API when the API is modified.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.
- 3. (Optional)Select a product from the Select Product drop-down list. Enter an API name in the search box.

Enter a full or partial API name to search for APIs.

4. Click the ∠ icon in the Actions column corresponding to the API whose information is to be modified.

Edit API:GabDescribeE	csDisk			
Basic Information	Edit Parameters			
		API Name:		
		Type:	1000	V
		Namespace:	apaliti	
		Endpoint:		
			Save	

5. In the Edit API dialog box that appears, click the **Basic Information** tab, modify the basic information, and click **Save**.

Parameter	Description	
API Name	The name of the API.	
Туре	 The type of the API. Different products have different API types. The types are as follows: asAPI: the API for operations and maintenance OpenAPI: the API for product operations customAPI: the API customized through the API designer 	
namespace	The namespace of the API. It corresponds to the product name.	
Endpoint	The domain name of the product that corresponds to the API.	

The following table describes the parameters of an API.

- 6. Click the Edit Parameters tab to modify configuration information such as the request parameters, response parameters, and error handling mechanism. Follow these steps: Set the request parameters in the Parameters section, response parameters in the ResultMapp ing section, and error handling mechanism in the ErrorMapping section.
- 7. After the modification is complete, click Save.

11.3.3. Test APIs

The Apsara Asapi Management system allows you to test APIs online to check whether an API is available. During the test, you can save input parameters as a test case for subsequent execution.

Method 1

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.
- 3. On the APIs page, click the 🔚 icon in the Actions column corresponding to the API to be

tested.

4. In the Test API dialog box, set Request Parameters.

Request parameters may vary with APIs. The following table describes the typical request parameters.

Parameter	Required	Description
regionId	Yes	The region ID of the test environment.
accessKeyld and accessKeySecret	Yes	The identification of the visitor. You can obtain them from the Apsara Stack console.

5. After request parameters are configured, click Send.

The Apsara Asapi Management system sends a test request to the configured domain name. The response appears in the **Responses** section.

 (Optional)After the test is complete, click Save As Test Case for subsequent execution of this test case. You can do this by choosing Testing Platform > Test Cases on the Test Cases page.

Method 2

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.
- 3. Select the product to be tested and search for the product API.
- 4. Configure request parameters of the API to be tested.
- 5. After the request parameters are configured, click Run. The Apsara Asapi Management system sends a test request to the configured domain name. The response appears in the Responses section.
- 6. After the test is complete, click **Save**. You can repeat the test case on the right of the Test Cases page.

11.3.4. Remove information about APIs

You can remove information about an API that you no longer need.

Procedure

> Document Version:20200918

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.
- 3. (Optional)Select a product from the Select Product drop-down list. Enter an API name in the search box.

Enter a full or partial API name to search for APIs.

- 4. Click the 🕆 icon in the Actions column corresponding to the API to be removed.
- 5. In the message that appears, click **OK**.

11.3.5. API design

The Apsara Asapi Management system provides an API designer to help you customize APIs.

11.3.5.1. Designers

This topic describes API designers.

If Apsara Stack Asapis do not match the APIs you are using, or if you need to customize APIs to meet the requirements of specific projects, you can use an API designer to assemble and create desired APIs in the flow design process.

The API designer is built based on the open-source project Node-RED. Node-RED is a powerful tool that is launched by IBM to build Internet of Things (IoT) applications. It uses the visual programming method that allows developers to connect predefined code blocks (nodes) to perform tasks. Connected nodes are a combination of input nodes, processing nodes, and output nodes. When they are connected to form a flow, they are able to process requests such as HTTP requests.

Node-RED is highly capable of customizing flows and processing HTTP messages. These capabilities can be easily expanded.

To design an API, follow these steps:

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose API Platform > API Design. The API Designer page appears.
- 3. Drag and drop the components on the left to the flow chart section. You can combine these components to complete a specific flow chart and design an API.

11.3.5.2. Designer nodes

This topic describes typical nodes that are used in the designer.

To customize Asapis, the Apsara Asapi Management system adds some nodes through the mechanism provided by Node-RED.

Typical nodes are described as follows:

- api-request: used to create a request to access an Asapi. There is a small icon before each node. When you click the icon, a request is sent to the flow that contains the node.
- api response: used to provide responses and format the returned data.
- api selector: used to select and execute an existing API.

- db exec: used to execute a specified SQL operation.
- new api: used to create an API that contains a specified endpoint and specific input parameters.
- sync msg: used to merge multiple responses into one response and send the response.
- py function: Python is used in some modules during the API design process.
- Input components: the detailed operations involved in a request process. For example, set the request protocol types such as HTTP, TCP, and UDP, status code of the request, and the created request link.
- Output components: the returned data, status code, and protocol in the response. Output components are used to describe fixed output modes such as request and response methods and formats of returned data.

11.3.5.3. Design an API flow

This topic describes how to design an API flow.

Each customized API has its own API flow. Each API flow consists of multiple connected nodes, including one input node, several processing nodes, and one return node (or output node).

Among the nodes:

- Typical input nodes are api request and http in.
- Typical return nodes are api response and http response.
- Typical processing nodes are function, api selector, and db exec. The function node is used to convert parameters and process simple logic.

An API flow is designed as follows:

- 1. Select an input node and an output node, and add processing nodes to the flow.
- 2. Define the name and configurations of each node, such as the endpoint of the input node.
- 3. Connect the nodes as needed to form a flow.
- 4. In the upper-right corner of API Designer, click **Deploy** to publish the flow.
- 5. Access this flow in the browser. You can obtain the response.

11.4. Version management

11.4.1. Apsara Stack version management

Apsara Stack has multiple versions that vary with projects.

11.4.1.1. Add information about versions

You can add information about Apsara Stack versions as needed to manage the relationships among Apsara Stack versions, products, and product versions.

Context

Each Apsara Stack version can have either one release version or one snapshot of the onpremises environment or deployment environment. It can be distinguished by its version name and description.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Apsara Stack Versions, The Apsara Stack Versions page appears.
- 3. In the upper-right corner of the page, click Add Version.
- 4. In the Add Version dialog box that appears, set Apsara Stack Version, Version, and Release Notes.

We recommend that you enter information about the current version for Release Notes.

5. Click Submit.

11.4.1.2. Select products for an Apsara Stack version

After adding information about an Apsara Stack version, you can select products that are supported in the Apsara Stack version based on version output conditions.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Apsara Stack Versions. The Apsara Stack Versions page appears.
- 3. In the version list, click **Configure Products** in the Product column corresponding to an Apsara Stack version.
- 4. In the Configure Products dialog box that appears, select a version from the Version dropdown list and select the check box in the **Output** column corresponding to the product.

Configure Apsara Stack Products: 3.3			Х
Product	Version	Output	*
Ecs	3.1.0 ~		
Rds	3.1.0 ~		
Oss	2.4.2 V		
Vpc	3.4.0 V		
SIb	3.4.0 V		.
	submit		

5. Click Submit to generate information of the products of the specified Apsara Stack version.

11.4.1.3. Compare versions

You can use the version comparison function to compare the product differences between two Apsara Stack versions. Based on these product differences, you can further learn about the differences of their APIs as well as of the definitions and parameters of these APIs.

Context

- Apsara Stack version: Each Apsara Stack version can have either one release version or one snapshot of the on-premises environment or deployment environment. Versions are distinguished by version names and descriptions.
- Product version: the specific version of a product when each Apsara Stack version is released, such as RDS 3.7.0. An Apsara Stack release version can have only one version of a specific product.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Apsara Stack Versions. The Apsara Stack Versions page appears.
- 3. In the upper-left corner of the page, click Compare Versions.
- 4. On the Select Version tab, select two versions to be compared. Click Next.

Compare Apsara Stack Versions			×
1 Select Version 2 Vers	sion Differences —	- 3 Product Differences -	(4) Compare API Versions
	Select Source Version	Select Target Version v3.3	
	v3.3		
	v3.4	v3.4	
	v3.5	v3.5	
		v3.6	
	v3.6	v3.7	
	v3.7	v3.8	
	v3.8	V5.0	
			Next

5. On the Version Difference tab, you can compare the product differences between the two versions. For example, you can view versions for which product information has been added or removed.

6. Click a product. Click Next to go to the Product Difference tab. You can compare the differences in product APIs between these two versions.You can view functions for which APIs are added, removed, or remain unchanged in these two versions.

ompare Apsara Stack Vers	ions					
Select Version	— 🕢 Version Difference	s — 🕻	Produc	t Differences	- (4	Compare API Versi
	rsion: N/A,Target Version: 80 eted:0 .,APIs Changed:0 .,A		ed:0 .			
Function Name	Source Version	Target V	ersion	Status		Mark
CountCloudInstances				2016-04- 28	new	• Added
DescribeNetworkQuotas				2016-04- 28	new	• Added
GaFillParams			2016-04- 28	new	• Added	
GaFillProduct			2016-04- 28	new	• Added	
GaNotifyPaid				2016-04- 28	new	• Added
						Previous Ne

7. Click an API. Click Next to view the changes that are made to this API.

Select Version ——— 🔗 Version Differences –	— 🔗 Product Differences — 4 Compare API Versio
aFillParamsCompare API Versions: 与 2016-04-28	
1	<pre>1 k?xml version="1.0" encoding="UTF-8"?></pre>
	2
	3 <api `<="" data"="" p="" tagname="data" version="2016-04</p></td></tr><tr><td></td><td>4 <Parameters></td></tr><tr><td></td><td>5 <Parameter name=" visibility="Private"></api>
	6 <parameter name="requestId" p="" stsaccesskeyid"="" tagname="R</p></td></tr><tr><td></td><td>7 <Parameter name=" tagnau<=""></parameter>
	8 <parameter <="" apiname"="" callerbid"="" name="apiName" p="" tagname="callerBid"></parameter>
	10 <parameter calleruidloginemail"<="" name="calleroid" ownerid"="" tagname="Own</td></tr><tr><td></td><td>12 <Parameter name=" td=""></parameter>
	13 <parameter name="callerBidLoginEmail" td="" ·<=""></parameter>
	14 <parameter name="ownerIdLoginEmail" ta:<="" td=""></parameter>
	15 <parameter <="" name="resourceOwnerAccount" td=""></parameter>
	<pre>16 <parameter name="resourceOwnerId" pre="" tagn;<=""></parameter></pre>
	17 <parameter <="" enable"="" name="clientIP" pre="" tagname="enable"></parameter>
	19 <parameter name="requestContent" tagna<="" td=""></parameter>
	20 <parameter <="" name="token" tagname="Clien" td=""></parameter>
	21 <parameter name="ownerAccount" p="" tagname<=""></parameter>
	22

11.4.1.4. Remove information about Apsara Stack versions

If you no longer need an Apsara Stack version, you can remove its version and output information.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Apsara Stack Versions page appears.
- 3. Click the 🗇 icon in the Actions column corresponding to the Apsara Stack version to be removed.
- 4. In the message that appears, click OK.

11.4.2. Product baseline

Product baselines are a set of configurations used by Apsara Stack products to define products, services, service roles, and applications. The Apsara Asapi Management system provides basic information about products, services, and service roles. During initialization, the Apsara Asapi Management system automatically scans baseline information of all products in the Apsara Stack environment. You can use the system to scan the metadatabases and servers of services and service roles.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. Follow these steps:
 - In the upper-right corner of the page, click Scan Apsara Stack Environment to scan metadatabases and servers that correspond to all products and update their information in the system.
 - Select a product from the drop-down list to query the services and service roles of the product.
 - Click the local in the Actions column corresponding to a service role to scan the metadatabases of the service role.
 - Click the 📇 icon in the Actions column corresponding to a service role to scan the server of the service role.

11.4.3. Products

Operations and maintenance engineers can manage information of current Apsara Stack versions and product versions in real time.

Context

• Apsara Stack version and product version: Each Apsara Stack version can have only one specific version of products.

• Product version and SDK version: Each product version can have one SDK version.

11.4.3.1. Add information about products

You can add information about a product you need to manage.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. In the upper-right corner of the page, click Add Product.
- 4. In the Add Prdouct dialog box that appears, set Product Name and Product Description.
- 5. Click Submit to add information about a product.

11.4.3.2. Add information about product versions

After you add information about a product, you need to add its product version and API version information for subsequent version management.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. On the Products page, click the 🗈 icon in the Actions column corresponding to the product of which the version information is to be added.
- 4. In the Add Version dialog box that appears, set Version and API Version.

Parameter	Description
Version	The version of the current product.
API Version	The API version of the current product.

5. Click Submit.

11.4.3.3. Import information about APIs

You can import information about a preset API to the Apsara Asapi Management system.

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. In the product list, click the 🗄 icon in the Actions column corresponding to the product to be managed.
- 4. In the dialog box that appears, select the product version from the left drop-down list.

If information about the API has been imported for the product version, this API is displayed in the APIs section.

- 5. In the upper-right corner, click Import API.
- 6. In the API Import dialog box that appears, set **Apsara Stack Version** and **API File** for the API to be imported.

API Import		×
Proprietary cloud v	select the propriet \vee	
API File:	select files	

7. Click OK to import the API to the system.

11.4.3.4. Set SDK versions

A product has multiple SDK versions. You can set the SDK version of a product to obtain the SDK version of the product in the Apsara Stack release version.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. In the product list, click the (a) icon in the Actions column corresponding to the product to be modified.
- 4. In the SDK Settings dialog box that appears, click Modify in the Actions column corresponding to the product version.

SDK settings		×
Product Version	SDK Version	operate
3.6		modify
3.5		modify
		< 1 >

5. Select the specified SDK version from the drop-down list. Click **Submit** in the **Actions** column corresponding to the SDK version. The SDK version is modified.

11.4.3.5. Modify product names and descriptions

You can modify the name and description of a product.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. On the Products page, click the ∠ icon in the Actions column corresponding to the product to be modified.
- 4. In the dialog box that appears, modify the product name or description. Click Submit.

11.4.3.6. View information about product versions

When you need to learn about how to use a product, you can view information about the product version and API version.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. On the Products page, click the *Q* icon in the Actions column corresponding to the target product. You can view information about the product version and API version.

11.4.3.7. Modify information about product versions

You can modify information about a product version or API version as needed.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. On the **Products** page, click the Q icon in the **Actions** column corresponding to the product about which the version information is to be modified.
- 4. In the View Version dialog box that appears, click the ∠ icon in the Actions column corresponding to the version to be modified.
- 5. In the dialog box that appears, modify information about the product version and API version.

11.4.3.8. Remove information about product versions

You can remove information about a product version that is not applicable.

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. On the **Products** page, click the Q icon in the **Actions** column corresponding to the product

whose version information is to be removed.

- 4. In the View Version dialog box that appears, click the 🕆 icon in the Actions column corresponding to the version to be removed.
- 5. In the message that appears, click OK.

11.4.3.9. Delete products

You can delete products that you no longer need.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. Click the 📋 icon in the Actions column corresponding to the product to be deleted.

11.4.3.10. Delete product APIs

You can delete APIs that are not applicable to a product.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
- 3. In the product list, click the ∺ icon in the Actions column corresponding to the product to be managed.
- 4. In the dialog box that appears, select the product version from the left drop-down list. Imported APIs are displayed in the APIs section.
- 5. Click the 📋 icon corresponding to the API to be deleted.
- 6. In the message that appears, click OK.

11.4.4. SDKs

The Apsara Asapi Management system enables you to customize SDKs. You can customize an SDK as needed to export APIs of Apsara Stack products of a specific version. You can also modify and delete the customized SDKs.

11.4.4.1. Customize SDKs

The Apsara Asapi Management system provides a tool to customize SDKs. The tool enables you to customize multiple combinations of SDKs for APIs within and across Apsara Stack products of specified versions.

Context

Each product has corresponding SDKs for different programming languages. The Apsara Asapi Management system supports only SDKs for Java and Python.

Each SDK consists of an SDK core and an SDK model. The SDK core is the framework of the SDK. It is used to generate HTTP requests or requests of other protocols. The SDK core is fixed. You do not need to generate it each time. The SDK model defines the request parameters and responses of each API.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > SDK Tools. The SDK Tools page appears.
- 3. In the upper-right corner of the page, click **Customize SDK**.
- 4. Set Apsara Stack Version, Product Name, Product Version, SDK Version, API Version, and Language. The corresponding APIs are displayed in the following APIs section.
- 5. Select APIs and click Create SDK.

After an SDK is created, you can view the SDK in the SDK list on the SDK Tools page.

6. (Optional)Click the link in the Download column corresponding to the product to download the SDK.

Customize SDK					
Product Name	Language	Apsara Stack Version	Generated At	Download	Actions
Drds	Java	v3.4	Jun 1, 2020, 10:03:26	drds-java-sdk_2020-06-01_100315.zip	∠ 0

(?) Note The SDK generated in the Apsara Asapi Management system is the SDK model. To use this SDK, you must download the SDK core. You can download the SDK core from the Alibaba Cloud official website or obtain the SDK core from the Apsara Stack aftersales service.

11.4.4.2. Modify SDKs

When you need to update an SDK, you can upload an SDK to replace the original one.

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > SDK Tools. The SDK Tools page appears.
- 3. In the product list, click the ∠ icon in the Actions column corresponding to the product for which the SDK is to be modified.
- 4. In the dialog box that appears, upload an SDK as prompted. Click Submit.

Edit SDK:drds-java-sdk_	2020-06-01_100315	×
SDK Version :	v3.4	
Re-upload SDK:	6	
	Click or drag the file to this area to upload the file.	
	Submit	

11.4.4.3. Delete SDKs

You can delete SDKs that you no longer need.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose Versions > SDK Tools. The SDK Tools page appears.
- 3. In the product list, click the 🗇 icon in the Actions column corresponding to the product for which the SDK is to be deleted.
- 4. In the message that appears, click Yes.

11.5. Test management

To facilitate API tests in the Apsara Asapi Management system, the system provides the test management function. Each API can be saved as a test case during the test. A test case contains the request parameters of an API. You can associate multiple test cases to create a test set. You can choose to run one test case and one test set at a time. You can view execution results on the Execution History page.

11.5.1. Test cases

A test case is used to test a specified API.

11.5.1.1. Modify test cases

You can modify the request parameters of a test case as needed.

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.
- 3. On the Test Cases page, click the 🖉 icon in the Actions column corresponding to the test

case to be modified.

Test Platform	st Platform > Test Case					
ID	name	API	product	APIAPI version	description	operation
30	Describelmages	Describelmages	Ecs	2014-05-26		Z • 10
29	DescribeBCCClusterInfo	DescribeBCCClusterInfo	opsAPI	2018-01-22		∠ ▶ 🗗
28	DescribeTemplate	DescribeTemplate	ROS	2015-09-01		∠ ▶ ⊡

4. In the Edit Test Case dialog box that appears, modify values of the regionId, accessKeyId, accessKeySecret, Product, apild, and apiVersion parameters.

Edit a Test Case	×
TestCase Name: Describelmages	
✓ Request Parameters	
regionId:	cn-qingdao-env12-d01
accessKeyld :	
accessKeySecret:	
apiName :	Describelmages
regionNo :	cn-qingdao-env12-d01
product:	Ecs
apild :	7824
apiVersion :	2014-05-26
L	save

5. Click Save.

11.5.1.2. Run test cases

You can run a test case as needed.

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.
- 3. On the Test Cases page, click the > icon in the Actions column corresponding to a test

case.After the test case is executed, you can view the request parameters and responses in the preview area. You can also choose **Testing Platform > Execution History** to view the execution results on the Execution History page.

11.5.1.3. Delete test cases

You can delete test cases that you no longer need.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.
- 3. On the Test Cases page, click the 🕆 icon in the Actions column corresponding to the test case to be deleted.
- 4. In the message that appears, click Yes.

11.5.2. Test sets

A test set consists of multiple associated test cases.

11.5.2.1. Create test sets

You can create a test set based on the test requirements.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
- 3. On the Test Sets page, click Create Test Set.
- 4. In the dialog box that appears, enter the test set name and description. Click **Save**.We recommend that you enter a test set name that can be identified easily.

11.5.2.2. Associate test cases

You can associate test cases with a test set to manage the test cases in a unified manner.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
- 3. On the Test Sets page, click the

icon in the Actions column corresponding to a test set.

4. (Optional)You can update the name and description of the test set as needed. Click Save.

- 5. Click Relate to Test Case.
- 6. In the dialog box that appears, search for and select the test case to associate. You can select multiple test cases and add them to the test set.
- 7. Click Save.

11.5.2.3. Run test sets

You can run a test set to check whether the APIs in the test set are available.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.

			+ Create Test Set
ID	Name	Description	Actions
1	osstestcase	test	∠ → ū

3. On the Test Sets page, click the **>** icon in the Actions column corresponding to a test case.

The test cases in the test set start to run.After the test set is run, you can choose **Testing Platform > Execution History** to view the execution results on the Execution History page.

11.5.2.4. Delete test sets

You can delete test sets that you no longer need.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
- 3. On the Test Sets page, click the 🗇 icon in the Actions column corresponding to the test case to be deleted.
- 4. In the message that appears, click Yes.

11.5.3. View execution history of test cases

You can view information about the API for which a test case was executed, including the corresponding product, version information, execution time, and execution status.

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **Testing Platform > Execution History**. The **Execution History** page appears.
- 3. On the Execution History page, click the details icon in the Actions column corresponding to

the API to view.

4. In the Execution Details dialog box that appears, view the execution details of the test case, including request parameters and responses.

11.6. System management

11.6.1. Metadatabase management

You can add or remove information about metadatabases in the Apsara Asapi Management system.

11.6.1.1. View information about added metadatabases

The Apsara Asapi Management system automatically scans all metadatabases in the Apsara Stack environment during initialization. The Apsara Asapi Management system allows you to scan all metadatabases to view information about the added metadatabases. You can also manually add information about the metadatabases.

Context

The metadatabase information contains the domain name, database name, port, and server that are used in Apsara Stack products.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **System Management > Metabase**. The **Metabase** page appears.
- 3. Use one of the following methods to view the connection information about metadatabases:
 - Scan metadatabases

Click Scan Metabase to scan all added metadatabases in the Aspara Stack environment.

• Add information about metadatabases

Click Add Metabase. In the Add Metabse dialog box that appears, set Product Name, Metabase Name, Metabase Server, Metabase Port, Username, and Password. Click Submit.

Add Metabase		×
* Product Name:		
* Metabase Name:		
* Metabase Server:	ecsdriver	
* Metabase Port:		
* Username :		
* Password :		
	Submit	

You can view information about the added metadatabases in the metadatabase list.

11.6.1.2. View connection information about

metadatabases

You can view connection information about a metadatabase on the Metabase page.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose System Management > Metabase. The Metabase page appears.
- 3. In the upper-left corner of the page, select the product to view from the drop-down list.
- 4. On the Metabase page, click the () icon in the Actions column corresponding to the

metadatabase. In the message that appears, you can view the metadatabase connection information.

Select a product. v		Scan Metabase	Add Metabase	
Product	Database Name	Server Name	Port	Actions
OSS	oss_chiji	OSS-	3896	Details
OSS	oss_chiji_slave	OSS-	3896	÷ 💽
OSS	mns_user_db	mns	3898	₽ ©

11.6.1.3. Remove information about metadatabases

To facilitate management, you can remove information about metadatabases that you no longer need.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose **System Management > Metabase**. The **Metabase** page appears.
- 3. On the Metabase page, click the 🕆 icon in the Actions column corresponding to the metadatabase whose information is to be removed.
- 4. In the message that appears, click Yes.

11.6.2. Server management

You can add or remove information about servers in the Apsara Asapi Management system.

11.6.2.1. View information about added servers

The Apsara Asapi Management system automatically scans all servers (including physical servers and VMs) in the Apsara Stack environment during initialization. When new servers are added to the Apsara Stack environment, you can scan servers to view information about the added servers. You can also add information about the added servers.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose System Management > Server Management. The Server Management page appears.
- 3. Use one of the following methods to view information about added servers:
 - Scan servers

Click **Scan Server** to scan the information about all servers in the Apsara Stack environment.

• Add information about new servers

Click Add Server. In the Add Server dialog box that appears, set Name, Server Name, Server IP, SSH Port, and SSH User. Upload the SSH private key. Click Submit.

* Name:	Alias. Typically it's the pr	
* Server Name	Server Name	
* Server IP:	IP Address	
* SSH Port:	Port	
* SSH User:	SSH User	
SSH Password:	SSH Password	
SSH Private Ke	上 Upload SSH Private Key	

You can view information about the added servers in the server list.

11.6.2.2. Remove server information

To facilitate management, you can remove information about servers that you no longer need.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose System Management > Server Management. The Server Management page appears.
- 3. On the Server Management page, click the 🗇 icon in the Actions column corresponding to the server whose information is to be removed.

11.6.3. Audit APIs

You can view the call records of all Asapis. The records contain the specific API, status, time, and result of each call.

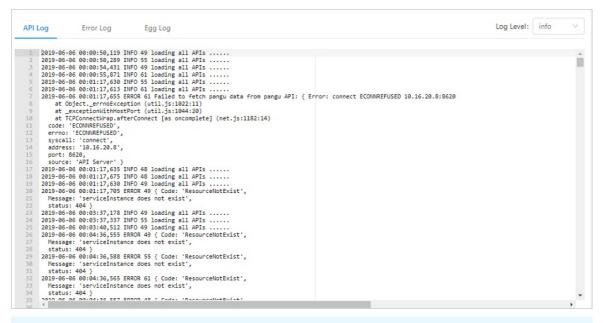
- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose System Management > API Audit. The API Audit page appears.
- 3. On the API Audit page, click the () icon in the Actions column corresponding to an API. You can view the call result of this API.

11.6.4. View logs

You can view API logs, error logs, and Egg logs to better maintain the backend.

Procedure

- 1. Log on to the Apsara Asapi Management platform.
- 2. In the left-side navigation pane, choose System Management > Log Management. The Log Management page appears.
- 3. View the details of all logs on the API Log, Error Log, and Egg Log tabs.



? Note You can modify log levels in the Apsara Asapi Management system in real time. The level you set is valid only during the active service operating period. If you restart the service, the default level remains.

12.Appendix 12.1. Operation Access Manager (OAM)

12.1.1. OAM introduction

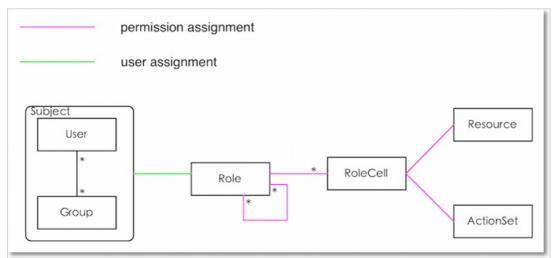
Overview

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

OAM permission model

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

See the OAM permission model as follows.



Permission model

12.1.2. Instructions

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

subject

Operators of the access control system. OAM has two types of subjects: users and groups.

user

Administrators and operators of operations systems.

group

A collection of users.

role

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

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

RoleHierarchy

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

RoleCell

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

resource

The description of an authorized object. For more information about resources of operations platforms, see Permission lists of operations platforms.

ActionSet

The description of authorized actions. An ActionSet can contain multiple actions. For more information about actions of operations platforms, see Permission lists of operations platforms.

available authorizations

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

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

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

12.1.3. Quick start

This topic describes how to add and assign roles quickly.

12.1.3.1. Log on to OAM

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

Prerequisites

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

The URL of the ASO console is in the following format: *region-id*.aso.*intranet-domain-id*.com.

• A browser is available. We recommend that you use the Google Chrome browser.

Procedure

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

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

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

3. Enter your username and password.

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

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

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

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

12.1.3.2. 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.

12.1.3.3. Add group members

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

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group whose name and description you want to modify and click Manage in the Actions column.
- 4. 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.

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

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group whose name and description you want to modify and click Manage in the Actions column.
- 4. 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.

12.1.3.5. 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.
 - i. Click Edit Tag.
 - ii. In the displayed Edit Tags dialog box, click Create.
 - iii. Enter the Key and the corresponding Value of the tag and then click Confirm.
 - iv. Repeat the preceding step to create more tags.

The created tags are displayed in the dotted box.

- v. Click Confirm to create the tags.
- 6. Click Confirm to create the role.

12.1.3.6. 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 Query roles.

Procedure

1. Log on to OAM.

- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role whose information you want to modify and click Manage in the Actions column.

- 4. 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.

12.1.3.7. Add resources to a role

You must add resources to a created role.

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role whose information you want to modify and click Manage in the Actions column.
- 4. Click the Resource List tab.
- 5. Click Add Resource.
- 6. Complete the configurations. For more information, see Configurations.

Configurations

Configuration item	Description
BID	The deployment region ID.
	The cloud product to be added, for example, rds.
Product	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 Permission lists of operations platforms.
Actions	An ActionSet, which can contain multiple actions. For more information about actions of operations platforms, see Permission lists of operations platforms.
Available Authorizations	The maximum number of authorizations in the cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.
Description	The description of the resource.

7. Click Add.

12.1.3.8. Add authorized users to a role

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

Prerequisites

The corresponding users or user groups are created. Users are created in the Apsara Stack console. For more information about how to create user groups, see Create groups.

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role whose information you want to modify and 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 groups.

- 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.

12.1.4. Manage groups

Group Management allows you to view, modify, or delete groups.

12.1.4.1. Modify the group information

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

- 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.

12.1.4.2. 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.

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Group Management > Owned Groups.
- 3. Find the group whose name and description you want to modify and click Manage in the Actions column.
- 4. In the 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.

12.1.4.3. 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.

12.1.4.4. View authorized groups

You can view the groups to which you are added on the Authorized Groups page.

Context

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

> Document Version:20200918

- 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.

12.1.5. Manage roles

Role Management allows you to view, modify, transfer, or delete roles.

12.1.5.1. 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.

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

12.1.5.2. Modify 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.

12.1.5.3. 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.

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. Find the role whose information you want to modify and click Manage in the Actions column.

4. 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.

12.1.5.4. 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.

12.1.5.5. Delete a role

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

Prerequisites

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

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, choose Role Management > Owned Roles.
- 3. At the right of the role to be deleted and then click Delete.

12.1.5.6. View authorized roles

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

- 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.

12.1.5.7. 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.

12.1.6. Search for resources

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

Procedure

- 1. Log on to OAM.
- 2. In the left-side navigation pane, 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.

12.1.7. 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.
 - i. Enter the resource information in the Resource field.

Note Use the English input method when entering values in the Resource and Action fields.

ii. Enter the permissions in the Action field, such as create, read, and write. Separate multiple permissions with commas (,).

12.1.8. Appendix

12.1.8.1. Default roles and their functions

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

12.1.8.1.1. Default role of OAM

This topic describes the default role of Operation Access Manager (OAM) and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Super administrator	An administrator with root permissions	*:*	*	10

12.1.8.1.2. Default roles of Apsara Infrastructure

Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations.

Apsara Infrastructure Management Framework is a distributed data center management system, used to manage applications on clusters containing multiple machines, and provides basic functions such as deployment, upgrade, expansion, contraction, and configuration change.

For more information about the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations, see the following table.

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

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

Role name	Role description	Resource	Actions	Available authorizations
Tianji_IDC administrator	Has all the permissions to Apsara Infrastructure Management Framework data centers, which allows you to view and modify the data center information	*:tianji:idcs	["*"]	0
Tianji administrator	Has all the permissions to Apsara Infrastructure Management Framework, which allows you to perform operations on all Apsara Infrastructure Management Framework configurations	*:tianji	["*"]	0

12.1.8.1.3. Default roles of Webapp-rule

This topic describes the default roles of Webapp-rule and the corresponding available authorizations.

Role name	Role description	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

Role name	Role description	Resource	Actions	Available authorizations
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

12.1.8.1.4. Default roles of the workflow console

This topic describes the default roles of the workflow console and the corresponding available authorizations.

The workflow console Grandcanal is an internally distributed process development framework. Developers can assemble, retry, roll back, and manually intervene the process based on this framework. Operations engineers can use the workflow console to manually intervene the corresponding processes.

For more information about the default roles of the workflow console and the corresponding available authorizations, see the following table.

Role name	Role description	Resource	Actions	Available authorizations
grandcanal.ADMI N	The workflow console administrator, who can query the workflow and activity details, and retry, roll back, terminate, and restart a workflow	26842:grandcana l	["write" ,"read"]	0
grandcanal.Read er	Has the read- only permission to the workflow console and can only perform the read operation	26842:grandcana l	["read"]	0

12.1.8.1.5. Default role of Tianjimon

This topic describes the default role of Tianjimon and the corresponding available authorizations.

Tianjimon, as the monitoring module of Apsara Infrastructure Management Framework, is used for the basic monitoring function of physical machines and services deployed based on Apsara Infrastructure Management Framework.

For more information about the default role of Tianjimon and the corresponding available authorizations, see the following table.

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

12.1.8.2. Operation permissions on O&M platforms

This topic describes the operation permissions on O&M platforms.

12.1.8.2.1. Permissions on Apsara Infrastructure

Management Framework

This topic describes the operation permissions on Apsara Infrastructure Management Framework.

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

Resource	Action	Description
*:tianji:services: [sname]:serviceinstances: [siname]:template	write	SetServiceInstanceTmpl
*:tianji:services: [sname]:serviceinstances: [siname]:template	delete	DeleteServiceInstanceTmpl
*:tianji:services: [sname]:serviceinstances: [siname]:template	read	GetServiceInstanceTmpl
*:tianji:services: [sname]:serviceinstances: [siname]:tags:[tag]	delete	DeleteServiceInstanceProductT agInService
*:tianji:services: [sname]:serviceinstances: [siname]:tags:[tag]	write	AddServiceInstanceProductTag InService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:resources	read	GetServerroleResourceInServic e
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]	write	OperateSRMachineInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]	read	GetMachineSRInfoInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]	delete	DeleteSRMachineActionInServic e
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	read	GetMachinesSRInfoInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	delete	DeleteSRMachinesActionInServ ice

Resource	Action	Description
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	write	OperateSRMachinesInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:apps: [app]:resources	read	GetAppResourceInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles: [serverrole]:apps: [app]:machines: [machine]:tianjilogs	read	TianjiLogsInService
*:tianji:services: [sname]:serviceinstances: [siname]:serverroles	read	GetServiceInstanceServerrolesI nService
*:tianji:services: [sname]:serviceinstances: [siname]:schema	write	SetServiceInstanceSchema
*:tianji:services: [sname]:serviceinstances: [siname]:schema	delete	DeleteServiceInstanceSchema
*:tianji:services: [sname]:serviceinstances: [siname]:rollings:[version]	write	OperateRollingJobInService
*:tianji:services: [sname]:serviceinstances: [siname]:rollings	read	ListRollingJobInService
*:tianji:services: [sname]:serviceinstances: [siname]:resources	read	GetInstanceResourceInService
*:tianji:services: [sname]:serviceinstances: [siname]:machines:[machine]	read	GetMachineAllSRInfoInService
*:tianji:services: [sname]:serviceinstances: [siname]	write	DeployServiceInstanceInServic e
*:tianji:services: [sname]:serviceinstances: [siname]	read	GetServiceInstanceConf

Resource	Action	Description
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:name	read	GetMachineAppFileListInServic e
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:download	read	GetMachineAppFileDownloadIn Service
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:content	read	GetMachineAppFileContentInSe rvice
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:filelist	read	GetMachineFileListInService
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:dockerlogs	read	DockerLogsInService
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:debuglog	read	GetMachineDebugLogInService
*:tianji:services: [sname]:serverroles: [serverrole]:machines: [machine]:apps	read	GetMachineAppListInService
*:tianji:services: [sname]:serverroles: [serverrole]:apps: [app]:dockerinspect	read	DockerInspect
*:tianji:services: [sname]:schemas: [schemaname]	write	PutServiceSchema
*:tianji:services: [sname]:schemas: [schemaname]	delete	DeleteServiceSchema

Resource	Action	Description
*:tianji:services: [sname]:resources	read	GetResourceInService
*:tianji:services:[sname]	delete	DeleteService
*:tianji:services:[sname]	write	CreateService
*:tianji:projects: [pname]:machinebuckets: [bname]:machines:[machine]	read	GetMachineBucketMachineInfo
*:tianji:projects: [pname]:machinebuckets: [bname]:machines	read	GetMachineBucketMachines
*:tianji:projects: [pname]:machinebuckets: [bname]	write	CreateMachineBucket
*:tianji:projects: [pname]:machinebuckets: [bname]	write	OperateMachineBucketMachine s
*:tianji:projects: [pname]:machinebuckets: [bname]	delete	DeleteMachineBucket
*:tianji:projects: [pname]:machinebuckets: [bname]	read	GetMachineBucketMachinesLeg acy
*:tianji:projects: [pname]:machinebuckets	read	GetMachineBucketList
*:tianji:projects: [pname]:projects: [pname]:clusters: [cname]:tssessions: [tssessionname]:tsses	terminal	UpdateTsSessionTssByCluster
*:tianji:projects: [pname]:projects: [pname]:clusters: [cname]:tssessions	terminal	CreateTsSessionByCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:tjmontemplate	read	GetServiceInstanceTjmonTmplI nCluster

Resource	Action	Description
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:template	delete	DeleteServiceInstanceTmplInCl uster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:template	write	SetServiceInstanceTmplInClust er
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:template	read	GetServiceInstanceTmplInClust er
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:tags:[tag]	write	AddServiceInstanceProductTag InCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:tags:[tag]	delete	DeleteServiceInstanceProductT agInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:resources	read	GetServerroleResourceInCluste r
<pre>*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:name</pre>	read	GetMachineAppFileList
<pre>*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:download</pre>	read	GetMachineAppFileDownload

Resource	Action	Description
<pre>*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:files:content</pre>	read	GetMachineAppFileContent
<pre>*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps:[app]:filelist</pre>	read	GetMachineFileList
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:dockerlogs	read	DockerLogsInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps: [app]:debuglog	read	GetMachineDebugLog
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]:apps	read	GetMachineAppList
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]	read	GetMachineSRInfoInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]	write	OperateSRMachineInCluster

Resource	Action	Description
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines: [machine]	delete	DeleteSRMachineActionInClust er
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	write	OperateSRMachinesInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	delete	DeleteSRMachinesActionInClus ter
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:serverroles: [serverrole]:machines	read	GetAllMachineSRInfoInCluster
*: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	GetServiceInstanceServerrolesI nCluster

Resource	Action	Description
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:schema	delete	DeleteServiceInstanceSchemaI nCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:schema	write	SetServiceInstanceSchemaInCl uster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]:resources	read	GetInstanceResourceInCluster
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]	delete	DeleteServiceInstance
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]	write	CreateServiceInstance
*:tianji:projects: [pname]:clusters: [cname]:serviceinstances: [siname]	read	GetServiceInstanceConfInClust er
*: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

Resource	Action	Description
*:tianji:projects: [pname]:clusters: [cname]:machines:[machine]	write	SetMachineAction
*:tianji:projects: [pname]:clusters: [cname]:machines:[machine]	delete	DeleteMachineAction
*:tianji:projects: [pname]:clusters: [cname]:machines	write	OperateClusterMachines
*:tianji:projects: [pname]:clusters: [cname]:difflist	read	GetVersionDiffList
*:tianji:projects: [pname]:clusters:[cname]:diff	read	GetVersionDiff
*:tianji:projects: [pname]:clusters: [cname]:deploylogs:[version]	read	GetDeployLogInCluster
*:tianji:projects: [pname]:clusters: [cname]:deploylogs	read	GetDeployLogListInCluster
*:tianji:projects: [pname]:clusters: [cname]:builds:[version]	read	GetBuildJob
*:tianji:projects: [pname]:clusters: [cname]:builds	read	ListBuildJob
*:tianji:projects: [pname]:clusters:[cname]	write	OperateCluster
*: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

Resource	Action	Description
*:tianji:idcs:[idc]:rooms: [room]:racks:[rack]:rackunits: [rackunit]	write	CreateRackunit
*:tianji:idcs:[idc]:rooms: [room]:racks:[rack]:rackunits: [rackunit]	write	SetRackunitAttr
*:tianji:idcs:[idc]:rooms: [room]:racks:[rack]:rackunits: [rackunit]	delete	DeleteRackunit
*:tianji:idcs:[idc]:rooms: [room]:racks:[rack]	write	SetRackAttr
*:tianji:idcs:[idc]:rooms: [room]:racks:[rack]	write	CreateRack
*:tianji:idcs:[idc]:rooms: [room]:racks:[rack]	delete	DeleteRack
*:tianji:idcs:[idc]:rooms:[room]	write	CreateRoom
*:tianji:idcs:[idc]:rooms:[room]	delete	DeleteRoom
*:tianji:idcs:[idc]:rooms:[room]	write	SetRoomAttr
*:tianji:idcs:[idc]	delete	DeleteIdc
*:tianji:idcs:[idc]	write	SetIdcAttr
*:tianji:idcs:[idc]	write	Createldc

12.1.8.2.2. 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

12.1.8.2.3. 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

12.1.8.2.4. Permissions on Monitoring System of Apsara

Infrastructure Management Framework

This topic describes the operation permissions on Monitoring System of Apsara Infrastructure Management Framework.

Resource	Action	Description
26842:tianjimon:monitor- manage	manage	Monitoring and O&M